

# Vecchia-Inducing-Points Full-Scale Approximations for Gaussian Processes

**Tim Gyger**

TIM.GYGER@HSLU.CH

*Lucerne University of Applied Sciences and Arts  
Department of Mathematical Modeling and Machine Learning  
University of Zurich  
SUURSTOFFI 1, 6343 ROTKREUZ, SWITZERLAND*

**Reinhard Furrer**

REINHARD.FURRER@UZH.CH

*Department of Mathematical Modeling and Machine Learning  
University of Zurich*

**Fabio Sigrist**

FABIO.SIGRIST@STAT.MATH.ETHZ.CH

*Seminar for Statistics, ETH Zurich  
Lucerne University of Applied Sciences and Arts*

**Editor:** Chris Oates

## Abstract

Gaussian processes are flexible, probabilistic, non-parametric models widely used in machine learning and statistics. However, their scalability to large data sets is limited by computational constraints. To overcome these challenges, we propose **Vecchia-inducing-points full-scale (VIF)** approximations combining the strengths of global inducing points and local Vecchia approximations. Vecchia approximations excel in settings with low-dimensional inputs and moderately smooth covariance functions, while inducing point methods are better suited to high-dimensional inputs and smoother covariance functions. Our VIF approach bridges these two regimes by using an efficient correlation-based neighbor-finding strategy for the Vecchia approximation of the residual process, implemented via a modified cover tree algorithm. We further extend our framework to non-Gaussian likelihoods by introducing iterative methods that substantially reduce computational costs for training and prediction by several orders of magnitude compared to Cholesky-based computations when using a Laplace approximation. In particular, we propose and compare novel preconditioners and provide theoretical convergence results. Extensive numerical experiments on simulated and real-world data sets show that VIF approximations are both computationally efficient as well as more accurate and numerically stable than state-of-the-art alternatives. All methods are implemented in the open-source C++ library **GPBoost** with high-level Python and R interfaces.

**Keywords:** Probabilistic modeling, uncertainty quantification, large data, iterative methods, Laplace approximation

## 1. Introduction

Gaussian process (GP) models are probabilistic non-parametric models which are used across various disciplines including machine learning (Rasmussen and Williams, 2006) and geostatistics (Cressie, 1993). However, their applicability to large data sets is limited by

computational costs, with  $\mathcal{O}(n^3)$  time and  $\mathcal{O}(n^2)$  memory complexity, where  $n$  is the number of observations. Various strategies have been proposed to mitigate this computational burden; see Liu et al. (2020) and Heaton et al. (2019) for reviews. Low-rank inducing-points approximations such as the fully independent training conditional (FITC) approximation (Quinero-Candela and Rasmussen, 2005), predictive process models (Banerjee et al., 2008; Finley et al., 2009), and variational approximations (Titsias, 2009; Hensman et al., 2013, 2015) are global approximations that focus on capturing the large-scale structures of GPs. Such inducing point methods are widely used in machine learning and are considered state-of-the-art methods. An alternative approach to improving computational efficiency is to leverage sparse linear algebra methods by enforcing sparsity in a Cholesky factor of precision matrices (Vecchia, 1988; Datta et al., 2016; Katzfuss and Guinness, 2021; Schäfer et al., 2021a; Cao et al., 2023). These local Vecchia approximations are considered as a “leader among the sea of approximation” (Guinness, 2021) in spatial statistics as they often yield superior approximation accuracy (Rambelli and Sigrist, 2026).

In this work, we propose an approach that integrates inducing point methods with Vecchia approximations and denote this as Vecchia-inducing-points full-scale (VIF) approximations. This method combines the advantages of global inducing points and local Vecchia approximations. Vecchia approximations excel in settings with low-dimensional inputs and moderately smooth covariance functions. The latter can be explained by the screening effect (Schäfer et al., 2021a, Section 3.3.3). Vecchia approximations are expected to work less well for high-dimensional data as they depend on the selection of neighbors, and this is more difficult in higher dimensions. Low-rank inducing point methods, on the other hand, are expected to work well for smoother covariance functions (Schäfer et al., 2021b, Section 2.4) and are better suited to high-dimensional inputs (see our experimental results in Section 7.1). VIF approximations bridge these two types of approximations. In addition, we propose a novel, efficient correlation-distance-based method for selecting conditioning sets for Vecchia approximations on the residual process using a modified cover tree algorithm.

We further extend our framework to non-Gaussian likelihoods to include, e.g., classification and count data regression tasks, by using a Laplace approximation. Although, depending on the likelihood, Laplace approximations can be inaccurate for small data sets, they are computationally very efficient (Nickisch and Rasmussen, 2008) and converge asymptotically to the correct quantity and are thus accurate for large data sets. To support this argument, Figure 1 shows the estimated marginal variance parameter obtained with a VIF-Laplace approximation for varying sample sizes  $n$ , based on simulated data with a Bernoulli likelihood as described in Section 7. Estimation is performed using the iterative methods introduced in this paper and repeated across 100 simulated data sets for each  $n$ . The results demonstrate that the downward bias in the variance parameter diminishes as  $n$  increases. However, standard Cholesky-based calculations do not scale well despite the use of VIF and Laplace approximations for non-Gaussian likelihoods since this requires computing the Cholesky factorization of a sparse  $n \times n$  matrix (see Section 3) whose computational complexity is not linear in  $n$  and depends on the graph structure as well as the effectiveness of the fill-reducing ordering (Davis, 2006). For two-dimensional spatial data, this complexity is approximately  $\mathcal{O}(n^{3/2})$ , while in higher dimensions, it is at least  $\mathcal{O}(n^2)$  (Lipton et al., 1979). To remedy this, we introduce iterative methods (Saad, 2003; Trefethen and Bau, 2022; Aune et al., 2014; Gardner et al., 2018a) that substantially reduce computational




Figure 1: Violin plots of the estimated variance parameter  $\sigma_1^2$  for different sample sizes  $n$  for binary data. Red dots indicate the mean estimates, and the dashed line marks the true parameter value  $\sigma_1^2 = 1$ .

costs for likelihood evaluations, gradient computations, and the calculation of predictive distributions when using a Laplace approximation. In contrast to the Cholesky decomposition, iterative methods require only matrix-vector multiplications which can be trivially parallelized. We propose and compare novel preconditioners, derive new convergence guarantees, introduce two simulation-based approaches to compute predictive variances, and provide both theoretical insights and empirical evaluations. Furthermore, we compare the prediction accuracy and runtime of our methods against existing state-of-the-art approaches on a large set of simulated and real-world data sets and find that VIF approximations are both computationally efficient and more accurate than state-of-the-art alternatives.

Full-scale approximations were first introduced in Sang et al. (2011) and Sang and Huang (2012), which combined inducing points with covariance tapering (Furrer et al., 2006) approximations. Zhang et al. (2019) extended this by partitioning the data into blocks and deriving a block-wise sparse precision matrix for the residual process. While this can also be interpreted as a special form of a Vecchia approximation, Zhang et al. (2019) do not explicitly consider the case when using single points as conditioning sets, they do not address the question of efficiently selecting neighbors, as this might be less of an issue when using blocks, and they do not cover non-Gaussian likelihoods, which entail additional computational challenges.

The remainder of the paper is structured as follows. Sections 2 and 3 present the VIF framework for Gaussian and non-Gaussian likelihoods, respectively. In Section 4, we introduce iterative methods for inference within the VIF framework, and in Section 5, we present convergence results for the preconditioned conjugate gradient (CG) methods. In Section 6, we propose efficient algorithms based on cover trees to identify the nearest Vecchia neighbors with respect to the correlation distance. In Section 7, we conduct simulation studies to analyze computational and statistical properties of VIF approximations. Finally, in Section 8, we apply our methodology to various real-world data sets, comparing runtime and prediction accuracy with several state-of-the-art methods. Additionally, we extend these

experiments by estimating the smoothness parameter of automatic relevance determination (ARD) Matérn kernels, contributing to the challenge of model selection - at least within the family of Matérn kernels.

### 1.1 Software implementation

The methods presented in this article are implemented in the `GPBoost` library written in C++ with Python and R interface packages, see <https://github.com/fabsig/GPBoost>.<sup>1</sup> In addition, to facilitate the practical use of the proposed methodology, we provide a self-contained generic Python example in the GitHub repository <https://github.com/TimGyger/VIF> that demonstrates how to fit a VIF GP model and obtain predictions using the `GPBoost` Python interface.

## 2. VIF approximations for Gaussian process regression

We consider the following GP model:

$$y = F(\mathbf{x}) + b(\mathbf{s}) + \epsilon(\mathbf{s}), \tag{1}$$

where  $y \in \mathbb{R}$  is the response variable,  $b(\mathbf{s})$  is a zero-mean GP with inputs  $\mathbf{s} \in \mathbb{R}^d$ , specified by a parametric covariance function  $c_\theta(\cdot, \cdot)$  with parameters  $\theta \in \mathbb{R}^q$ ,  $F(\cdot) : \mathbb{R}^p \rightarrow \mathbb{R}$  is a fixed effects prior mean function with predictor variables  $\mathbf{x} \in \mathbb{R}^p$ , and  $\epsilon(\mathbf{s})$  is a zero-mean independent Gaussian white noise process with variance  $\sigma^2$ . Note that the GP inputs  $\mathbf{s}$  and the fixed effects predictor variables  $\mathbf{x}$  may or may not be overlapping. For instance, in spatial statistics,  $\mathbf{s}$  often consists of spatial and/or temporal coordinates whereas  $\mathbf{x}$  contains additional predictor variables. But in machine learning applications, the GP and fixed effects inputs can also be equal; see, e.g., Sigrist (2022a) for an example. For notational simplicity, we will assume  $F(\mathbf{x}) = 0$  in the following, but  $F(\cdot)$  can be easily extended to a linear mean function  $F(\mathbf{x}) = \mathbf{x}^T \boldsymbol{\beta}$ ,  $\boldsymbol{\beta} \in \mathbb{R}^p$ , and to machine learning methods such as tree-boosting (Sigrist, 2022a,b) and neural networks (Simchoni and Rosset, 2023).

If  $n$  samples  $\mathbf{y} = (y(\mathbf{s}_1), \dots, y(\mathbf{s}_n))^T \in \mathbb{R}^n$  are observed at inputs  $\mathcal{S} = \{\mathbf{s}_1, \dots, \mathbf{s}_n\}$ ,  $\mathbf{y}$  has a multivariate distribution with mean zero and covariance matrix

$$\tilde{\boldsymbol{\Sigma}} = \boldsymbol{\Sigma} + \sigma^2 \mathbf{I}_n, \tag{2}$$

where  $\boldsymbol{\Sigma} \in \mathbb{R}^{n \times n}$  is the covariance matrix of  $\mathbf{b} = (b(\mathbf{s}_1), \dots, b(\mathbf{s}_n))^T \in \mathbb{R}^n$  with entries  $\Sigma_{ij} = \text{Cov}(b(\mathbf{s}_i), b(\mathbf{s}_j)) = c_\theta(\mathbf{s}_i, \mathbf{s}_j)$ ,  $\mathbf{s}_i, \mathbf{s}_j \in \mathcal{S}$ . Concerning the notation in this article, we use the symbol  $\tilde{\cdot}$  to distinguish matrices including the error variance  $\sigma^2 \mathbf{I}_n$  from those without this diagonal error variance matrix. E.g.,  $\tilde{\boldsymbol{\Sigma}}$  contains the error variance  $\sigma^2 \mathbf{I}_n$  whereas  $\boldsymbol{\Sigma}$  does not. For maximum likelihood estimation, we minimize the negative log-likelihood function

$$\mathcal{L}(\boldsymbol{\theta}; \mathbf{y}) = \frac{n}{2} \log(2\pi) + \frac{1}{2} \log \det(\tilde{\boldsymbol{\Sigma}}) + \frac{1}{2} \mathbf{y}^T \tilde{\boldsymbol{\Sigma}}^{-1} \mathbf{y}.$$

---

1. The VIF approximation can be enabled via the parameter `gp_approx = "vif"`, iterative methods can be enabled via the parameter `matrix_inversion_method = "iterative"` (=default), and the preconditioner is chosen via the parameter `cg_preconditioner_type`.

Evaluation of the log-likelihood and its derivatives involves the calculation of linear solves with the  $n \times n$  matrix  $\tilde{\Sigma}$  and calculating its determinant, which has  $\mathcal{O}(n^3)$  complexity when using a Cholesky decomposition.

For prediction at  $n_p$  new locations  $\mathcal{S}^p = \{\mathbf{s}_1^p, \dots, \mathbf{s}_{n_p}^p\}$ , the predictive distribution  $\mathbf{y}^p | \mathbf{y}$  is given by  $\mathcal{N}(\boldsymbol{\mu}^p, \boldsymbol{\Sigma}^p)$ , where  $\boldsymbol{\mu}^p = \boldsymbol{\Sigma}_{nn_p}^T \tilde{\Sigma}^{-1} \mathbf{y}$ ,  $\boldsymbol{\Sigma}^p = \boldsymbol{\Sigma}_{n_p} + \sigma^2 \mathbf{I}_{n_p} - \boldsymbol{\Sigma}_{nn_p}^T \tilde{\Sigma}^{-1} \boldsymbol{\Sigma}_{nn_p} \in \mathbb{R}^{n_p \times n_p}$ ,  $\boldsymbol{\Sigma}_{nn_p} = \left[ c_{\boldsymbol{\theta}}(\mathbf{s}_i, \mathbf{s}_j^p) \right]_{i=1:n, j=1:n_p}$  is a cross-covariance matrix, and  $\boldsymbol{\Sigma}_{n_p} = \left[ c_{\boldsymbol{\theta}}(\mathbf{s}_i^p, \mathbf{s}_j^p) \right]_{i,j=1:n_p}$ . For large  $n_p$ , it is usually infeasible to store the complete predictive covariance matrix  $\boldsymbol{\Sigma}^p$ , and the primary focus is solely on the predictive variances, i.e., the diagonal elements of  $\boldsymbol{\Sigma}^p$ . Nevertheless, Cholesky-based computation of these variances still entails a substantial computational burden involving  $\mathcal{O}(n^2 \cdot n_p)$  operations even when a Cholesky factor of  $\tilde{\Sigma}$  has been precomputed.

## 2.1 Vecchia-inducing-points full-scale (VIF) approximation

The idea of the Vecchia-inducing-points full-scale (VIF) approximation is to decompose  $b(\mathbf{s}) + \epsilon(\mathbf{s})$  into two parts  $b(\mathbf{s}) + \epsilon(\mathbf{s}) = b_1(\mathbf{s}) + b_s(\mathbf{s})$ , where  $b_1(\mathbf{s})$  is a low-rank predictive process modeling large-scale dependence and  $b_s(\mathbf{s}) = b(\mathbf{s}) + \epsilon(\mathbf{s}) - b_1(\mathbf{s})$  is a residual process capturing the residual variation that is unexplained by  $b_1(\mathbf{s})$ . Specifically, for a given set  $\mathcal{S}^* = \{\mathbf{s}_1^*, \dots, \mathbf{s}_m^*\}$  of  $m$  inducing points, or knots, the predictive process is given by

$$b_1(\mathbf{s}) = \mathbf{c}(\mathbf{s}, \mathcal{S}^*)^T \boldsymbol{\Sigma}_m^{-1} \mathbf{b}^*, \quad (3)$$

where  $\mathbf{b}^* = (b(\mathbf{s}_1^*), \dots, b(\mathbf{s}_m^*))^T$ . Its covariance is

$$\text{Cov}(b_1(\mathbf{s}_i), b_1(\mathbf{s}_j)) = \mathbf{c}(\mathbf{s}_i, \mathcal{S}^*)^T \boldsymbol{\Sigma}_m^{-1} \mathbf{c}(\mathbf{s}_j, \mathcal{S}^*),$$

where  $\mathbf{c}(\mathbf{s}_i, \mathcal{S}^*) = (c_{\boldsymbol{\theta}}(\mathbf{s}_i, \mathbf{s}_1^*), \dots, c_{\boldsymbol{\theta}}(\mathbf{s}_i, \mathbf{s}_m^*))^T$  and  $\boldsymbol{\Sigma}_m = [c_{\boldsymbol{\theta}}(\mathbf{s}_i^*, \mathbf{s}_j^*)]_{i=1:m, j=1:m} \in \mathbb{R}^{m \times m}$  is a covariance matrix of the inducing points. The covariance matrix of  $\mathbf{b}_1 = (b_1(\mathbf{s}_1), \dots, b_1(\mathbf{s}_n))^T$  is thus given by

$$\boldsymbol{\Sigma}^1 = \text{Cov}(\mathbf{b}_1) = \boldsymbol{\Sigma}_{mn}^T \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn},$$

where  $\boldsymbol{\Sigma}_{mn} = [c_{\boldsymbol{\theta}}(\mathbf{s}_i^*, \mathbf{s}_j)]_{i=1:m, j=1:n} \in \mathbb{R}^{m \times n}$  is a cross-covariance matrix between the inducing and data points. Concerning the residual process  $b_s(\mathbf{s})$ ,  $\text{Cov}(\mathbf{b}_s)$  is given by

$$\text{Cov}(\mathbf{b}_s) = \tilde{\Sigma} - \boldsymbol{\Sigma}_{mn}^T \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn},$$

where  $\mathbf{b}_s = (b_s(\mathbf{s}_1), \dots, b_s(\mathbf{s}_n))^T$ . In a VIF approximation, this residual covariance matrix is approximated using a Vecchia approximation, which represents the full joint density  $p(\mathbf{b}_s | \boldsymbol{\theta})$  through a product of low-dimensional conditional densities,

$$p(\mathbf{b}_s | \boldsymbol{\theta}) \approx \prod_{i=1}^n p(b_s(\mathbf{s}_i) | \mathbf{b}_{\mathbf{s}_{N(i)}}, \boldsymbol{\theta}),$$

where  $N(i) \subseteq \{1, \dots, i-1\}$  are sets of conditioning points, often called neighbors in Vecchia approximations (Katzfuss and Guinness, 2021; Datta et al., 2016). This leads to a sparse approximate Cholesky factorization of the precision matrix  $\text{Cov}(\mathbf{b}_s)^{-1}$ :

$$(\tilde{\Sigma}^s)^{-1} = \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \approx \text{Cov}(\mathbf{b}_s)^{-1},$$

where  $\mathbf{D} = \text{diag}(D_i)$  and  $\mathbf{B}$  is a sparse lower triangular matrix with 1's on the diagonal and non-zero off-diagonal entries  $\mathbf{B}_{iN(i)} = -\mathbf{A}_i \in \mathbb{R}^{|N(i)|}$ , and  $D_i$  and  $\mathbf{A}_i$  are defined as

$$\begin{aligned} D_i &= \tilde{\Sigma}_{ii} - \Sigma_{mi}^T \Sigma_m^{-1} \Sigma_{mi} - \left( \mathbf{A}_i (\Sigma_{iN(i)}^T - \Sigma_{mN(i)}^T \Sigma_m^{-1} \Sigma_{mi}) \right), \\ \mathbf{A}_i &= (\Sigma_{iN(i)} - \Sigma_{mi}^T \Sigma_m^{-1} \Sigma_{mN(i)}) (\tilde{\Sigma}_{N(i)} - \Sigma_{mN(i)}^T \Sigma_m^{-1} \Sigma_{mN(i)})^{-1}, \end{aligned} \quad (4)$$

where  $\Sigma_{mN(i)} = [c\theta(\mathbf{s}_l^*, \mathbf{s}_k^{N_i})]_{l=1:m, k=1:|N(i)|} \in \mathbb{R}^{m \times |N(i)|}$ ,  $\mathbf{s}_k^{N_i}$  are the Vecchia neighbor inputs of  $\mathbf{s}_i$ , and  $\Sigma_{mi} = [c\theta(\mathbf{s}_l^*, \mathbf{s}_i)]_{l=1:m} \in \mathbb{R}^m$ . As is commonly done, we choose  $N(i)$  as the indices of the  $m_v$  nearest (in a sense to be defined below) neighbors of  $\mathbf{s}_i$  among  $\mathbf{s}_1, \dots, \mathbf{s}_{i-1}$  if  $i > m_v + 1$ , and  $N(i) = \{1, \dots, i-1\}$  if  $i \leq m_v + 1$ . Calculating a Vecchia approximation for the residual process has  $\mathcal{O}(n \cdot (m_v^3 + m_v^2 \cdot m))$  computational cost and requires  $\mathcal{O}(n \cdot (m_v + m))$  memory storage.

In summary, a VIF approximation for the covariance matrix  $\tilde{\Sigma}$  in (2) is given by

$$\tilde{\Sigma}_{\dagger} = \Sigma^{\dagger} + \tilde{\Sigma}^s \approx \tilde{\Sigma}$$

and the corresponding negative log-likelihood is

$$\mathcal{L}_{\dagger}(\boldsymbol{\theta}; \mathbf{y}) = \frac{n}{2} \log(2\pi) + \frac{1}{2} \log \det(\tilde{\Sigma}_{\dagger}) + \frac{1}{2} \mathbf{y}^T \tilde{\Sigma}_{\dagger}^{-1} \mathbf{y}.$$

Note that if the number of Vecchia neighbors is zero, the VIF approximation reduces to the FITC approximation as a special case, and if there are no inducing points, the VIF approximation coincides with a classical Vecchia approximation. In addition, the VIF approximation can be interpreted as a general Vecchia approximation (Katzfuss and Guinness, 2021), in which each conditional distribution conditions on all inducing points in addition to the local Vecchia neighbors. However, the inverse Cholesky factor of the VIF approximation is, in general, not sparse. Sparsity would only arise in the special case where the inducing points coincide with observation locations.

## 2.2 Computational complexity and calculation of gradients

For computational efficiency, we can apply the Sherman-Woodbury-Morrison formula:

$$\begin{aligned} \tilde{\Sigma}_{\dagger}^{-1} &= (\tilde{\Sigma}^s + \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn})^{-1} \\ &= \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} - \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \Sigma_{mn}^T (\Sigma_m + \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \Sigma_{mn}^T)^{-1} \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}, \end{aligned}$$

where we denote

$$\mathbf{M} = \Sigma_m + \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \Sigma_{mn}^T. \quad (5)$$

Moreover, by Sylvester's determinant theorem, the determinant of  $\tilde{\Sigma}_{\dagger}$  can be calculated as

$$\begin{aligned} \det(\tilde{\Sigma}_{\dagger}) &= \det(\tilde{\Sigma}^s + \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn}) \\ &= \det(\Sigma_m + \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \Sigma_{mn}^T) \cdot \det(\Sigma_m)^{-1} \cdot \det(\mathbf{D}). \end{aligned}$$

If a first- or second-order method for convex optimization, such as the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm (Fletcher, 2000) is used, gradients of the negative log-likelihood function are required. For the VIF approximation, the gradient with respect to  $\boldsymbol{\theta}$  is given by

$$\frac{\partial}{\partial \boldsymbol{\theta}} \mathcal{L}_{\dagger}(\boldsymbol{\theta}; \mathbf{y}) = \frac{1}{2} \text{Tr} \left( \tilde{\boldsymbol{\Sigma}}_{\dagger}^{-1} \frac{\partial \tilde{\boldsymbol{\Sigma}}_{\dagger}}{\partial \boldsymbol{\theta}} \right) - \frac{1}{2} \mathbf{y}^{\text{T}} \tilde{\boldsymbol{\Sigma}}_{\dagger}^{-1} \frac{\partial \tilde{\boldsymbol{\Sigma}}_{\dagger}}{\partial \boldsymbol{\theta}} \tilde{\boldsymbol{\Sigma}}_{\dagger}^{-1} \mathbf{y},$$

where  $\frac{\partial \tilde{\boldsymbol{\Sigma}}_{\dagger}}{\partial \boldsymbol{\theta}} = \frac{\partial \tilde{\boldsymbol{\Sigma}}^{\text{s}}}{\partial \boldsymbol{\theta}} + \frac{\partial \boldsymbol{\Sigma}^{\text{l}}}{\partial \boldsymbol{\theta}}$  with

$$\begin{aligned} \frac{\partial \tilde{\boldsymbol{\Sigma}}^{\text{s}}}{\partial \boldsymbol{\theta}} &= \mathbf{B}^{-1} \frac{\partial \mathbf{D}}{\partial \boldsymbol{\theta}} \mathbf{B}^{-\text{T}} - \mathbf{B}^{-1} \frac{\partial \mathbf{B}}{\partial \boldsymbol{\theta}} \mathbf{B}^{-1} \mathbf{D} \mathbf{B}^{-\text{T}} - \mathbf{B}^{-1} \mathbf{D} \mathbf{B}^{-\text{T}} \frac{\partial \mathbf{B}^{\text{T}}}{\partial \boldsymbol{\theta}} \mathbf{B}^{-\text{T}}, \\ \frac{\partial \boldsymbol{\Sigma}^{\text{l}}}{\partial \boldsymbol{\theta}} &= \frac{\partial \boldsymbol{\Sigma}_{mn}^{\text{T}}}{\partial \boldsymbol{\theta}} \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn} + \boldsymbol{\Sigma}_{mn}^{\text{T}} \boldsymbol{\Sigma}_m^{-1} \frac{\partial \boldsymbol{\Sigma}_{mn}}{\partial \boldsymbol{\theta}} - \boldsymbol{\Sigma}_{mn}^{\text{T}} \boldsymbol{\Sigma}_m^{-1} \frac{\partial \boldsymbol{\Sigma}_m}{\partial \boldsymbol{\theta}} \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn}. \end{aligned}$$

In Appendix A, we derive the gradients of  $\mathbf{B}$  and  $\mathbf{D}$ .

The computational complexity associated with the computation of the negative log-likelihood and its derivatives is of the order  $\mathcal{O}(n \cdot (m_v^3 + m_v^2 \cdot m + m^2))$  and the required storage of order  $\mathcal{O}(n \cdot (m + m_v))$ .

### 2.3 Prediction with VIF approximations

We propose to do predictions for the response  $\mathbf{y}^p$  at  $n_p$  prediction inputs  $\mathcal{S}^p = \{\mathbf{s}_1^p, \dots, \mathbf{s}_{n_p}^p\}$  with VIF approximations using the following result.

**Proposition 1** *A VIF-approximated posterior predictive distribution for  $p(\mathbf{y}^p | \mathbf{y}, \boldsymbol{\theta})$ ,  $\mathbf{y}^p = (y(\mathbf{s}_1^p), \dots, y(\mathbf{s}_{n_p}^p))^{\text{T}} \in \mathbb{R}^{n_p}$ , is given by  $\mathcal{N}(\boldsymbol{\mu}_{\dagger}^p, \boldsymbol{\Sigma}_{\dagger}^p)$ , where*

$$\begin{aligned} \boldsymbol{\mu}_{\dagger}^p &= -\mathbf{B}_p^{-1} \mathbf{B}_{p0} \mathbf{y} + \boldsymbol{\Sigma}_{mn_p}^{\text{T}} \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn} \mathbf{B}^{\text{T}} \mathbf{D}^{-1} \mathbf{B} \mathbf{y} \\ &\quad - \boldsymbol{\Sigma}_{mn_p}^{\text{T}} \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn} \mathbf{B}^{\text{T}} \mathbf{D}^{-1} \mathbf{B} \boldsymbol{\Sigma}_{mn}^{\text{T}} \mathbf{M}^{-1} \boldsymbol{\Sigma}_{mn} \mathbf{B}^{\text{T}} \mathbf{D}^{-1} \mathbf{B} \mathbf{y} \end{aligned} \quad (6)$$

$$+ \mathbf{B}_p^{-1} \mathbf{B}_{p0} \boldsymbol{\Sigma}_{mn}^{\text{T}} \mathbf{M}^{-1} \boldsymbol{\Sigma}_{mn} \mathbf{B}^{\text{T}} \mathbf{D}^{-1} \mathbf{B} \mathbf{y}$$

$$\begin{aligned} \boldsymbol{\Sigma}_{\dagger}^p &= \mathbf{B}_p^{-1} \mathbf{D}_p \mathbf{B}_p^{-\text{T}} + \mathbf{B}_p^{-1} \mathbf{B}_{p0} (\mathbf{B}^{\text{T}} \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{B}_{p0}^{\text{T}} \mathbf{B}_p^{-\text{T}} + \boldsymbol{\Sigma}_{mn_p}^{\text{T}} \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn_p} \\ &\quad - (\boldsymbol{\Sigma}_{mn_p}^{\text{T}} \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn} - \mathbf{B}_p^{-1} \mathbf{B}_{p0} (\mathbf{B}^{\text{T}} \mathbf{D}^{-1} \mathbf{B})^{-1}) \tilde{\boldsymbol{\Sigma}}_{\dagger}^{-1} \\ &\quad \cdot (\boldsymbol{\Sigma}_{mn}^{\text{T}} \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn_p} - (\mathbf{B}^{\text{T}} \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{B}_{p0}^{\text{T}} \mathbf{B}_p^{-\text{T}}), \end{aligned} \quad (7)$$

where  $\boldsymbol{\Sigma}_{mn_p} = [c_{\boldsymbol{\theta}}(\mathbf{s}_i^*, \mathbf{s}_j^p)]_{i=1:m, j=1:n_p} \in \mathbb{R}^{m \times n_p}$ , and  $\mathbf{B}_{p0} \in \mathbb{R}^{n_p \times n}$ ,  $\mathbf{B}_p \in \mathbb{R}^{n_p \times n_p}$ , and  $\mathbf{D}_p^{-1} \in \mathbb{R}^{n_p \times n_p}$  are defined below in (8).

**Proof** [Proof of Proposition 1] First, note that  $(\mathbf{y}^{\text{T}}, \mathbf{y}^{p\text{T}})^{\text{T}} = (\mathbf{b}_l^{\text{T}}, \mathbf{b}_l^{p\text{T}})^{\text{T}} + (\mathbf{b}_s^{\text{T}}, \mathbf{b}_s^{p\text{T}})^{\text{T}}$ , where  $\mathbf{b}_l^p$  and  $\mathbf{b}_s^p$  denote the low-rank and residual processes, respectively, at the prediction points  $\mathcal{S}^p$ . Following common practice for predictions with Vecchia approximations (Katzfuss et al., 2020), we apply a Vecchia approximation to the joint distribution of the residual process  $(\mathbf{b}_s^{\text{T}}, \mathbf{b}_s^{p\text{T}})^{\text{T}}$  at the training and prediction inputs whose covariance matrix is

$$\text{Cov} \left( (\mathbf{b}_s^{\text{T}}, \mathbf{b}_s^{p\text{T}})^{\text{T}} \right) = \begin{pmatrix} \tilde{\boldsymbol{\Sigma}} & \boldsymbol{\Sigma}_{mn_p} \\ \boldsymbol{\Sigma}_{mn_p}^{\text{T}} & \tilde{\boldsymbol{\Sigma}}_p \end{pmatrix} - \begin{pmatrix} \boldsymbol{\Sigma}_{mn}^{\text{T}} \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn} & \boldsymbol{\Sigma}_{mn}^{\text{T}} \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn_p} \\ \boldsymbol{\Sigma}_{mn_p}^{\text{T}} \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn} & \boldsymbol{\Sigma}_{mn_p}^{\text{T}} \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn_p} \end{pmatrix}.$$

This gives

$$\text{Cov}\left(\left(\mathbf{b}_s^\top, \mathbf{b}_s^{p\top}\right)^\top\right)^{-1} \approx \begin{pmatrix} \mathbf{B} & \mathbf{0} \\ \mathbf{B}_{po} & \mathbf{B}_p \end{pmatrix}^\top \begin{pmatrix} \mathbf{D}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{D}_p^{-1} \end{pmatrix} \begin{pmatrix} \mathbf{B} & \mathbf{0} \\ \mathbf{B}_{po} & \mathbf{B}_p \end{pmatrix}. \quad (8)$$

It follows that

$$\text{Cov}\left(\left(\mathbf{b}_s^\top, \mathbf{b}_s^{p\top}\right)^\top\right) \approx \begin{pmatrix} (\mathbf{B}^\top \mathbf{D}^{-1} \mathbf{B})^{-1} & -(\mathbf{B}^\top \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{B}_{po}^\top \mathbf{B}_p^{-\top} \\ -\mathbf{B}_p^{-1} \mathbf{B}_{po} (\mathbf{B}^\top \mathbf{D}^{-1} \mathbf{B})^{-1} & (\mathbf{B}_p^\top \mathbf{D}_p^{-1} \mathbf{B}_p)^{-1} + \mathbf{B}_p^{-1} \mathbf{B}_{po} (\mathbf{B}^\top \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{B}_{po}^\top \mathbf{B}_p^{-\top} \end{pmatrix}$$

and thus

$$\text{Cov}\left(\left(\mathbf{y}^\top, \mathbf{y}^{p\top}\right)^\top\right) \approx \begin{pmatrix} \Sigma_{mn}^\top \Sigma_m^{-1} \Sigma_{mn} & \Sigma_{mn}^\top \Sigma_m^{-1} \Sigma_{mn_p} \\ \Sigma_{mn_p}^\top \Sigma_m^{-1} \Sigma_{mn} & \Sigma_{mn_p}^\top \Sigma_m^{-1} \Sigma_{mn_p} \end{pmatrix} + \begin{pmatrix} (\mathbf{B}^\top \mathbf{D}^{-1} \mathbf{B})^{-1} & -(\mathbf{B}^\top \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{B}_{po}^\top \mathbf{B}_p^{-\top} \\ -\mathbf{B}_p^{-1} \mathbf{B}_{po} (\mathbf{B}^\top \mathbf{D}^{-1} \mathbf{B})^{-1} & (\mathbf{B}_p^\top \mathbf{D}_p^{-1} \mathbf{B}_p)^{-1} + \mathbf{B}_p^{-1} \mathbf{B}_{po} (\mathbf{B}^\top \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{B}_{po}^\top \mathbf{B}_p^{-\top} \end{pmatrix}.$$

By standard arguments for conditional probabilities of multivariate Gaussian distributions, we obtain  $\mathbf{y}^p \mid \mathbf{y} \sim \mathcal{N}\left(\boldsymbol{\mu}_\dagger^p, \boldsymbol{\Sigma}_\dagger^p\right)$  in Proposition 1.  $\blacksquare$

In Appendix C, we additionally derive an alternative and equivalent expression for  $\boldsymbol{\Sigma}_\dagger^p$ , which allows for a more efficient calculation and is used in our software implementation. The predictive distribution  $\mathbf{b}^p \mid \mathbf{y}$  of the latent GP  $\mathbf{b}^p$  is obtained analogously by subtracting  $\sigma^2 \mathbf{I}_{n_p}$  from  $\boldsymbol{\Sigma}_\dagger^p$  in (7). The computational costs for the predictive means  $\boldsymbol{\mu}_\dagger^p$  and variances  $\text{diag}(\boldsymbol{\Sigma}_\dagger^p)$  are  $\mathcal{O}(n_p \cdot (m_v^3 + m_v^2 \cdot m) + n \cdot (m_v + m))$  and  $\mathcal{O}(n_p \cdot (m_v + m \cdot m_v) + n \cdot (m \cdot m_v + m^2))$ , respectively.

### 3. VIF-Laplace approximations for latent Gaussian process models

In the following, we assume that the response variable  $\mathbf{y} = (y(\mathbf{s}_1), \dots, y(\mathbf{s}_n))^\top \in \mathbb{R}^n$  follows a parametric distribution with a density of  $p(\mathbf{y} \mid \boldsymbol{\mu}, \boldsymbol{\xi}) = \prod_{i=1}^n p(y_i \mid \mu_i, \xi)$  with respect to a sigma finite product measure conditional on  $\boldsymbol{\mu} = F(\mathbf{X}) + \mathbf{b}$ . This density has, potentially link function-transformed, parameters  $\boldsymbol{\mu} \in \mathbb{R}^n$  and additional auxiliary parameters  $\boldsymbol{\xi} \in \Xi \subset \mathbb{R}^r$ . For instance,  $\mu_i$  and  $\xi$  can be the mean and variance of a Gaussian distribution, the log-mean and the shape parameter of a gamma likelihood, or  $\mu_i$  can be the logit-transformed success probability of a Bernoulli distribution for which there is no additional parameter  $\xi$ . Similarly as in (1), the parameter  $\boldsymbol{\mu}$  is modeled as the sum of fixed effects  $F(\mathbf{X})$  and a finite-dimensional version of a GP  $\mathbf{b}$ ,  $\boldsymbol{\mu} = F(\mathbf{X}) + \mathbf{b}$ , where  $\mathbf{X} \in \mathbb{R}^{n \times p}$  contains predictor variables and  $\mathbf{b} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$ . We again assume for simplicity that  $F(\mathbf{X}) = \mathbf{0}$ , but this assumption can be easily relaxed. Estimation of the covariance  $\boldsymbol{\theta} \in \mathbb{R}^q$  and auxiliary parameters  $\boldsymbol{\xi}$  is done by minimizing the negative log-marginal likelihood

$$-\log(p(\mathbf{y} \mid \boldsymbol{\theta}, \boldsymbol{\xi})) = -\log\left(\int p(\mathbf{y} \mid \mathbf{b}, \boldsymbol{\xi}) p(\mathbf{b} \mid \boldsymbol{\theta}) d\mathbf{b}\right). \quad (9)$$

For non-Gaussian likelihoods, there is typically no analytic expression for  $p(\mathbf{y} \mid \mathbf{b}, \boldsymbol{\xi})$  and an approximation has to be used. In this article, we use the Laplace approximation (Tierney and Kadane, 1986), additionally assuming that  $p(\mathbf{y} \mid \boldsymbol{\mu}, \boldsymbol{\xi})$  is log-concave in  $\boldsymbol{\mu}$ . The Laplace approximation for (9) is given by

$$L^{LA}(\mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\xi}) = -\log p(\mathbf{y} \mid \tilde{\mathbf{b}}, \boldsymbol{\xi}) + \frac{1}{2} \tilde{\mathbf{b}}^T \boldsymbol{\Sigma}^{-1} \tilde{\mathbf{b}} + \frac{1}{2} \log \det(\boldsymbol{\Sigma} \mathbf{W} + \mathbf{I}_n), \quad (10)$$

where  $\tilde{\mathbf{b}} = \operatorname{argmax}_{\mathbf{b}} \log p(\mathbf{y} \mid \mathbf{b}, \boldsymbol{\xi}) - \frac{1}{2} \mathbf{b}^T \boldsymbol{\Sigma}^{-1} \mathbf{b}$  is the mode of  $p(\mathbf{y} \mid \mathbf{b}, \boldsymbol{\xi}) p(\mathbf{b} \mid \boldsymbol{\theta})$ , and  $\mathbf{W} \in \mathbb{R}^{n \times n}$  is diagonal with

$$W_{ii} = - \left. \frac{\partial^2 \log p(y(\mathbf{s}_i) \mid b(\mathbf{s}_i), \boldsymbol{\xi})}{\partial b(\mathbf{s}_i)^2} \right|_{\mathbf{b}=\tilde{\mathbf{b}}}. \quad (11)$$

The mode  $\tilde{\mathbf{b}}$  is usually found with Newton's method (Rasmussen and Williams, 2006). Gradients of  $L^{LA}(\mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\xi})$  for parameter estimation can be found, e.g., in Rasmussen and Williams (2006) and Sigrist (2022b). We apply a VIF approximation to the latent process  $\mathbf{b}$  by again decomposing it into a low-rank and a residual process,  $\mathbf{b} = \mathbf{b}_l + \mathbf{b}_s$ , where  $\mathbf{b}_l$  and  $\mathbf{b}_s$  are defined as in Section 2.1 except that the residual process  $\mathbf{b}_s$  does not contain any error variance. Specifically, the precision matrix  $\operatorname{Cov}(\mathbf{b}_s)^{-1} = (\boldsymbol{\Sigma} - \boldsymbol{\Sigma}_{mn}^T \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn})^{-1}$  is approximated as

$$(\boldsymbol{\Sigma}^s)^{-1} = \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \approx \operatorname{Cov}(\mathbf{b}_s)^{-1},$$

where  $\mathbf{B}$  and  $\mathbf{D}$  are defined analogously as in Section 2.1, but  $\tilde{\boldsymbol{\Sigma}}_{ii}$  and  $\tilde{\boldsymbol{\Sigma}}_{N(i)}$  are replaced by  $\boldsymbol{\Sigma}_{ii}$  and  $\boldsymbol{\Sigma}_{N(i)}$  not including the error variance in (4). A VIF-Laplace approximation (VIFLA) is then given by combining the VIF and Laplace approximations to obtain the following approximate negative log-marginal likelihood  $-\log(p(\mathbf{y} \mid \mathbf{b}, \boldsymbol{\xi}))$ :

$$L^{VIFLA}(\mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\xi}) = -\log p(\mathbf{y} \mid \tilde{\mathbf{b}}, \boldsymbol{\xi}) + \frac{1}{2} \tilde{\mathbf{b}}^T \boldsymbol{\Sigma}_\dagger^{-1} \tilde{\mathbf{b}} + \frac{1}{2} \log \det(\boldsymbol{\Sigma}_\dagger \mathbf{W} + \mathbf{I}_n), \quad (12)$$

where  $\tilde{\mathbf{b}} = \operatorname{argmax}_{\mathbf{b}} \log p(\mathbf{y} \mid \mathbf{b}, \boldsymbol{\xi}) - \frac{1}{2} \mathbf{b}^T \boldsymbol{\Sigma}_\dagger^{-1} \mathbf{b}$  is the mode of  $p(\mathbf{y} \mid \mathbf{b}, \boldsymbol{\xi}) p_\dagger(\mathbf{b} \mid \boldsymbol{\theta})$  with  $p_\dagger(\mathbf{b} \mid \boldsymbol{\theta})$  being the density of  $\mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_\dagger)$  and

$$\boldsymbol{\Sigma}_\dagger^{-1} = (\mathbf{B}^{-1} \mathbf{D} \mathbf{B}^{-T} + \boldsymbol{\Sigma}_{mn}^T \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn})^{-1} = \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} - \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \boldsymbol{\Sigma}_{mn}^T \mathbf{M}^{-1} \boldsymbol{\Sigma}_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B},$$

where we have applied the Sherman-Woodbury-Morrison formula, and we recall that  $\mathbf{M}$  is defined in (5). By properties of the determinant of matrix products and Sylvester's determinant theorem, we have

$$\begin{aligned} \log \det(\boldsymbol{\Sigma}_\dagger \mathbf{W} + \mathbf{I}_n) &= -\log \det(\boldsymbol{\Sigma}_m) - \log \det(\mathbf{D}^{-1}) + \log \det(\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}) \\ &\quad + \log \det(\mathbf{M} - \boldsymbol{\Sigma}_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} (\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{B}^T \mathbf{D}^{-1} \boldsymbol{\Sigma}_{mn}^T), \end{aligned}$$

see Appendix B for the detailed derivation. Furthermore, one iteration of Newton's method is given by

$$\tilde{\mathbf{b}}^{t+1} = \left( \mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1} \right)^{-1} \left( \mathbf{W} \tilde{\mathbf{b}}^t + \frac{\partial \log p(\mathbf{y} \mid \tilde{\mathbf{b}}^t, \boldsymbol{\xi})}{\partial \mathbf{b}} \right), t = 0, 1, \dots \quad (13)$$

Applying the Sherman-Woodbury-Morrison formula, we can do linear solves with  $\mathbf{W} + \Sigma_{\dagger}^{-1}$  using the following relationship:

$$\begin{aligned}
 (\mathbf{W} + \Sigma_{\dagger}^{-1})^{-1} &= \mathbf{W}^{-1}(\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}(\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{W} - \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}(\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \\
 &\quad \cdot \mathbf{W} \Sigma_{mn}^T (\Sigma_m + \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}(\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{W} \Sigma_{mn}^T)^{-1} \\
 &\quad \cdot \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}(\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{W}) \Sigma_{\dagger},
 \end{aligned} \tag{14}$$

see Appendix B for a detailed derivation. The gradients of  $L^{VIFLA}(\mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\xi})$ , with respect to  $\boldsymbol{\theta}$  and  $\boldsymbol{\xi}$  are also derived in Appendix B. As the above results show, the computational complexity of evaluating the negative log-marginal likelihood and its derivatives is primarily determined by the Cholesky factorization of the matrix  $\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}$ . Despite that  $\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}$  is sparse, this is computationally expensive for large  $n$  as mentioned in the introduction. For two-dimensional spatial data, the Cholesky decomposition for  $\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}$  has complexity of approximately  $\mathcal{O}(n^{3/2})$ , while in higher dimensions, it can be  $\mathcal{O}(n^2)$  or larger (Lipton et al., 1979).

### 3.1 Prediction with VIF-Laplace approximations

The goal is to predict either the latent GP  $\mathbf{b}^p \in \mathbb{R}^{n_p}$  or the response variable  $\mathbf{y}^p \in \mathbb{R}^{n_p}$  at  $n_p$  new inputs  $\mathcal{S}^p = \{\mathbf{s}_1^p, \dots, \mathbf{s}_{n_p}^p\}$  using the posterior predictive distributions  $p(\mathbf{b}^p | \mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\xi})$  and  $p(\mathbf{y}^p | \mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\xi})$ , respectively. For the former, we have the following result.

**Proposition 2** *A VIF-Laplace-approximated posterior predictive distribution for  $p(\mathbf{b}^p | \mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\xi})$ ,  $\mathbf{b}^p = (b(\mathbf{s}_1^p), \dots, b(\mathbf{s}_{n_p}^p))^T \in \mathbb{R}^{n_p}$ , is given by  $\mathcal{N}(\boldsymbol{\omega}_p, \boldsymbol{\Omega}_p)$ , where*

$$\begin{aligned}
 \boldsymbol{\omega}_p &= -\mathbf{B}_p^{-1} \mathbf{B}_{po} \tilde{\mathbf{b}} + \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \tilde{\mathbf{b}} \\
 &\quad - \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \Sigma_{mn}^T \mathbf{M}^{-1} \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \tilde{\mathbf{b}} \\
 &\quad + \mathbf{B}_p^{-1} \mathbf{B}_{po} \Sigma_{mn}^T \mathbf{M}^{-1} \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \tilde{\mathbf{b}} \\
 \boldsymbol{\Omega}_p &= \mathbf{B}_p^{-1} \mathbf{D}_p \mathbf{B}_p^{-T} + \mathbf{B}_p^{-1} \mathbf{B}_{po} (\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{B}_{po}^T \mathbf{B}_p^{-T} + \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn_p} \\
 &\quad - (\Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} - \mathbf{B}_p^{-1} \mathbf{B}_{po} (\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1}) (\Sigma_{\dagger}^{-1} - \Sigma_{\dagger}^{-1} (\mathbf{W} + \Sigma_{\dagger}^{-1})^{-1} \Sigma_{\dagger}^{-1}) \\
 &\quad \cdot (\Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} - (\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{B}_{po}^T \mathbf{B}_p^{-T}),
 \end{aligned} \tag{15}$$

and  $\mathbf{B}_{po} \in \mathbb{R}_p^{n_p \times n}$ ,  $\mathbf{B}_p \in \mathbb{R}_p^{n_p \times n_p}$ , and  $\mathbf{D}_p^{-1} \in \mathbb{R}_p^{n_p \times n_p}$  are defined analogously as in (8).

**Proof** [Proof of Proposition 2]

This can be derived as follows. As in Section 2.3, a Vecchia approximation is applied to the joint distribution of the residual GP at the training and prediction inputs, and we obtain  $\mathbf{b}^p | \mathbf{b} \sim \mathcal{N}(\mathbf{v}_p, \boldsymbol{\Xi}_p)$  with

$$\begin{aligned}
 \mathbf{v}_p &= -\mathbf{B}_p^{-1} \mathbf{B}_{po} \mathbf{b} + \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \mathbf{b} \\
 &\quad - \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \Sigma_{mn}^T \mathbf{M}^{-1} \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \mathbf{b} \\
 &\quad + \mathbf{B}_p^{-1} \mathbf{B}_{po} \Sigma_{mn}^T \mathbf{M}^{-1} \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \mathbf{b} \\
 \boldsymbol{\Xi}_p &= \mathbf{B}_p^{-1} \mathbf{D}_p \mathbf{B}_p^{-T} + \mathbf{B}_p^{-1} \mathbf{B}_{po} (\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{B}_{po}^T \mathbf{B}_p^{-T} + \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn_p}.
 \end{aligned}$$

By the law of total probability, we have

$$p(\mathbf{b}^p \mid \mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\xi}) = \int p(\mathbf{b}^p \mid \mathbf{b}, \boldsymbol{\theta}) p(\mathbf{b} \mid \mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\xi}) d\mathbf{b}.$$

The Laplace approximation  $p(\mathbf{b} \mid \mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\xi}) \approx \mathcal{N}(\tilde{\mathbf{b}}, (\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})^{-1})$  and  $\mathbf{b}^p \mid \mathbf{b} \sim \mathcal{N}(\mathbf{v}_p, \boldsymbol{\Xi}_p)$  then give the result in (15).  $\blacksquare$

See Appendix C for an alternative and equivalent expression for  $\boldsymbol{\Omega}_p$ , which allows for a more efficient calculation and is used in our software implementation. For predicting the response variable, the integral  $p(\mathbf{y}^p \mid \mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\xi}) = \int p(\mathbf{y}^p \mid \mathbf{b}^p, \boldsymbol{\xi}) p(\mathbf{b}^p \mid \mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\xi}) d\mathbf{b}^p$  is analytically intractable for most likelihoods and must be approximated using numerical integration or by simulating from  $p(\mathbf{b}^p \mid \mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\xi}) \approx \mathcal{N}(\boldsymbol{\omega}_p, \boldsymbol{\Omega}_p)$ .

#### 4. Iterative methods for VIF-Laplace approximations

While VIF-Laplace approximations improve scalability, a sparse Cholesky factorization of the matrix  $\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}$  remains computationally expensive as mentioned above. To address this, we next show how iterative methods and stochastic approximations can be used for fast computations. Parameter estimation and prediction with VIF-Laplace approximations involve several computationally expensive operations outlined in the following. First, calculating linear solves  $(\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})\mathbf{u} = \mathbf{v}$ ,  $\mathbf{v} \in \mathbb{R}^n$ , (i) in Newton's method for finding the mode, see (13), (ii) for implicit derivatives of the log-marginal likelihood such as  $(\frac{\partial L^{\text{VIFLA}}(\mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\xi})}{\partial \tilde{\mathbf{b}}})^T \frac{\partial \tilde{\mathbf{b}}}{\partial \theta_k}$ , see Appendix B, and (iii) for predictive distributions in (15). The latter is particularly challenging as the number of prediction points  $n_p$  is typically large, and the linear systems thus contain many right-hand sides, see Section 4.2 for more information. Second, calculating log-determinants  $\log \det(\boldsymbol{\Sigma}_\dagger \mathbf{W} + \mathbf{I}_n)$  in the log-marginal likelihood given in (12). And third, calculating trace terms such as  $\text{Tr}((\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})^{-1} \frac{\partial(\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})}{\partial \theta_k}) = \frac{\partial \log \det(\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})}{\partial \theta_k}$  for the derivatives of log-determinants, see Appendix B.

##### 4.1 Likelihood evaluation and gradient calculation

To compute linear solves  $(\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})\mathbf{u} = \mathbf{v}$ , we apply the preconditioned conjugate gradient (CG) method (Saad, 2003). Preconditioning increases the convergence speed of the CG method, and it also reduces the variance of stochastic log-determinant and gradient approximations. Below in Section 4.3, we propose two preconditioners. Preconditioning means that we solve the equivalent system

$$\mathbf{P}^{-\frac{1}{2}}(\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})\mathbf{P}^{-\frac{T}{2}}\mathbf{P}^{\frac{T}{2}}\mathbf{u} = \mathbf{P}^{-\frac{1}{2}}\mathbf{v}, \quad (16)$$

where  $\mathbf{P}$  is a symmetric positive definite preconditioner matrix. Alternatively, since  $\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1} = \boldsymbol{\Sigma}_\dagger^{-1}(\mathbf{W}^{-1} + \boldsymbol{\Sigma}_\dagger)\mathbf{W}$ , we can equivalently solve

$$\mathbf{P}^{-\frac{1}{2}}(\mathbf{W}^{-1} + \boldsymbol{\Sigma}_\dagger)\mathbf{P}^{-\frac{T}{2}}\mathbf{P}^{\frac{T}{2}}\mathbf{W}\mathbf{u} = \mathbf{P}^{-\frac{1}{2}}\boldsymbol{\Sigma}_\dagger\mathbf{v}. \quad (17)$$

The version in (16) is used for the VIFDU preconditioner defined in Section 4.3.1, while (17) is used for the FITC preconditioner defined in Section 4.3.2. In each iteration of the

preconditioned CG method for (16) and (17), we calculate matrix-vector products of the form  $(\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} - \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \boldsymbol{\Sigma}_{mn}^T \mathbf{M}^{-1} \boldsymbol{\Sigma}_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}) \mathbf{w}$  and  $(\mathbf{W}^{-1} + \mathbf{B}^{-1} \mathbf{D} \mathbf{B}^{-T} + \boldsymbol{\Sigma}_{mn}^T \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn}) \mathbf{w}$ , respectively,  $\mathbf{w} \in \mathbb{R}^n$ , which require  $\mathcal{O}(n \cdot (m + m_v))$  operations. The computational complexity for computing and inverting the matrix  $\mathbf{M}$  is  $\mathcal{O}(n \cdot (m^2 + m \cdot m_v))$ .

To calculate  $\log \det(\boldsymbol{\Sigma}_\dagger \mathbf{W} + \mathbf{I}_n)$ , we employ the stochastic Lanczos quadrature (SLQ) method (Ubaru et al., 2017) which combines a quadrature technique based on the Lanczos algorithm (Lanczos, 1950) with Hutchinson’s stochastic trace estimator (Hutchinson, 1989). Dong et al. (2017) analyzed different approaches to estimate log-determinants and found that the SLQ method is superior to other methods. The preconditioned SLQ method gives the following approximation:

$$\log \det(\boldsymbol{\Sigma}_\dagger \mathbf{W} + \mathbf{I}_n) \approx \log \det(\boldsymbol{\Sigma}_\dagger) + \frac{n}{\ell} \sum_{i=1}^{\ell} \mathbf{e}_1^T \log(\tilde{\mathbf{T}}_i) \mathbf{e}_1 + \log \det(\mathbf{P}), \quad (18)$$

where  $\tilde{\mathbf{T}}_i \in \mathbb{R}^{k \times k}$  is the tridiagonal matrix of the Lanczos tridiagonalization  $\tilde{\mathbf{Q}}_i \tilde{\mathbf{T}}_i \tilde{\mathbf{Q}}_i^T \approx \mathbf{P}^{-\frac{1}{2}} (\mathbf{W} + \tilde{\boldsymbol{\Sigma}}_\dagger^{-1}) \mathbf{P}^{-\frac{1}{2}}$  obtained by running the Lanczos algorithm for  $k$  steps with the matrix  $\mathbf{P}^{-\frac{1}{2}} (\mathbf{W} + \tilde{\boldsymbol{\Sigma}}_\dagger^{-1}) \mathbf{P}^{-\frac{1}{2}}$  and initial vector  $\mathbf{P}^{-\frac{1}{2}} \mathbf{z}_i / \sqrt{n}$ , where  $\mathbf{z}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{P})$  and  $\sqrt{n}$  approximates the normalization factor  $\|\mathbf{P}^{-\frac{1}{2}} \mathbf{z}_i\|_2$ . See Appendix D for a detailed derivation. Alternatively, we can obtain the following SLQ approximation:

$$\log \det(\boldsymbol{\Sigma}_\dagger \mathbf{W} + \mathbf{I}_n) \approx \log \det(\mathbf{W}) + \frac{n}{\ell} \sum_{i=1}^{\ell} \mathbf{e}_1^T \log(\tilde{\mathbf{T}}_i) \mathbf{e}_1 + \log \det(\mathbf{P}), \quad (19)$$

where now  $\tilde{\mathbf{T}}_i \in \mathbb{R}^{k \times k}$  is the tridiagonal matrix of the partial Lanczos tridiagonalization  $\tilde{\mathbf{Q}}_i \tilde{\mathbf{T}}_i \tilde{\mathbf{Q}}_i^T \approx \mathbf{P}^{-\frac{1}{2}} (\mathbf{W}^{-1} + \boldsymbol{\Sigma}_\dagger) \mathbf{P}^{-\frac{1}{2}}$  obtained by running the Lanczos algorithm for  $k$  steps with the matrix  $\mathbf{P}^{-\frac{1}{2}} (\mathbf{W}^{-1} + \boldsymbol{\Sigma}_\dagger) \mathbf{P}^{-\frac{1}{2}}$  and initial vector  $\mathbf{P}^{-\frac{1}{2}} \mathbf{z}_i / \|\mathbf{P}^{-\frac{1}{2}} \mathbf{z}_i\|_2$ ,  $\mathbf{z}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{P})$ . The version in (18) is used for the VIFDU preconditioner defined in Section 4.3.1, whereas (19) is used for the FITC preconditioner defined in Section 4.3.2.

As in Gardner et al. (2018a), we use a technique from Saad (2003) to calculate the partial Lanczos tridiagonal matrices  $\tilde{\mathbf{T}}_1, \dots, \tilde{\mathbf{T}}_\ell$  from the coefficients of the preconditioned CG algorithm when solving  $(\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})^{-1} \mathbf{z}_i$ , or  $(\mathbf{W}^{-1} + \boldsymbol{\Sigma}_\dagger)^{-1} \mathbf{z}_i$ ,  $i = 1, \dots, \ell$ . In doing so, we avoid running the Lanczos algorithm, which brings multiple advantages: Numerical instabilities are not an issue and storing  $\tilde{\mathbf{Q}}_i$  is not necessary. In addition to the  $\ell$  partial Lanczos tridiagonal matrices, the modified CG method computes the  $\ell$  linear solves  $(\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})^{-1} \mathbf{z}_i$ , or  $(\mathbf{W}^{-1} + \boldsymbol{\Sigma}_\dagger)^{-1} \mathbf{z}_i$ , respectively, for the probe vectors  $\mathbf{z}_i$ . This brings the advantage that once the log-determinant is calculated, its gradients can be calculated with minimal computational overhead. Specifically, the trace terms for calculating derivatives of log-determinants given in Appendix B can be computed using stochastic trace estimation

(STE) as follows:

$$\begin{aligned}
 \text{Tr} \left( (\mathbf{W} + \boldsymbol{\Sigma}_{\dagger}^{-1})^{-1} \frac{\partial(\mathbf{W} + \boldsymbol{\Sigma}_{\dagger}^{-1})}{\partial b^*(\mathbf{s}_i)} \right) &= \text{Tr} \left( (\mathbf{W} + \boldsymbol{\Sigma}_{\dagger}^{-1})^{-1} \frac{\partial(\mathbf{W} + \boldsymbol{\Sigma}_{\dagger}^{-1})}{\partial b^*(\mathbf{s}_i)} \mathbb{E}_{\mathbf{z}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{P})} [\mathbf{P}^{-1} \mathbf{z}_i \mathbf{z}_i^{\text{T}}] \right) \\
 &\approx \frac{1}{\ell} \sum_{i=1}^{\ell} (\mathbf{z}_i^{\text{T}} (\mathbf{W} + \boldsymbol{\Sigma}_{\dagger}^{-1})^{-1}) \left( \frac{\partial \mathbf{W}}{\partial b^*(\mathbf{s}_i)} \mathbf{P}^{-1} \mathbf{z}_i \right) \quad i = 1, \dots, p, \\
 \text{Tr} \left( (\mathbf{W} + \boldsymbol{\Sigma}_{\dagger}^{-1})^{-1} \frac{\partial(\mathbf{W} + \boldsymbol{\Sigma}_{\dagger}^{-1})}{\partial \theta_k} \right) &\approx \frac{1}{\ell} \sum_{i=1}^{\ell} (\mathbf{z}_i^{\text{T}} (\mathbf{W} + \boldsymbol{\Sigma}_{\dagger}^{-1})^{-1}) \left( \frac{\partial \boldsymbol{\Sigma}_{\dagger}^{-1}}{\partial \theta_k} \mathbf{P}^{-1} \mathbf{z}_i \right) \quad k = 1, \dots, q, \\
 \text{Tr} \left( (\mathbf{W} + \boldsymbol{\Sigma}_{\dagger}^{-1})^{-1} \frac{\partial(\mathbf{W} + \boldsymbol{\Sigma}_{\dagger}^{-1})}{\partial \xi_l} \right) &\approx \frac{1}{\ell} \sum_{i=1}^{\ell} (\mathbf{z}_i^{\text{T}} (\mathbf{W} + \boldsymbol{\Sigma}_{\dagger}^{-1})^{-1}) \left( \frac{\partial \mathbf{W}}{\partial \xi_l} \mathbf{P}^{-1} \mathbf{z}_i \right) \quad l = 1, \dots, r,
 \end{aligned}$$

see Appendix D for a detailed derivation. We additionally apply variance reduction by using the preconditioner  $\mathbf{P}$  to construct a control variate (Lemieux, 2014), see Appendix D for more information. We choose the Lanczos rank  $k$  adaptively by running the preconditioned CG algorithm until it has converged.

In each iteration of the preconditioned CG method, we calculate matrix-vector products, which involve  $\mathcal{O}(n \cdot (m + m_v))$  operations. In total, we have computational cost for calculating the negative log-marginal likelihood and its derivatives of  $\mathcal{O}(n \cdot (m_v^3 + m^2 + m \cdot t + m_v \cdot t + m \cdot m_v^2))$ , where  $t$  is the number of iterations of the CG algorithm.

## 4.2 Predictive variances

Calculating predictive (co-)variances given in (15) is computationally expensive when  $n$  and  $n_p$  are large, even with iterative methods due to  $n_p$  right-hand sides. In the following, we propose two computationally efficient simulation-based approaches both relying on iterative methods to compute  $\text{diag}(\boldsymbol{\Omega}_p)$ . For this, we split  $\boldsymbol{\Omega}_p$  given in (15) into two parts:

$$\begin{aligned}
 &\mathbf{B}_p^{-1} \mathbf{D}_p \mathbf{B}_p^{-\text{T}} + \mathbf{B}_p^{-1} \mathbf{B}_{po} (\mathbf{B}^{\text{T}} \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{B}_{po}^{\text{T}} \mathbf{B}_p^{-\text{T}} + \boldsymbol{\Sigma}_{mn_p}^{\text{T}} \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn_p} \\
 &\quad - (\boldsymbol{\Sigma}_{mn_p}^{\text{T}} \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn} - \mathbf{B}_p^{-1} \mathbf{B}_{po} (\mathbf{B}^{\text{T}} \mathbf{D}^{-1} \mathbf{B})^{-1}) \boldsymbol{\Sigma}_{\dagger}^{-1} \\
 &\quad \cdot (\boldsymbol{\Sigma}_{mn}^{\text{T}} \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn_p} - (\mathbf{B}^{\text{T}} \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{B}_{po}^{\text{T}} \mathbf{B}_p^{-\text{T}})
 \end{aligned} \tag{20}$$

and

$$\begin{aligned}
 &(\boldsymbol{\Sigma}_{mn_p}^{\text{T}} \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn} - \mathbf{B}_p^{-1} \mathbf{B}_{po} (\mathbf{B}^{\text{T}} \mathbf{D}^{-1} \mathbf{B})^{-1}) \boldsymbol{\Sigma}_{\dagger}^{-1} (\mathbf{W} + \boldsymbol{\Sigma}_{\dagger}^{-1})^{-1} \boldsymbol{\Sigma}_{\dagger}^{-1} \\
 &\quad \cdot (\boldsymbol{\Sigma}_{mn}^{\text{T}} \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn_p} - (\mathbf{B}^{\text{T}} \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{B}_{po}^{\text{T}} \mathbf{B}_p^{-\text{T}}),
 \end{aligned} \tag{21}$$

where the diagonal of the first part in (20) can be computed efficiently in a deterministic manner without using iterative methods, because the term  $\mathbf{B}_p^{-1} \mathbf{B}_{po} (\mathbf{B}^{\text{T}} \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{B}_{po}^{\text{T}} \mathbf{B}_p^{-\text{T}}$  cancels out.

In Algorithm 1, the second term in (21) is approximated stochastically by sampling from a normal distribution with the matrix in (21) as its covariance matrix. This algorithm provides an unbiased and consistent approximation for  $\boldsymbol{\Omega}_p$ ; see Appendix C for the proof of Proposition 3. It can also be adapted to compute only the predictive variances  $\text{diag}(\boldsymbol{\Omega}_p)$  by summing  $\mathbf{z}_i^{(8)} \circ \mathbf{z}_i^{(8)}$  in Line 9 and extracting the diagonal of each term in Line 11. To

compute the linear solve  $(\Sigma_{\dagger}^{-1} + \mathbf{W})^{-1} \mathbf{z}_i^{(6)}$ , or  $\mathbf{W}^{-1}(\Sigma_{\dagger} + \mathbf{W}^{-1})^{-1} \Sigma_{\dagger} \mathbf{z}_i^{(6)}$ , in Line 7, we use the preconditioned CG method.

---

**Algorithm 1** Simulation-based Predictive (Co-)Variance Estimator (SBPV)

---

**Require:** Matrices  $\mathbf{B}_{po}$ ,  $\mathbf{D}_p$ ,  $\mathbf{D}^{-1}$ ,  $\mathbf{B}$ ,  $\mathbf{B}_p^{-1}$ ,  $\Sigma_{mn}$ ,  $\Sigma_m$ ,  $\Sigma_{mn_p}$ ,  $\mathbf{M}$ ,  $\mathbf{W}$

**Ensure:** Approximated predictive covariances  $\Omega_p$

- 1: Initialize:  $\mathbf{Z} \leftarrow \mathbf{0} \in \mathbb{R}^{n_p \times n_p}$
  - 2: **for**  $i = 1$  to  $\ell$  **do**
  - 3:   Let  $\mathbf{z}_i^{(1)}$ ,  $\mathbf{z}_i^{(2)} \sim \mathcal{N}(0, \mathbf{I}_n)$  and  $\mathbf{z}_i^{(3)} \sim \mathcal{N}(0, \mathbf{I}_m)$
  - 4:    $\mathbf{z}_i^{(4)} \leftarrow \Sigma_{mn}^T \Sigma_m^{-\frac{1}{2}} \mathbf{z}_i^{(3)} + \mathbf{B}^{-1} \mathbf{D}^{\frac{1}{2}} \mathbf{z}_i^{(1)} \quad \triangleright \mathbf{z}_i^{(4)} \sim \mathcal{N}(0, \Sigma_{\dagger})$
  - 5:    $\mathbf{z}_i^{(5)} \leftarrow (\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} - \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \Sigma_{mn}^T \mathbf{M}^{-1} \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}) \mathbf{z}_i^{(4)} \quad \triangleright \mathbf{z}_i^{(5)} \sim \mathcal{N}(0, \Sigma_{\dagger}^{-1})$
  - 6:    $\mathbf{z}_i^{(6)} \leftarrow \mathbf{z}_i^{(5)} + \mathbf{W}^{\frac{1}{2}} \mathbf{z}_i^{(2)} \quad \triangleright \mathbf{z}_i^{(6)} \sim \mathcal{N}(0, \Sigma_{\dagger}^{-1} + \mathbf{W})$
  - 7:    $\mathbf{z}_i^{(7)} \leftarrow \begin{cases} (\Sigma_{\dagger}^{-1} + \mathbf{W})^{-1} \mathbf{z}_i^{(6)} \\ \text{or } \mathbf{W}^{-1}(\Sigma_{\dagger} + \mathbf{W}^{-1})^{-1} \Sigma_{\dagger} \mathbf{z}_i^{(6)} \end{cases} \quad \triangleright \mathbf{z}_i^{(7)} \sim \mathcal{N}(0, (\Sigma_{\dagger}^{-1} + \mathbf{W})^{-1})$
  - 8:    $\mathbf{z}_i^{(8)} \leftarrow (\Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} - \mathbf{B}_p^{-1} \mathbf{B}_{po} (\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1}) \Sigma_{\dagger}^{-1} \mathbf{z}_i^{(7)}$   
 $\quad \triangleright \mathbf{z}_i^{(8)} \sim \mathcal{N}\left(0, (\Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} - \mathbf{B}_p^{-1} \mathbf{B}_{po} (\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1}) \Sigma_{\dagger}^{-1} (\Sigma_{\dagger}^{-1} + \mathbf{W})^{-1} \Sigma_{\dagger}^{-1} \cdot (\Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} - (\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{B}_{po}^T \mathbf{B}_p^{-T})\right)$
  - 9:    $\mathbf{Z} \leftarrow \mathbf{Z} + \mathbf{z}_i^{(8)} (\mathbf{z}_i^{(8)})^T$
  - 10: **end for**
  - 11:  $\Omega_p \leftarrow \mathbf{B}_p^{-1} \mathbf{D}_p \mathbf{B}_p^{-T} + \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn_p} - \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p}$   
 $+ \mathbf{B}_p^{-1} \mathbf{B}_{po} \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} - \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \Sigma_{mn}^T \mathbf{M}^{-1} \Sigma_{mn} \mathbf{B}_{po}^T \mathbf{B}_p^{-T}$   
 $+ (\mathbf{B}_p^{-1} \mathbf{B}_{po} \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p})^T - (\Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \Sigma_{mn}^T \mathbf{M}^{-1} \Sigma_{mn} \mathbf{B}_{po}^T \mathbf{B}_p^{-T})^T$   
 $+ \mathbf{B}_p^{-1} \mathbf{B}_{po} \Sigma_{mn}^T \mathbf{M}^{-1} \Sigma_{mn} \mathbf{B}_{po}^T \mathbf{B}_p^{-T}$   
 $+ \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \Sigma_{mn}^T \mathbf{M}^{-1} \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} + \frac{1}{\ell} \mathbf{Z}$
- 

**Proposition 3** *Algorithm 1 produces an unbiased and consistent estimator for the predictive covariance matrix  $\Omega_p$ .*

In Algorithm 2, we estimate the diagonal of matrix in (21) using the stochastic approach proposed by Bekas et al. (2007), which approximates the diagonal of a matrix  $\mathbf{A} \in \mathbb{R}^{n \times n}$  as  $\text{diag}(\mathbf{A}) \approx \sum_{i=1}^{\ell} \mathbf{z}_i \circ \mathbf{A} \mathbf{z}_i$ , where  $\mathbf{z}_i$  are Rademacher random vectors with entries in  $\{\pm 1\}$ . To compute the linear solves involving  $(\Sigma_{\dagger}^{-1} + \mathbf{W})^{-1}$ , or  $(\Sigma_{\dagger} + \mathbf{W}^{-1})^{-1}$ , respectively, in line 3, we again apply the preconditioned CG method. This algorithm also results in an unbiased and consistent approximation for  $\text{diag}(\Omega_p)$ , see Appendix C for a proof of Proposition 4.

**Proposition 4** *Algorithm 2 produces an unbiased and consistent estimator of the predictive variance  $\text{diag}(\Omega_p)$ .*

Both Algorithms 1 and 2 have a computational complexity of  $\mathcal{O}((n + n_p) \cdot (\ell \cdot m + \ell \cdot m_v + m \cdot m_v))$  for the calculation of predictive variances.

---

**Algorithm 2** Stochastic Predictive Variance Estimator (SPV)
 

---

**Require:** Matrices  $B_{po}$ ,  $D_p$ ,  $D^{-1}$ ,  $B$ ,  $B_p^{-1}$ ,  $\Sigma_{mn}$ ,  $\Sigma_m$ ,  $\Sigma_{mn_p}$ ,  $M$ ,  $W$ 
**Ensure:** Approximated predictive variances  $\text{diag}(\Omega_p)$ 

 1: Initialize:  $Z \leftarrow \mathbf{0} \in \mathbb{R}^{n_p}$ 

 2: **for**  $i = 1$  to  $\ell$  **do**

$$3: \quad z_i^{(2)} \leftarrow \begin{cases} (\Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} - B_p^{-1} B_{po} (B^T D^{-1} B)^{-1}) \Sigma_{\dagger}^{-1} (W + \Sigma_{\dagger}^{-1})^{-1} \\ \cdot \Sigma_{\dagger}^{-1} (\Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} - (B^T D^{-1} B)^{-1} B_{po}^T B_p^{-T}) z_i^{(1)} \\ \text{or } (\Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} - B_p^{-1} B_{po} (B^T D^{-1} B)^{-1}) \Sigma_{\dagger}^{-1} W^{-1} \\ \cdot (W^{-1} + \Sigma_{\dagger})^{-1} (\Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} - (B^T D^{-1} B)^{-1} B_{po}^T B_p^{-T}) z_i^{(1)} \end{cases}$$

 where  $z_i^{(1)} \sim \text{Rademacher}$ 

 4:  $Z \leftarrow Z + z_i^{(1)} \circ z_i^{(2)}$ 

 5: **end for**

$$6: \text{diag}(\Omega_p) \leftarrow \text{diag}(B_p^{-1} D_p B_p^{-T}) + \text{diag}(\Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn_p}) \\ - \text{diag}(\Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} B^T D^{-1} B \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p}) + 2 \cdot \text{diag}(B_p^{-1} B_{po} \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p}) \\ - 2 \cdot \text{diag}(\Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} B^T D^{-1} B \Sigma_{mn}^T M^{-1} \Sigma_{mn} B_{po}^T B_p^{-T}) \\ + \text{diag}(B_p^{-1} B_{po} \Sigma_{mn}^T M^{-1} \Sigma_{mn} B_{po}^T B_p^{-T}) \\ + \text{diag}(\Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} B^T D^{-1} B \Sigma_{mn}^T M^{-1} \Sigma_{mn} B^T D^{-1} B \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p}) + \frac{1}{\ell} Z$$


---

Alternatively, Pleiss et al. (2018) propose to use the Lanczos algorithm to approximate predictive variances. However, recent research (Kündig and Sigrüst, 2025; Gyger et al., 2026) has found that this Lanczos tridiagonalization-based approach is considerably more inaccurate compared to simulation-based methods. Intuitively, the Lanczos algorithm can work relatively well for approximating covariance matrices since their eigenvalue distribution often consists of a few large and many small eigenvalues, making them amenable to low-rank approximations. However, when applied to invert a covariance matrix, as required for computing predictive covariances, the eigenvalue distribution is reversed, resulting in many large eigenvalues and necessitating a very high rank for moderate accuracy.

### 4.3 Preconditioners

In the following, we present two preconditioners for iterative methods with VIF-Laplace approximations.

#### 4.3.1 THE VIF WITH DIAGONAL UPDATE PRECONDITIONER

The “**VIF** with **diagonal update**” (VIFDU) preconditioner is inspired by the “**Vecchia** approximation with **diagonal update**” (VADU) preconditioner introduced by Küdig and Sigrüst (2025) and exploits the structure of the Vecchia approximation to approximate

$\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} + \mathbf{W} \approx \mathbf{B}^T (\mathbf{D}^{-1} + \mathbf{W}) \mathbf{B}$ . Specifically, the VIFDU preconditioner  $\widehat{\mathbf{P}}$  is given by

$$\begin{aligned} \widehat{\mathbf{P}} &= \mathbf{B}^T (\mathbf{D}^{-1} - \mathbf{D}^{-1} \mathbf{B} \boldsymbol{\Sigma}_{mn}^T \mathbf{M}^{-1} \boldsymbol{\Sigma}_{mn} \mathbf{B}^T \mathbf{D}^{-1} + \mathbf{W}) \mathbf{B} \\ &\approx \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} - \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \boldsymbol{\Sigma}_{mn}^T \mathbf{M}^{-1} \boldsymbol{\Sigma}_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} + \mathbf{W} \\ &= (\boldsymbol{\Sigma}^s + \boldsymbol{\Sigma}_{mn}^T \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn})^{-1} + \mathbf{W} \\ &= \boldsymbol{\Sigma}_{\dagger}^{-1} + \mathbf{W}. \end{aligned}$$

This VIFDU preconditioner is applied when using the versions in (16) and (18) for the CG and SLQ methods, respectively. Linear solves with  $\widehat{\mathbf{P}}$ , the log-determinant  $\log \det(\widehat{\mathbf{P}})$  and its derivatives are computed in  $\mathcal{O}(n \cdot (m_v \cdot m + m^2))$  time. Sampling from  $\mathcal{N}(\mathbf{0}, \widehat{\mathbf{P}})$  is not straightforward since we can not apply the reparameterization trick used by Gardner et al. (2018a, Appendix C.1). If  $\boldsymbol{\epsilon}_1 \in \mathbb{R}^m$  and  $\boldsymbol{\epsilon}_2 \in \mathbb{R}^n$  are standard normal vectors, then  $(\mathbf{B}^T (\mathbf{W} + \mathbf{D}^{-1})^{\frac{1}{2}} \boldsymbol{\epsilon}_2 - \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \boldsymbol{\Sigma}_{mn}^T \mathbf{M}^{-\frac{T}{2}} \boldsymbol{\epsilon}_1)$ , is a sample from the distribution  $\mathcal{N}(\mathbf{0}, \mathbf{B}^T (\mathbf{W} + \mathbf{D}^{-1} + \mathbf{D}^{-1} \mathbf{B} \boldsymbol{\Sigma}_{mn}^T \mathbf{M}^{-1} \boldsymbol{\Sigma}_{mn} \mathbf{B}^T \mathbf{D}^{-1}) \mathbf{B})$  and not from  $\mathcal{N}(\mathbf{0}, \mathbf{B}^T (\mathbf{W} + \mathbf{D}^{-1} - \mathbf{D}^{-1} \mathbf{B} \boldsymbol{\Sigma}_{mn}^T \mathbf{M}^{-1} \boldsymbol{\Sigma}_{mn} \mathbf{B}^T \mathbf{D}^{-1}) \mathbf{B})$ , where  $\mathbf{M}^{\frac{1}{2}}$  is the Cholesky factor of  $\mathbf{M}$  and  $(\mathbf{W} + \mathbf{D}^{-1})^{\frac{1}{2}}$  is the elementwise square-root of  $(\mathbf{W} + \mathbf{D}^{-1})$ . However, sampling from  $\mathcal{N}(\mathbf{0}, \widehat{\mathbf{P}})$  can be done as follows. First, for sampling from  $\mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{\dagger}^{-1})$ , we compute a sample from  $\mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{\dagger})$  as  $\mathbf{B}^{-1} \mathbf{D}^{\frac{1}{2}} \boldsymbol{\epsilon}_2 + \boldsymbol{\Sigma}_{mn}^T \boldsymbol{\Sigma}_m^{-\frac{1}{2}} \boldsymbol{\epsilon}_1$ , where  $\boldsymbol{\Sigma}_m^{\frac{1}{2}}$  is the Cholesky factor of  $\boldsymbol{\Sigma}_m$ . Then we multiply by  $\boldsymbol{\Sigma}_{\dagger}^{-1}$  to obtain a sample from  $\mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{\dagger}^{-1})$  and add  $\mathbf{B}^T \mathbf{W}^{\frac{1}{2}} \boldsymbol{\epsilon}_3$ , where  $\boldsymbol{\epsilon}_3 \in \mathbb{R}^n$  is a standard normal vector, such that  $\mathbf{B}^T \mathbf{W}^{\frac{1}{2}} \boldsymbol{\epsilon}_3 + \boldsymbol{\Sigma}_{\dagger}^{-1} (\mathbf{B}^{-1} \mathbf{D}^{\frac{1}{2}} \boldsymbol{\epsilon}_2 + \boldsymbol{\Sigma}_{mn}^T \boldsymbol{\Sigma}_m^{-\frac{1}{2}} \boldsymbol{\epsilon}_1) \sim \mathcal{N}(\mathbf{0}, \mathbf{B}^T (\mathbf{W} + \mathbf{D}^{-1} - \mathbf{D}^{-1} \mathbf{B} \boldsymbol{\Sigma}_{mn}^T \mathbf{M}^{-1} \boldsymbol{\Sigma}_{mn} \mathbf{B}^T \mathbf{D}^{-1}) \mathbf{B}) = \mathcal{N}(\mathbf{0}, \widehat{\mathbf{P}})$ . The computational overhead for this procedure is  $\mathcal{O}(n \cdot (m_v + m))$ . For further information on the VIFDU preconditioner including how to do linear solves, calculate log-determinants and their derivatives, see Appendix E.1.

#### 4.3.2 THE FITC PRECONDITIONER

The **FITC** preconditioner is given by

$$\widehat{\mathbf{P}} = \boldsymbol{\Sigma}_{kn}^T \boldsymbol{\Sigma}_k^{-1} \boldsymbol{\Sigma}_{kn} + \text{diag}(\boldsymbol{\Sigma} - \boldsymbol{\Sigma}_{kn}^T \boldsymbol{\Sigma}_k^{-1} \boldsymbol{\Sigma}_{kn}) + \mathbf{W}^{-1}$$

where  $\boldsymbol{\Sigma}_k = [c_{\theta}(\hat{\mathbf{s}}_i, \hat{\mathbf{s}}_j)]_{i=1:k, j=1:k} \in \mathbb{R}^{k \times k}$  and  $\boldsymbol{\Sigma}_{kn} = [c_{\theta}(\hat{\mathbf{s}}_i, \mathbf{s}_j)]_{i=1:k, j=1:n} \in \mathbb{R}^{k \times n}$  are (cross-) covariance matrices with respect to a set of inducing points  $\widehat{\mathcal{S}} = \{\hat{\mathbf{s}}_1, \dots, \hat{\mathbf{s}}_k\}$ . This preconditioner is applied when using the versions in (17) and (19) for the CG and SLQ methods, respectively. Note that the FITC preconditioner can use a different set of inducing points than those used in the VIF approximation. For example, using a larger number of inducing points can lead to a more effective preconditioner compared to relying solely on the inducing points used by the VIF approximation. The construction of the FITC preconditioner, linear solves involving  $\widehat{\mathbf{P}}$ , the computation of the log-determinant  $\log \det(\widehat{\mathbf{P}})$ , and its derivatives all require  $\mathcal{O}(n \cdot k^2)$  operations. Sampling from  $\mathcal{N}(\mathbf{0}, \widehat{\mathbf{P}})$  requires  $\mathcal{O}(n \cdot k)$  operations. For additional details, see Appendix E.2.

An alternative low-rank preconditioner for Gaussian processes is the pivoted Cholesky decomposition (Harbrecht et al., 2012; Gardner et al., 2018a). However, the construction

of the pivoted Cholesky preconditioner is computationally more expensive than the FITC preconditioner. This means that the FITC preconditioner can use a higher rank for the same computational budget compared to the pivoted Cholesky preconditioner. Previous research (Gyger et al., 2026) has shown that the FITC preconditioner is more accurate than the pivoted Cholesky preconditioner, also for GP models that do not use inducing points.

## 5. Convergence analysis

In the following, we analyze the convergence properties of the preconditioned CG method with the VIFDU and FITC preconditioners. We show that the convergence speed is influenced by the VIF approximation tuning parameters, the number of inducing points  $m$  and the number of Vecchia neighbors  $m_v$ , and the eigenvalue structure of the covariance matrix  $\Sigma$ . We denote the Frobenius and the 2-norm (spectral norm) by  $\|\cdot\|_F$  and  $\|\cdot\|_2$ , respectively, and define the vector norm  $\|\mathbf{v}\|_{\mathbf{A}} = \sqrt{\mathbf{v}^T \mathbf{A} \mathbf{v}}$  for  $\mathbf{v} \in \mathbb{R}^n$  and a positive semi-definite matrix  $\mathbf{A} \in \mathbb{R}^{n \times n}$ . In the subsequent theorems, we make the following assumptions:

**Assumption 1**  $n \geq 2$  and  $m \in \{1, 2, \dots, n\}$ .

**Assumption 2** The covariance matrix  $\Sigma$  is of the form  $\Sigma_{ij} = \sigma_1^2 \cdot r(\mathbf{s}_i, \mathbf{s}_j)$ , where  $r(\cdot, \cdot)$  is positive and continuous, and  $r(\mathbf{s}_i, \mathbf{s}_i) = 1$ . Additionally, the matrix  $\Sigma$  has eigenvalues  $\lambda_1 \geq \dots \geq \lambda_n > 0$ .

**Assumption 3** The response variable  $\mathbf{y} \in \{0, 1\}^n$  is binary, following a Bernoulli likelihood with a logit link function.

Assumption 3 implies that the diagonal entries of  $\mathbf{W}$ , defined in equation (11), satisfy  $0 \leq \mathbf{W}_{ii} \leq 1$  for  $i \in \{1, 2, \dots, n\}$ .

First, we analyze the convergence speed of the preconditioned CG method for linear solves with  $\Sigma_{\dagger}^{-1} + \mathbf{W}$  using the VIFDU preconditioner.

### Theorem 5 Convergence of the CG method with the VIFDU preconditioner

Let  $\Sigma_{\dagger} \in \mathbb{R}^{n \times n}$  be the VIF approximation of a covariance matrix  $\Sigma \in \mathbb{R}^{n \times n}$  with  $m$  inducing points and  $m_v$  Vecchia neighbors. Consider the linear system  $(\Sigma_{\dagger}^{-1} + \mathbf{W})\mathbf{u}^* = \mathbf{v}$ , where  $\mathbf{v} \in \mathbb{R}^n$ . Let  $\mathbf{u}_k$  be the approximation in the  $k$ -th iteration of the preconditioned CG method with the VIFDU preconditioner  $\hat{\mathbf{P}} \in \mathbb{R}^{n \times n}$ . Under Assumptions 1–3, the following holds for the relative error:

$$\frac{\|\mathbf{u}^* - \mathbf{u}_k\|_{\Sigma_{\dagger}^{-1} + \mathbf{W}}}{\|\mathbf{u}^* - \mathbf{u}_0\|_{\Sigma_{\dagger}^{-1} + \mathbf{W}}} \leq 2 \cdot \left( \frac{1}{1 + \left( \alpha^n \cdot (\lambda_1 + (\lambda_{m+1} \cdot m_v)^n) \cdot (\sqrt{n} \cdot m_v + 1) \right)^{-1}} \right)^k,$$

where  $\alpha > 1$  is a constant.

For the proof, see Appendix F. Theorem 5 shows that selecting less Vecchia neighbors  $m_v$  leads to faster convergence of the CG method. However, the relationship is more complicated regarding the number of inducing points  $m$ . While the term  $\lambda_{m+1}$  decreases with

larger  $m$  leading to faster convergence, we bound  $\|\Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn}\|_2$  in the proof of Theorem 5 by  $\|\Sigma\|_2 = \lambda_1$  (see Appendix F), and  $\|\Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn}\|_2$  grows with  $m$ . Additionally, we observe that the convergence is slower with increasing values of  $\lambda_1$  and  $\lambda_{m+1}$ .

Next, we analyze the convergence speed of the preconditioned CG method for linear solves with  $\Sigma_{\dagger} + \mathbf{W}^{-1}$  using the FITC preconditioner.

**Theorem 6** *Convergence of the CG method with the FITC preconditioner*

Let  $\Sigma_{\dagger} \in \mathbb{R}^{n \times n}$  be the VIF approximation of a covariance matrix  $\Sigma \in \mathbb{R}^{n \times n}$  with  $m$  inducing points and  $m_v$  Vecchia neighbors. Consider the linear system  $(\Sigma_{\dagger} + \mathbf{W}^{-1})\mathbf{u}^* = \mathbf{v}$ , where  $\mathbf{v} \in \mathbb{R}^n$ . Let  $\mathbf{u}_k$  be the approximation in the  $k$ -th iteration of the preconditioned CG method with the FITC preconditioner  $\hat{\mathbf{P}} \in \mathbb{R}^{n \times n}$ . The preconditioner is constructed with the same set of inducing points as those used for  $\Sigma_{\dagger}$ . Under Assumptions 1–3, the following holds for the relative error:

$$\frac{\|\mathbf{u}^* - \mathbf{u}_k\|_{\Sigma_{\dagger} + \mathbf{W}^{-1}}}{\|\mathbf{u}^* - \mathbf{u}_0\|_{\Sigma_{\dagger} + \mathbf{W}^{-1}}} \leq 2 \cdot \left( \frac{1}{1 + (\alpha \cdot \lambda_{m+1} \cdot m_v)^{-n}} \right)^k,$$

where  $\alpha > 1$  is a constant.

For the proof, see Appendix F. Theorem 6 shows that selecting fewer Vecchia neighbors  $m_v$  and a higher number of inducing points  $m$  leads to faster convergence of the preconditioned CG method. In contrast to the results in Theorem 5, the convergence rate of the CG method with the FITC preconditioner does not depend on the largest eigenvalue  $\lambda_1$ . Consequently, using the FITC preconditioner exhibits less sensitivity to the specific covariance matrix, the sample size, and the parametric covariance function.

## 6. Selecting inducing points and Vecchia neighbors

There are various ways of choosing inducing points. We use the kMeans++ algorithm (Arthur and Vassilvitskii, 2007) since Gyger et al. (2026) have found that it is an effective method for selecting inducing points for Gaussian processes with full-scale and FITC approximations. Its computational complexity is  $\mathcal{O}(n \cdot m)$ . If the covariance function is of the form

$$c_{\theta}(\mathbf{s}_i, \mathbf{s}_j) = c_{\theta}^o(\|q_{\lambda}(\mathbf{s}_i) - q_{\lambda}(\mathbf{s}_j)\|), \tag{22}$$

where  $q_{\lambda}(\cdot)$  is a transformation and  $c_{\theta}^o(\cdot)$  is an isotropic covariance function such as the Matérn kernel, then the distance  $\|q_{\lambda}(\mathbf{s}_i) - q_{\lambda}(\mathbf{s}_j)\|$  depends on scaled inputs  $\mathbf{s}^{\lambda} = q_{\lambda}(\mathbf{s})$ . Automatic relevance determination (ARD) kernels with input dimension-specific length scales  $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_d)$ ,  $q_{\lambda}(\mathbf{s}) = (s_1/\lambda_1, \dots, s_d/\lambda_d)$ , are examples of such covariance functions. Since distances are computed in a transformed space, we also determine the inducing points using the transformed input locations  $q_{\lambda}(\mathbf{s})$ . Intuitively, if a dimension  $s_l$  has a relatively large length scale  $\lambda_l$  for an ARD kernel, it is less relevant for the covariance function and should thus also have less impact on choosing inducing points. This is what is accomplished when choosing the inducing points using the transformed locations  $q_{\lambda}(\mathbf{s}) = (s_1/\lambda_1, \dots, s_d/\lambda_d)$ . However, since the inducing points depend on  $\boldsymbol{\lambda}$ , the inducing points must be updated dynamically during the parameter optimization. An advantage of

the kMeans++ algorithm in terms of computational efficiency is that it can use the inducing points from a previous optimization iteration as initialization.

For selecting Vecchia neighbors, a common approach is to choose the  $m_v$  nearest input points that appear earlier in a given ordering of the data points. Usually, the Euclidean distance is used as a metric. However, previous research (Kang and Katzfuss, 2023) has shown that a correlation-based selection of neighbors can lead to more accurate approximations. For ARD covariance functions, for example, the intuitive argument why this choice of neighbors is generally beneficial is analogous to the one made above for choosing the inducing points in the transformed space: less relevant input dimensions should have less impact on the choice of neighbors. Often, correlation-based selection of neighbors is equivalent to simply using the Euclidean distance in a transformed space (Kang and Katzfuss, 2023). In our case, though, simply using the Euclidean metric in a transformed space is not applicable as we compute a Vecchia approximation for the residual process with covariance  $\Sigma - \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn}$ , which corresponds to a non-stationary covariance function that cannot be written in the form of (22). In the following, we present a computationally efficient method for the correlation-based selection of Vecchia neighbors for VIF approximations.

We define the following correlation-based distance  $d_c(\cdot, \cdot)$  to determine the nearest neighbors for the residual Vecchia approximation:

$$d_c(\mathbf{s}_i, \mathbf{s}_j) = \sqrt{1 - \left| \frac{c_r(\mathbf{s}_i, \mathbf{s}_j)}{\sqrt{c_r(\mathbf{s}_i, \mathbf{s}_i) \cdot c_r(\mathbf{s}_j, \mathbf{s}_j)}} \right|}, \quad c_r(\mathbf{s}_i, \mathbf{s}_j) = [\Sigma]_{ij} - \Sigma_{mi}^T \Sigma_m^{-1} \Sigma_{mj},$$

where  $\Sigma_{mj} = [c_{\theta}(\mathbf{s}_i^*, \mathbf{s}_j)]_{i=1:m} \in \mathbb{R}^{m \times 1}$ . There are efficient algorithms for finding the  $m_v$  nearest neighbors for the Euclidean distance, such as k-d trees and ball trees, enabling fast searches even in high-dimensional spaces. However, when employing an arbitrary metric, such as the proposed correlation-based metric, nearest neighbor search becomes significantly more challenging. Unlike the Euclidean distance, these specialized metrics often lack geometric properties that facilitate efficient search algorithms, leading to increased computational complexity. To address this, we utilize cover trees (Beygelzimer et al., 2006) which enable  $m_v$ -nearest neighbor search for a set of  $n$  points with a complexity of  $\mathcal{O}(C_d \cdot n \cdot \log(m_v) \cdot (m_v + \log(n)))$ , where  $C_d > 0$  is a hidden dimensionality factor according to Beygelzimer et al. (2006).<sup>2</sup> For computational efficiency in our setting, we modify both the cover tree construction and the nearest neighbor search algorithm as follows. For the tree construction, instead of randomly selecting knots from the remaining data, we systematically insert the point with the smallest index into the cover tree. The full algorithm is given in Algorithm 3. This adaptation simplifies the nearest neighbor search by restricting potential neighbors to those with smaller indices, see Algorithm 4, resulting

2. Elkin and Kurlin (2023) identified potential issues in the original computational complexity analysis of Beygelzimer et al. (2006) and introduced new algorithms for constructing compressed cover trees and performing neighbor searches, which provably achieve the stated time complexity. Despite these refinements, the original and simpler algorithms proposed by Beygelzimer et al. (2006) have been shown to outperform other nearest neighbor search methods (Beygelzimer et al., 2006), and they are expected to achieve the claimed time complexity on well-behaved data sets in practice (Curtin, 2016). We use the original version of the cover tree algorithm throughout this paper, including for our computational complexity analysis. We conjecture that our adaptations could also be applied to the compressed versions, and we leave this as future work.



---

**Algorithm 4** Find  $m_v$  nearest Vecchia neighbors with respect to the metric  $d_c$ 


---

**Require:** Query point  $\mathbf{s}_i \in \mathcal{S}$ , cover tree  $\mathcal{T}$ , number of neighbors  $m_v$ **Ensure:**  $\mathcal{N}_{m_v}$  set of  $m_v$  nearest Vecchia neighbors of  $\mathbf{s}_i$ 

- 1: Initialize: Maximal distance:  $R_{\max} = 1$   
 Number of levels in tree  $\mathcal{T}$ :  $l \leftarrow \text{depth}(\mathcal{T})$   
 Set of potential nearest neighbors:  $\mathcal{Q} \leftarrow \{\mathcal{T}_{1,1}\}$   
 Distance to  $m_v$  closest point in  $\mathcal{Q}$ :  $D_{m_v} \leftarrow 1$  ▷ Set to 1 if  $|\mathcal{Q}| < m_v$
  - 2: **for**  $j = 1$  to  $l$  **do**
  - 3:    $\mathcal{C} \leftarrow \{\mathbf{s}_k \in \text{Children}(\mathbf{s}) \mid \mathbf{s} \in \mathcal{Q} \cap k < i\} \cup \mathcal{Q}$  ▷ Children of  $\mathcal{Q}$  with index  $< i$
  - 4:    $D_{m_v} \leftarrow m_v$ -th smallest distance  $d_c(\mathbf{s}, \mathbf{s}_i)$  for  $\mathbf{s} \in \mathcal{C}$  ▷ Set to 1 if  $|\mathcal{C}| < m_v$
  - 5:    $\mathcal{Q} \leftarrow \{\mathbf{s} \in \mathcal{C} \mid d_c(\mathbf{s}, \mathbf{s}_i) \leq D_{m_v} + 1/2^{j-1}\}$
  - 6: **end for**
  - 7:  $\mathcal{N}_{m_v} \leftarrow$  Find  $m_v$  nearest points to  $\mathbf{s}_i$  in  $\mathcal{Q}$  by brute-force search
- 

rameters  $\boldsymbol{\lambda} \in \mathbb{R}^d$ . See the following subsections for the specific choices of  $\boldsymbol{\lambda}$  and the sample size  $n$ . Sample locations are drawn uniformly from the unit hypercube  $[0, 1]^d$ . We simulate data for both Gaussian and non-Gaussian likelihoods. For the Gaussian likelihood, we use a variance of 0.001 for the error term, and in the non-Gaussian case, the response variable  $\mathbf{y} \in \{0, 1\}^n$  follows a Bernoulli likelihood with a logit link function. Unless stated otherwise, we use  $m = 200$  inducing points and  $m_v = 30$  Vecchia neighbors for VIF, Vecchia, and FITC approximations, but other choices are also considered below. Except as otherwise indicated, Vecchia neighbors for VIF and Vecchia approximations are selected using the correlation-based distance described in Section 6, and inducing points for VIF and FITC approximations are selected in the transformed space as outlined in Section 6. Furthermore, unless specified otherwise, we use the FITC preconditioner with  $k = 200$  inducing points for the iterative methods for VIF approximations, Algorithm 1 (SBPV) with  $\ell = 100$  sample vectors for predictive variances, 50 sample vectors for the SLQ method, and a convergence tolerance of  $\delta = 0.01$  for the CG method. Estimation is done by minimizing the negative log-marginal likelihood using the limited-memory BFGS algorithm. All calculations are performed on an AMD EPYC 7742 processor with 64 CPU cores and 512 GB of RAM, using the GPBoost library version 1.5.8 compiled with GCC 11.2.0. To ensure reproducibility, the code for the experiments is available at <https://github.com/TimGyger/VIF>.

### 7.1 Accuracy for varying input dimensions and smoothness of covariance functions

We begin by comparing the accuracy of VIF approximations with FITC and Vecchia standalone approximations for varying input dimensions and covariance functions with different smoothness parameters. Specifically, we consider input dimensions  $d \in \{2, 5, 10, 20, 50, 100\}$  and the following ARD covariance functions: 1/2-Matérn, 3/2-Matérn, 5/2-Matérn, and Gaussian ( $\infty$ -Matérn). Since these covariance functions have different effective ranges and the average distance among two randomly chosen points changes with the dimension  $d$ , we adapt the data-generating length scale parameters for different input dimensions  $d$  and covariance functions for better comparability. See Table 5 in Appendix 7 for the specific

length scale parameters. Standalone Vecchia approximations are applied to the observable response variable. We have also run the experiments applying a Vecchia approximation to the latent GP and obtained almost equal prediction accuracy measures (results not shown). For each setting, we generate 10 independent data sets consisting of 20,000 samples using a Gaussian likelihood. To assess prediction accuracy, we randomly select 10,000 samples from each data set to serve as test data. Prediction accuracy is measured using the RMSE, the log-score (LS),  $-\frac{1}{n_p} \sum_{i=1}^{n_p} \log \left( \phi \left( \frac{y_i^* - \mu_{\dagger,i}^p}{\sigma_{\dagger,i}^p} \right) \right)$ , and the continuous ranked probability score (CRPS)  $-\frac{1}{n_p} \sum_{i=1}^{n_p} \sigma_{\dagger,i}^p \left( \frac{1}{\sqrt{\pi}} - 2 \cdot \phi \left( \frac{y_i^* - \mu_{\dagger,i}^p}{\sigma_{\dagger,i}^p} \right) - \frac{y_i^* - \mu_{\dagger,i}^p}{\sigma_{\dagger,i}^p} (2 \cdot \Phi \left( \frac{y_i^* - \mu_{\dagger,i}^p}{\sigma_{\dagger,i}^p} \right) - 1) \right)$ , where  $\phi(x) = \frac{1}{\sqrt{2\pi}} \exp \left( -\frac{x^2}{2} \right)$  and  $\Phi(x)$  are the density and cumulative distribution functions, respectively, of a standard normal distribution,  $y^*$  is the test response, and  $\mu_{\dagger,i}^p$  and  $\sigma_{\dagger,i}^p$  are the predictive means and variances, respectively.

Figure 2 compares the VIF, FITC, and Vecchia approximations and an exact GP model without an approximation for varying input dimensions  $d$  for a 3/2-Matérn kernel. We observe that the VIF approximation consistently outperforms both the Vecchia and FITC approximations across all input dimensions. As expected, for low-dimensional inputs, the Vecchia approximation is very accurate, and the FITC approximation is substantially less accurate. However, the accuracy of the Vecchia approximation declines relatively quickly with increasing dimension  $d$ , and the FITC approximation is considerably more accurate for large dimensions. Moreover, for higher dimensions  $d \geq 10$ , none of the approximations achieves the accuracy of an exact GP model, highlighting that modeling high-dimensional functions is challenging. We also repeat these experiments for a different choice of parameters. Specifically, in Figure 14 in Appendix 7, we present the results for a 3/2-Matérn kernel using a larger error variance of  $\sigma^2 = 0.01$ , and the length scale parameters are chosen such that the covariance remains approximately equal at the average distance for all  $d$  (to the one of a Gaussian kernel with length scales  $\boldsymbol{\lambda} = (0.35, 0.4, 0.45, 0.5, 0.55)^T$ ). The results are similar to those shown in Figure 2.

In Table 4 in Appendix 7, we additionally report the corresponding average runtimes for training and prediction for each method. As expected, the training runtimes for ARD kernels increase with the input dimension  $d$  across all methods as the corresponding optimization problems become more high-dimensional. Overall, VIF and FITC approximations have similar runtimes for estimation, and Vecchia approximations are faster. All three methods have similar prediction runtimes. As mentioned above, we use  $m = 200$  inducing points and  $m_v = 30$  Vecchia neighbors for VIF, Vecchia, and FITC approximations. In Figure 10 in Appendix 7, we additionally show results obtained by increasing the number of Vecchia neighbors to  $m_v = 60$  for the Vecchia approximation to ensure that the computational time for estimation is approximately equal for all three methods. Although the Vecchia approximation slightly benefits from the increased approximation complexity, the results are qualitatively very similar. Moreover, Figures 11 and 12 in Appendix 7 compare the prediction accuracy versus the total training and prediction runtime for the three approximations across varying numbers of inducing points and Vecchia neighbors, for input dimensions  $d = 10$  and  $d = 100$ . For  $d = 100$ , the VIF approximation achieves the best accuracy-runtime trade-off, while for  $d = 10$ , Vecchia and VIF approximations exhibit similar accuracy-runtime trade-offs. Furthermore, we also compare the VIF approximation for two inducing-point-to-neighbor ratios  $m/m_v$  in Figures 11 and 12 in Appendix 7. In line




Figure 2: Comparison of the VIF, FITC, and Vecchia approximations and an exact GP model (without an approximation) for varying dimensions  $d$  using an ARD 3/2-Matérn kernel. The plots show the RMSE (log-scale), log-score (LS), and CRPS (log-scale) (mean  $\pm$  2 standard errors).

with the above findings, these results indicate that in high dimensions, it is beneficial for the VIF approximations to rather increase the number of inducing points than the number of Vecchia neighbors.

The results when comparing VIF, FITC, and Vecchia approximations and an exact GP model without an approximation for the different covariance functions with varying smoothness and  $d = 10$  are shown in Figure 3. The VIF approximation is again consistently more accurate compared to both FITC and Vecchia approximations. All approximations are more accurate for smoother covariance functions. However, the relative difference in accuracy between the Vecchia and both the VIF and FITC approximations increases with increasing smoothness levels. This is in line with the prior expectation that Vecchia approximations are more accurate for moderately smooth covariance functions, whereas FITC approximations become more accurate for smoother covariance functions. In Figure 13 in Appendix 7, we additionally report the results for varying smoothness when  $d = 2$ , i.e., for spatial data. As expected, the VIF approximation outperforms the Vecchia approximation only marginally in this low-dimensional setting. For the Gaussian kernel, all approximations are almost equally accurate.

## 7.2 Comparison of preconditioners

For all subsequent experiments and unless stated otherwise, we generate 100,000 samples from a zero-mean Gaussian process with five-dimensional inputs and an ARD Gaussian covariance function with length scale parameters  $\boldsymbol{\lambda} = (0.15, 0.30, 0.45, 0.60, 0.75)$ .

In the following, we analyze the accuracy and runtimes of the iterative methods for non-Gaussian likelihoods introduced in Section 4. First, we compare the VIFDU and FITC preconditioners introduced in Section 4.3 with respect to the runtime and accuracy of log-likelihood approximations for a binary likelihood. In Figure 4, we report the wall-clock times and the RMSE between log-marginal likelihoods computed using iterative methods




Figure 3: Comparison of the VIF, FITC, and Vecchia approximations and an exact GP model (without an approximation) for varying smoothness parameters of an ARD Matérn kernel for  $d = 10$  (1/2-Matérn, 3/2-Matérn, 5/2-Matérn, and  $\infty$ -Matérn = Gaussian kernel). The plots show the RMSE (log-scale), log-score (LS), and CRPS (log-scale) (mean  $\pm$  2 standard errors).

and those computed using a Cholesky decomposition for three different VIF approximations with different numbers of inducing points  $m$  and Vecchia neighbors  $m_v$ . The figure also shows the runtime for the calculations based on a Cholesky decomposition. The marginal likelihood is evaluated at the data-generating parameters and repeated 100 times. Overall, both preconditioners lead to very accurate approximations. For instance, a log-marginal likelihood difference in 10 corresponds to a relative error of approximately  $10^{-4}$ . However, we find that the FITC preconditioner consistently outperforms the VIFDU preconditioner, having both faster runtimes and more accurate log-marginal likelihood approximations. Moreover, the iterative methods are substantially faster than traditional Cholesky-based computations with a speed-up of approximately three orders of magnitude for both preconditioners. Note that both preconditioners yield unbiased stochastic log-likelihood approximations (results not shown).

The accuracy and runtimes of the iterative methods depend on the CG convergence tolerance  $\delta$  and the number of sample vectors  $\ell$  for the SLQ method. Table 6 in Appendix 7 reports the accuracy and runtimes for different tolerances  $\delta \in \{0.0001, 0.001, 0.01, 0.1, 1\}$  and numbers of sample vectors  $\ell \in \{10, 50, 100\}$ . The results show that decreasing  $\delta$  below 0.01 yields negligible improvements in accuracy while increasing computational cost, indicating that  $\delta = 0.01$  provides a good balance between accuracy and efficiency. In contrast, the number of sample vectors  $\ell$  has a more pronounced impact on the accuracy of SLQ-approximated log-determinants than the CG tolerance.

The FITC preconditioner has one tuning parameter, the number of inducing points, which governs a trade-off between computational efficiency and the accuracy of the SLQ approximation. To analyze this parameter, we compute the log-marginal likelihood using varying numbers of inducing points,  $k \in \{10, 50, 100, 200, 300, 400, 500\}$ . In Figure 15 in Appendix 7, we report the runtime and the log-marginal likelihood compared to exact Cholesky-based calculations for three different VIF approximations. Our analysis shows




Figure 4: Accuracy-runtime comparison of preconditioners: RMSE between log-marginal likelihoods computed using iterative methods and those computed using a Cholesky decomposition versus runtime for the VIFDU and FITC preconditioners and varying numbers of sample vectors  $\ell$  (annotated in the plot). A binary likelihood,  $n = 100,000$  samples, and three different VIF approximations are used.

that the FITC preconditioner overall yields the fastest runtime for  $k = 200$ . Note that the experiments in this and the next section are conducted on only one simulated data set, since we do not want to mix sampling variability and randomness of the methods that are analyzed. However, the results do not change when using other samples (results not shown).

### 7.3 Predictive variances

Next, we compare the two simulation- and iterative-methods-based approaches introduced in Section 4.2 for calculating predictive variances, namely Algorithm 1 (SBPV) and Algorithm 2 (SPV). Both algorithms are run twice with the VIFDU and the FITC preconditioners. We use simulated data with  $n = n_p = 100,000$  training and test points, a binary likelihood for the response variable, and predictive distributions are calculated using the true data-generating covariance parameters. To measure the accuracy of the methods, we compute the RMSE of the simulation-based predictive variances relative to the “exact” Cholesky-based results. In addition, we measure the total runtime, which includes computing the mode, the latent predictive means, and the latent predictive variances. Figure 5 shows the RMSE as a function of the wall-clock time for varying numbers of random vectors  $\ell$ . We find that all algorithms yield very accurate predictive variances already for small runtimes. Furthermore, the results show that the SBPV algorithm is more accurate than the SPV algorithm, and that the FITC preconditioner consistently leads to lower runtimes compared to the VIFDU preconditioner.

### 7.4 Runtime analysis of VIF approximations

In the following, we analyze how the runtime of a VIF approximation scales with the sample size, and how it depends on the two tuning parameters of a VIF approximation: the number

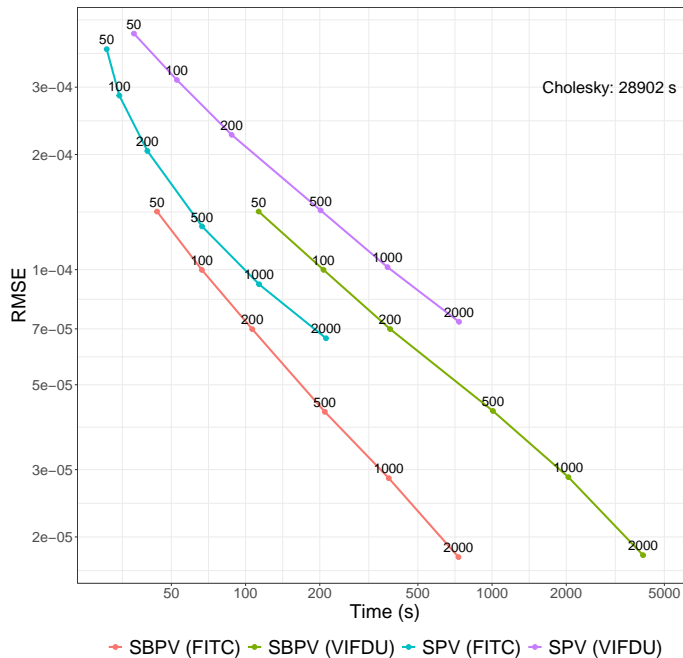


Figure 5: Accuracy-runtime comparison of simulation- and iterative-methods-based predictive variances (SBPV and SPV algorithms with the FITC and VIFDU preconditioners): RMSE between predictive variances computed using simulation-based methods and those computed using a Cholesky decomposition versus runtime for a binary likelihood. The number of random vectors  $\ell$  is annotated in the plot.

of inducing points  $m$  and the number of Vecchia neighbors  $m_v$ . Specifically, we vary each of the three quantities  $m$ ,  $m_v$ , and  $n$  in turn while keeping the others fixed at  $m = 200$ ,  $m_v = 30$ , and  $n = 100,000$ . We consider both a Gaussian and a Bernoulli likelihood, and we additionally compare VIF approximations to the FITC and Vecchia approximations. For the Bernoulli likelihood, we apply the VIF approximations using our proposed iterative methods with the FITC and VIFDU preconditioners. Similarly, the standalone Vecchia approximation uses iterative methods with the Vecchia approximation with diagonal update (VADU) preconditioner (Kündig and Sigrist, 2025). Figure 6 illustrates the runtime for evaluating the log-marginal likelihood for Gaussian and Bernoulli likelihoods for different approximation parameter values  $m \in \{10, 20, 50, 100, 200, 500\}$ ,  $m_v \in \{5, 10, 15, 20, 30, 50\}$ , and sample sizes  $n \in \{10,000, 20,000, 30,000, 50,000, 80,000, 100,000\}$ . As expected, the runtime increases with the sample size  $n$ , the number of inducing points  $m$ , and the number of Vecchia neighbors  $m_v$  for both likelihoods. For the non-Gaussian likelihood, we find that the FITC preconditioner consistently outperforms the VIFDU preconditioner across all parameter settings. Notably, the VIF approximation with the FITC preconditioner has similar runtimes as the Vecchia approximation.

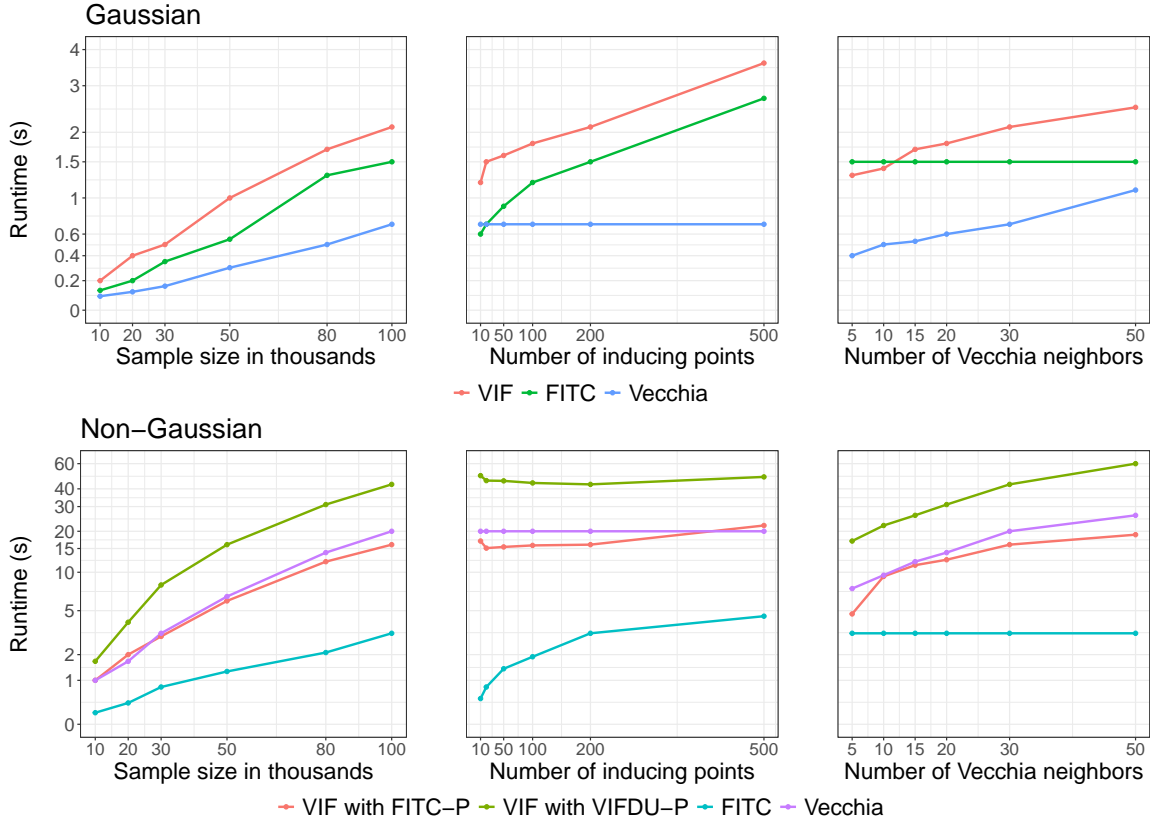


Figure 6: Time (s) for computing the marginal likelihood with VIF, FITC, and Vecchia approximations for varying sample sizes  $n$ , numbers of inducing points  $m$ , and numbers of Vecchia neighbors  $m_v$  for both a Gaussian (top row) and a non-Gaussian likelihood (bottom row).

For the experiments presented in Figure 6, we exclude the time required to determine the  $m_v$  Vecchia neighbors, as these are not necessarily recomputed for each log-likelihood evaluation in the optimization process. Figure 7 illustrates the runtime for constructing the cover tree and identifying the  $m_v$  nearest neighbors with respect to the correlation distance for varying sample sizes, feature dimensions  $d \in \{2, 5, 10, 20, 50, 100\}$ , numbers of inducing points, and number of Vecchia neighbors. We again vary each of these quantities in turn while keeping the others fixed at  $m = 200$ ,  $m_v = 30$ ,  $n = 100,000$ , and  $d = 5$ . These results show that the runtime primarily depends on the sample size  $n$  and the input dimension  $d$ . The runtime scales approximately linearly with the number of inducing points, which is expected, as computing the correlation distance requires  $\mathcal{O}(m)$  operations. The number of Vecchia neighbors  $m_v$  has only a minor impact on the overall runtime.

In Figure 16 in Appendix 7, we additionally analyze the runtime required to compute the predictive means and variances for Gaussian and Bernoulli likelihoods, respectively, for different approximation parameters, sample sizes, number of prediction points, and comparing VIF, FITC, and Vecchia approximations. For the VIF approximation and the non-Gaussian




Figure 7: Time (s) for constructing the cover tree and finding the  $m_v$  nearest neighbors with respect to the correlation distance for varying sample sizes  $n$ , input dimensions  $d$ , numbers of inducing points  $m$ , and numbers of Vecchia neighbors  $m_v$ .

likelihood, iterative methods are employed with the FITC and VIFDU preconditioners. As the results show, the FITC preconditioner consistently outperforms the VIFDU preconditioner across all parameter settings and even surpasses the Vecchia approximation with the VADU preconditioner in terms of runtime.

## 8. Real-world data experiments

In the following, we apply our proposed methods to multiple real-world data sets comparing the VIF approximation to the following state-of-the-art GP approximations: sparse Gaussian process regression (SGPR) (Titsias, 2009), stochastic variational Gaussian processes (SVGP) (Hensman et al., 2013, 2015), scalable kernel interpolation for product kernels (SKIP) (Gardner et al., 2018b), and double-Kullback-Leibler-optimal Gaussian process approximation (DKLGP) (Cao et al., 2023). SGPR and SVGP are both variational low-rank inducing-point approximations which differ in the way the evidence lower bound (ELBO) is maximized. SKIP, an extension of structured kernel interpolation (SKI) (Wilson and Nickisch, 2015), is an inducing points-based approximation that relies on local kernel interpolation, the fast Fourier transform (FFT), and iterative methods. DKLGP uses a sparse inverse Cholesky factor approximation whose parameters are trained in a variational framework. In other words, DKLGP is a variational form of a Vecchia approximation.

We use GPyTorch (Gardner et al., 2018a) version 1.13, for SKIP, SGPR, and SVGP. For DKLGP, we use the implementation of Cao et al. (2023) available in the repository <https://github.com/katzfuss-group/DKL-GP>. Furthermore, for SKIP, SGPR, SVGP, and DKLGP, training is done using the ADAM optimizer, following the recommendations of Gardner et al. (2018a) and Cao et al. (2023). For SVGP, the inducing points are optimized jointly with the variational parameters, and for SGPR, the inducing points are chosen by random subsampling from the training data. Unless otherwise specified, a 3/2-Matérn ARD kernel is used. For the non-Gaussian data sets, we employ the VIF approximation with iterative methods and the same settings as in the simulated experiments; see the beginning of Section 7. We use 5-fold cross-validation to assess the prediction accuracy. All data sets and code to reproduce our experiments are available on <https://github.com/TimGyger/VIF>.

## 8.1 Data sets

We consider a variety of data sets from the UCI and OpenML repositories, which are widely used for benchmarking GP approximations in the machine learning literature, as well as additional spatial data sets. Appropriate likelihoods are selected based on the characteristics of each data set; see below for the specific choices. We follow Cao et al. (2023) and apply the following pre-processing steps. First, the input features are transformed to the interval  $[0, 1]$  based on the training data. Second, input features with a standard deviation below 0.01 after standardization are excluded from the analysis to prevent numerical issues for some of the methods. For the same reason, input features with a standard deviation below 0.1 are additionally excluded for the binary classification data sets. Third, the data sets are filtered to ensure a minimum Euclidean distance  $\|\mathbf{s}_i - \mathbf{s}_j\|$  of 0.001 between any two data points, retaining only the first point in cases where multiple points are too close, to prevent numerical singularities for some methods. With the exception of the 3dRoad data set, this procedure has a negligible impact on the sample size. For the data sets modeled using a Gaussian likelihood, the response variable is normalized based on the training data sets to have zero mean and unit variance for better comparability across different data sets.

## 8.2 Results

We first present the results for the Gaussian likelihood in Table 1. Entries marked as NA indicate numerical problems (SKIP did not converge and yields predictions that are very inaccurate). Additionally, we illustrate the log-score (LS) results in Figure 8 (left plot). We find that the VIF approximation consistently outperforms all other methods in terms of all accuracy measures across all data sets. Furthermore, for data sets with low input dimensions (‘3dRoad’, ‘Protein’, ‘Kin40K’), we observe that the inducing point methods SKIP, SGPR, and SVGP are very inaccurate, and the variational Vecchia approximation DKLGP is more accurate compared to these inducing point-based methods. The situation is reversed for the higher-dimensional data sets (‘KEGGU’, ‘KEGG’, ‘Elevators’, ‘Ailerons’), where the best inducing-points approximation, SVGP, is more accurate than the variational Vecchia approximation DKLGP. These findings are in line with the results from the simulated experiments in Section 7.1, where we found that Vecchia approximations are accurate for low-dimensional inputs but less accurate for higher dimensions. The VIF approximation achieves the highest prediction accuracy for all data sets with both low and higher-dimensional inputs as it combines low-rank inducing points and residual process Vecchia approximations. Concerning computational time, we find that the VIF approximation is, overall, faster than all other approximations.

Next, we present the results for the binary data sets in Table 2 reporting the area under the receiver operating characteristic curve (AUC), the root mean squared error (RMSE), which in this case corresponds to the square root of the Brier score, the accuracy (ACC), and the log-score (LS). Figure 8 (right plot) additionally visualizes the log-score (LS) for all data sets modeled using non-Gaussian likelihoods. The differences in prediction accuracy among all approximations are very small for the binary classification data sets. The likely explanation for this is that binary response variable data provide less information compared to continuous and count data response variables. Furthermore, the VIF approximation has

| Data   | Accuracy measure | VIF<br>$m_v = 30$<br>$m = 200$ | SKIP<br>$m = 1000$ | SGPR<br>$m = 1000$ | SVGP<br>$m = 1000$   | DKLGP<br>$\rho = 1.5$ |
|--|------------------|--------------------------------|--------------------|--------------------|----------------------|-----------------------|
| <b>3dRoad</b><br>$n = 434,874$<br>$d = 3$    | RMSE             | <b>0.145 ± 0.002</b>           | 0.603 ± 0.010      | 0.734 ± 0.014      | 0.550 ± 0.008        | 0.255 ± 0.004         |
|  | CRPS             | <b>0.068 ± 0.002</b>           | 0.294 ± 0.014      | 0.402 ± 0.008      | 0.297 ± 0.004        | 0.121 ± 0.002         |
|  | LS               | <b>-0.625 ± 0.010</b>          | 2.371 ± 0.022      | 1.114 ± 0.020      | 0.825 ± 0.006        | 0.009 ± 0.004         |
|  | Time             | 376 s                          | 1861 s             | 759 s              | 966 s                | 890 s                 |
| <b>KEGGU</b><br>$n = 63,608$<br>$d = 26$     | RMSE             | <b>0.094 ± 0.008</b>           | NA                 | 0.117 ± 0.024      | <b>0.094 ± 0.004</b> | 0.146 ± 0.014         |
|  | CRPS             | <b>0.030 ± 0.002</b>           | convergence issues | 0.068 ± 0.008      | 0.034 ± 0.002        | 0.054 ± 0.002         |
|  | LS               | <b>-1.081 ± 0.072</b>          |                    | -0.441 ± 0.081     | -0.955 ± 0.042       | -0.845 ± 0.038        |
|  | Time             | 738 s                          |                    | 502 s              | 617 s                | 1362 s                |
| <b>KEGG</b><br>$n = 48,827$<br>$d = 18$      | RMSE             | <b>0.100 ± 0.010</b>           | 0.672 ± 0.166      | 0.699 ± 0.170      | <b>0.103 ± 0.004</b> | 0.142 ± 0.008         |
|  | CRPS             | <b>0.041 ± 0.004</b>           | 0.407 ± 0.086      | 0.450 ± 0.088      | 0.048 ± 0.002        | 0.067 ± 0.002         |
|  | LS               | <b>-1.082 ± 0.070</b>          | 1.158 ± 0.134      | 1.194 ± 0.144      | -0.893 ± 0.024       | -0.593 ± 0.040        |
|  | Time             | 383 s                          | 1202 s             | 863 s              | 951 s                | 1858 s                |
| <b>Elevators</b><br>$n = 16,599$<br>$d = 17$ | RMSE             | <b>0.355 ± 0.004</b>           | NA                 | 0.692 ± 0.224      | 0.373 ± 0.008        | 0.418 ± 0.006         |
|  | CRPS             | <b>0.196 ± 0.004</b>           | convergence issues | 0.458 ± 0.202      | 0.205 ± 0.004        | 0.230 ± 0.004         |
|  | LS               | <b>0.388 ± 0.012</b>           |                    | 2.045 ± 1.534      | 0.435 ± 0.020        | 0.907 ± 0.054         |
|  | Time             | 711 s                          |                    | 520 s              | 1195 s               | 2583 s                |
| <b>Protein</b><br>$n = 45,730$<br>$d = 8$    | RMSE             | <b>0.516 ± 0.006</b>           | 0.831 ± 0.008      | 0.822 ± 0.048      | 0.729 ± 0.006        | 0.621 ± 0.006         |
|  | CRPS             | <b>0.259 ± 0.002</b>           | 0.610 ± 0.008      | 0.467 ± 0.028      | 0.411 ± 0.004        | 0.327 ± 0.006         |
|  | LS               | <b>0.667 ± 0.016</b>           | 6.310 ± 0.216      | 1.234 ± 0.058      | 1.105 ± 0.008        | 0.900 ± 0.052         |
|  | Time             | 547 s                          | 1807 s             | 676 s              | 1309 s               | 1531 s                |
| <b>Kin40K</b><br>$n = 40,000$<br>$d = 8$     | RMSE             | <b>0.114 ± 0.002</b>           | 1.401 ± 0.024      | 0.321 ± 0.008      | 0.175 ± 0.004        | 0.284 ± 0.002         |
|  | CRPS             | <b>0.059 ± 0.002</b>           | 0.862 ± 0.020      | 0.271 ± 0.010      | 0.099 ± 0.002        | 0.151 ± 0.002         |
|  | LS               | <b>-0.808 ± 0.012</b>          | 3.192 ± 0.034      | 0.131 ± 0.008      | -0.230 ± 0.012       | 0.082 ± 0.008         |
|  | Time             | 333 s                          | 1328 s             | 947 s              | 1179 s               | 2732 s                |
| <b>Ailerons</b><br>$n = 13,750$<br>$d = 33$  | RMSE             | <b>0.399 ± 0.019</b>           | NA                 | 0.428 ± 0.062      | <b>0.386 ± 0.015</b> | 0.424 ± 0.028         |
|  | CRPS             | <b>0.208 ± 0.008</b>           | convergence issues | 0.238 ± 0.043      | <b>0.208 ± 0.003</b> | 0.230 ± 0.009         |
|  | LS               | <b>0.436 ± 0.030</b>           |                    | 0.676 ± 0.251      | <b>0.459 ± 0.017</b> | 0.642 ± 0.068         |
|  | Time             | 528 s                          |                    | 209 s              | 510 s                | 1918 s                |

Table 1: RMSE, CRPS, log-score (LS) (mean ± 2 standard errors), and runtime in seconds (s) for the regression data sets modeled using a Gaussian likelihood. The score of the best-performing method is shown in bold. If the mean score of another method lies within two standard errors of the best mean, it is also in bold.

a faster runtime than the most accurate inducing points method (SVGP) and the variational Vecchia approximation DKLGP.

In Table 3, we present the results for data sets modeled using Poisson, Student-t, and Gamma likelihoods. Specifically, the House data set is modeled with a Student-t likelihood, the Bike data set with a Poisson likelihood, and the Power and WaterVapor (WV) data sets with a Gamma likelihood. The current implementation of DKLGP by Cao et al. (2023) available on <https://github.com/katzfuss-group/DKL-GP/> does not support Poisson or Gamma likelihoods. The NA entry indicates that SGPR crashed. We find that the VIF




Figure 8: Log-score (LS) with error bars (mean  $\pm$  2 standard errors) for the data sets modeled with a Gaussian (left plot) and non-Gaussian likelihoods (right plot).

approximation consistently outperforms all other methods in prediction accuracy across all data sets. The strong performance of the VIF-Laplace approximation compared to the variational benchmark methods is likely due to its more accurate covariance approximation. Concerning computational time, the VIF approximation has a similar runtime as SVGP for two data sets (‘Power’, ‘WaterVapor’), but it is slower on the two other data sets (‘Bike’, ‘House’).

In Tables 8 and 9 in Appendix 7, we additionally report results comparing the VIF approximation to classical Vecchia and FITC approximations. Overall, the VIF approximation is more accurate than both Vecchia and FITC approximations. The Vecchia approximation also yields competitive prediction accuracy, often approaching that of the VIF approximation. However, the VIF approximation achieves higher accuracy than the Vecchia approximations on multiple data sets such as the ‘Kin40K’, ‘Elevators’, and ‘Ailerons’ data sets.

### 8.3 Matérn kernel smoothness selection and non-zero mean functions

Assuming a fixed value for the smoothness parameter  $\nu$  of the Matérn covariance function and a zero prior mean function  $F(\cdot) = 0$  might result in misspecified models. In the following, we analyze these two issues. We first extend our comparison of VIF approximations for the regression tasks with Gaussian and non-Gaussian likelihoods from the previous section by additionally estimating the shape (or smoothness) parameter  $\nu$  of the Matérn ARD kernel, rather than assuming a fixed  $\nu = 3/2$  Matérn ARD kernel. This approach enables model selection within the Matérn kernel family. Note that estimating the smoothness parameter is possible for VIF approximations implemented in the `GPBoost` library but not for the other approximations implemented in `GPYtorch`. The results are shown in Figure 9 (left plot) for the log-score and in more detail in Table 7 in Appendix 7. We find that estimating  $\nu$  improves the prediction accuracy across nearly all data sets, but the differences

| Data                                      | Accuracy measure | VIF<br>$m_\nu = 30$<br>$m = 200$ | SGPR<br>$m = 1000$    | SVGP<br>$m = 1000$    | DKLGP<br>$\rho = 1.5$ |
|---|------------------|----------------------------------|-----------------------|-----------------------|-----------------------|
| <b>Bank</b><br>$n = 45,211$<br>$d = 16$   | AUC              | <b>0.806 ± 0.008</b>             | <b>0.800 ± 0.006</b>  | <b>0.804 ± 0.008</b>  | 0.789 ± 0.006         |
|   | RMSE             | <b>0.284 ± 0.004</b>             | <b>0.286 ± 0.004</b>  | <b>0.284 ± 0.002</b>  | 0.292 ± 0.004         |
|   | ACC              | <b>0.895 ± 0.002</b>             | <b>0.896 ± 0.002</b>  | <b>0.896 ± 0.002</b>  | 0.892 ± 0.004         |
|   | LS               | <b>0.035 ± 0.016</b>             | <b>0.048 ± 0.012</b>  | <b>0.040 ± 0.014</b>  | 0.754 ± 0.094         |
|   | Time             | 1643 s                           | 1270 s                | 2603 s                | 5012 s                |
| <b>Adult</b><br>$n = 48,842$<br>$d = 14$  | AUC              | <b>0.880 ± 0.008</b>             | <b>0.887 ± 0.007</b>  | <b>0.883 ± 0.006</b>  | 0.869 ± 0.004         |
|   | RMSE             | <b>0.336 ± 0.006</b>             | <b>0.332 ± 0.004</b>  | <b>0.335 ± 0.004</b>  | 0.342 ± 0.004         |
|   | ACC              | <b>0.838 ± 0.006</b>             | <b>0.840 ± 0.007</b>  | <b>0.835 ± 0.006</b>  | 0.833 ± 0.004         |
|   | LS               | <b>0.003 ± 0.012</b>             | <b>-0.023 ± 0.104</b> | <b>-0.040 ± 0.062</b> | 0.844 ± 0.048         |
|   | Time             | 2854 s                           | 1561 s                | 5661 s                | 2815 s                |
| <b>Credit</b><br>$n = 30,000$<br>$d = 22$ | AUC              | 0.768 ± 0.006                    | 0.767 ± 0.006         | <b>0.773 ± 0.004</b>  | 0.752 ± 0.004         |
|   | RMSE             | <b>0.378 ± 0.004</b>             | <b>0.379 ± 0.004</b>  | <b>0.377 ± 0.004</b>  | 0.385 ± 0.004         |
|   | ACC              | <b>0.808 ± 0.006</b>             | <b>0.808 ± 0.006</b>  | <b>0.809 ± 0.006</b>  | 0.799 ± 0.006         |
|   | LS               | <b>0.407 ± 0.012</b>             | <b>0.413 ± 0.012</b>  | <b>0.402 ± 0.012</b>  | 0.996 ± 0.070         |
|   | Time             | 1161 s                           | 1235 s                | 2786 s                | 2816 s                |
| <b>MAGIC</b><br>$n = 19,020$<br>$d = 9$   | AUC              | <b>0.920 ± 0.004</b>             | 0.902 ± 0.006         | 0.915 ± 0.004         | 0.900 ± 0.004         |
|   | RMSE             | <b>0.316 ± 0.004</b>             | 0.334 ± 0.004         | 0.321 ± 0.004         | 0.340 ± 0.004         |
|   | ACC              | <b>0.866 ± 0.006</b>             | 0.850 ± 0.006         | <b>0.860 ± 0.006</b>  | 0.842 ± 0.004         |
|   | LS               | <b>0.006 ± 0.022</b>             | 0.144 ± 0.072         | 0.166 ± 0.352         | 0.374 ± 0.030         |
|   | Time             | 236 s                            | 437 s                 | 663 s                 | 459 s                 |

Table 2: AUC, RMSE (Brier score), ACC, LS (mean ± 2 standard errors), and runtime in seconds (s) for the binary classification data sets. Bold indicates the best mean; other means are in bold if within two standard errors.

are relatively small for most data sets except for the ‘3dRoad’, ‘Kin40K’, ‘House’, ‘Power’, and ‘Ailerons’ data sets, for which estimated  $\nu$ ’s yield substantially higher accuracy. In addition, estimating the smoothness parameter increases the runtime due to the additional evaluations of Bessel functions.

Next, we extend the GP model (1) by allowing for non-zero prior mean (fixed effects) functions  $F(\cdot)$ . Specifically, we consider a linear regression function  $F(\mathbf{x}) = \mathbf{x}^T \boldsymbol{\beta}$  as well as a function that is modeled using tree-boosting (Sigrist, 2022a,b). The latter is denoted as ‘GPBoost model’ in the following. In both cases, we use the GP input variables  $\mathbf{s}$  as covariates  $\mathbf{x}$  for the mean function, i.e.,  $\mathbf{s} = \mathbf{x}$ . In addition, we use a VIF approximation and a 3/2-Matérn kernel. The results are shown in Figure 9 (right plot) for the log-score and in more detail in Table 10 in Appendix 7. We find that while the GPBoost model yields slightly more accurate predictions in terms of the log-score for some data sets (‘Elevators’, ‘Protein’, ‘Kin40K’), all three models have overall similar prediction accuracy measures except for the linear model on the ‘KEGGU’ data set.

| Data   | Accuracy measure           | VIF<br>$m_v = 30$<br>$m = 200$  | SGPR<br>$m = 1000$  | SVGP<br>$m = 1000$   | DKLGP<br>$\rho = 1.5$  |
|--|----------------------------|---|---|--|--|
| <b>Bike</b><br>$n = 17,379$<br>$d = 12$<br>(Poisson)     | RMSE<br>CRPS<br>LS<br>Time | <b><math>38.904 \pm 1.332</math></b><br><b><math>17.162 \pm 0.454</math></b><br><b><math>4.676 \pm 0.030</math></b><br>1534 s | NA<br>crashed   | $42.256 \pm 0.850$<br>$20.746 \pm 0.194$<br>$7.130 \pm 0.208$<br>817 s                 | not implemented  |
| <b>House</b><br>$n = 20,640$<br>$d = 8$<br>(Student-t)   | RMSE<br>CRPS<br>LS<br>Time | <b><math>0.214 \pm 0.008</math></b><br><b><math>0.103 \pm 0.002</math></b><br>$0.092 \pm 0.076$<br>1932 s                     | $0.271 \pm 0.004$<br>$0.259 \pm 0.002$<br>$0.942 \pm 0.002$<br>349 s  | $0.258 \pm 0.002$<br>$0.137 \pm 0.002$<br><b><math>0.067 \pm 0.012</math></b><br>568 s | $0.267 \pm 0.004$<br>$0.144 \pm 0.002$<br>$0.523 \pm 0.058$<br>525 s |
| <b>Power</b><br>$n = 52,417$<br>$d = 5$<br>(Gamma)       | RMSE<br>CRPS<br>LS<br>Time | <b><math>0.218 \pm 0.002</math></b><br><b><math>0.121 \pm 0.001</math></b><br><b><math>-0.084 \pm 0.010</math></b><br>5035 s  | $0.826 \pm 0.456$<br>$0.837 \pm 0.090$<br>$2.142 \pm 0.088$<br>1597 s | $0.567 \pm 0.030$<br>$0.787 \pm 0.006$<br>$2.096 \pm 0.004$<br>4063 s                  | not implemented  |
| <b>WaterVapor</b><br>$n = 100,000$<br>$d = 2$<br>(Gamma) | RMSE<br>CRPS<br>LS<br>Time | <b><math>0.102 \pm 0.001</math></b><br><b><math>0.056 \pm 0.001</math></b><br><b><math>-0.728 \pm 0.022</math></b><br>3975 s  | $0.423 \pm 0.002$<br>$0.293 \pm 0.001$<br>$0.941 \pm 0.001$<br>2027 s | $0.424 \pm 0.004$<br>$0.294 \pm 0.002$<br>$0.943 \pm 0.004$<br>4747 s                  | not implemented  |

Table 3: RMSE, CRPS, log-score (LS) (mean  $\pm$  2 standard errors), and runtime in seconds (s) for the regression data sets modeled using non-Gaussian likelihoods. Bold indicates the best mean; other means are in bold if within two standard errors.

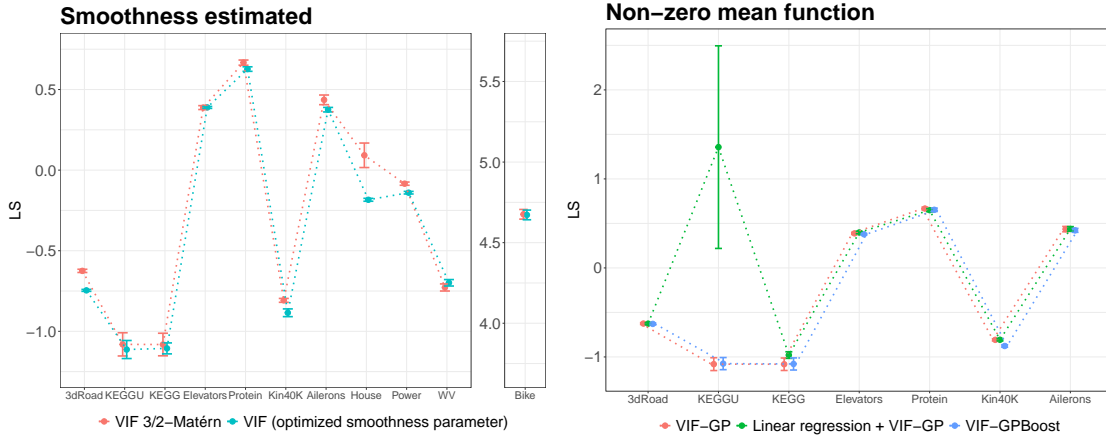


Figure 9: Log-score (LS) with error bars (mean  $\pm$  2 standard errors) when estimating the smoothness parameter for the regression data sets (left plot) and when using non-zero prior mean functions (right plot).

## 9. Conclusion

In this work, we introduce Vecchia-inducing-points full-scale (VIF) approximations, a novel approach that leverages the complementary strengths of global inducing points and local Vecchia approximations to enhance the scalability and accuracy of Gaussian processes. We develop an efficient correlation-based neighbor selection strategy for the residual process using a modified cover tree algorithm, and we introduce iterative methods and preconditioners for VIF-Laplace approximation for non-Gaussian likelihoods, significantly reducing computational costs by several orders of magnitude compared to Cholesky-based computations. Our theoretical analysis provides insights into the convergence properties of the preconditioned conjugate gradient method. We systematically evaluate our methods concerning both runtime and accuracy through extensive numerical experiments on simulated and real-world data sets. We find that VIF approximations consistently outperform state-of-the-art alternatives in both accuracy and computational efficiency across various real-world data sets. Future research can analyze alternative methods for choosing both inducing points and Vecchia neighbors in VIF approximations, e.g., using a variational framework.

## Acknowledgments

This research was partially supported by the Swiss Innovation Agency - Innosuisse (grant number ‘57667.1 IP-ICT’).

## Appendix

### Appendix A. Gradients of the residual process Vecchia approximation

The gradients of  $\mathbf{B}$  and  $\mathbf{D}$  defined in Section 2.1 with respect to the covariance parameters  $\boldsymbol{\theta}$  are given by

$$\begin{aligned}
 \left(\frac{\partial \mathbf{B}}{\partial \boldsymbol{\theta}}\right)_{iN(i)} &= -\left(\frac{\partial \boldsymbol{\Sigma}_{iN(i)}}{\partial \boldsymbol{\theta}} - \boldsymbol{\Sigma}_{mi}^T \boldsymbol{\Sigma}_m^{-1} \frac{\partial \boldsymbol{\Sigma}_{mN(i)}}{\partial \boldsymbol{\theta}} - \frac{\partial \boldsymbol{\Sigma}_{mi}^T}{\partial \boldsymbol{\theta}} \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mN(i)}\right. \\
 &\quad + \boldsymbol{\Sigma}_{mi}^T \boldsymbol{\Sigma}_m^{-1} \frac{\partial \boldsymbol{\Sigma}_m}{\partial \boldsymbol{\theta}} \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mN(i)} \left.\right) (\boldsymbol{\Sigma}_{N(i)} - \boldsymbol{\Sigma}_{mN(i)}^T \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mN(i)})^{-1} \\
 &\quad + (\boldsymbol{\Sigma}_{iN(i)} - \boldsymbol{\Sigma}_{mi}^T \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mN(i)}) (\boldsymbol{\Sigma}_{N(i)} - \boldsymbol{\Sigma}_{mN(i)}^T \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mN(i)})^{-1} \\
 &\quad \cdot \left(\frac{\partial \boldsymbol{\Sigma}_{N(i)}}{\partial \boldsymbol{\theta}} - \boldsymbol{\Sigma}_{mN(i)}^T \boldsymbol{\Sigma}_m^{-1} \frac{\partial \boldsymbol{\Sigma}_{mN(i)}}{\partial \boldsymbol{\theta}} - \frac{\partial \boldsymbol{\Sigma}_{mN(i)}^T}{\partial \boldsymbol{\theta}} \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mN(i)}\right. \\
 &\quad \left.+ \boldsymbol{\Sigma}_{mN(i)}^T \boldsymbol{\Sigma}_m^{-1} \frac{\partial \boldsymbol{\Sigma}_m}{\partial \boldsymbol{\theta}} \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mN(i)}\right) (\boldsymbol{\Sigma}_{N(i)} - \boldsymbol{\Sigma}_{mN(i)}^T \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mN(i)})^{-1} \\
 \frac{\partial \mathbf{D}}{\partial \boldsymbol{\theta}} &= \frac{\partial \boldsymbol{\Sigma}_i}{\partial \boldsymbol{\theta}} - \boldsymbol{\Sigma}_{mi}^T \boldsymbol{\Sigma}_m^{-1} \frac{\partial \boldsymbol{\Sigma}_{mi}}{\partial \boldsymbol{\theta}} - \frac{\partial \boldsymbol{\Sigma}_{mi}^T}{\partial \boldsymbol{\theta}} \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mi} + \boldsymbol{\Sigma}_{mi}^T \boldsymbol{\Sigma}_m^{-1} \frac{\partial \boldsymbol{\Sigma}_m}{\partial \boldsymbol{\theta}} \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mi} \\
 &\quad - \frac{\partial \mathbf{A}_i}{\partial \boldsymbol{\theta}} (\boldsymbol{\Sigma}_i - \boldsymbol{\Sigma}_{mi}^T \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mi}) \\
 &\quad - \mathbf{A}_i \left(\frac{\partial \boldsymbol{\Sigma}_i}{\partial \boldsymbol{\theta}} - \boldsymbol{\Sigma}_{mi}^T \boldsymbol{\Sigma}_m^{-1} \frac{\partial \boldsymbol{\Sigma}_{mi}}{\partial \boldsymbol{\theta}} - \frac{\partial \boldsymbol{\Sigma}_{mi}^T}{\partial \boldsymbol{\theta}} \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mi} + \boldsymbol{\Sigma}_{mi}^T \boldsymbol{\Sigma}_m^{-1} \frac{\partial \boldsymbol{\Sigma}_m}{\partial \boldsymbol{\theta}} \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mi}\right)
 \end{aligned}$$

and cost also  $\mathcal{O}(n \cdot (m_v^3 + m_v^2 \cdot m))$ .

### Appendix B. VIF and Laplace Approximation

#### Log-determinant:

$$\begin{aligned}
 \log \det(\boldsymbol{\Sigma}_\dagger \mathbf{W} + \mathbf{I}_n) &= \log \det(\boldsymbol{\Sigma}_\dagger) + \log \det(\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1}) \\
 &= \log \det(\mathbf{M}) - \log \det(\boldsymbol{\Sigma}_m) - \log \det(\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}) + \log \det(\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1}) \\
 &= \log \det(\mathbf{M}) - \log \det(\boldsymbol{\Sigma}_m) - \log \det(\mathbf{D}^{-1}) \\
 &\quad + \log \det(\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} - \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \boldsymbol{\Sigma}_{mn}^T \mathbf{M}^{-1} \boldsymbol{\Sigma}_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}) \\
 &= \log \det(\mathbf{M}) - \log \det(\boldsymbol{\Sigma}_m) - \log \det(\mathbf{D}^{-1}) \\
 &\quad + \log \det(\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}) - \log \det(\mathbf{M}) + \log \det(\mathbf{M}_1) \\
 &= -\log \det(\boldsymbol{\Sigma}_m) - \log \det(\mathbf{D}^{-1}) + \log \det(\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}) \\
 &\quad + \log \det(\mathbf{M}_1),
 \end{aligned}$$

where  $\mathbf{M}_1 = \mathbf{M} - \boldsymbol{\Sigma}_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} (\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \boldsymbol{\Sigma}_{mn}^T$ .

**Linear solves:**

$$\begin{aligned}
 (\mathbf{W} + \Sigma_{\dagger}^{-1})^{-1} &= \mathbf{W}^{-1}(\mathbf{W}^{-1} + \Sigma_{\dagger})^{-1}\Sigma_{\dagger} = \mathbf{W}^{-1}(\mathbf{W}^{-1} + \mathbf{B}^{-1}\mathbf{D}\mathbf{B}^{-T} + \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn})^{-1}\Sigma_{\dagger} \\
 &= \mathbf{W}^{-1}((\mathbf{W}^{-1} + \mathbf{B}^{-1}\mathbf{D}\mathbf{B}^{-T})^{-1} \\
 &\quad - (\mathbf{W}^{-1} + \mathbf{B}^{-1}\mathbf{D}\mathbf{B}^{-T})^{-1}\Sigma_{mn}^T M_2^{-1}\Sigma_{mn}(\mathbf{W}^{-1} + \mathbf{B}^{-1}\mathbf{D}\mathbf{B}^{-T})^{-1})\Sigma_{\dagger} \\
 &= \mathbf{W}^{-1}(\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}(\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{W} \\
 &\quad - \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}(\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{W} \Sigma_{mn}^T M_3^{-1} \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \\
 &\quad \cdot (\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{W}) \Sigma_{\dagger},
 \end{aligned}$$

where  $M_2 = \Sigma_m + \Sigma_{mn}(\mathbf{W}^{-1} + \mathbf{B}^{-1}\mathbf{D}\mathbf{B}^{-T})^{-1}\Sigma_{mn}^T$  and  $M_3 = \Sigma_m + \Sigma_{mn}\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}(\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1}\mathbf{W}\Sigma_{mn}^T$ .

**Gradients:**

The gradients with respect to  $\boldsymbol{\theta}$  and  $\boldsymbol{\xi}$  are given by

$$\begin{aligned}
 \frac{\partial L^{VIFLA}(\mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\xi})}{\partial \theta_k} &= \frac{1}{2} \tilde{\mathbf{b}}^T \frac{\partial \Sigma_{\dagger}^{-1}}{\partial \theta_k} \tilde{\mathbf{b}} + \frac{1}{2} \frac{\partial \log \det(\Sigma_{\dagger} \mathbf{W} + \mathbf{I}_n)}{\partial \theta_k} \\
 &\quad + \left( \frac{\partial L^{VIFLA}(\mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\xi})}{\partial \tilde{\mathbf{b}}} \right)^T \frac{\partial \tilde{\mathbf{b}}}{\partial \theta_k}, \quad k = 1, \dots, q, \\
 \frac{\partial L^{VIFLA}(\mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\xi})}{\partial \xi_l} &= - \frac{\partial \log p(\mathbf{y} | \mathbf{b}^*, \boldsymbol{\xi})}{\partial \xi_l} + \frac{1}{2} \frac{\partial \log \det(\Sigma_{\dagger} \mathbf{W} + \mathbf{I}_n)}{\partial \xi_l} \\
 &\quad + \left( \frac{\partial L^{LA}(\mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\xi})}{\partial \tilde{\mathbf{b}}} \right)^T \frac{\partial \tilde{\mathbf{b}}}{\partial \xi_l}, \quad l = 1, \dots, r,
 \end{aligned}$$

where

$$\begin{aligned}
 \frac{\partial L^{VIFLA}(\mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\xi})}{\partial b^*(s_i)} &= \frac{1}{2} \frac{\partial \log \det(\Sigma_{\dagger} \mathbf{W} + \mathbf{I}_n)}{\partial b^*(s_i)} = \frac{1}{2} \text{Tr} \left( (\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \frac{\partial \mathbf{W}}{\partial b^*(s_i)} \right) \\
 &\quad + \frac{1}{2} \text{Tr} \left( M_1^{-1} \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} (\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \right. \\
 &\quad \cdot \left. \frac{\partial \mathbf{W}}{\partial b^*(s_i)} (\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \Sigma_{mn}^T \right),
 \end{aligned}$$

$$\begin{aligned}
 \frac{\partial \tilde{\mathbf{b}}}{\partial F_i} &= - \left( \mathbf{W} + \Sigma_{\dagger}^{-1} \right)^{-1} \mathbf{W}_{\cdot i} = -\mathbf{W}^{-1}(\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}(\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{W} \\
 &\quad - \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}(\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{W} \Sigma_{mn}^T \\
 &\quad \cdot M_3^{-1} \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}(\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{W}) \Sigma_{\dagger} \mathbf{W}_{\cdot i},
 \end{aligned}$$

$$\begin{aligned}
 \frac{\partial \tilde{\mathbf{b}}}{\partial \theta_k} &= - \left( \mathbf{W} + \Sigma_{\dagger}^{-1} \right)^{-1} \frac{\partial \Sigma_{\dagger}^{-1}}{\partial \theta_k} \tilde{\mathbf{b}} = -\mathbf{W}^{-1}(\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}(\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{W} \\
 &\quad - \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}(\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{W} \Sigma_{mn}^T \\
 &\quad \cdot M_3^{-1} \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}(\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{W}) \Sigma_{\dagger} \frac{\partial \Sigma_{\dagger}^{-1}}{\partial \theta_k} \tilde{\mathbf{b}},
 \end{aligned}$$

$$\begin{aligned}
 \frac{\partial \tilde{\mathbf{b}}}{\partial \xi_l} &= \left( \mathbf{W} + \Sigma_{\dagger}^{-1} \right)^{-1} \frac{\partial^2 \log p(\mathbf{y} \mid \mathbf{b}^*, \boldsymbol{\xi})}{\partial \xi_l \partial \tilde{\mathbf{b}}} \\
 &= \mathbf{W}^{-1} \left( \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} (\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{W} - \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} (\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{W} \Sigma_{mn}^T \right. \\
 &\quad \left. \cdot \mathbf{M}_3^{-1} \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} (\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{W} \right) \Sigma_{\dagger} \frac{\partial^2 \log p(\mathbf{y} \mid \mathbf{b}^*, \boldsymbol{\xi})}{\partial \xi_l \partial \tilde{\mathbf{b}}}, \\
 \frac{\partial \log \det(\Sigma_{\dagger} \mathbf{W} + \mathbf{I}_n)}{\partial b^*(s_i)} &= \frac{\partial(\log \det(\Sigma_{\dagger}) + \log \det(\mathbf{W} + \Sigma_{\dagger}^{-1}))}{\partial b^*(s_i)} = \frac{\partial(\log \det(\mathbf{W} + \Sigma_{\dagger}^{-1}))}{\partial b^*(s_i)} \\
 &= \text{Tr} \left( (\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \frac{\partial \mathbf{W}}{\partial b^*(s_i)} \right) \\
 &\quad + \text{Tr} \left( \mathbf{M}_1^{-1} \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} (\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \right. \\
 &\quad \left. \cdot \frac{\partial \mathbf{W}}{\partial b^*(s_i)} (\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \Sigma_{mn}^T \right) \quad i = 1, \dots, p, \\
 \frac{\partial \log \det(\Sigma_{\dagger} \mathbf{W} + \mathbf{I}_n)}{\partial \theta_k} &= -\text{Tr} \left( \Sigma_m^{-1} \frac{\partial \Sigma_m}{\partial \theta_k} \right) + \text{Tr} \left( \mathbf{M}_1^{-1} \frac{\partial \mathbf{M}_1}{\partial \theta_k} \right) + \text{Tr} \left( \mathbf{D}^{-1} \frac{\partial \mathbf{D}}{\partial \theta_k} \right) \\
 &\quad + \text{Tr} \left( (\mathbf{W} + \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \frac{\partial (\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})}{\partial \theta_k} \right) \quad k = 1, \dots, q, \\
 \frac{\partial \log \det(\Sigma_{\dagger} \mathbf{W} + \mathbf{I}_n)}{\partial \xi_l} &= \text{Tr} \left( (\mathbf{W} + \Sigma_{\dagger}^{-1})^{-1} \frac{\partial \mathbf{W}}{\partial \xi_l} \right) \quad l = 1, \dots, r.
 \end{aligned}$$

## Appendix C. Predictive (co-)variances

### C.1 Gaussian

$$\begin{aligned}
 \Sigma_{\dagger}^p &= \mathbf{B}_p^{-1} \mathbf{D}_p \mathbf{B}_p^{-T} + \mathbf{B}_p^{-1} \mathbf{B}_{po} (\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{B}_{po}^T \mathbf{B}_p^{-T} + \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn_p} \\
 &\quad - (\Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} - \mathbf{B}_p^{-1} \mathbf{B}_{po} (\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1}) \tilde{\Sigma}_{\dagger}^{-1} (\Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} \\
 &\quad - (\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{B}_{po}^T \mathbf{B}_p^{-T}) \\
 &= \mathbf{B}_p^{-1} \mathbf{D}_p \mathbf{B}_p^{-T} + \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn_p} - \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} \\
 &\quad + \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} \mathbf{B}_{po}^T \mathbf{B}_p^{-T} + \mathbf{B}_p^{-1} \mathbf{B}_{po} \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} + \mathbf{B}_p^{-1} \mathbf{B}_{po} \Sigma_{mn}^T \mathbf{M}^{-1} \Sigma_{mn} \mathbf{B}_{po}^T \mathbf{B}_p^{-T} \\
 &\quad - \mathbf{B}_p^{-1} \mathbf{B}_{po} \Sigma_{mn}^T \mathbf{M}^{-1} \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} \\
 &\quad - \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \Sigma_{mn}^T \mathbf{M}^{-1} \Sigma_{mn} \mathbf{B}_{po}^T \mathbf{B}_p^{-T} \\
 &\quad + \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \Sigma_{mn}^T \mathbf{M}^{-1} \Sigma_{mn} \mathbf{B}^T \mathbf{D}^{-1} \mathbf{B} \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p}
 \end{aligned}$$

## C.2 Non-Gaussian

$$\begin{aligned}
 \Omega_p &= B_p^{-1} D_p B_p^{-T} + B_p^{-1} B_{po} (B^T D^{-1} B)^{-1} B_{po}^T B_p^{-T} + \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn_p} \\
 &\quad - (\Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} - B_p^{-1} B_{po} (B^T D^{-1} B)^{-1}) (\Sigma_{\dagger}^{-1} - \Sigma_{\dagger}^{-1} (W + \Sigma_{\dagger}^{-1})^{-1} \Sigma_{\dagger}^{-1}) \\
 &\quad \cdot (\Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} - (B^T D^{-1} B)^{-1} B_{po}^T B_p^{-T}) \\
 &\stackrel{(1)}{=} B_p^{-1} D_p B_p^{-T} + B_p^{-1} B_{po} (B^T D^{-1} B)^{-1} B_{po}^T B_p^{-T} + \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn_p} \\
 &\quad - (\Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} - B_p^{-1} B_{po} (B^T D^{-1} B)^{-1}) (W^{-1} + \Sigma_{\dagger}^{-1})^{-1} (\Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} \\
 &\quad - (B^T D^{-1} B)^{-1} B_{po}^T B_p^{-T}) \\
 &= B_p^{-1} D_p B_p^{-T} + B_p^{-1} B_{po} (B^T D^{-1} B)^{-1} B_{po}^T B_p^{-T} + \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn_p} \\
 &\quad - (\Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} - B_p^{-1} B_{po} (B^T D^{-1} B)^{-1}) \\
 &\quad \cdot ((W^{-1} + B^{-1} D B^{-T})^{-1} \\
 &\quad - (W^{-1} + B^{-1} D B^{-T})^{-1} \Sigma_{mn}^T M_2^{-1} \Sigma_{mn} (W^{-1} + B^{-1} D B^{-T})^{-1}) \\
 &\quad \cdot (\Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} - (B^T D^{-1} B)^{-1} B_{po}^T B_p^{-T}) \\
 &= B_p^{-1} D_p B_p^{-T} + B_p^{-1} B_{po} (B^T D^{-1} B)^{-1} B_{po}^T B_p^{-T} + \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn_p} \\
 &\quad - (\Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} - B_p^{-1} B_{po} (B^T D^{-1} B)^{-1}) (W^{-1} + B^{-1} D B^{-T})^{-1} \\
 &\quad \cdot (\Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} - (B^T D^{-1} B)^{-1} B_{po}^T B_p^{-T}) \tag{23}
 \end{aligned}$$

$$\begin{aligned}
 &+ (\Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} - B_p^{-1} B_{po} (B^T D^{-1} B)^{-1}) \\
 &\quad \cdot (W^{-1} + B^{-1} D B^{-T})^{-1} \Sigma_{mn}^T M_2^{-1} \Sigma_{mn} (W^{-1} + B^{-1} D B^{-T})^{-1} \\
 &\quad \cdot (\Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} - (B^T D^{-1} B)^{-1} B_{po}^T B_p^{-T}), \tag{24}
 \end{aligned}$$

where in (1) we used the Sherman-Woodbury-Morrison formula. We compute the terms in (23) and (24) separately:

$$\begin{aligned}
 &(\Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} - B_p^{-1} B_{po} (B^T D^{-1} B)^{-1}) (W^{-1} + B^{-1} D B^{-T})^{-1} \\
 &\quad \cdot (\Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} - (B^T D^{-1} B)^{-1} B_{po}^T B_p^{-T}) \\
 &\stackrel{(1)}{=} (\Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} - B_p^{-1} B_{po} (B^T D^{-1} B)^{-1}) \\
 &\quad \cdot (B^T D^{-1} B - B^T D^{-1} B (W + B^T D^{-1} B)^{-1} B^T D^{-1} B) \\
 &\quad \cdot (\Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} - (B^T D^{-1} B)^{-1} B_{po}^T B_p^{-T}) \\
 &= B_p^{-1} B_{po} (B^T D^{-1} B)^{-1} B_{po}^T B_p^{-T} + \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} B^T D^{-1} B \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} \\
 &\quad - B_p^{-1} B_{po} \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} - \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} B_{po}^T B_p^{-T} \\
 &\quad - B_{po}^T B_p^{-T} (W + B^T D^{-1} B)^{-1} B_{po}^T B_p^{-T} \\
 &\quad - \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} B^T D^{-1} B (W + B^T D^{-1} B)^{-1} B^T D^{-1} B \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} \\
 &\quad + \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} B^T D^{-1} B (W + B^T D^{-1} B)^{-1} B_{po}^T B_p^{-T} \\
 &\quad + B_{po}^T B_p^{-T} (W + B^T D^{-1} B)^{-1} B^T D^{-1} B \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p},
 \end{aligned}$$

where in (1) we used again the Sherman-Woodbury-Morrison formula.

$$\begin{aligned}
 & (\Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} - B_p^{-1} B_{po} (B^T D^{-1} B)^{-1}) \\
 & \cdot (W^{-1} + B^{-1} D B^{-T})^{-1} \Sigma_{mn}^T M_2^{-1} \Sigma_{mn} (W^{-1} + B^{-1} D B^{-T})^{-1} \\
 & \cdot (\Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} - (B^T D^{-1} B)^{-1} B_{po}^T B_p^{-T}) \\
 & \stackrel{(2)}{=} (\Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} - B_p^{-1} B_{po} (B^T D^{-1} B)^{-1}) \\
 & \cdot B^T D^{-1} B (W + B^T D^{-1} B)^{-1} W^{-1} \Sigma_{mn}^T M_3^{-1} \Sigma_{mn} W^{-1} (W + B^T D^{-1} B)^{-1} B^T D^{-1} B \\
 & \cdot (\Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} - (B^T D^{-1} B)^{-1} B_{po}^T B_p^{-T}) \\
 & = B_p^{-1} B_{po} (W + B^T D^{-1} B)^{-1} W^{-1} \Sigma_{mn}^T M_3^{-1} \Sigma_{mn} W^{-1} (W + B^T D^{-1} B)^{-1} B_{po}^T B_p^{-T} \\
 & + \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} \\
 & \cdot B^T D^{-1} B (W + B^T D^{-1} B)^{-1} W^{-1} \Sigma_{mn}^T M_3^{-1} \Sigma_{mn} W^{-1} (W + B^T D^{-1} B)^{-1} B^T D^{-1} B \\
 & \cdot \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} \\
 & - B_p^{-1} B_{po} (W + B^T D^{-1} B)^{-1} W^{-1} \Sigma_{mn}^T M_3^{-1} \Sigma_{mn} W^{-1} \\
 & \cdot (W + B^T D^{-1} B)^{-1} B^T D^{-1} B \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} \\
 & - \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} B^T D^{-1} B (W + B^T D^{-1} B)^{-1} W^{-1} \Sigma_{mn}^T M_3^{-1} \Sigma_{mn} W^{-1} \\
 & \cdot (W + B^T D^{-1} B)^{-1} B_{po}^T B_p^{-T},
 \end{aligned}$$

where in (2) we used that  $(W^{-1} + B^{-1} D B^{-T})^{-1} = W^{-1} (W + B^T D^{-1} B)^{-1} B^T D^{-1} B = B^T D^{-1} B (W + B^T D^{-1} B)^{-1} W^{-1}$ . Finally, we obtain:

$$\begin{aligned}
 \Omega_p & = B_p^{-1} D_p B_p^{-T} + \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn_p} - \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} B^T D^{-1} B \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} \\
 & + B_p^{-1} B_{po} \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} + \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} B_{po}^T B_p^{-T} \\
 & + B_{po}^T B_p^{-T} (W + B^T D^{-1} B)^{-1} B_{po}^T B_p^{-T} \\
 & + \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} B^T D^{-1} B (W + B^T D^{-1} B)^{-1} B^T D^{-1} B \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} \\
 & - \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} B^T D^{-1} B (W + B^T D^{-1} B)^{-1} B_{po}^T B_p^{-T} \\
 & - B_{po}^T B_p^{-T} (W + B^T D^{-1} B)^{-1} B^T D^{-1} B \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} \\
 & + B_p^{-1} B_{po} (W + B^T D^{-1} B)^{-1} W^{-1} \Sigma_{mn}^T M_3^{-1} \Sigma_{mn} W^{-1} (W + B^T D^{-1} B)^{-1} B_{po}^T B_p^{-T} \\
 & + \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} \\
 & \cdot B^T D^{-1} B (W + B^T D^{-1} B)^{-1} W^{-1} \Sigma_{mn}^T M_3^{-1} \Sigma_{mn} W^{-1} \\
 & \cdot (W + B^T D^{-1} B)^{-1} B^T D^{-1} B \\
 & \cdot \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} \\
 & - B_p^{-1} B_{po} (W + B^T D^{-1} B)^{-1} W^{-1} \Sigma_{mn}^T M_3^{-1} \Sigma_{mn} W^{-1} \\
 & \cdot (W + B^T D^{-1} B)^{-1} B^T D^{-1} B \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn_p} \\
 & - \Sigma_{mn_p}^T \Sigma_m^{-1} \Sigma_{mn} B^T D^{-1} B (W + B^T D^{-1} B)^{-1} W^{-1} \Sigma_{mn}^T M_3^{-1} \Sigma_{mn} W^{-1} \\
 & \cdot (W + B^T D^{-1} B)^{-1} B_{po}^T B_p^{-T}.
 \end{aligned}$$

While Algorithms 1 and 2 produce estimators based on sampled realizations, the following proofs are formulated in terms of the underlying random vectors  $\mathbf{z}_i$ .

**Proof** [Proof of Proposition 3] By standard results, in Algorithm 1, the only non-deterministic term  $1/\ell \sum_{i=1}^{\ell} \mathbf{z}_i^{(8)} (\mathbf{z}_i^{(8)})^T$  is an unbiased and consistent estimator for  $(\boldsymbol{\Sigma}_{mn_p}^T \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn} - \mathbf{B}_p^{-1} \mathbf{B}_{po} (\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1}) \boldsymbol{\Sigma}_{\dagger}^{-1} (\boldsymbol{\Sigma}_{\dagger}^{-1} + \mathbf{W})^{-1} \boldsymbol{\Sigma}_{\dagger}^{-1} (\boldsymbol{\Sigma}_{mn}^T \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn_p} - (\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{B}_{po}^T \mathbf{B}_p^{-T})$ , and the claim in Proposition 3 thus follows.  $\blacksquare$

**Proof** [Proof of Proposition 4] By standard results, in Algorithm 2, the only non-deterministic term  $1/\ell \sum_{i=1}^{\ell} \mathbf{z}_i^{(1)} \circ (\boldsymbol{\Sigma}_{mn_p}^T \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn} - \mathbf{B}_p^{-1} \mathbf{B}_{po} (\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1}) \boldsymbol{\Sigma}_{\dagger}^{-1} (\mathbf{W} + \boldsymbol{\Sigma}_{\dagger}^{-1})^{-1} \boldsymbol{\Sigma}_{\dagger}^{-1} (\boldsymbol{\Sigma}_{mn}^T \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn_p} - (\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{B}_{po}^T \mathbf{B}_p^{-T}) \mathbf{z}_i^{(1)}$  is an unbiased and consistent estimate of  $\text{diag}(\boldsymbol{\Sigma}_{mn_p}^T \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn} - \mathbf{B}_p^{-1} \mathbf{B}_{po} (\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1}) \boldsymbol{\Sigma}_{\dagger}^{-1} (\mathbf{W} + \boldsymbol{\Sigma}_{\dagger}^{-1})^{-1} \boldsymbol{\Sigma}_{\dagger}^{-1} (\boldsymbol{\Sigma}_{mn}^T \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn_p} - (\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} \mathbf{B}_{po}^T \mathbf{B}_p^{-T})$ , and the claim in Proposition 4 thus follows.  $\blacksquare$

## Appendix D. Derivations of SLQ-approximated log-determinants and stochastic trace estimation for gradients

The stochastic Lanczos quadrature (SLQ) method can be derived as follows:

$$\begin{aligned} \log \det (\boldsymbol{\Sigma}_{\dagger} \mathbf{W} + \mathbf{I}_n) &= \log \det (\boldsymbol{\Sigma}_{\dagger}) + \log \det (\mathbf{W} + \tilde{\boldsymbol{\Sigma}}_{\dagger}^{-1}) \\ &= \log \det (\boldsymbol{\Sigma}_{\dagger}) + \log \det (\mathbf{P}^{-\frac{1}{2}} (\mathbf{W} + \tilde{\boldsymbol{\Sigma}}_{\dagger}^{-1}) \mathbf{P}^{-\frac{1}{2}}) + \log \det (\mathbf{P}) \\ &\approx \log \det (\boldsymbol{\Sigma}_{\dagger}) + \frac{1}{\ell} \sum_{i=1}^{\ell} \mathbf{z}_i^T \mathbf{P}^{-\frac{1}{2}} \tilde{\mathbf{Q}}_i \log (\tilde{\mathbf{T}}_i) \tilde{\mathbf{Q}}_i^T \mathbf{P}^{-\frac{1}{2}} \mathbf{z}_i + \log \det (\mathbf{P}) \\ &\approx \log \det (\boldsymbol{\Sigma}_{\dagger}) + \frac{n}{\ell} \sum_{i=1}^{\ell} \mathbf{e}_1^T \log (\tilde{\mathbf{T}}_i) \mathbf{e}_1 + \log \det (\mathbf{P}), \end{aligned}$$

where  $\tilde{\mathbf{Q}}_i \in \mathbb{R}^{n \times k}$  has orthonormal columns,  $\tilde{\mathbf{T}}_i \in \mathbb{R}^{k \times k}$  is the tridiagonal matrix of the partial Lanczos tridiagonalization  $\tilde{\mathbf{Q}}_i \tilde{\mathbf{T}}_i \tilde{\mathbf{Q}}_i^T \approx \mathbf{P}^{-\frac{1}{2}} (\mathbf{W} + \tilde{\boldsymbol{\Sigma}}_{\dagger}^{-1}) \mathbf{P}^{-\frac{1}{2}}$  obtained by running the Lanczos algorithm for  $k$  steps with  $\mathbf{P}^{-\frac{1}{2}} (\mathbf{W} + \tilde{\boldsymbol{\Sigma}}_{\dagger}^{-1}) \mathbf{P}^{-\frac{1}{2}}$  and initial vector  $\mathbf{P}^{-\frac{1}{2}} \mathbf{z}_i / \sqrt{n} \approx \mathbf{P}^{-\frac{1}{2}} \mathbf{z}_i / \|\mathbf{P}^{-\frac{1}{2}} \mathbf{z}_i\|_2$ , and  $\mathbf{z}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{P})$ . Alternatively, we have

$$\begin{aligned} \log \det (\boldsymbol{\Sigma}_{\dagger} \mathbf{W} + \mathbf{I}_n) &= \log \det (\mathbf{W}) + \log \det (\mathbf{W}^{-1} + \tilde{\boldsymbol{\Sigma}}_{\dagger}) = \log \det (\mathbf{W}) \\ &\quad + \log \det (\mathbf{W}^{-1} + \tilde{\boldsymbol{\Sigma}}_{\dagger}) \\ &\approx \log \det (\mathbf{W}) + \frac{n}{\ell} \sum_{i=1}^{\ell} \mathbf{e}_1^T \log (\tilde{\mathbf{T}}_i) \mathbf{e}_1 + \log \det (\mathbf{P}), \end{aligned}$$

where  $\tilde{\mathbf{T}}_i \in \mathbb{R}^{k \times k}$  is the tridiagonal matrix of the partial Lanczos tridiagonalization  $\tilde{\mathbf{Q}}_i \tilde{\mathbf{T}}_i \tilde{\mathbf{Q}}_i^T \approx \mathbf{P}^{-\frac{1}{2}} (\mathbf{W}^{-1} + \tilde{\boldsymbol{\Sigma}}_{\dagger}) \mathbf{P}^{-\frac{1}{2}}$  obtained by running the Lanczos algorithm for  $k$  steps with  $\mathbf{P}^{-\frac{1}{2}} (\mathbf{W}^{-1} + \tilde{\boldsymbol{\Sigma}}_{\dagger}) \mathbf{P}^{-\frac{1}{2}}$  and initial vector  $\mathbf{P}^{-\frac{1}{2}} \mathbf{z}_i / \|\mathbf{P}^{-\frac{1}{2}} \mathbf{z}_i\|_2$ , and  $\mathbf{z}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{P})$ .

Furthermore, the stochastic trace estimation for gradients can be computed as follows:

$$\begin{aligned}
 \text{Tr} \left( (\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})^{-1} \frac{\partial(\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})}{\partial b^*(\mathbf{s}_i)} \right) &= \text{Tr} \left( (\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})^{-1} \frac{\partial(\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})}{\partial b^*(\mathbf{s}_i)} \mathbb{E}_{\mathbf{z}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{P})} [\mathbf{P}^{-1} \mathbf{z}_i \mathbf{z}_i^\top] \right) \\
 &\approx \frac{1}{\ell} \sum_{i=1}^{\ell} (\mathbf{z}_i^\top (\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})^{-1}) \left( \frac{\partial \mathbf{W}}{\partial b^*(\mathbf{s}_i)} \mathbf{P}^{-1} \mathbf{z}_i \right) \\
 &= \tilde{T}_\ell^{b^*(\mathbf{s}_i)} \quad i = 1, \dots, p, \\
 \text{Tr} \left( (\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})^{-1} \frac{\partial(\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})}{\partial \theta_k} \right) &\approx \frac{1}{\ell} \sum_{i=1}^{\ell} (\mathbf{z}_i^\top (\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})^{-1}) \left( \frac{\partial(\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})}{\partial \theta_k} \mathbf{P}^{-1} \mathbf{z}_i \right) \\
 &= \tilde{T}_\ell^{\theta_k} \quad k = 1, \dots, q, \\
 \text{Tr} \left( (\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})^{-1} \frac{\partial(\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})}{\partial \xi_l} \right) &\approx \frac{1}{\ell} \sum_{i=1}^{\ell} (\mathbf{z}_i^\top (\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})^{-1}) \left( \frac{\partial \mathbf{W}}{\partial \xi_l} \mathbf{P}^{-1} \mathbf{z}_i \right) \\
 &= \tilde{T}_\ell^{\xi_l} \quad l = 1, \dots, r,
 \end{aligned}$$

or e.g.

$$\begin{aligned}
 \text{Tr} \left( (\mathbf{W}^{-1} + \boldsymbol{\Sigma}_\dagger)^{-1} \frac{\partial(\mathbf{W}^{-1} + \boldsymbol{\Sigma}_\dagger)}{\partial \theta_k} \right) &\approx \frac{1}{\ell} \sum_{i=1}^{\ell} (\mathbf{z}_i^\top (\mathbf{W}^{-1} + \boldsymbol{\Sigma}_\dagger)^{-1}) \left( \frac{\partial(\mathbf{W}^{-1} + \boldsymbol{\Sigma}_\dagger)}{\partial \theta_k} \mathbf{P}^{-1} \mathbf{z}_i \right) \\
 &= \tilde{T}_\ell^{\theta_k} \quad k = 1, \dots, q.
 \end{aligned}$$

Further, we can reduce the variance of this stochastic estimate by using the preconditioner  $\mathbf{P}$  to build a control variate, e.g.,

$$\begin{aligned}
 \text{Tr} \left( (\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})^{-1} \frac{\partial(\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})}{\partial \theta_k} \right) &\approx \frac{1}{\ell} \sum_{i=1}^{\ell} \left( (\mathbf{z}_i^\top (\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})^{-1}) \left( \frac{\partial(\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})}{\partial \theta_k} \mathbf{P}^{-1} \mathbf{z}_i \right) \right. \\
 &\quad \left. - \hat{c}_{\text{opt}} \cdot (\mathbf{z}_i^\top \mathbf{P}^{-1}) \left( \frac{\partial \mathbf{P}}{\partial \theta_k} \mathbf{P}^{-1} \mathbf{z}_i \right) \right) + \hat{c}_{\text{opt}} \cdot \text{Tr} \left( \mathbf{P}^{-1} \frac{\partial \mathbf{P}}{\partial \theta_k} \right), \\
 \hat{c}_{\text{opt}} &= \frac{\sum_{i=1}^{\ell} \left( (\mathbf{z}_i^\top (\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})^{-1}) \left( \frac{\partial(\mathbf{W} + \boldsymbol{\Sigma}_\dagger^{-1})}{\partial \theta_k} \mathbf{P}^{-1} \mathbf{z}_i \right) - \tilde{T}_\ell \right) \left( (\mathbf{z}_i^\top \mathbf{P}^{-1}) \left( \frac{\partial \mathbf{P}}{\partial \theta_k} \mathbf{P}^{-1} \mathbf{z}_i \right) - \text{Tr} \left( \mathbf{P}^{-1} \frac{\partial \mathbf{P}}{\partial \theta_k} \right) \right)}{\sum_{i=1}^{\ell} \left[ (\mathbf{z}_i^\top \mathbf{P}^{-1}) \left( \frac{\partial \mathbf{P}}{\partial \theta_k} \mathbf{P}^{-1} \mathbf{z}_i \right) - \text{Tr} \left( \mathbf{P}^{-1} \frac{\partial \mathbf{P}}{\partial \theta_k} \right) \right]^2}.
 \end{aligned}$$

Or,

$$\begin{aligned}
 \text{Tr} \left( (\mathbf{W}^{-1} + \boldsymbol{\Sigma}_\dagger)^{-1} \frac{\partial(\mathbf{W}^{-1} + \boldsymbol{\Sigma}_\dagger)}{\partial \theta_k} \right) &\approx \frac{1}{\ell} \sum_{i=1}^{\ell} \left( (\mathbf{z}_i^\top (\mathbf{W}^{-1} + \boldsymbol{\Sigma}_\dagger)^{-1}) \left( \frac{\partial(\mathbf{W}^{-1} + \boldsymbol{\Sigma}_\dagger)}{\partial \theta_k} \mathbf{P}^{-1} \mathbf{z}_i \right) \right. \\
 &\quad \left. - \hat{c}_{\text{opt}} \cdot (\mathbf{z}_i^\top \mathbf{P}^{-1}) \left( \frac{\partial \mathbf{P}}{\partial \theta_k} \mathbf{P}^{-1} \mathbf{z}_i \right) \right) + \hat{c}_{\text{opt}} \cdot \text{Tr} \left( \mathbf{P}^{-1} \frac{\partial \mathbf{P}}{\partial \theta_k} \right), \\
 \hat{c}_{\text{opt}} &= \frac{\sum_{i=1}^{\ell} \left( (\mathbf{z}_i^\top (\mathbf{W}^{-1} + \boldsymbol{\Sigma}_\dagger)^{-1}) \left( \frac{\partial(\mathbf{W}^{-1} + \boldsymbol{\Sigma}_\dagger)}{\partial \theta_k} \mathbf{P}^{-1} \mathbf{z}_i \right) - \tilde{T}_\ell \right) \left( (\mathbf{z}_i^\top \mathbf{P}^{-1}) \left( \frac{\partial \mathbf{P}}{\partial \theta_k} \mathbf{P}^{-1} \mathbf{z}_i \right) - \text{Tr} \left( \mathbf{P}^{-1} \frac{\partial \mathbf{P}}{\partial \theta_k} \right) \right)}{\sum_{i=1}^{\ell} \left[ (\mathbf{z}_i^\top \mathbf{P}^{-1}) \left( \frac{\partial \mathbf{P}}{\partial \theta_k} \mathbf{P}^{-1} \mathbf{z}_i \right) - \text{Tr} \left( \mathbf{P}^{-1} \frac{\partial \mathbf{P}}{\partial \theta_k} \right) \right]^2}.
 \end{aligned}$$

## Appendix E. Additional information on the preconditioners

### E.1 VIF with diagonal update (VIFDU) preconditioner

The “**VIF** with **diagonal update**” (VIFDU) preconditioner is given by

$$\hat{\mathbf{P}} = \mathbf{B}^\top \mathbf{W} \mathbf{B} + \boldsymbol{\Sigma}_\dagger^{-1} = \mathbf{B}^\top (\mathbf{W} + \mathbf{D}^{-1} - \mathbf{D}^{-1} \mathbf{B} \boldsymbol{\Sigma}_{mn}^\top \mathbf{M}^{-1} \boldsymbol{\Sigma}_{mn} \mathbf{B}^\top \mathbf{D}^{-1}) \mathbf{B}.$$

Linear solves with  $\hat{\mathbf{P}}$

$$\hat{\mathbf{P}}^{-1} \mathbf{w} =$$

$\mathbf{B}^{-1} ((\mathbf{W} + \mathbf{D}^{-1})^{-1} + (\mathbf{W} + \mathbf{D}^{-1})^{-1} \mathbf{D}^{-1} \mathbf{B} \boldsymbol{\Sigma}_{mn}^\top \mathbf{M}_3^{-1} \boldsymbol{\Sigma}_{mn} \mathbf{B}^\top \mathbf{D}^{-1} (\mathbf{W} + \mathbf{D}^{-1})^{-1}) \mathbf{B}^{-\top} \mathbf{w}$ , where  $\mathbf{w} \in \mathbb{R}^n$  and  $\mathbf{M}_3 = \mathbf{M} - \boldsymbol{\Sigma}_{mn} \mathbf{B}^\top \mathbf{D}^{-1} (\mathbf{W} + \mathbf{D}^{-1})^{-1} \mathbf{D}^{-1} \mathbf{B} \boldsymbol{\Sigma}_{mn}^\top$ , are computed in  $\mathcal{O}(n \cdot (m_v + m))$  when  $\mathbf{M}_3$  is precomputed in  $\mathcal{O}(n \cdot (m_v \cdot m + m^2))$ . We can compute the log-determinant and its derivatives using Sylvester’s determinant theorem

$$\begin{aligned} \log \det (\hat{\mathbf{P}}) &= \log \det \left( \mathbf{B}^\top (\mathbf{W} + \mathbf{D}^{-1} - \mathbf{D}^{-1} \mathbf{B} \boldsymbol{\Sigma}_{mn}^\top \mathbf{M}^{-1} \boldsymbol{\Sigma}_{mn} \mathbf{B}^\top \mathbf{D}^{-1}) \mathbf{B} \right) \\ &= \log \det (\mathbf{W} + \mathbf{D}^{-1} - \mathbf{D}^{-1} \mathbf{B} \boldsymbol{\Sigma}_{mn}^\top \mathbf{M}^{-1} \boldsymbol{\Sigma}_{mn} \mathbf{B}^\top \mathbf{D}^{-1}) \\ &= \log \det (\mathbf{W} + \mathbf{D}^{-1}) - \log \det (\mathbf{M}) + \log \det (\mathbf{M}_3) \end{aligned}$$

$$\begin{aligned} \text{Tr} \left( \hat{\mathbf{P}}^{-1} \frac{\partial \hat{\mathbf{P}}}{\partial b^*(\mathbf{s}_i)} \right) &= \text{Tr} \left( (\mathbf{W} + \mathbf{D}^{-1})^{-1} \frac{\partial \mathbf{W}}{\partial b^*(\mathbf{s}_i)} \right) \\ &\quad + \text{Tr} \left( \mathbf{M}_3^{-1} \boldsymbol{\Sigma}_{mn} \mathbf{B}^\top \mathbf{D}^{-1} (\mathbf{W} + \mathbf{D}^{-1})^{-1} \frac{\partial \mathbf{W}}{\partial b^*(\mathbf{s}_i)} (\mathbf{W} + \mathbf{D}^{-1})^{-1} \mathbf{D}^{-1} \mathbf{B} \boldsymbol{\Sigma}_{mn}^\top \right) \end{aligned}$$

$$\text{Tr} \left( \hat{\mathbf{P}}^{-1} \frac{\partial \hat{\mathbf{P}}}{\partial \theta_k} \right) = \text{Tr} \left( (\mathbf{W} + \mathbf{D}^{-1})^{-1} \frac{\partial (\mathbf{W} + \mathbf{D}^{-1})}{\partial \theta_k} \right) - \text{Tr} \left( \mathbf{M}^{-1} \frac{\partial \mathbf{M}}{\partial \theta_k} \right) + \text{Tr} \left( \mathbf{M}_3^{-1} \frac{\partial \mathbf{M}_3}{\partial \theta_k} \right)$$

$$\begin{aligned} \text{Tr} \left( \hat{\mathbf{P}}^{-1} \frac{\partial \hat{\mathbf{P}}}{\partial \xi_l} \right) &= \text{Tr} \left( (\mathbf{W} + \mathbf{D}^{-1})^{-1} \frac{\partial \mathbf{W}}{\partial \xi_l} \right) \\ &\quad + \text{Tr} \left( \mathbf{M}_3^{-1} \boldsymbol{\Sigma}_{mn} \mathbf{B}^\top \mathbf{D}^{-1} (\mathbf{W} + \mathbf{D}^{-1})^{-1} \frac{\partial \mathbf{W}}{\partial \xi_l} (\mathbf{W} + \mathbf{D}^{-1})^{-1} \mathbf{D}^{-1} \mathbf{B} \boldsymbol{\Sigma}_{mn}^\top \right) \end{aligned}$$

in  $\mathcal{O}(n \cdot (m_v \cdot m + m^2))$ .

### E.2 FITC preconditioner

The **FITC** preconditioner is given by

$$\hat{\mathbf{P}} = \text{diag}(\boldsymbol{\Sigma} - \boldsymbol{\Sigma}_{kn}^\top \boldsymbol{\Sigma}_k^{-1} \boldsymbol{\Sigma}_{kn}) + \mathbf{W}^{-1} + \boldsymbol{\Sigma}_{kn}^\top \boldsymbol{\Sigma}_k^{-1} \boldsymbol{\Sigma}_{kn} = \mathbf{D}_V + \boldsymbol{\Sigma}_{kn}^\top \boldsymbol{\Sigma}_k^{-1} \boldsymbol{\Sigma}_{kn},$$

where  $\mathbf{D}_V = \text{diag}(\boldsymbol{\Sigma} - \boldsymbol{\Sigma}_{kn}^\top \boldsymbol{\Sigma}_k^{-1} \boldsymbol{\Sigma}_{kn}) + \mathbf{W}^{-1} \in \mathbb{R}^{n \times n}$ ,  $\boldsymbol{\Sigma}_k = [c_\theta(\hat{\mathbf{s}}_i, \hat{\mathbf{s}}_j)]_{i=1:k, j=1:k} \in \mathbb{R}^{k \times k}$  is a covariance matrix and  $\boldsymbol{\Sigma}_{kn} = [c_\theta(\hat{\mathbf{s}}_i, \mathbf{s}_j)]_{i=1:k, j=1:n} \in \mathbb{R}^{k \times n}$  a cross-covariance matrix with respect to the set of inducing points  $\hat{\mathcal{S}} = \{\hat{\mathbf{s}}_1, \dots, \hat{\mathbf{s}}_k\}$ . This formulation allows for flexibility in selecting inducing points, enabling the use of a different or larger set of inducing points than those employed in the VIF. Such flexibility can lead to a more effective preconditioner compared to relying solely on the inducing points defined by the VIF structure. The construction of the FITC preconditioner requires  $\mathcal{O}(n \cdot k^2)$  operations and linear solves  $\hat{\mathbf{P}}^{-1} \mathbf{w} = \mathbf{D}_V^{-1} \mathbf{w} - \mathbf{D}_V^{-1} \boldsymbol{\Sigma}_{kn}^\top \mathbf{M}_V^{-1} \boldsymbol{\Sigma}_{kn} \mathbf{D}_V^{-1} \mathbf{w}$ , where  $\mathbf{w} \in \mathbb{R}^n$  and  $\mathbf{M}_V = \boldsymbol{\Sigma}_k + \boldsymbol{\Sigma}_{kn} \mathbf{D}_V^{-1} \boldsymbol{\Sigma}_{kn}^\top$ , are computed in  $\mathcal{O}(n \cdot k)$  when  $\mathbf{M}_V$  is precomputed in  $\mathcal{O}(n \cdot k^2)$ . We can compute the

log-determinant and its derivatives using Sylvester's determinant theorem

$$\begin{aligned}\log \det (\widehat{\mathbf{P}}) &= \log \det (\mathbf{D}_V) - \log \det (\boldsymbol{\Sigma}_k) + \log \det (\mathbf{M}_V) \\ \text{Tr} \left( \widehat{\mathbf{P}}^{-1} \frac{\partial \widehat{\mathbf{P}}}{\partial b^*(\mathbf{s}_i)} \right) &= \text{Tr} \left( \mathbf{D}_V^{-1} \frac{\partial \mathbf{W}^{-1}}{\partial b^*(\mathbf{s}_i)} \right) + \text{Tr} \left( \mathbf{M}_V^{-1} \boldsymbol{\Sigma}_{kn} \mathbf{D}_V^{-1} \frac{\partial \mathbf{W}^{-1}}{\partial b^*(\mathbf{s}_i)} \mathbf{D}_V^{-1} \boldsymbol{\Sigma}_{kn}^T \right) \\ \text{Tr} \left( \widehat{\mathbf{P}}^{-1} \frac{\partial \widehat{\mathbf{P}}}{\partial \theta_k} \right) &= \text{Tr} \left( \mathbf{D}_V^{-1} \frac{\partial \mathbf{D}_V}{\partial \theta_k} \right) - \text{Tr} \left( \boldsymbol{\Sigma}_k^{-1} \frac{\partial \boldsymbol{\Sigma}_k}{\partial \theta_k} \right) + \text{Tr} \left( \mathbf{M}_V^{-1} \frac{\partial \mathbf{M}_V}{\partial \theta_k} \right) \\ \text{Tr} \left( \widehat{\mathbf{P}}^{-1} \frac{\partial \widehat{\mathbf{P}}}{\partial \xi_l} \right) &= \text{Tr} \left( \mathbf{D}_V^{-1} \frac{\partial \mathbf{W}^{-1}}{\partial \xi_l} \right) + \text{Tr} \left( \mathbf{M}_V^{-1} \boldsymbol{\Sigma}_{kn} \mathbf{D}_V^{-1} \frac{\partial \mathbf{W}^{-1}}{\partial \xi_l} \mathbf{D}_V^{-1} \boldsymbol{\Sigma}_{kn}^T \right)\end{aligned}$$

in  $\mathcal{O}(n \cdot k^2)$ . Sampling from a normal distribution with mean  $\mathbf{0}$  and covariance  $\widehat{\mathbf{P}}$  can be done by using the reparameterization trick described in Gardner et al. (2018a, Appendix C.1). If  $\boldsymbol{\epsilon}_1$  is a standard normal vector in  $\mathbb{R}^k$  and  $\boldsymbol{\epsilon}_2$  is a standard normal vector in  $\mathbb{R}^n$ , then the expression  $(\mathbf{D}_V^{\frac{1}{2}} \boldsymbol{\epsilon}_2 + \boldsymbol{\Sigma}_{kn}^T \boldsymbol{\Sigma}_k^{-\frac{1}{2}} \boldsymbol{\epsilon}_1)$  is a sample from  $\mathcal{N}(\mathbf{0}, \mathbf{D}_V + \boldsymbol{\Sigma}_{kn}^T \boldsymbol{\Sigma}_k^{-1} \boldsymbol{\Sigma}_{kn})$ . The computational complexity of this procedure is  $\mathcal{O}(n \cdot k)$ .

## Appendix F. Convergence analysis of the preconditioned CG method

**Lemma 7 Error bound for Vecchia approximation of  $\boldsymbol{\Sigma} - \boldsymbol{\Sigma}_{mn}^T \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn}$**

Under Assumptions 1–3, let  $(\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1}$  be a Vecchia approximation with  $m_v$  neighbors for  $\boldsymbol{\Sigma} - \boldsymbol{\Sigma}_{mn}^T \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn}$ , where  $\boldsymbol{\Sigma}_{mn}^T \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn}$  is the rank- $m$  approximation of a covariance matrix  $\boldsymbol{\Sigma}$  with eigenvalues  $\lambda_1 \geq \dots \geq \lambda_n > 0$  as described in Section 2.1. Then, the following inequality holds

$$\|(\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1}\|_2 \leq (\alpha \cdot \lambda_{m+1} \cdot m_v)^n,$$

where  $\alpha > 1$  is a constant.

### Proof

Since  $\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}$  is symmetric positive definite, its singular values coincide with its eigenvalues. Therefore, for the spectral norm,

$$\|(\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1}\|_2 = \frac{1}{\lambda_{\min}(\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})} = \frac{1}{\sigma_{\min}(\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})},$$

where  $\lambda_{\min}$  and  $\sigma_{\min}$  denote the smallest eigenvalue and smallest singular value, respectively.

Furthermore, we apply the lower bound for the smallest singular value of a matrix  $\mathbf{A} \in \mathbb{R}^{n \times n}$  as established by Shun (2021); Yi-Sheng and Dun-He (1997)

$$\sigma_{\min}(\mathbf{A}) \geq |\det \mathbf{A}| \cdot \left( \frac{n-1}{\|\mathbf{A}\|_F^2} \right)^{(n-1)/2},$$

where  $\|\cdot\|_F$  is the Frobenius-norm. Applying this bound, we obtain

$$\begin{aligned}\frac{1}{\sigma_{\min}(\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})} &\leq |\det(\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1}| \cdot \left( \frac{\|\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}\|_F^2}{n-1} \right)^{(n-1)/2} \\ &= |\det(\mathbf{D})| \cdot \left( \frac{\|\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}\|_F^2}{n-1} \right)^{(n-1)/2}.\end{aligned}$$

By Katzfuss and Guinness (2021), we know that, for some constant  $\alpha_0 > 1$ , the matrix  $\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}$  contains at most  $\alpha_0 \cdot n \cdot m_v^2$  non-zero elements and therefore, we obtain

$$\begin{aligned} |\det(\mathbf{D})| \cdot \left( \frac{\|\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}\|_F^2}{n-1} \right)^{(n-1)/2} &\leq |\det(\mathbf{D})| \cdot \left( \frac{\alpha_1 \cdot n \cdot m_v^2}{n-1} \right)^{(n-1)/2} \\ &\leq \alpha_1^n \cdot |\det(\mathbf{D})| \cdot m_v^{n-1} \cdot \left( 1 + \frac{1}{n-1} \right)^{(n-1)/2} \\ &\leq \alpha_1^n \cdot |\det(\mathbf{D})| \cdot m_v^{n-1} \cdot e^1 = \alpha_2^n \cdot |\det(\mathbf{D})| \cdot m_v^n, \end{aligned}$$

where  $\alpha_1 = \alpha_0 \cdot (\max_{i,j=1,\dots,n} [\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B}]_{i,j})^2$  and  $\alpha_2 = \alpha_1 \cdot e^{1/n}$ . Furthermore, we have by Sun et al. (2015)

$$\begin{aligned} |\det(\mathbf{D})| &\leq (\max_{i=1,\dots,n} |D_i|)^n \leq (\max_{i=1,\dots,n} |\Sigma_{ii} - \Sigma_{mi}^T \Sigma_m^{-1} \Sigma_{mi}|)^n \\ &\leq (\lambda_{\max}(\Sigma - \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn}))^n \leq \alpha_3^n \cdot \lambda_{\max}(\Sigma - \widehat{\Sigma}^m)^n \leq \alpha_3^n \cdot \lambda_{m+1}^n, \end{aligned}$$

where  $\alpha_3 > 1$  is a constant,  $\widehat{\Sigma}^m \in \mathbb{R}^{n \times n}$  is the best rank- $m$  approximation of  $\Sigma$  given by the  $m$ -truncated eigenvalue decomposition (see the Eckart-Young-Mirsky theorem for the spectral norm (Eckart and Young, 1936)), and  $\lambda_{\max}$  the largest eigenvalue. Hence, we obtain

$$\|(\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1}\|_2 \leq \alpha_2^n \cdot |\det(\mathbf{D})| \cdot m_v^n \leq (\alpha \cdot \lambda_{m+1} \cdot m_v)^n,$$

where  $\alpha = \alpha_2 \cdot \alpha_3$ . ■

### Lemma 8 Bound for the Condition Number of $\widehat{\mathbf{P}}^{-1}(\Sigma_{\dagger}^{-1} + \mathbf{W})$

Let  $\Sigma_{\dagger} \in \mathbb{R}^{n \times n}$  be a VIF approximation with  $m$  inducing points and  $m_v$  Vecchia neighbors of a covariance matrix  $\Sigma \in \mathbb{R}^{n \times n}$ , where  $\Sigma$  has eigenvalues  $\lambda_1 \geq \dots \geq \lambda_n > 0$ . Furthermore, let  $\widehat{\mathbf{P}} \in \mathbb{R}^{n \times n}$  be the VIFDU preconditioner for  $\Sigma_{\dagger}^{-1} + \mathbf{W}$ . Under Assumptions 1–3, the condition number  $\kappa(\widehat{\mathbf{P}}^{-1}(\Sigma_{\dagger}^{-1} + \mathbf{W}))$  is bounded by

$$\kappa(\widehat{\mathbf{P}}^{-1}(\Sigma_{\dagger}^{-1} + \mathbf{W})) \leq \left( 1 + \alpha^n \cdot (\lambda_1 + (\lambda_{m+1} \cdot m_v)^n) \cdot (\sqrt{n} \cdot m_v + \max(\mathbf{W})) \right)^2,$$

where  $\alpha > 1$  is a constant.

**Proof** We use the identity

$$\kappa(\widehat{\mathbf{P}}^{-1}(\Sigma_{\dagger}^{-1} + \mathbf{W})) = \|\widehat{\mathbf{P}}^{-1}(\Sigma_{\dagger}^{-1} + \mathbf{W})\|_2 \cdot \|(\Sigma_{\dagger}^{-1} + \mathbf{W})^{-1} \widehat{\mathbf{P}}\|_2,$$

For  $\mathbf{E} = \widehat{\mathbf{P}} - (\Sigma_{\dagger}^{-1} + \mathbf{W})$ , we obtain

$$\begin{aligned} \kappa(\widehat{\mathbf{P}}^{-1}(\Sigma_{\dagger}^{-1} + \mathbf{W})) &= \|\widehat{\mathbf{P}}^{-1}(\Sigma_{\dagger}^{-1} + \mathbf{W})\|_2 \cdot \|(\Sigma_{\dagger}^{-1} + \mathbf{W})^{-1} \widehat{\mathbf{P}}\|_2 \\ &= \|\widehat{\mathbf{P}}^{-1}(\widehat{\mathbf{P}} - \mathbf{E})\|_2 \cdot \|(\Sigma_{\dagger}^{-1} + \mathbf{W})^{-1}((\Sigma_{\dagger}^{-1} + \mathbf{W}) + \mathbf{E})\|_2 \\ &= \|\mathbf{I}_n - \widehat{\mathbf{P}}^{-1} \mathbf{E}\|_2 \cdot \|\mathbf{I}_n + (\Sigma_{\dagger}^{-1} + \mathbf{W})^{-1} \mathbf{E}\|_2. \end{aligned}$$

Applying Cauchy-Schwarz and the triangle inequality, we have

$$\kappa(\widehat{\mathbf{P}}^{-1}(\Sigma_{\dagger}^{-1} + \mathbf{W})) \leq (1 + \|\widehat{\mathbf{P}}^{-1}\|_2 \cdot \|\mathbf{E}\|_2) \cdot (1 + \|(\Sigma_{\dagger}^{-1} + \mathbf{W})^{-1}\|_2 \cdot \|\mathbf{E}\|_2).$$

Furthermore, we obtain

$$\|(\Sigma_{\dagger}^{-1} + \mathbf{W})^{-1}\|_2 = \frac{1}{\sigma_{\min}(\Sigma_{\dagger}^{-1} + \mathbf{W})} \leq \frac{1}{\sigma_{\min}(\Sigma_{\dagger}^{-1}) + \min(\mathbf{W})} \leq \frac{1}{\sigma_{\min}(\Sigma_{\dagger}^{-1})} = \|\Sigma_{\dagger}\|_2$$

and

$$\|\widehat{\mathbf{P}}^{-1}\|_2 = \frac{1}{\sigma_{\min}(\boldsymbol{\Sigma}_{\dagger}^{-1} + \mathbf{B}^T \mathbf{W} \mathbf{B})} \leq \frac{1}{\sigma_{\min}(\boldsymbol{\Sigma}_{\dagger}^{-1}) + \sigma_{\min}(\mathbf{B}^T \mathbf{W} \mathbf{B})} \leq \frac{1}{\sigma_{\min}(\boldsymbol{\Sigma}_{\dagger}^{-1})} = \|\boldsymbol{\Sigma}_{\dagger}\|_2.$$

Therefore, we obtain

$$\begin{aligned} \kappa(\widehat{\mathbf{P}}^{-1}(\boldsymbol{\Sigma}_{\dagger}^{-1} + \mathbf{W})) &\leq (1 + \|\boldsymbol{\Sigma}_{\dagger}\|_2 \cdot \|\mathbf{E}\|_2)^2 = (1 + \|\boldsymbol{\Sigma}_{\dagger}\|_2 \cdot \|\mathbf{B}^T \mathbf{W} \mathbf{B} - \mathbf{W}\|_2)^2 \\ &\leq (1 + \|\boldsymbol{\Sigma}_{\dagger}\|_2 \cdot (\|\mathbf{B}^T \mathbf{W} \mathbf{B}\|_2 + \max(\mathbf{W})))^2 \\ &\leq (1 + \|\boldsymbol{\Sigma}_{\dagger}\|_2 \cdot (\|\mathbf{B}^T \mathbf{W} \mathbf{B}\|_F + \max(\mathbf{W})))^2 \\ &\leq \left(1 + \|\boldsymbol{\Sigma}_{\dagger}\|_2 \cdot (\alpha_1 \cdot \sqrt{n} \cdot m_v + \max(\mathbf{W}))\right)^2 \\ &= \left(1 + \|\boldsymbol{\Sigma}_{mn}^T \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn} + (\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1}\|_2 \cdot (\alpha_1 \cdot \sqrt{n} \cdot m_v + \max(\mathbf{W}))\right)^2 \\ &\leq \left(1 + (\|\boldsymbol{\Sigma}_{mn}^T \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn}\|_2 + \|(\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1}\|_2) \cdot (\alpha_1 \cdot \sqrt{n} \cdot m_v + \max(\mathbf{W}))\right)^2 \\ &\leq \left(1 + \left(\|\boldsymbol{\Sigma}\|_2 + (\alpha_2 \cdot \lambda_{m+1} \cdot m_v)^n\right) \cdot (\alpha_1 \cdot \sqrt{n} \cdot m_v + \max(\mathbf{W}))\right)^2 \\ &= \left(1 + \left(\lambda_1 + (\alpha_2 \cdot \lambda_{m+1} \cdot m_v)^n\right) \cdot (\alpha_1 \cdot \sqrt{n} \cdot m_v + \max(\mathbf{W}))\right)^2 \\ &\leq \left(1 + \alpha^n \cdot (\lambda_1 + (\lambda_{m+1} \cdot m_v)^n) \cdot (\sqrt{n} \cdot m_v + \max(\mathbf{W}))\right)^2, \end{aligned}$$

where, in the fourth inequality, we again invoke the result of Katzfuss and Guinness (2021), which guarantees the existence of a constant  $\alpha_0 > 1$  such that the matrix  $\mathbf{B}^T \mathbf{W}^{-1} \mathbf{B}$  contains at most  $\alpha_0 \cdot n \cdot m_v^2$  non-zero entries. This implies  $\alpha_1 = \sqrt{\alpha_0} \max_{i,j=1,\dots,n} |[\mathbf{B}^T \mathbf{W} \mathbf{B}]_{i,j}|$ . Furthermore, in the sixth inequality we apply Lemma 7 with some constant  $\alpha_2 > 1$ . Finally, the constant  $\alpha > 1$  is chosen such that  $\alpha = \max(\alpha_1^{1/n}, \alpha_2, \alpha_1^{1/n} \cdot \alpha_2)$ . ■

**Proof** [Proof of Theorem 5] By using a well-known CG convergence result (Trefethen and Bau, 2022), which bounds the error in terms of the conditioning number, and Lemma 8, we

obtain

$$\begin{aligned}
 \frac{\|\mathbf{u}^* - \mathbf{u}_k\|_{\Sigma_{\dagger}^{-1} + \mathbf{W}}}{\|\mathbf{u}^* - \mathbf{u}_0\|_{\Sigma_{\dagger}^{-1} + \mathbf{W}}} &\leq 2 \cdot \left( \frac{\sqrt{\kappa(\widehat{\mathbf{P}}^{-1}(\Sigma_{\dagger}^{-1} + \mathbf{W}))} - 1}{\sqrt{\kappa(\widehat{\mathbf{P}}^{-1}(\Sigma_{\dagger}^{-1} + \mathbf{W}))} + 1} \right)^k \\
 &\leq 2 \cdot \left( \frac{1 + \alpha^n \cdot (\lambda_1 + (\lambda_{m+1} \cdot m_v)^n) \cdot (\sqrt{n} \cdot m_v + \max(\mathbf{W})) - 1}{1 + \alpha^n \cdot (\lambda_1 + (\lambda_{m+1} \cdot m_v)^n) \cdot (\sqrt{n} \cdot m_v + \max(\mathbf{W})) + 1} \right)^k \\
 &= 2 \cdot \left( \frac{\alpha^n \cdot (\lambda_1 + (\lambda_{m+1} \cdot m_v)^n) \cdot (\sqrt{n} \cdot m_v + \max(\mathbf{W}))}{\alpha^n \cdot (\lambda_1 + (\lambda_{m+1} \cdot m_v)^n) \cdot (\sqrt{n} \cdot m_v + \max(\mathbf{W})) + 2} \right)^k \\
 &\leq 2 \cdot \left( \frac{1}{1 + \left( \alpha^n \cdot (\lambda_1 + (\lambda_{m+1} \cdot m_v)^n) \cdot (\sqrt{n} \cdot m_v + \max(\mathbf{W})) \right)^{-1}} \right)^k. \quad \blacksquare
 \end{aligned}$$

**Lemma 9** *Bound for the Condition Number of  $\widehat{\mathbf{P}}^{-1}(\Sigma_{\dagger} + \mathbf{W}^{-1})$*

Let  $\Sigma_{\dagger} \in \mathbb{R}^{n \times n}$  be a VIF approximation with  $m$  inducing points and  $m_v$  Vecchia neighbors of a covariance matrix  $\Sigma \in \mathbb{R}^{n \times n}$ , where  $\Sigma$  has eigenvalues  $\lambda_1 \geq \dots \geq \lambda_n > 0$ . Furthermore, let  $\widehat{\mathbf{P}} \in \mathbb{R}^{n \times n}$  be the FITC preconditioner for  $\Sigma_{\dagger} + \mathbf{W}^{-1}$  with the same set of inducing points as those used for  $\Sigma_{\dagger}$ . Under Assumptions 1–3, the condition number  $\kappa(\widehat{\mathbf{P}}^{-1}(\Sigma_{\dagger} + \mathbf{W}^{-1}))$  is bounded by

$$\kappa(\widehat{\mathbf{P}}^{-1}(\Sigma_{\dagger} + \mathbf{W}^{-1})) \leq \left( 1 + \max(\mathbf{W}) \cdot (\alpha \cdot \lambda_{m+1} \cdot m_v)^n \right)^2,$$

where  $\alpha > 1$  is a constant.

**Proof** For  $\mathbf{E} = (\Sigma_{\dagger} + \mathbf{W}^{-1}) - \widehat{\mathbf{P}}$ , we obtain

$$\begin{aligned}
 \kappa(\widehat{\mathbf{P}}^{-1}(\Sigma_{\dagger} + \mathbf{W}^{-1})) &= \|\widehat{\mathbf{P}}^{-1}(\Sigma_{\dagger} + \mathbf{W}^{-1})\|_2 \cdot \|(\Sigma_{\dagger} + \mathbf{W}^{-1})^{-1} \widehat{\mathbf{P}}\|_2 \\
 &= \|\widehat{\mathbf{P}}^{-1}(\widehat{\mathbf{P}} + \mathbf{E})\|_2 \cdot \|(\Sigma_{\dagger} + \mathbf{W}^{-1})^{-1}((\Sigma_{\dagger} + \mathbf{W}^{-1}) - \mathbf{E})\|_2 \\
 &= \|\mathbf{I}_n + \widehat{\mathbf{P}}^{-1} \mathbf{E}\|_2 \cdot \|\mathbf{I}_n - (\Sigma_{\dagger} + \mathbf{W}^{-1})^{-1} \mathbf{E}\|_2.
 \end{aligned}$$

Applying Cauchy-Schwarz and the triangle inequality, we have

$$\kappa(\widehat{\mathbf{P}}^{-1}(\Sigma_{\dagger} + \mathbf{W}^{-1})) \leq (1 + \|\widehat{\mathbf{P}}^{-1}\|_2 \cdot \|\mathbf{E}\|_2) \cdot (1 + \|(\Sigma_{\dagger} + \mathbf{W}^{-1})^{-1}\|_2 \cdot \|\mathbf{E}\|_2).$$

Furthermore, we obtain

$$\|(\Sigma_{\dagger} + \mathbf{W}^{-1})^{-1}\|_2 = \frac{1}{\sigma_{\min}(\Sigma_{\dagger} + \mathbf{W}^{-1})} \leq \frac{1}{\sigma_{\min}(\Sigma_{\dagger}) + \min(\mathbf{W}^{-1})} \leq \frac{1}{\min(\mathbf{W}^{-1})} = \max(\mathbf{W})$$

and

$$\begin{aligned}
 \|\widehat{\mathbf{P}}^{-1}\|_2 &= \frac{1}{\sigma_{\min}(\mathbf{D}_I + \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn} + \mathbf{W}^{-1})} \leq \frac{1}{\sigma_{\min}(\mathbf{D}_I + \Sigma_{mn}^T \Sigma_m^{-1} \Sigma_{mn}) + \min(\mathbf{W}^{-1})} \\
 &\leq \frac{1}{\min(\mathbf{W}^{-1})} = \max(\mathbf{W}),
 \end{aligned}$$

where  $\mathbf{D}_I = \text{diag}(\boldsymbol{\Sigma} - \boldsymbol{\Sigma}_{nn}^T \boldsymbol{\Sigma}_m^{-1} \boldsymbol{\Sigma}_{mn})$ . Therefore, we obtain

$$\begin{aligned} \kappa(\widehat{\mathbf{P}}^{-1}(\boldsymbol{\Sigma}_\dagger + \mathbf{W}^{-1})) &\leq (1 + \max(\mathbf{W}) \cdot \|\mathbf{E}\|_2)^2 = (1 + \max(\mathbf{W}) \cdot \|(\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1} - \mathbf{D}_I\|_2)^2 \\ &\leq (1 + \max(\mathbf{W}) \cdot \|(\mathbf{B}^T \mathbf{D}^{-1} \mathbf{B})^{-1}\|_2)^2 \\ &\leq \left(1 + \max(\mathbf{W}) \cdot (\alpha \cdot \lambda_{m+1} \cdot m_v)^n\right)^2, \end{aligned}$$

where we used Lemma 7 with a constant  $\alpha > 1$ . ■

**Proof** [Proof of Theorem 6] By proceeding analogue to the proof of Theorem 5 and by using Lemma 9, we obtain

$$\begin{aligned} \frac{\|\mathbf{u}^* - \mathbf{u}_k\|_{\boldsymbol{\Sigma}_\dagger + \mathbf{W}^{-1}}}{\|\mathbf{u}^* - \mathbf{u}_0\|_{\boldsymbol{\Sigma}_\dagger + \mathbf{W}^{-1}}} &\leq 2 \cdot \left( \frac{\sqrt{\kappa(\widehat{\mathbf{P}}^{-1}(\boldsymbol{\Sigma}_\dagger + \mathbf{W}^{-1}))} - 1}{\sqrt{\kappa(\widehat{\mathbf{P}}^{-1}(\boldsymbol{\Sigma}_\dagger + \mathbf{W}^{-1}))} + 1} \right)^k \\ &\leq 2 \cdot \left( \frac{1 + \max(\mathbf{W}) \cdot (\alpha \cdot \lambda_{m+1} \cdot m_v)^n - 1}{1 + \max(\mathbf{W}) \cdot (\alpha \cdot \lambda_{m+1} \cdot m_v)^n + 1} \right)^k \\ &= 2 \cdot \left( \frac{\max(\mathbf{W}) \cdot (\alpha \cdot \lambda_{m+1} \cdot m_v)^n}{\max(\mathbf{W}) \cdot (\alpha \cdot \lambda_{m+1} \cdot m_v)^n + 2} \right)^k \\ &\leq 2 \cdot \left( \frac{1}{1 + \max(\mathbf{W})^{-1} \cdot (\alpha \cdot \lambda_{m+1} \cdot m_v)^{-n}} \right)^k. \end{aligned}$$
■

## 7. Additional results

|                 | Vecchia   | FITC | VIF  |      | Vecchia           | FITC      | VIF |     |     |
|-----------------|-----------|------|------|------|-------------------|-----------|-----|-----|-----|
| <b>Training</b> | $d = 2$   | 4    | 13   | 24   | <b>Prediction</b> | $d = 2$   | 0.6 | 0.7 | 0.9 |
|                 | $d = 5$   | 8    | 38   | 39   |                   | $d = 5$   | 0.7 | 0.8 | 1.2 |
|                 | $d = 10$  | 16   | 91   | 86   |                   | $d = 10$  | 0.7 | 0.9 | 1.4 |
|                 | $d = 20$  | 35   | 272  | 244  |                   | $d = 20$  | 1.0 | 1.3 | 1.6 |
|                 | $d = 50$  | 129  | 1081 | 811  |                   | $d = 50$  | 1.3 | 1.8 | 2.0 |
|                 | $d = 100$ | 396  | 3179 | 2840 |                   | $d = 100$ | 1.6 | 2.2 | 2.4 |

Table 4: Average runtimes in seconds for the experiments reported in Figure 2 in Section 7.1




Figure 10: RMSE, log-score (LS), and CRPS (mean  $\pm$  2 standard errors) for VIF ( $m_v = 30$  &  $m = 200$ ), FITC ( $m = 200$ ), and Vecchia ( $m_v = 60$ ) approximation for varying dimensions  $d$  using a 3/2-Matérn kernel.




Figure 11: RMSE, log-score (LS), and CRPS (mean) versus runtime (training + prediction) in seconds (s) for VIF, FITC, and Vecchia approximations across varying numbers of inducing points  $m \in \{50, 100, 150, 200, 250, 300, 500, 1000\}$  and Vecchia neighbors  $m_v \in \{10, 20, 30, 50, 100\}$ . For VIF, two parameter ratios are illustrated:  $\frac{m}{m_v} = 5$  (solid) and  $\frac{m}{m_v} = 10$  (dashed). Experiments are conducted with input dimension  $d = 10$  using a Matérn-3/2 kernel, and the corresponding values of  $m$  and  $m_v$  are annotated in the figure.




Figure 12: RMSE, log-score (LS), and CRPS (mean) versus runtime (training + prediction) in seconds (s) for VIF, FITC, and Vecchia approximations across varying numbers of inducing points  $m \in \{50, 100, 150, 200, 250, 300, 500, 1000\}$  and Vecchia neighbors  $m_v \in \{10, 20, 30, 50, 100\}$ . For VIF, two parameter ratios are illustrated:  $\frac{m}{m_v} = 5$  (solid) and  $\frac{m}{m_v} = 10$  (dashed). Experiments are conducted with input dimension  $d = 100$  using a Matérn-3/2 kernel, and the corresponding values of  $m$  and  $m_v$  are annotated in the figure.




Figure 13: RMSE (log-scale), log-score (LS), and CRPS (log-scale) (mean  $\pm 2$  standard errors) for VIF ( $m_v = 30$  &  $m = 200$ ), FITC ( $m = 200$ ), and Vecchia ( $m_v = 30$ ) approximations for 1/2-Matérn, 3/2-Matérn, 5/2-Matérn, and Gaussian ( $\infty$ -Matérn) ARD kernels when  $d = 2$ .

| Kernel:   | Figures 2, 3, and 13          |                               |                               |                               | Figure 14                      |
|-----------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|--------------------------------|
|           | 1/2-Matern                    | 3/2-Matern                    | 5/2-Matern                    | Gaussian                      | 3/2-Matern                     |
| $d = 2$   | $(0.07, 0.30)^T$              | $(0.10, 0.22)^T$              | $(0.12, 0.21)^T$              | $(0.13, 0.19)^T$              | $(0.20, 0.36)^T$               |
| $d = 5$   |                               | $(0.13, 0.473, \dots, 1.5)^T$ |                               |                               | $(0.23, 0.41, \dots, 0.96)^T$  |
| $d = 10$  | $(0.15, 0.389, \dots, 2.3)^T$ | $(0.25, 0.467, \dots, 2.2)^T$ | $(0.27, 0.473, \dots, 2.1)^T$ | $(0.28, 0.471, \dots, 2.0)^T$ | $(0.24, 0.43, \dots, 1.96)^T$  |
| $d = 20$  |                               | $(0.50, 0.763, \dots, 5.5)^T$ |                               |                               | $(0.25, 0.45, \dots, 4.00)^T$  |
| $d = 50$  |                               | $(0.55, 0.661, \dots, 6.0)^T$ |                               |                               | $(0.25, 0.45, \dots, 10.16)^T$ |
| $d = 100$ |                               | $(0.60, 0.665, \dots, 7.0)^T$ |                               |                               | $(0.25, 0.46, \dots, 20.45)^T$ |

Table 5: Length scale parameters  $\lambda$  used in Figures 2, 3, 13, and 14 for various dimensions  $d$  and kernels (“...” means linearly interpolated).




Figure 14: RMSE, log-score (LS), and CRPS (mean  $\pm$  2 standard errors) for VIF ( $m_v = 30$  &  $m = 200$ ), FITC ( $m = 200$ ), and Vecchia ( $m_v = 30$ ) approximations for various dimensions  $d$  using an error variance of 0.01 and length scale parameters chosen such that the covariance remains approximately equal (to the one of a Gaussian kernel with length scales  $\lambda = (0.35, 0.4, 0.45, 0.5, 0.55)^T$ ) at the average distance among two randomly chosen points.

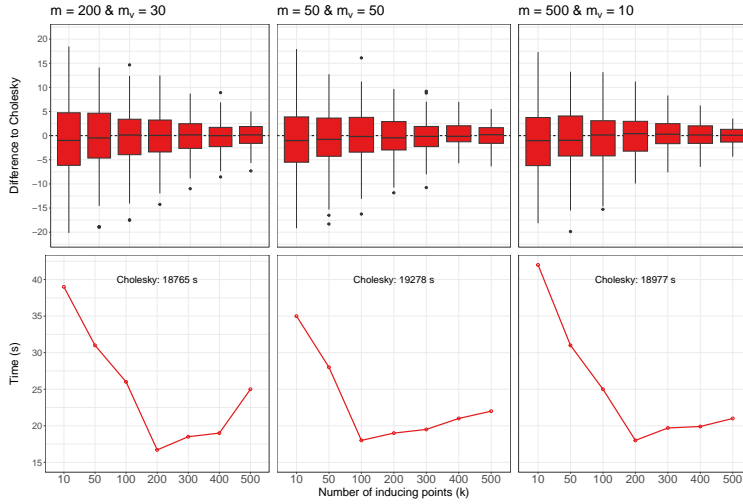


Figure 15: Log-marginal likelihood differences relative to Cholesky-based computations and runtime of the FITC preconditioner for varying numbers of inducing points  $k$  for three different VIF approximations using  $\ell = 50$  sample vectors.

| Preconditioner | CG convergence tolerance ( $\delta$ ) | $\ell = 10$   | $\ell = 50$  | $\ell = 100$ |
|----------------|---------------------------------------|---------------|--------------|--------------|
| <b>FITC</b>    | $\delta = 1$                          | 8.813 (10 s)  | 3.456 (16 s) | 2.626 (20 s) |
|                | $\delta = 0.1$                        | 8.792 (16 s)  | 3.410 (20 s) | 2.610 (23 s) |
|                | $\delta = 0.01$                       | 8.723 (19 s)  | 3.389 (23 s) | 2.602 (26 s) |
|                | $\delta = 0.001$                      | 8.723 (33 s)  | 3.389 (38 s) | 2.600 (41 s) |
|                | $\delta = 0.0001$                     | 8.723 (40 s)  | 3.389 (46 s) | 2.600 (49 s) |
| <b>VIFDU</b>   | $\delta = 1$                          | 12.295 (22 s) | 5.301 (40 s) | 4.887 (72 s) |
|                | $\delta = 0.1$                        | 12.269 (24 s) | 5.263 (45 s) | 4.849 (76 s) |
|                | $\delta = 0.01$                       | 12.234 (27 s) | 5.218 (48 s) | 4.804 (80 s) |
|                | $\delta = 0.001$                      | 12.234 (31 s) | 5.217 (52 s) | 4.803 (85 s) |
|                | $\delta = 0.0001$                     | 12.233 (35 s) | 5.217 (56 s) | 4.803 (89 s) |

Table 6: Accuracy-runtime comparison of preconditioners: RMSE between log-marginal likelihoods computed using iterative methods and those computed using a Cholesky decomposition and runtime for the VIFDU and FITC preconditioners and varying numbers of sample vectors  $\ell \in \{10, 50, 100\}$  and CG convergence tolerances  $\delta \in \{0.0001, 0.001, 0.01, 0.1, 1\}$ . We used a binary likelihood,  $n = 100,000$  samples, and a VIF approximation with  $m = 200$  &  $m_v = 30$ .

| Data   | Accuracy measure | VIF<br>$m_v = 30$<br>$m = 200$ |                       | Data  | Accuracy measure | VIF<br>$m_v = 30$<br>$m = 200$ |                       |
|--|------------------|--------------------------------|-----------------------|---|------------------|--------------------------------|-----------------------|
|  |                  |                                |                       |   |                  |                                |                       |
| <b>3dRoad</b><br>$n = 434,874$<br>$d = 3$    | RMSE             | <b>0.145 ± 0.002</b>           | <b>0.143 ± 0.002</b>  | <b>Kin40K</b><br>$n = 40,000$<br>$d = 8$                  | RMSE             | 0.114 ± 0.002                  | <b>0.111 ± 0.002</b>  |
|  | CRPS             | <b>0.068 ± 0.002</b>           | <b>0.067 ± 0.002</b>  |   | CRPS             | 0.059 ± 0.002                  | <b>0.055 ± 0.002</b>  |
|  | LS               | -0.625 ± 0.010                 | <b>-0.746 ± 0.006</b> |   | LS               | -0.808 ± 0.012                 | <b>-0.885 ± 0.024</b> |
|  | $\nu$            | 1.5                            | 0.696 ± 0.004         |   | $\nu$            | 1.5                            | 7.426 ± 0.842         |
|  | Time             | 376 s                          | 865 s                 |   | Time             | 333 s                          | 1009 s                |
| <b>KEGGU</b><br>$n = 63,608$<br>$d = 26$     | RMSE             | <b>0.094 ± 0.008</b>           | <b>0.097 ± 0.002</b>  | <b>Bike</b><br>$n = 17,379$<br>$d = 12$<br>(Poisson)      | RMSE             | <b>38.904 ± 1.332</b>          | <b>39.065 ± 1.346</b> |
|  | CRPS             | <b>0.030 ± 0.002</b>           | <b>0.028 ± 0.002</b>  |   | CRPS             | <b>17.162 ± 0.454</b>          | <b>17.477 ± 0.381</b> |
|  | LS               | <b>-1.081 ± 0.072</b>          | <b>-1.113 ± 0.056</b> |   | LS               | <b>4.676 ± 0.030</b>           | <b>4.672 ± 0.030</b>  |
|  | $\nu$            | 1.5                            | 2.137 ± 0.950         |   | $\nu$            | 1.5                            | 0.959 ± 0.026         |
|  | Time             | 738 s                          | 1214 s                |   | Time             | 1534 s                         | 2017 s                |
| <b>KEGG</b><br>$n = 48,827$<br>$d = 18$      | RMSE             | <b>0.100 ± 0.010</b>           | <b>0.104 ± 0.012</b>  | <b>House</b><br>$n = 20,640$<br>$d = 8$<br>(Student-t)    | RMSE             | 0.214 ± 0.008                  | <b>0.203 ± 0.006</b>  |
|  | CRPS             | <b>0.041 ± 0.004</b>           | <b>0.040 ± 0.002</b>  |   | CRPS             | 0.103 ± 0.002                  | <b>0.099 ± 0.003</b>  |
|  | LS               | <b>-1.082 ± 0.070</b>          | <b>-1.106 ± 0.034</b> |   | LS               | 0.092 ± 0.076                  | <b>-0.184 ± 0.009</b> |
|  | $\nu$            | 1.5                            | 2.769 ± 0.622         |   | $\nu$            | 1.5                            | 0.469 ± 0.023         |
|  | Time             | 383 s                          | 687 s                 |   | Time             | 1932 s                         | 3021 s                |
| <b>Elevators</b><br>$n = 16,599$<br>$d = 17$ | RMSE             | <b>0.355 ± 0.004</b>           | <b>0.351 ± 0.004</b>  | <b>Power</b><br>$n = 52,417$<br>$d = 5$<br>(Gamma)        | RMSE             | 0.218 ± 0.002                  | <b>0.201 ± 0.002</b>  |
|  | CRPS             | <b>0.196 ± 0.004</b>           | <b>0.194 ± 0.002</b>  |   | CRPS             | 0.121 ± 0.001                  | <b>0.111 ± 0.001</b>  |
|  | LS               | <b>0.388 ± 0.012</b>           | <b>0.389 ± 0.006</b>  |   | LS               | -0.084 ± 0.010                 | <b>-0.140 ± 0.008</b> |
|  | $\nu$            | 1.5                            | 0.912 ± 0.218         |   | $\nu$            | 1.5                            | 0.240 ± 0.018         |
|  | Time             | 711 s                          | 935 s                 |   | Time             | 5035 s                         | 4452 s                |
| <b>Protein</b><br>$n = 45,730$<br>$d = 8$    | RMSE             | 0.516 ± 0.006                  | <b>0.499 ± 0.006</b>  | <b>Water-Vapor</b><br>$n = 100,000$<br>$d = 2$<br>(Gamma) | RMSE             | <b>0.102 ± 0.001</b>           | <b>0.101 ± 0.001</b>  |
|  | CRPS             | 0.259 ± 0.002                  | <b>0.250 ± 0.002</b>  |   | CRPS             | <b>0.056 ± 0.001</b>           | <b>0.055 ± 0.001</b>  |
|  | LS               | 0.667 ± 0.016                  | <b>0.627 ± 0.014</b>  |   | LS               | <b>-0.728 ± 0.022</b>          | -0.699 ± 0.020        |
|  | $\nu$            | 1.5                            | 0.567 ± 0.022         |   | $\nu$            | 1.5                            | 0.383 ± 0.005         |
|  | Time             | 547 s                          | 835 s                 |   | Time             | 3975 s                         | 5084 s                |
| <b>Ailerons</b><br>$n = 13,750$<br>$d = 33$  | RMSE             | 0.399 ± 0.019                  | <b>0.377 ± 0.007</b>  |   |                  |                                |                       |
|  | CRPS             | 0.208 ± 0.008                  | <b>0.200 ± 0.003</b>  |   |                  |                                |                       |
|  | LS               | 0.436 ± 0.030                  | <b>0.375 ± 0.014</b>  |   |                  |                                |                       |
|  | $\nu$            | 1.5                            | 0.674 ± 0.033         |   |                  |                                |                       |
|  | Time             | 528 s                          | 602 s                 |   |                  |                                |                       |

Table 7: RMSE, CRPS, log-score (LS),  $\nu$  (mean  $\pm$  2 standard errors) and runtime in seconds (s) for a fixed 3/2-Matérn kernel (left columns) and Matérn kernel with the smoothness parameter estimated (right columns). Bold indicates the best mean; other means are in bold if within two standard errors.

| Data   | Performance Measure | VIF<br>$m_v = 30$<br>$m = 200$ | Vecchia<br>$m_v = 30$ | FITC<br>$m = 200$ |
|--|---------------------|--------------------------------|-----------------------|-------------------|
| <b>3dRoad</b><br>$n = 434,874$<br>$d = 3$                | RMSE                | <b>0.145 ± 0.002</b>           | <b>0.145 ± 0.001</b>  | 0.554 ± 0.006     |
|  | CRPS                | <b>0.068 ± 0.002</b>           | <b>0.068 ± 0.002</b>  | 0.300 ± 0.003     |
|  | LS                  | <b>-0.625 ± 0.010</b>          | <b>-0.625 ± 0.010</b> | 0.831 ± 0.011     |
|  | Time                | 376 s                          | 45 s                  | 144 s             |
| <b>KEGGU</b><br>$n = 63,608$<br>$d = 26$                 | RMSE                | <b>0.094 ± 0.008</b>           | <b>0.093 ± 0.006</b>  | 0.120 ± 0.023     |
|  | CRPS                | <b>0.030 ± 0.002</b>           | <b>0.030 ± 0.001</b>  | 0.046 ± 0.019     |
|  | LS                  | <b>-1.081 ± 0.072</b>          | <b>-1.060 ± 0.063</b> | -1.052 ± 0.123    |
|  | Time                | 738 s                          | 233 s                 | 628 s             |
| <b>KEGG</b><br>$n = 48,827$<br>$d = 18$                  | RMSE                | <b>0.100 ± 0.010</b>           | <b>0.102 ± 0.005</b>  | NA                |
|  | CRPS                | <b>0.041 ± 0.004</b>           | <b>0.042 ± 0.001</b>  | convergence       |
|  | LS                  | <b>-1.082 ± 0.070</b>          | <b>-1.050 ± 0.021</b> | issues            |
|  | Time                | 383 s                          | 60 s                  |                   |
| <b>Elevators</b><br>$n = 16,599$<br>$d = 17$             | RMSE                | <b>0.355 ± 0.004</b>           | 0.382 ± 0.009         | 0.453 ± 0.011     |
|  | CRPS                | <b>0.196 ± 0.004</b>           | 0.215 ± 0.004         | 0.248 ± 0.012     |
|  | LS                  | <b>0.388 ± 0.012</b>           | 0.431 ± 0.017         | 0.542 ± 0.028     |
|  | Time                | 711 s                          | 37 s                  | 347 s             |
| <b>Protein</b><br>$n = 45,730$<br>$d = 8$                | RMSE                | <b>0.516 ± 0.006</b>           | <b>0.517 ± 0.006</b>  | 0.720 ± 0.004     |
|  | CRPS                | <b>0.259 ± 0.002</b>           | <b>0.259 ± 0.003</b>  | 0.402 ± 0.002     |
|  | LS                  | <b>0.667 ± 0.016</b>           | <b>0.670 ± 0.017</b>  | 1.085 ± 0.005     |
|  | Time                | 547 s                          | 64 s                  | 392 s             |
| <b>Kin40K</b><br>$n = 40,000$<br>$d = 8$                 | RMSE                | <b>0.114 ± 0.002</b>           | 0.140 ± 0.002         | 0.422 ± 0.004     |
|  | CRPS                | <b>0.059 ± 0.002</b>           | 0.076 ± 0.001         | 0.232 ± 0.002     |
|  | LS                  | <b>-0.808 ± 0.012</b>          | -0.552 ± 0.011        | 0.523 ± 0.010     |
|  | Time                | 333 s                          | 153 s                 | 852 s             |
| <b>Ailerons</b><br>$n = 13,750$<br>$d = 33$              | RMSE                | <b>0.399 ± 0.019</b>           | 0.429 ± 0.016         | 0.879 ± 0.240     |
|  | CRPS                | <b>0.208 ± 0.008</b>           | 0.218 ± 0.004         | 0.478 ± 0.131     |
|  | LS                  | <b>0.436 ± 0.030</b>           | 0.472 ± 0.026         | 1.228 ± 0.381     |
|  | Time                | 528 s                          | 108 s                 | 259 s             |
| <b>Bike</b><br>$n = 17,379$<br>$d = 12$<br>(Poisson)     | RMSE                | <b>38.904 ± 1.332</b>          | <b>38.768 ± 1.096</b> | NA                |
|  | CRPS                | <b>17.162 ± 0.454</b>          | <b>17.073 ± 0.390</b> | convergence       |
|  | LS                  | <b>4.676 ± 0.030</b>           | <b>4.686 ± 0.029</b>  | issues            |
|  | Time                | 1534 s                         | 1296 s                |                   |
| <b>House</b><br>$n = 20,640$<br>$d = 8$<br>(Student-t)   | RMSE                | <b>0.214 ± 0.008</b>           | <b>0.213 ± 0.007</b>  | 0.292 ± 0.014     |
|  | CRPS                | <b>0.103 ± 0.002</b>           | <b>0.102 ± 0.002</b>  | 0.159 ± 0.008     |
|  | LS                  | 0.092 ± 0.076                  | <b>0.081 ± 0.074</b>  | 0.395 ± 0.124     |
|  | Time                | 1932 s                         | 2205 s                | 1052 s            |
| <b>Power</b><br>$n = 52,417$<br>$d = 5$<br>(Gamma)       | RMSE                | <b>0.218 ± 0.002</b>           | <b>0.218 ± 0.003</b>  | NA                |
|  | CRPS                | <b>0.121 ± 0.001</b>           | <b>0.122 ± 0.002</b>  | convergence       |
|  | LS                  | <b>-0.084 ± 0.010</b>          | <b>-0.086 ± 0.009</b> | issues            |
|  | Time                | 5035 s                         | 4806 s                |                   |
| <b>WaterVapor</b><br>$n = 100,000$<br>$d = 2$<br>(Gamma) | RMSE                | <b>0.102 ± 0.001</b>           | <b>0.102 ± 0.001</b>  | 0.636 ± 0.011     |
|  | CRPS                | <b>0.056 ± 0.001</b>           | <b>0.056 ± 0.001</b>  | 0.525 ± 0.009     |
|  | LS                  | <b>-0.728 ± 0.022</b>          | <b>-0.728 ± 0.021</b> | 0.123 ± 0.007     |
|  | Time                | 3975 s                         | 2696 s                | 716 s             |

Table 8: RMSE, CRPS, log-score (LS) (mean  $\pm$  2 standard errors), and runtime in seconds (s) for the regression data sets. Bold indicates the best mean; other means are in bold if within two standard errors.

| Data                                      | Performance Measure | VIF<br>$m_v = 30$<br>$m = 200$ | Vecchia<br>$m_v = 30$ | FITC<br>$m = 200$           |
|---|---------------------|--------------------------------|-----------------------|-----------------------------|
| <b>Bank</b><br>$n = 45,211$<br>$d = 16$   | AUC                 | <b>0.806 ± 0.008</b>           | <b>0.800 ± 0.007</b>  | 0.724 ± 0.010               |
|   | RMSE                | <b>0.284 ± 0.004</b>           | <b>0.282 ± 0.005</b>  | 0.340 ± 0.009               |
|   | ACC                 | <b>0.895 ± 0.002</b>           | <b>0.898 ± 0.003</b>  | 0.801 ± 0.009               |
|   | LS                  | <b>0.035 ± 0.016</b>           | <b>0.040 ± 0.011</b>  | 0.067 ± 0.019               |
|   | Time                | 1643 s                         | 798 s                 | 692 s                       |
| <b>Adult</b><br>$n = 48,842$<br>$d = 14$  | AUC                 | <b>0.880 ± 0.008</b>           | <b>0.888 ± 0.009</b>  | NA<br>convergence<br>issues |
|   | RMSE                | <b>0.336 ± 0.006</b>           | <b>0.329 ± 0.007</b>  |                             |
|   | ACC                 | <b>0.838 ± 0.006</b>           | <b>0.843 ± 0.006</b>  |                             |
|   | LS                  | <b>0.003 ± 0.012</b>           | <b>-0.007 ± 0.012</b> |                             |
|   | Time                | 2854 s                         | 805 s                 |                             |
| <b>Credit</b><br>$n = 30,000$<br>$d = 22$ | AUC                 | <b>0.768 ± 0.006</b>           | <b>0.770 ± 0.007</b>  | <b>0.772 ± 0.005</b>        |
|   | RMSE                | <b>0.378 ± 0.004</b>           | <b>0.378 ± 0.005</b>  | <b>0.378 ± 0.005</b>        |
|   | ACC                 | <b>0.808 ± 0.006</b>           | <b>0.807 ± 0.006</b>  | <b>0.807 ± 0.007</b>        |
|   | LS                  | <b>0.407 ± 0.012</b>           | <b>0.402 ± 0.013</b>  | <b>0.400 ± 0.014</b>        |
|   | Time                | 1161 s                         | 518 s                 | 378 s                       |
| <b>MAGIC</b><br>$n = 19,020$<br>$d = 9$   | AUC                 | <b>0.920 ± 0.004</b>           | <b>0.918 ± 0.005</b>  | 0.895 ± 0.008               |
|   | RMSE                | <b>0.316 ± 0.004</b>           | <b>0.313 ± 0.005</b>  | 0.329 ± 0.008               |
|   | ACC                 | <b>0.866 ± 0.006</b>           | <b>0.871 ± 0.006</b>  | 0.830 ± 0.009               |
|   | LS                  | <b>0.006 ± 0.022</b>           | <b>0.012 ± 0.025</b>  | 0.092 ± 0.032               |
|   | Time                | 236 s                          | 170 s                 | 113 s                       |

Table 9: AUC, RMSE (Brier score), ACC, LS (mean ± 2 standard errors), and runtime in seconds (s) for the binary classification data sets. Bold indicates the best mean; other means are in bold if within two standard errors.

| Data   | Accuracy measure | Pure VIF-GP           | Linear regression + VIF-GP | VIF-GPBoost           |
|--|------------------|-----------------------|----------------------------|-----------------------|
| <b>3dRoad</b><br>$n = 434,874$<br>$d = 3$    | RMSE             | <b>0.145 ± 0.002</b>  | <b>0.145 ± 0.002</b>       | <b>0.145 ± 0.002</b>  |
|  | CRPS             | <b>0.068 ± 0.002</b>  | <b>0.068 ± 0.001</b>       | <b>0.068 ± 0.001</b>  |
|  | LS               | <b>-0.625 ± 0.010</b> | <b>-0.626 ± 0.010</b>      | <b>-0.629 ± 0.011</b> |
| <b>KEGGU</b><br>$n = 63,608$<br>$d = 26$     | RMSE             | <b>0.094 ± 0.008</b>  | 0.131 ± 0.009              | <b>0.098 ± 0.009</b>  |
|  | CRPS             | <b>0.030 ± 0.002</b>  | 0.041 ± 0.002              | <b>0.032 ± 0.001</b>  |
|  | LS               | <b>-1.081 ± 0.072</b> | 1.357 ± 1.138              | <b>-1.075 ± 0.068</b> |
| <b>KEGG</b><br>$n = 48,827$<br>$d = 18$      | RMSE             | <b>0.100 ± 0.010</b>  | 0.118 ± 0.020              | <b>0.099 ± 0.009</b>  |
|  | CRPS             | <b>0.041 ± 0.004</b>  | <b>0.045 ± 0.001</b>       | <b>0.043 ± 0.002</b>  |
|  | LS               | <b>-1.082 ± 0.070</b> | -0.979 ± 0.037             | <b>-1.079 ± 0.068</b> |
| <b>Elevators</b><br>$n = 16,599$<br>$d = 17$ | RMSE             | 0.355 ± 0.004         | 0.359 ± 0.005              | <b>0.347 ± 0.004</b>  |
|  | CRPS             | <b>0.196 ± 0.004</b>  | <b>0.199 ± 0.005</b>       | <b>0.196 ± 0.003</b>  |
|  | LS               | 0.388 ± 0.012         | 0.398 ± 0.014              | <b>0.378 ± 0.008</b>  |
| <b>Protein</b><br>$n = 45,730$<br>$d = 8$    | RMSE             | <b>0.516 ± 0.006</b>  | <b>0.512 ± 0.007</b>       | <b>0.513 ± 0.006</b>  |
|  | CRPS             | <b>0.259 ± 0.002</b>  | <b>0.256 ± 0.003</b>       | <b>0.256 ± 0.003</b>  |
|  | LS               | 0.667 ± 0.016         | <b>0.649 ± 0.017</b>       | <b>0.653 ± 0.019</b>  |
| <b>Kin40K</b><br>$n = 40,000$<br>$d = 8$     | RMSE             | <b>0.114 ± 0.002</b>  | 0.115 ± 0.002              | <b>0.112 ± 0.002</b>  |
|  | CRPS             | 0.059 ± 0.002         | 0.059 ± 0.001              | <b>0.057 ± 0.001</b>  |
|  | LS               | -0.808 ± 0.012        | -0.807 ± 0.013             | <b>-0.877 ± 0.014</b> |
| <b>Ailerons</b><br>$n = 13,750$<br>$d = 33$  | RMSE             | <b>0.399 ± 0.019</b>  | <b>0.401 ± 0.017</b>       | <b>0.385 ± 0.016</b>  |
|  | CRPS             | <b>0.208 ± 0.008</b>  | <b>0.209 ± 0.007</b>       | <b>0.205 ± 0.005</b>  |
|  | LS               | <b>0.436 ± 0.030</b>  | <b>0.438 ± 0.026</b>       | <b>0.422 ± 0.022</b>  |

Table 10: RMSE, CRPS, log-score (LS) (mean  $\pm$  2 standard errors) when using non-zero prior mean functions (linear and tree-boosting). Bold indicates the best mean; other means are in bold if within two standard errors.

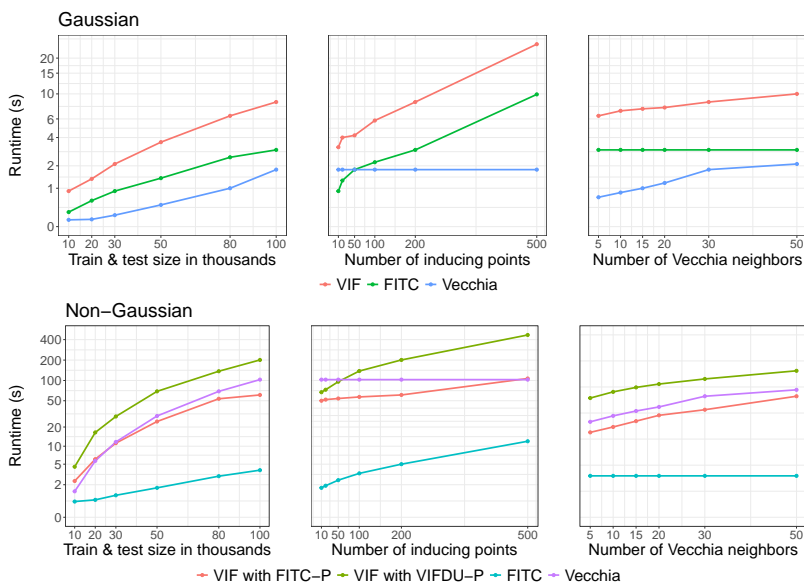


Figure 16: Time (s) for computing predictive distributions with VIF, FITC, and Vecchia approximations on simulated data for varying sample sizes  $n$ , numbers of inducing points  $m$ , and numbers of Vecchia neighbors  $m_v$ .

## References

- David Arthur and Sergei Vassilvitskii. K-means++ the advantages of careful seeding. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1027–1035, 2007.
- Erlend Aune, Daniel P Simpson, and Jo Eidsvik. Parameter estimation in high dimensional Gaussian distributions. *Statistics and Computing*, 24:247–263, 2014.
- Sudipto Banerjee, Alan E Gelfand, Andrew O Finley, and Huiyan Sang. Gaussian predictive process models for large spatial data sets. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 70(4):825–848, 2008.
- Costas Bekas, Effrosyni Kokiopoulou, and Yousef Saad. An estimator for the diagonal of a matrix. *Applied Numerical Mathematics*, 57(11-12):1214–1229, 2007.
- Alina Beygelzimer, Sham Kakade, and John Langford. Cover trees for nearest neighbor. In *Proceedings of the 23rd international conference on Machine learning*, pages 97–104, 2006.
- Jian Cao, Myeongjong Kang, Felix Jimenez, Huiyan Sang, Florian Tobias Schaefer, and Matthias Katzfuss. Variational sparse inverse Cholesky approximation for latent Gaussian processes via double Kullback-Leibler minimization. In *International Conference on Machine Learning*, pages 3559–3576. PMLR, 2023.
- Noel Cressie. *Statistics for spatial data*. John Wiley & Sons, 1993.

- Ryan R Curtin. *Improving dual-tree algorithms*. PhD thesis, Georgia Institute of Technology, Atlanta, GA, USA, 2016.
- Abhirup Datta, Sudipto Banerjee, Andrew O Finley, and Alan E Gelfand. Hierarchical nearest-neighbor Gaussian process models for large geostatistical datasets. *Journal of the American Statistical Association*, 111(514):800–812, 2016.
- Timothy A Davis. *Direct methods for sparse linear systems*. SIAM, 2006.
- Kun Dong, David Eriksson, Hannes Nickisch, David Bindel, and Andrew G Wilson. Scalable log determinants for Gaussian process kernel learning. *Advances in Neural Information Processing Systems*, 30, 2017.
- Carl Eckart and Gale Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936.
- Yury Elkin and Vitaliy Kurlin. A new near-linear time algorithm for k-nearest neighbor search using a compressed cover tree. In *International Conference on Machine Learning*, pages 9267–9311. PMLR, 2023.
- Andrew O Finley, Huiyan Sang, Sudipto Banerjee, and Alan E Gelfand. Improving the performance of predictive process modeling for large datasets. *Computational Statistics & Data Analysis*, 53(8):2873–2884, 2009.
- Roger Fletcher. *Practical methods of optimization*. John Wiley & Sons, 2000.
- Reinhard Furrer, Marc G Genton, and Douglas Nychka. Covariance tapering for interpolation of large spatial datasets. *Journal of Computational and Graphical Statistics*, 15(3):502–523, 2006.
- Jacob Gardner, Geoff Pleiss, Kilian Q Weinberger, David Bindel, and Andrew G Wilson. Gpytorch: Blackbox matrix-matrix Gaussian process inference with gpu acceleration. *Advances in Neural Information Processing Systems*, 31, 2018a.
- Jacob Gardner, Geoff Pleiss, Ruihan Wu, Kilian Weinberger, and Andrew Wilson. Product kernel interpolation for scalable Gaussian processes. In *International Conference on Artificial Intelligence and Statistics*, pages 1407–1416. PMLR, 2018b.
- Joseph Guinness. Gaussian process learning via Fisher scoring of Vecchia’s approximation. *Statistics and Computing*, 31(3):1–8, 2021.
- Tim Gyger, Reinhard Furrer, and Fabio Sigrist. Iterative Methods for Full-Scale Gaussian Process Approximations for Large Spatial Data. *SIAM/ASA Journal on Uncertainty Quantification*, 14(1):142–167, 2026.
- Helmut Harbrecht, Michael Peters, and Reinhold Schneider. On the low-rank approximation by the pivoted Cholesky decomposition. *Applied Numerical Mathematics*, 62(4):428–440, 2012.

- Matthew J Heaton, Abhirup Datta, Andrew O Finley, Reinhard Furrer, Joseph Guinness, Rajarshi Guhaniyogi, Florian Gerber, Robert B Gramacy, Dorit Hammerling, Matthias Katzfuss, et al. A case study competition among methods for analyzing large spatial data. *Journal of Agricultural, Biological and Environmental Statistics*, 24:398–425, 2019.
- James Hensman, Nicolò Fusi, and Neil D Lawrence. Gaussian processes for big data. In *Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence*, pages 282–290, 2013.
- James Hensman, Alexander Matthews, and Zoubin Ghahramani. Scalable variational Gaussian process classification. In *Artificial intelligence and statistics*, pages 351–360. PMLR, 2015.
- Michael F Hutchinson. 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.
- Myeongjong Kang and Matthias Katzfuss. Correlation-based sparse inverse Cholesky factorization for fast Gaussian-process inference. *Statistics and Computing*, 33(3):56, 2023.
- Matthias Katzfuss and Joseph Guinness. A general framework for Vecchia approximations of Gaussian processes. *Statistical Science*, 36(1):124–141, 2021.
- Matthias Katzfuss, Joseph Guinness, Wenlong Gong, and Daniel Zilber. Vecchia approximations of Gaussian-process predictions. *Journal of Agricultural, Biological and Environmental Statistics*, 25:383–414, 2020.
- Pascal Kündig and Fabio Sigrist. Iterative methods for Vecchia-Laplace approximations for latent Gaussian process models. *Journal of the American Statistical Association*, 120(550):1267–1280, 2025.
- Cornelius Lanczos. An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *Journal of research of the National Bureau of Standards*, 45(4):255–282, 1950.
- Christiane Lemieux. Control variates. *Wiley StatsRef: Statistics Reference Online*, pages 1–8, 2014.
- Richard J Lipton, Donald J Rose, and Robert Endre Tarjan. Generalized nested dissection. *SIAM Journal on Numerical Analysis*, 16(2):346–358, 1979.
- Haitao Liu, Yew-Soon Ong, Xiaobo Shen, and Jianfei Cai. When Gaussian process meets big data: A review of scalable GPs. *IEEE Transactions on Neural Networks and Learning Systems*, 31(11):4405–4423, 2020.
- Hannes Nickisch and Carl Edward Rasmussen. Approximations for binary Gaussian process classification. *Journal of Machine Learning Research*, 9(Oct):2035–2078, 2008.
- Geoff Pleiss, Jacob Gardner, Kilian Weinberger, and Andrew Gordon Wilson. Constant-time predictive distributions for Gaussian processes. In *International Conference on Machine Learning*, pages 4114–4123. PMLR, 2018.

- Joaquin Quinonero-Candela and Carl Edward Rasmussen. A unifying view of sparse approximate Gaussian process regression. *The Journal of Machine Learning Research*, 6: 1939–1959, 2005.
- Filippo Rambelli and Fabio Sigrist. An accuracy-runtime trade-off comparison of scalable Gaussian process approximations for spatial data. *Journal of Agricultural, Biological, and Environmental Statistics (in press)*, 2026.
- Carl Edward Rasmussen and Christopher K.I. Williams. *Gaussian processes for machine learning*. MIT Press Cambridge, MA, 2006.
- Yousef Saad. *Iterative methods for sparse linear systems*. SIAM, 2003.
- Huiyan Sang and Jianhua Z Huang. A full scale approximation of covariance functions for large spatial data sets. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 74(1):111–132, 2012.
- Huiyan Sang, Mikyoung Jun, and Jianhua Z Huang. Covariance approximation for large multivariate spatial data sets with an application to multiple climate model errors. *The Annals of Applied Statistics*, pages 2519–2548, 2011.
- Florian Schäfer, Matthias Katzfuss, and Houman Owhadi. Sparse Cholesky factorization by Kullback–Leibler minimization. *SIAM Journal on Scientific Computing*, 43(3):A2019–A2046, 2021a.
- Florian Schäfer, Timothy John Sullivan, and Houman Owhadi. Compression, inversion, and approximate pca of dense kernel matrices at near-linear computational complexity. *Multiscale Modeling & Simulation*, 19(2):688–730, 2021b.
- Xu Shun. Two new lower bounds for the smallest singular value. *arXiv preprint arXiv:2108.01221*, 2021.
- Fabio Sigrist. Gaussian process boosting. *The Journal of Machine Learning Research*, 23(1):10565–10610, 2022a.
- Fabio Sigrist. Latent Gaussian model boosting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2):1894–1905, 2022b.
- Giora Simchoni and Saharon Rosset. Integrating random effects in deep neural networks. *Journal of Machine Learning Research*, 24(156):1–57, 2023.
- Shiliang Sun, Jing Zhao, and Jiang Zhu. A review of nyström methods for large-scale machine learning. *Information Fusion*, 26:36–48, 2015.
- Luke Tierney and Joseph B Kadane. Accurate approximations for posterior moments and marginal densities. *Journal of the American Statistical Association*, 81(393):82–86, 1986.
- Michalis Titsias. Variational learning of inducing variables in sparse Gaussian processes. In *Artificial intelligence and statistics*, pages 567–574. PMLR, 2009.
- Lloyd N Trefethen and David Bau. *Numerical linear algebra*, volume 181. Siam, 2022.

- Shashanka Ubaru, Jie Chen, and Yousef Saad. Fast estimation of  $\text{tr}(f(a))$  via stochastic lanczos quadrature. *SIAM Journal on Matrix Analysis and Applications*, 38(4):1075–1099, 2017.
- Aldo V Vecchia. Estimation and model identification for continuous spatial processes. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 50(2):297–312, 1988.
- Andrew Wilson and Hannes Nickisch. Kernel interpolation for scalable structured Gaussian processes (KISS-GP). In *International conference on machine learning*, pages 1775–1784. PMLR, 2015.
- Yu Yi-Sheng and Gu Dun-He. A note on a lower bound for the smallest singular value. *Linear algebra and its Applications*, 253(1-3):25–38, 1997.
- Bohai Zhang, Huiyan Sang, and Jianhua Z Huang. Smoothed full-scale approximation of Gaussian process models for computation of large spatial data sets. *Statistica Sinica*, 29(4):1711–1737, 2019.