



Vecchia Approximations of Gaussian-Process Predictions

Matthias KATZFUSS, Joseph GUINNESS, Wenlong GONG, and Daniel ZILBER

Gaussian processes (GPs) are highly flexible function estimators used for geospatial analysis, nonparametric regression, and machine learning, but they are computationally infeasible for large datasets. Vecchia approximations of GPs have been used to enable fast evaluation of the likelihood for parameter inference. Here, we study Vecchia approximations of spatial predictions at observed and unobserved locations, including obtaining joint predictive distributions at large sets of locations. We consider a general Vecchia framework for GP predictions, which contains some novel and some existing special cases. We study the accuracy and computational properties of these approaches theoretically and numerically, proving that our new methods exhibit linear computational complexity in the total number of spatial locations. We show that certain choices within the framework can have a strong effect on uncertainty quantification and computational cost, which leads to specific recommendations on which methods are most suitable for various settings. We also apply our methods to a satellite dataset of chlorophyll fluorescence, showing that the new methods are faster or more accurate than existing methods and reduce unrealistic artifacts in prediction maps. Supplementary materials accompanying this paper appear on-line.

Key Words: Computational complexity; Kriging; Large datasets; Sparsity; Spatial statistics.

1. INTRODUCTION

Gaussian processes (GPs) are popular models for functions, time series, and spatial fields, with many application areas such as geospatial analysis (e.g., Banerjee et al. 2004; Cressie and Wikle 2011), nonparametric regression and machine learning (e.g., Rasmussen and Williams 2006), the analysis of computer experiments (e.g., Kennedy and O'Hagan 2001), and Bayesian optimization of expensive functions (Jones et al. 1998) and of the tuning parameters in neural networks (e.g., Snoek et al. 2012). Here, we focus on spatial prediction using GPs. GPs are flexible, interpretable, allow natural probabilistic quantification of

M. Katzfuss (✉) · W. Gong · D. Zilber Department of Statistics, Texas A&M University, College Station, USA (E-mail: katzfuss@gmail.com). J. Guinness Department of Statistics and Data Science, Cornell University, Ithaca, USA. W. Gong Department of Statistics, NC State University, Raleigh, USA.

© 2020 International Biometric Society
Journal of Agricultural, Biological, and Environmental Statistics
<https://doi.org/10.1007/s13253-020-00401-7>

uncertainty, and are thus well-suited for big-data applications in principle. However, direct application of GPs incurs computational cost that is cubic in the data size, which is too expensive for many modern datasets of interest.

To deal with this computational problem, numerous GP approximations or simplifying assumptions have been proposed. These include imposing sparsity on covariance matrices (Furrer et al. 2006; Kaufman et al. 2008; Du et al. 2009), sparsity on precision matrices (Rue and Held 2005; Lindgren et al. 2011; Nychka et al. 2015), composite likelihoods (e.g., Curriero and Lele 1999; Stein et al. 2004; Eidsvik et al. 2014), and low-rank structure (e.g., Higdon 1998; Wikle and Cressie 1999; Quiñonero-Candela and Rasmussen 2005; Banerjee et al. 2008; Cressie and Johannesson 2008; Katzfuss and Cressie 2011; Tzeng and Huang 2018). While low-rank approaches are poorly suited for capturing fine-scale dependence, sparsity-based approaches can generally not guarantee linear scaling in the data or grid size, especially in higher dimensions. Local GP approximations (e.g., Gramacy and Apley 2015) are fast but do not scale well to joint predictions at many locations.

We focus on Vecchia approximations, which obtain a sparse Cholesky factor of the precision matrix by removing conditioning variables in a factorization of the joint density of the GP observations into a product of conditional distributions (Vecchia 1988). This approach has become very popular for likelihood approximations for parameter inference (e.g., Stein et al. 2004; Sun and Stein 2016; Guinness 2018). For the typical setting of GP observations that include additive noise, Katzfuss and Guinness (2019) consider a general Vecchia framework that applies the Vecchia approximation to a vector consisting of both the latent GP realizations and the noisy data. This framework contains many other popular GP approximations as special cases (e.g., Snelson and Ghahramani 2007; Finley et al. 2009; Sang et al. 2011; Datta et al. 2016; Katzfuss 2017; Katzfuss and Gong 2019).

Several authors have also proposed the use of Vecchia approximations for the important task of GP prediction, also referred to as kriging. The approach in Vecchia (1992) for one-at-a-time Vecchia predictions has squared time complexity in the number of observations. Datta et al. (2016) and Finley et al. (2019) proposed Bayesian inference and prediction based on Vecchia-type approximations, which we will discuss and compare to in detail in the present paper. Guinness (2018) considered prediction using conditional expectation and uncertainty quantification using conditional simulations in a Vecchia approach based solely on conditioning on observed variables. This is relatively computationally cheap, but uncertainty measures contain random simulation error, and the observed conditioning might not provide accurate approximations in the presence of noise (cf. Katzfuss and Guinness 2019). Vecchia approximations have also been employed as preconditioners in iterative solvers that are used in prediction (Stroud et al. 2017), but this approach is feasible only if prediction is desired at a small number of locations, or if additional approximations are made (Guinness 2019). The multi-resolution approximation (Katzfuss 2017; Katzfuss and Gong 2019) and related approaches relying on domain partitioning (e.g., Sang et al. 2011; Zhang et al. 2019), shown in Katzfuss and Guinness (2019) to be special cases of Vecchia approximations, also provide fast GP prediction, but they can lead to artifacts along partition boundaries. We will discuss these connections and provide numerical comparisons.

Our article synthesizes and extends the literature on Vecchia approximations of spatial GP predictions, in particular the use of the general Vecchia approximation (Katzfuss and

Guinness 2019) for marginal and joint predictive distributions. Extension of the general Vecchia framework to GP prediction was not considered in Katzfuss and Guinness (2019) and requires consideration of complex issues, including how to order variables and choose conditioning sets to achieve accurate joint predictions at observed and unobserved locations and how to guarantee fast computation of relevant summaries of the joint predictive distribution. Here, we systematically study these issues with regard to the accuracy of the approximations and their computational burden. We introduce novel approaches within the framework, for which we can guarantee sparsity of the matrices necessary for inference, resulting in linear memory and time complexity in the number of data points and predictions for fixed conditioning-set size. Our framework enables systematic discussion, study, and comparison of our new methods and existing approaches, based on which we make specific recommendations about which methods are most suitable in various situations. Our framework is agnostic with respect to the inferential framework, allowing both frequentist and Bayesian inference on potential hyperparameters. We focus on the approximation of predictive distributions conditional on hyperparameters, to avoid confounding with the choice of hyperparameter priors or inference approaches.

Answering scientific questions sometimes requires quantifying the uncertainty of linear combinations or other functions of multiple predictions. For example, climate scientists are interested in global average temperature, hydrologists consider the total rainfall in a catchment area, and carbon-cycle scientists want to infer CO_2 surface fluxes from kriged maps of atmospheric CO_2 concentrations. Joint predictive distributions at a set of prediction locations are required to quantify uncertainties of these spatial averages, totals, or other follow-up or “downstream” analyses. Our article details how to compute prediction variances for linear combinations of predictions under the general Vecchia approximation. We also consider prediction of the latent process at observed locations, which is useful in spatial smoothing and in a Vecchia–Laplace approximation of generalized GPs for non-Gaussian spatial data (Zilber and Katzfuss 2019).

This article is organized as follows. Section 2 reviews GP prediction. In Sect. 3, we introduce general Vecchia approximations of GP prediction. In Sect. 4, we discuss specific methods and study their properties. Sections 5 and 6 provide numerical comparisons using simulated and real data, respectively. We conclude in Sect. 7. Appendices A–E contain details and proofs. A separate Supplementary Material document contains Sections S1–S4 with additional details, plots, comparisons, and a description of another Vecchia prediction method. The proposed methods are implemented in the R package `GPvecchia` (Katzfuss et al. 2020b). Code to reproduce our results is provided with this article.

2. EXACT GAUSSIAN-PROCESS PREDICTION

The process of interest is denoted by $\{y(\mathbf{s}) : \mathbf{s} \in \mathbb{D}\}$, or $y(\cdot)$, on a continuous (i.e., non-gridded) domain $\mathbb{D} \subset \mathbb{R}^d$, $d \in \mathbb{N}^+$. We assume that $y(\cdot) \sim GP(0, K)$ is a Gaussian process (GP) with mean zero and covariance function $K : \mathbb{D} \times \mathbb{D} \rightarrow \mathbb{R}$, which is assumed known up to some parameters. Let $\mathbf{s}_i \in \mathbb{D}$ for $i = 1, \dots, n$ and define the location vector $\mathcal{S} = (\mathbf{s}_1, \dots, \mathbf{s}_n)$. For simplicity, we assume throughout that the locations in \mathcal{S} are unique.

Define $y_i = y(\mathbf{s}_i)$ and the vectors $\mathbf{y} = (y_1, \dots, y_n)$ and $\mathbf{z} = (z_1, \dots, z_n)$. The response variables z_i are noisy versions of latent y_i : $z_i | \mathbf{y} \sim \mathcal{N}(y_i, \tau_i^2)$ independently for all i . Thus, the covariance matrix of \mathbf{y} is $\mathbf{K} = K(\mathcal{S}, \mathcal{S})$, and the covariance matrix of \mathbf{z} is $\mathbf{C} = \mathbf{K} + \mathbf{D}$, where \mathbf{D} is a diagonal matrix containing the noise or nugget variances, $\mathbf{D}_{ii} = \tau_i^2$. Define the index vector $o \subset (1, \dots, n)$ of length $n_O = |o|$ such that the subvector \mathbf{z}_o contains all observed response variables (i.e., the data), and \mathcal{S}_o represents the vector of observed locations. (We use the vector and indexing notation described in Appendix A.) We also define $p = (1, \dots, n) \setminus o$ to be an index vector of length $n_P = |p| = n - n_O$, such that \mathcal{S}_p is the vector of unobserved (prediction) locations (cf. [Le and Zidek 2006](#)). In summary, we have:

| Notation | Terminology |
|--|---|
| $o \subset (1, \dots, n)$ | vector of indices of observed locations |
| $p = (1, \dots, n) \setminus o$ | vector of indices of (unobserved) prediction locations |
| $\mathbf{y}, \mathbf{y}_o, \mathbf{y}_p$ | vectors of latent variables |
| $\mathbf{z}, \mathbf{z}_o, \mathbf{z}_p$ | vectors of response variables |

Inference on unknown parameters $\boldsymbol{\theta}$ in K and τ_i^2 can be carried out based on the multivariate normal likelihood, $f(\mathbf{z}_o) = \mathcal{N}_{n_O}(\mathbf{z}_o | \boldsymbol{\theta}, \mathbf{C}_{oo})$, or approximations thereof.

The goal for prediction is to obtain the posterior predictive distribution of \mathbf{y} via

$$f(\mathbf{y} | \mathbf{z}_o) = \int f(\mathbf{y} | \mathbf{z}_o, \boldsymbol{\theta}) dF(\boldsymbol{\theta} | \mathbf{z}_o). \quad (1)$$

The density $f(\mathbf{y} | \mathbf{z}_o, \boldsymbol{\theta})$ is normal with mean $\boldsymbol{\mu}(\boldsymbol{\theta}) = \mathbf{K}_{\bullet o} \mathbf{C}_{oo}^{-1} \mathbf{z}_o$ and covariance matrix

$$\boldsymbol{\Sigma}(\boldsymbol{\theta}) = \mathbf{K} - \mathbf{K}_{\bullet o} \mathbf{C}_{oo}^{-1} \mathbf{K}_{o \bullet}, \quad (2)$$

where \mathbf{K} and \mathbf{C} implicitly depend on $\boldsymbol{\theta}$ and \bullet denotes the vector of all indices. When using maximum-likelihood estimation, the posterior distribution $F(\boldsymbol{\theta} | \mathbf{z}_o)$ of the parameters is effectively approximated by a point mass at $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}$ in (1), and so $f(\mathbf{y} | \mathbf{z}_o) = \mathcal{N}(\mathbf{y} | \boldsymbol{\mu}(\hat{\boldsymbol{\theta}}), \boldsymbol{\Sigma}(\hat{\boldsymbol{\theta}}))$. For Bayesian inference using MCMC, the parameter posterior in (1) is approximated as discrete uniform on, say, $\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(L)}$, and so $f(\mathbf{y} | \mathbf{z}_o) = (1/L) \sum_l \mathcal{N}(\mathbf{y} | \boldsymbol{\mu}(\boldsymbol{\theta}^{(l)}), \boldsymbol{\Sigma}(\boldsymbol{\theta}^{(l)}))$, which is often further approximated by samples from the summands. Therefore, for both inferential paradigms, GP prediction requires obtaining $f(\mathbf{y} | \mathbf{z}_o, \boldsymbol{\theta}) = \mathcal{N}(\mathbf{y} | \boldsymbol{\mu}(\boldsymbol{\theta}), \boldsymbol{\Sigma}(\boldsymbol{\theta}))$ for particular fixed values of $\boldsymbol{\theta}$. Here and in the following, we will thus suppress dependence on $\boldsymbol{\theta}$ and regard it as fixed, unless stated otherwise. Sometimes (e.g., for cross validation), interest might also be in predicting \mathbf{z}_p , but this is a trivial extension of predicting \mathbf{y}_p , in that $\mathbf{z}_p | \mathbf{z}_o \sim \mathcal{N}(\boldsymbol{\mu}_p, \boldsymbol{\Sigma}_{pp} + \mathbf{D}_{pp})$.

Point predictions at individual locations are functions of the marginal distributions $y_i | \mathbf{z}_o$ only, but quantifying the uncertainty of linear combinations (e.g., spatial averages) and generating posterior simulations require the joint posterior distribution of \mathbf{y} given \mathbf{z}_o . While GP prediction is mathematically straightforward, it can be computationally expensive. The time complexity for obtaining the entire matrix $\boldsymbol{\Sigma}$ in (2) is $\mathcal{O}(n_O^3 + nn_O^2 + n^2 n_O)$, and even just

obtaining its diagonal elements (i.e., the prediction variances) requires $\mathcal{O}(n_O^3 + nn_O^2)$ time. Thus, GP prediction is computationally infeasible for large n_O or n , and approximations or simplifying assumptions are necessary.

3. THE GENERAL VECCHIA FRAMEWORK FOR GP PREDICTION

3.1. DEFINITION OF THE FRAMEWORK

The density of any random vector \mathbf{x} can be factored exactly as $f(\mathbf{x}) = \prod_i f(x_i | x_1, \dots, x_{i-1})$. This motivates a general Vecchia approximation (Katzfuss and Guinness 2019) for GP prediction, which applies Vecchia’s approximation (Vecchia 1988) to the vector $\mathbf{x} = \mathbf{z}_o \cup \mathbf{y}$:

$$\hat{f}(\mathbf{x}) = \prod_{i=1}^{n+n_O} f(x_i | \mathbf{x}_{g(i)}), \quad (3)$$

where $g(i) \subset (1, \dots, i-1)$ is a conditioning index vector of size $|g(i)|$, often formed based on variables with locations nearby in space to x_i . If $g(i) = (1, \dots, i-1)$ for every i , then the exact distribution is recovered: $\hat{f}(\mathbf{x}) = f(\mathbf{x})$. If $|g(i)|$ is bounded by some small integer $m \ll n$, the approximation can lead to enormous computational savings, because only matrices of size $m \times m$ need to be decomposed to evaluate (3). Recent results (Schäfer et al. 2020) indicate that in some settings, the approximation error can be bounded with m increasing only polylogarithmically in n . Because the general Vecchia approximation $\hat{f}(\mathbf{x})$ is a valid probability distribution (e.g., Datta et al., 2016, App. A; Katzfuss and Guinness, 2019, Prop. 1), it can be used for approximating the posterior predictive distribution $f(\mathbf{y} | \mathbf{z}_o)$ by applying the rules of probability to $\hat{f}(\mathbf{x})$ to obtain $\hat{f}(\mathbf{y} | \mathbf{z}_o)$. If predictions at additional locations are desired later, the corresponding realizations of $y(\cdot)$ can be appended to the end of \mathbf{x} for consistency (see Sect. 4.3.4).

The accuracy of a general Vecchia approximation depends on the choice of the ordering of the variables in \mathbf{x} and the specification of the conditioning index vectors $g(i)$. The computational efficiency of the approximation is governed by several factors, including the sparsity of the precision matrix for $\mathbf{y} | \mathbf{z}_o$ and its Cholesky factor. For a given m , Katzfuss and Guinness (2019) showed that there is often a trade-off in conditioning on latent versus response variables, in that it can be more accurate but also more computationally expensive to condition on y_k rather than on z_k . In Sect. 4, we study how ordering and conditioning choices affect both the quality of the approximation and its computational burden, for the purpose of providing practical guidelines for using Vecchia’s approximation for spatial prediction.

3.2. MATRIX REPRESENTATIONS

In this subsection, we introduce matrix notation and recapitulate existing results (e.g., Katzfuss and Guinness 2019). Let \mathbf{Q} be the precision matrix for \mathbf{x} under $\hat{f}(\mathbf{x})$. The joint distribution implied by the approximation in (3) is multivariate normal, $\hat{f}(\mathbf{x}) = \mathcal{N}(\mathbf{0}, \mathbf{Q}^{-1})$. Define $\text{chol}(\mathbf{M})$ to return the (lower-triangular) Cholesky factor of \mathbf{M} , $\text{rev}(\mathbf{M})$ to return the

reverse row–column reordering of \mathbf{M} and $\text{rchol}(\mathbf{M}) = \text{rev}(\text{chol}(\text{rev}(\mathbf{M})))$, which gives the (upper-triangular) upper–lower decomposition for \mathbf{M} . Our notation for the relevant matrices is the following:

| Notation | Note |
|--|--|
| $\ell = \#(\mathbf{y}, \mathbf{x})$ | indices in \mathbf{x} occupied by latent variables \mathbf{y} (i.e., $\mathbf{y} = \mathbf{x}_\ell$) |
| $r = \#(\mathbf{z}_o, \mathbf{x})$ | indices in \mathbf{x} occupied by response variables \mathbf{z}_o (i.e., $\mathbf{z}_o = \mathbf{x}_r$) |
| $\mathbf{U} = \text{rchol}(\mathbf{Q})$ | upper-lower Cholesky decomposition of \mathbf{Q} (i.e., $\mathbf{Q} = \mathbf{U}\mathbf{U}'$) |
| $\mathbf{W} = \mathbf{Q}_{\ell\ell} = \mathbf{U}_{\ell,\bullet}\mathbf{U}'_{\ell,\bullet}$ | posterior precision matrix of \mathbf{y} given \mathbf{z}_o |
| $\mathbf{V} = \text{rchol}(\mathbf{W})$ | upper-lower Cholesky decomposition of \mathbf{W} (i.e., $\mathbf{W} = \mathbf{V}\mathbf{V}'$) |

In practice, there is no need to construct the matrix \mathbf{Q} ; rather, we compute the nonzero entries of \mathbf{U} directly via the methods outlined in Appendix B. From the expressions in Appendix B, it is easy to see that \mathbf{U} is sparse with at most m off-diagonal nonzero entries per column, and \mathbf{U} can be computed in $\mathcal{O}(nm^3)$ time.

3.3. GENERAL VECCHIA PREDICTIONS

The goal for GP prediction is to obtain the posterior predictive distribution of \mathbf{y} given the response \mathbf{z}_o , or desired summaries of this distribution. As explained in Appendix C, it suffices to consider this distribution for certain values of the parameters $\boldsymbol{\theta}$, which we again suppress for notational simplicity. General Vecchia prediction approximates the exact conditional distribution $f(\mathbf{y}|\mathbf{z}_o)$ as implied by the joint distribution $\hat{f}(\mathbf{x})$ in (3) with $\mathbf{x} = \mathbf{y} \cup \mathbf{z}_o$:

$$\hat{f}(\mathbf{y}|\mathbf{z}_o) = \frac{\hat{f}(\mathbf{x})}{\int \hat{f}(\mathbf{x}) d\mathbf{y}} =: \mathcal{N}_n(\boldsymbol{\mu}, \boldsymbol{\Sigma}).$$

Since $\mathbf{W} = \mathbf{Q}_{\ell\ell}$ is the submatrix corresponding to \mathbf{y} of the full precision matrix \mathbf{Q} of $\mathbf{x} = \mathbf{y} \cup \mathbf{z}_o$, it is a well-known property of precision matrices that $\boldsymbol{\Sigma} = \mathbf{W}^{-1}$. While the posterior precision matrix \mathbf{W} is sparse, the covariance matrix $\boldsymbol{\Sigma}$ will generally be a dense $n \times n$ matrix. Thus, it is infeasible to actually compute and store this entire matrix when n is large. However, quantities of interest in the context of prediction can be computed using the general Vecchia approximation as follows:

1. The posterior mean or kriging predictor $\boldsymbol{\mu} = \mathbb{E}(\mathbf{y}|\mathbf{z}_o)$: It is straightforward to show that $\boldsymbol{\mu} = -(\mathbf{V}')^{-1}\mathbf{V}^{-1}\mathbf{U}_{\ell,\bullet}\mathbf{U}'_{r,\bullet}\mathbf{z}_o$ (Katzfuss and Guinness 2019, proof of Prop. 2).
2. The prediction variances $\text{diag}(\boldsymbol{\Sigma}) = (\text{var}(y_1|\mathbf{z}_o), \dots, \text{var}(y_n|\mathbf{z}_o))$: Based on \mathbf{V} , a selected inversion algorithm, also referred to as the Takahashi recursions (Erisman and Tinney 1975; Li et al. 2008; Lin et al. 2011), can be used to compute Σ_{ij} for all pairs i, j with $\mathbf{W}_{ij} \neq 0$. Thus, it also returns the prediction variances Σ_{ii} .
3. The joint posterior distribution of linear combinations (e.g., spatial averages): $\mathbf{H}\mathbf{y}|\mathbf{z}_o \sim \mathcal{N}_k(\mathbf{H}\boldsymbol{\mu}, (\mathbf{V}^{-1}\mathbf{H}')'(\mathbf{V}^{-1}\mathbf{H}'))$, where \mathbf{H} is $k \times n$. As $\mathbf{V}^{-1}\mathbf{H}'$ is generally

Table 1. Summary of considered methods, with details on response-first ordering (RF) given in Sect. 4.1 and on latent-first ordering (LF) in Sect. 4.2.

| Method | Related to | References | cnvrg. | Linear | Recommended |
|----------|----------------------------------|----------------------|--------|--------|----------------------|
| RF-full | | | ✓ | ✓ | 2D, large n_O, n_P |
| RF-stand | Standard Vecchia | Guinness (2018) | ✓ | ✓ | 2D, $n_P \ll n_O$ |
| RF-ind | NNGPR, local kriging | Finley et al. (2019) | ✗ | ✓ | |
| LF-full | NNGP with ref. set \mathcal{S} | Datta et al. (2016) | ✓ | ✗ | 2D, small n_O |
| LF-ind | NNGPC | Finley et al. (2019) | ✗ | ✗ | |
| LF-auto | | | ✓ | ✓ | 1D |

cnvrg.: convergence to exact joint distribution, $\hat{f}(\mathbf{y}|\mathbf{z}_o) \rightarrow f(\mathbf{y}|\mathbf{z}_o)$ as conditioning-set size $m \rightarrow n$; linear: computational complexity guaranteed to be linear in n for fixed m ; NNGP: nearest neighbor Gaussian process; NNGPR: NNGP-response; NNGPC: NNGP-collapsed.

dense, only a moderate k is computationally feasible. The variances of linear combinations can be computed faster, as $\text{diag}(\text{var}(\mathbf{H}\mathbf{y}|\mathbf{z}_o)) = ((\mathbf{V}^{-1}\mathbf{H}') \circ (\mathbf{V}^{-1}\mathbf{H}'))'\mathbf{1}_n$, where \circ denotes element-wise multiplication and $\mathbf{1}_n$ is an n -vector of ones.

- Conditional simulation from the posterior predictive distribution $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$: Draw n i.i.d. samples $a_i \sim \mathcal{N}(0, 1)$ from the standard normal distribution, set $\mathbf{a} = (a_1, \dots, a_n)'$ and $\mathbf{y}^* = \boldsymbol{\mu} + (\mathbf{V}')^{-1}\mathbf{a}$. Then, $\mathbf{y}^* \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$.

All of these tasks require computation of $\mathbf{V} = \text{rchol}(\mathbf{W})$ from \mathbf{U} . The cost of this Cholesky factorization depends on the number of nonzero entries per column in \mathbf{V} . In general, it is crucial for fast predictions for large n that \mathbf{V} is sufficiently sparse. Computational complexity is discussed in Sect. 4.3.3 in more detail.

4. SPECIFIC METHODS AND THEIR PROPERTIES

We now consider two different ordering schemes, response-first and latent-first ordering, within the general Vecchia framework for fast GP prediction, along with several special cases. The methods are summarized in Table 1 and illustrated in Fig. 1. In Section S2, we consider an additional approach based on a third ordering scheme, which is an extension of the sparse general Vecchia likelihood approximation of Katzfuss and Guinness (2019), but this approach is less suitable for GP prediction. The general Vecchia framework allows a coherent, systematic discussion and comparison of the different methods, which has so far been lacking in the literature on scalable GP predictions.

In our framework, the vector \mathcal{S} is obtained based on an ordering of the unordered set of locations $\{\mathbf{s}_1, \dots, \mathbf{s}_n\}$. For most of the methods below, we assume an observed-prediction (OP) restriction, meaning that the observed locations are ordered first and prediction locations last; that is, $o = (1, \dots, n_O)$ and $p = (n_O + 1, \dots, n)$. Unless stated otherwise, we recommend and use a maximum–minimum distance (maxmin) ordering (Guinness 2018; Schäfer et al. 2017) constrained to order the prediction locations last. The latent and response variables are then ordered according to how the locations are ordered, with y_i and z_i corresponding to s_i . After ordering \mathbf{y} and \mathbf{z} according to the locations, we must consider the

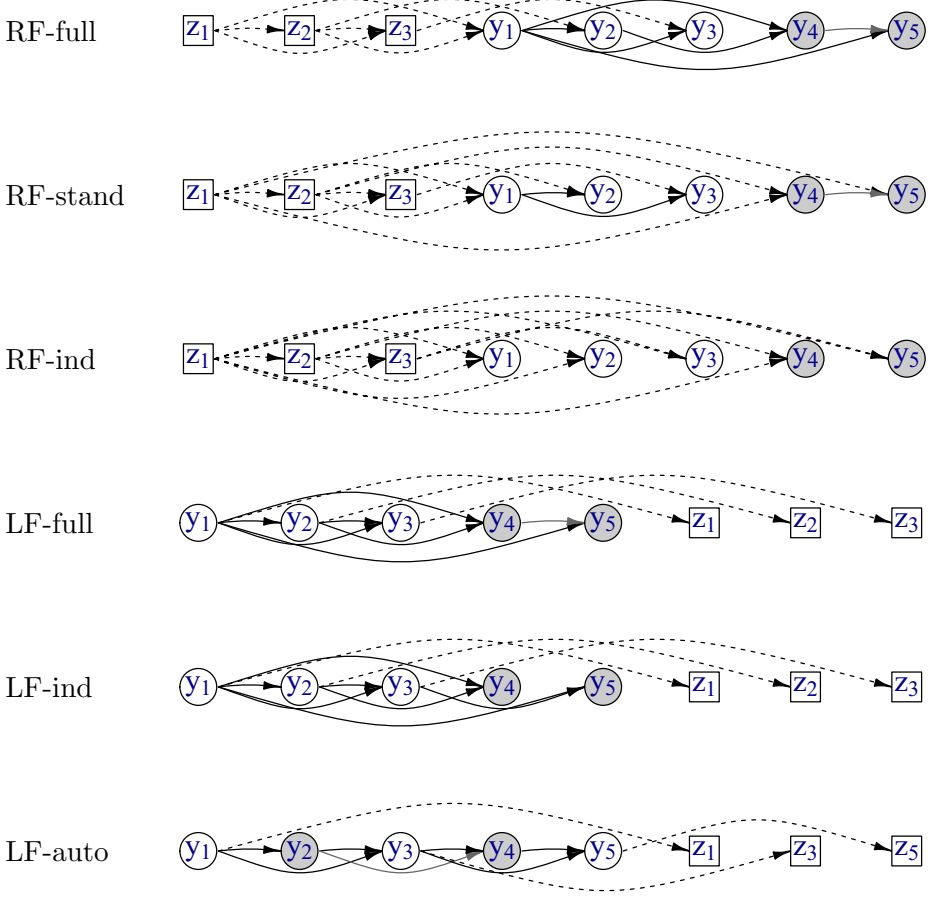


Figure 1. Illustration of the methods in Table 1 as directed acyclic graphs (DAGs) for a toy example with $n = 5$, $n_O = 3$, $n_P = 2$, and conditioning sets of size $m = 2$. Variable ordering is from left to right, and conditioning indicated by arrows. Response variables z_i with $i \in o$ are in squares, with dashed arrows to and from response variables. Latent variables y_i are in circles; y_i with $i \in o$ are unshaded; y_i with $i \in p$ are in grey, as are arrows between them.

separate issue of how to order \mathbf{y} and \mathbf{z}_o within \mathbf{x} . We study the impact that this choice has on approximation accuracy and computational cost.

This section also studies the impact of how the conditioning sets are chosen. The Vecchia approximation assumes that each variable x_i only conditions on variables that are ordered previously in \mathbf{x} . Of those previously ordered variables, we recommend considering those corresponding to the nearest m locations, potentially under additional restrictions. However, if \mathbf{s}_j is one of the nearest m locations, and both y_j and z_j are ordered before x_i , then we only consider one of them, because of conditional independence of x_i and z_j given y_j . Similarly, if y_i is ordered before z_i , then z_i will only condition on y_i , because z_i is conditionally independent of all other variables in \mathbf{x} given y_i .

4.1. RESPONSE-FIRST ORDERING

Response-first ordering means that \mathbf{x} is ordered as $\mathbf{x} = (\mathbf{x}_r, \mathbf{x}_\ell) = (\mathbf{z}_o, \mathbf{y})$, allowing us to rewrite (3) as

$$\hat{f}(\mathbf{x}) = \left(\prod_{i=1}^n f(y_i | \mathbf{y}_{q_y(i)}, \mathbf{z}_{q_z(i)}) \right) \left(\prod_{i \in o} f(z_i | \mathbf{z}_{g(i)}) \right), \quad (4)$$

where $q_y(i)$ and $q_z(i)$ are index sets implied by $g(j)$ for $y_i = \mathbf{x}_j$; for example, if $k \in g(j)$ and $\mathbf{x}_k = z_l$, then $l \in q_z(i)$. Under response-first ordering, $\mathbf{V} = \mathbf{U}_{\ell\ell}$ is simply a submatrix of \mathbf{U} . To see this, note that \mathbf{U} can be written in block form as

$$\mathbf{U} = \begin{bmatrix} \mathbf{U}_{rr} & \mathbf{U}_{r\ell} \\ \mathbf{U}_{\ell r} & \mathbf{U}_{\ell\ell} \end{bmatrix} = \begin{bmatrix} \mathbf{U}_{rr} & \mathbf{U}_{r\ell} \\ \mathbf{0} & \mathbf{U}_{\ell\ell} \end{bmatrix}, \quad (5)$$

where \mathbf{U} , and hence, $\mathbf{U}_{\ell\ell}$ are upper-triangular. Therefore, $\mathbf{W} = \mathbf{U}_{\ell,\bullet} \mathbf{U}'_{\ell,\bullet} = \mathbf{U}_{\ell\ell} \mathbf{U}'_{\ell\ell}$, and so $\mathbf{V} = \text{rchol}(\mathbf{W}) = \mathbf{U}_{\ell\ell}$. Thus, after constructing \mathbf{U} , no additional computation is required to obtain \mathbf{V} , so predictions can be computed in linear time. It is important to use this result to fill the entries of \mathbf{V} directly, rather than forming \mathbf{W} and factoring it. The latter approach can lead to a large number of numerical nonzeros in \mathbf{V} , which are symbolic nonzero entries in \mathbf{V} that are zero in theory but nonzero in practice due to numerical errors. These numerical nonzeros are illustrated in Figure S1 and described in detail in Appendix D.

For the following three methods, we consider response-first ordering under OP restriction (see beginning of Sect. 4), meaning that $\mathbf{y} = (\mathbf{y}_o, \mathbf{y}_p)$, and so $\mathbf{x} = (\mathbf{z}_o, \mathbf{y}_o, \mathbf{y}_p)$.

Response-first ordering, full conditioning (RF-full) This scheme is labeled as “full” because we allow every variable to condition on any variables ordered previously in \mathbf{x} . In (4), the conditioning vectors $\mathbf{y}_{q_y(i)}$ and $\mathbf{z}_{q_z(i)}$ are chosen as the m variables closest in space to y_i , among those that are previously ordered in \mathbf{x} , conditioning on the latent y_j instead of the response z_j whenever possible. Specifically, we set $q(i)$ to consist of the indices corresponding to the m nearest locations to \mathbf{s}_i , including i for $i \in o$, and not including i for $i \in p$. Then, for any $j \in q(i)$, we let y_i condition on y_j if it is ordered previously in \mathbf{x} , and condition on z_j otherwise. More precisely, we set $q_y(i) = \{j \in q(i) : j < i\}$ and $q_z(i) = \{j \in q(i) : j \geq i\}$.

Response-first ordering, standard conditioning (RF-stand) This scheme is identical to RF-full except that \mathbf{y}_p conditions only on \mathbf{y}_p and \mathbf{z}_o , not on \mathbf{y}_o . More precisely, we use the same $q(i)$ as in RF-full, but then set $q_y(i) = \{j \in q(i) : j < i; j \in p\}$ and $q_z(i) = q(i) \setminus q_y(i)$. This approach is labeled as “standard” because the posterior mean and conditional simulations of \mathbf{y}_p from this model are equivalent to those obtained in Guinness (2018), which used the standard Vecchia approximation.

RF-stand is computationally useful if only predictions of \mathbf{y}_p (not of \mathbf{y}_o) are desired, because \mathbf{y}_o can be removed from \mathbf{x} without changing the predictions of \mathbf{y}_p . Specifically, we then have $\mathbf{U}_{\ell_o \ell_p} = \mathbf{0}$ from (9). It can be shown that $\mathbf{W} = \text{blockdiag}(\mathbf{W}_{oo}, \mathbf{W}_{pp})$ and $\mathbf{V} = \text{blockdiag}(\mathbf{V}_{oo}, \mathbf{V}_{pp})$ are both block diagonal with $\mathbf{V}_{pp} = \mathbf{U}_{\ell_p \ell_p}$, and $\boldsymbol{\mu}_p = (\mathbf{U}'_{\ell_p \ell_p})^{-1} (\mathbf{U}_{\ell_o \ell_p})' \mathbf{z}_o$. Thus, when only prediction at unobserved locations \mathcal{S}_p is desired,

the prediction tasks laid out in Sect. 3.3 can be carried out solely based on $\mathbf{U}_{\bullet\ell_p}$, which is the submatrix formed by the last n_p columns of \mathbf{U} corresponding to \mathbf{y}_p . That is, the first $2n_o$ columns of \mathbf{U} corresponding to \mathbf{y}_o and \mathbf{z}_o would then not be required for prediction, resulting in a prediction complexity that depends on n_p , not on $n = n_o + n_p$ (once the $q(i)$ have been determined). This computational simplification comes at the price of some loss of accuracy relative to RF-full because of the restriction of the conditioning sets (see Proposition 2 and Section S3).

Response-first ordering, independent conditioning (RF-ind) RF-ind also uses the ordering $(\mathbf{z}_o, \mathbf{y}_o, \mathbf{y}_p)$, but we enforce that latent variables condition only on \mathbf{z}_o , and so $q_y(i) = \emptyset$ for all $i = 1, \dots, n$ in (4). Then, $q_z(i)$ consists of the indices corresponding to the m nearest observed locations to \mathbf{s}_i among S_o , including \mathbf{s}_i if $i \in o$. RF-ind ignores any posterior dependence between entries of \mathbf{y} . More precisely, note that, from (9), $\mathbf{U}_{\ell\ell}$ and hence \mathbf{W} are now diagonal, and so:

$$\Sigma = \mathbf{W}^{-1} = \text{diag}(\{\mathbf{U}_{\ell_i\ell_i}^2 : i = 1, \dots, n\})^{-1} = \text{diag}(\{\text{var}(\mathbf{y}_i | \mathbf{z}_{q_z(i)}) : i = 1, \dots, n\}).$$

Similarly, it is straightforward to show that $\mu_i = \mathbb{E}(y_i | \mathbf{z}_{q_z(i)})$. RF-ind is equivalent to local kriging, in which each marginal predictive distribution is obtained by considering the conditional distribution of y_i given the neighboring observations from \mathbf{z}_o . This is implicitly the same predictions as in the NNGP-response model in Finley et al. (2019), except that we here predict \mathbf{y}_p instead of predicting \mathbf{z}_p as in the NNGPR. The latter can be easily achieved by adding the noise or nugget variance to each prediction variance (see end of Sect. 2). RF-ind can, in principle, be extremely fast in parallel computing environments because each conditional mean and variance can be calculated completely in parallel.

4.2. LATENT-FIRST ORDERING

Latent-first ordering means that \mathbf{x} is ordered as $\mathbf{x} = (\mathbf{y}, \mathbf{z}_o)$, resulting in the approximation

$$\hat{f}(\mathbf{x}) = \left(\prod_{i=1}^n f(y_i | \mathbf{y}_{q_y(i)}) \right) \left(\prod_{i \in o} f(z_i | y_i) \right), \quad (6)$$

where z_i conditions only on y_i because of conditional independence from all other variables in the model.

Under latent-first ordering, \mathbf{V} cannot be obtained directly as a submatrix of \mathbf{U} as in response-first ordering. However, under the OP restriction that \mathbf{y} is ordered as $\mathbf{y} = (\mathbf{y}_o, \mathbf{y}_p)$, \mathbf{U} does contain some exploitable structure. Note that under this ordering, we have $\ell_o = o$, $\ell_p = p$, and $\mathbf{U}_{pr} = \mathbf{0}$. Hence,

$$\mathbf{U} = \begin{bmatrix} \mathbf{U}_{oo} & \mathbf{U}_{op} & \mathbf{U}_{or} \\ \mathbf{0} & \mathbf{U}_{pp} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{U}_{rr} \end{bmatrix}, \quad \mathbf{W} = \begin{bmatrix} \mathbf{W}_{oo} & \mathbf{U}_{op}\mathbf{U}'_{pp} \\ \mathbf{U}_{pp}\mathbf{U}'_{op} & \mathbf{U}_{pp}\mathbf{U}'_{pp} \end{bmatrix}, \quad \mathbf{V} = \begin{bmatrix} \mathbf{V}_{oo} & \mathbf{U}_{op} \\ \mathbf{0} & \mathbf{U}_{pp} \end{bmatrix}, \quad (7)$$

where only the entries of $\mathbf{V}_{oo} = \text{rchol}(\mathbf{W}_{oo})$ must be computed from $\mathbf{W}_{oo} = \mathbf{U}_{oo}\mathbf{U}'_{oo} + \mathbf{U}_{or}\mathbf{U}'_{or}$, and the last n_p columns of \mathbf{V} corresponding to \mathbf{y}_p can simply be “copied” from

U. This result can make latent-first orderings computationally competitive when there are many more prediction locations than observation locations, for example when predictions are required on a fine grid from a small number of observations.

Latent-first ordering, full conditioning (LF-full) We consider latent-first ordering under OP restriction: $\mathbf{x} = (\mathbf{y}_o, \mathbf{y}_p, \mathbf{z}_o)$. As in RF-full, LF-full scheme is labeled as “full” because we allow every variable to condition on any variables ordered previously in \mathbf{x} . Specifically, in (6), the latent conditioning vector $\mathbf{y}_{q_y(i)}$ simply consists of the m latent variables y_j with $j < i$ whose locations are closest in space to \mathbf{s}_i .

We have found that LF-full is usually the most accurate scheme in the examples we have tried; however, it can also be the most computationally demanding. As shown in (7), only parts of \mathbf{V} can be recovered directly from \mathbf{U} , and linear computational complexity cannot be guaranteed in general; see Sect. 4.3.3 for more details.

LF-full can be viewed as a special version of the NNGP in Datta et al. (2016), in which their reference set is chosen as \mathcal{S} , the vector of observed and prediction locations.

Latent-first ordering, independent conditioning (LF-ind) This scheme is a special case of LF-full with $q_y(i) \subset o$, and hence, there is no conditioning on variables at prediction locations: $q_y(i) \cap p = \emptyset$. This assumption of conditional independence of the entries of \mathbf{y}_p given \mathbf{y}_o can lead to inaccurate approximations of the joint predictive distribution at a set of locations (see, e.g., top row in Fig. 5). As with LF-full, linear computational complexity cannot be guaranteed because the submatrix of \mathbf{V} corresponding to \mathbf{y}_o cannot be obtained directly from \mathbf{U} . LF-ind is implicitly the same approximation as used in the NNGP-collapsed model in Finley et al. (2019).

Multi-resolution approximation (MRA) Predictions using the MRA (Katzfuss 2017; Katzfuss and Gong 2019) can be viewed as a version of LF-full, except that the $q_y(i)$ are based on iterative domain partitioning (cf. Katzfuss and Guinness 2019, Sect. 3.7). In contrast to the nearest-neighbor conditioning in LF-full, MRA conditioning ensures sparsity and linear complexity, which can be shown using Katzfuss and Guinness (2019, Prop. 6). While RF-full can be more accurate than MRA for a given m (Section S3), the special conditioning structure of the MRA has many other benefits (Jurek and Katzfuss 2018). For example, \mathbf{V}^{-1} has the same sparsity as \mathbf{V} for the MRA, and so $\mathbf{V}^{-1}\mathbf{H}'$ is generally sparse if \mathbf{H} is sparse. This allows computing the posterior covariance matrix of a large number of linear combinations $\mathbf{H}\mathbf{y}$.

Latent-first and coordinate ordering, autoregressive conditioning (LF-auto) Finally, we consider a method that is based on a general ordering of the locations, without OP restriction. We consider latent-first ordering, $\mathbf{x} = (\mathbf{y}, \mathbf{z}_o)$, and the approximation in (6), where each y_i simply conditions on $(y_{i-m}, \dots, y_{i-1})$; more precisely, we set $q_y(i) = \{\max(1, i - m), \dots, i - 1\}$. This amounts to a latent autoregressive process of order m . It is easy to verify that \mathbf{W} is banded with bandwidth m , and so its Cholesky factor \mathbf{V} can be obtained in $\mathcal{O}(nm^2)$ time.

LF-auto is most appropriate when successive locations in \mathcal{S} (and hence variables in \mathbf{x}) tend to be close in space, so that $\mathbf{y}_{q_y(i)}$ has strong correlation with y_i . This is often the case with coordinate-based ordering, especially in one-dimensional space. Therefore, we only consider LF-auto based on left-to-right ordering in one-dimensional domains—see Fig. 2 for an illustration. GPs are often used to model functions in one dimension, for example in astronomy (e.g., Wang et al. 2012; Kelly et al. 2014; Foreman-Mackey et al. 2017).

4.3. SUMMARY AND PROPERTIES

We now discuss some of the properties of the methods summarized in Table 1.

4.3.1. Zero Noise

In the case of zero noise or nugget (i.e., $\tau_i = 0$ for all $i = 1, \dots, n$), we have $\mathbf{y}_o = \mathbf{z}_o$, and so there is no real distinction between RF-full, RF-stand, and LF-full anymore. Similarly, LF-ind and RF-ind then become equivalent.

4.3.2. Approximation Accuracy

Each of our methods produces a prediction $\hat{f}(\mathbf{y}|\mathbf{z}_o)$, which can be viewed as an approximation of the exact GP prediction $f(\mathbf{y}|\mathbf{z}_o)$, or as a valid and exact conditional distribution implied by the multivariate normal Vecchia density $\hat{f}(\mathbf{x})$ in (3). We now discuss the Vecchia approximation error from the former perspective, in terms of the (conditional) Kullback–Leibler (KL) divergence to the exact distribution f . The KL divergence is the expected difference in log-likelihood:

$$\text{KL}(f(\mathbf{x})\|\hat{f}(\mathbf{x})) = \mathbb{E}^{\mathbf{x}} \log f(\mathbf{x}) - \mathbb{E}^{\mathbf{x}} \log \hat{f}(\mathbf{x}) = \int f(\mathbf{x}) \log (f(\mathbf{x})/\hat{f}(\mathbf{x})) d\mathbf{x}.$$

Similarly, for generic random vectors \mathbf{z} and \mathbf{y} , we define the conditional KL (CKL) divergence as

$$\begin{aligned} \text{CKL}(f(\mathbf{y}|\mathbf{z})\|\hat{f}(\mathbf{y}|\mathbf{z})) &= \mathbb{E}^{\mathbf{z}} \text{KL}(f(\mathbf{y}|\mathbf{z})\|\hat{f}(\mathbf{y}|\mathbf{z})) \\ &= \int f(\mathbf{z}) \int f(\mathbf{y}|\mathbf{z}) \log (f(\mathbf{y}|\mathbf{z})/\hat{f}(\mathbf{y}|\mathbf{z})) d\mathbf{y} d\mathbf{z}, \end{aligned}$$

which is the KL divergence for the conditional distribution of \mathbf{y} given \mathbf{z} , averaged over possible realizations of \mathbf{z} .

Proposition 1. *Let $\mathbf{x} = (\mathbf{x}^{(1)}, \mathbf{x}^{(2)})$ be a multivariate normal random vector, with $\mathbf{x}^{(1)} = (x_1, \dots, x_{n_1})$ and $\mathbf{x}^{(2)} = (x_{n_1+1}, \dots, x_{n_1+n_2})$. Further, let $\hat{f}_1(\mathbf{x})$ and $\hat{f}_2(\mathbf{x})$ be two Vecchia approximations of the form (3), with conditioning vectors $g_1(i)$ and $g_2(i)$, respectively.*

1. *If $g_1(i) \subset g_2(i)$ for all $i = 1, \dots, n_1 + n_2$, then $\text{KL}(f(\mathbf{x})\|\hat{f}_1(\mathbf{x})) \geq \text{KL}(f(\mathbf{x})\|\hat{f}_2(\mathbf{x}))$.*
2. *If $g_1(i) \subset g_2(i)$ for all $i = n_1 + 1, \dots, n_1 + n_2$, then $\text{CKL}(f(\mathbf{x}^{(2)}|\mathbf{x}^{(1)})\|\hat{f}_1(\mathbf{x}^{(2)}|\mathbf{x}^{(1)})) \geq \text{CKL}(f(\mathbf{x}^{(2)}|\mathbf{x}^{(1)})\|\hat{f}_2(\mathbf{x}^{(2)}|\mathbf{x}^{(1)}))$.*

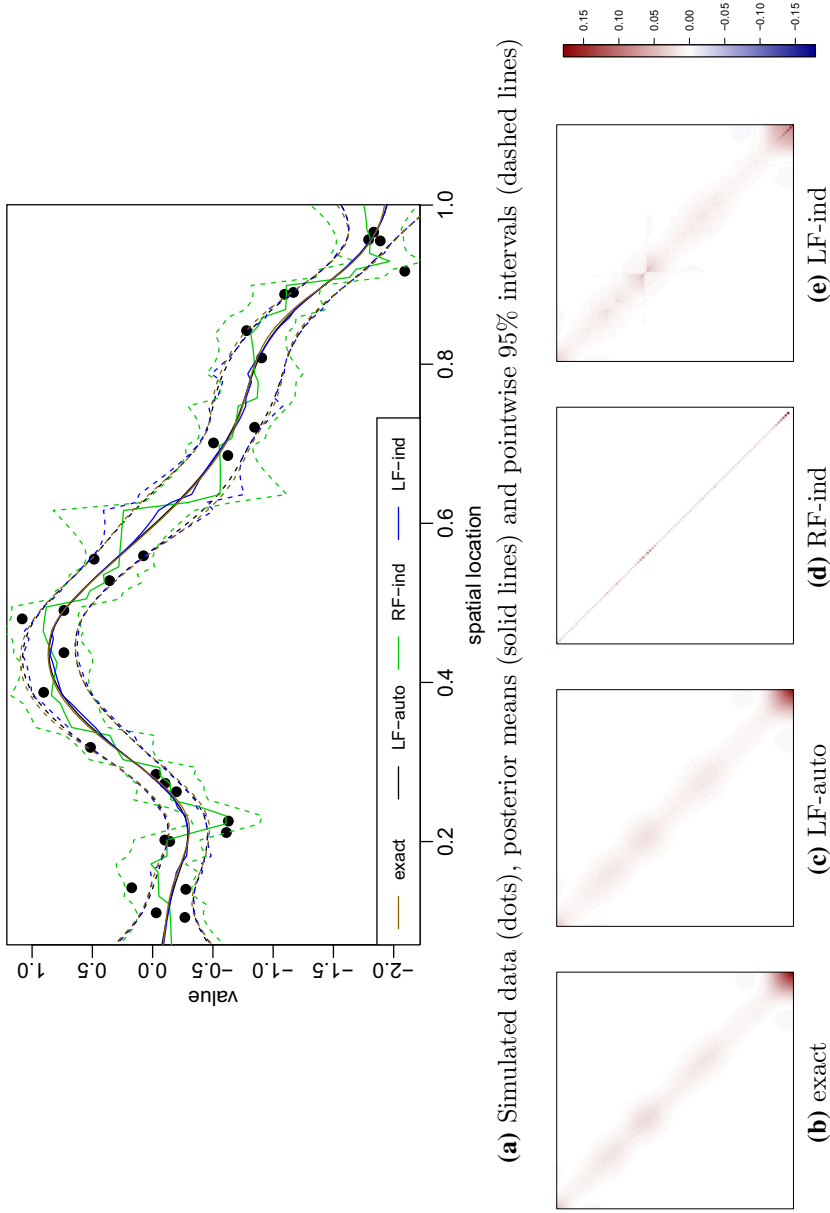


Figure 2. Illustration of predictions at $n_P = 100$ locations based on $n_O = 30$ simulated data on $\mathbb{D} = [0, 1]$ for a GP with Matérn covariance with smoothness 1.5 using $m = 2$ and coordinate (left-to-right) ordering. **a** LF-auto lines are largely covered up by exact lines. **b–e** Posterior predictive covariance matrices Σ_{pp} .

Using these properties, the following can be said about the approximation accuracy of the general Vecchia prediction methods in Table 1:

Proposition 2. *Let $\text{KL}_m^A(\mathbf{x}) = \text{KL}(f(\mathbf{x}) \parallel \hat{f}(\mathbf{x}))$, where $\hat{f}(\mathbf{x})$ is the approximate density obtained using Vecchia approach A (e.g., $A = \text{RF-full}$) with conditioning vectors of size m , and similarly for CKL_m^A .*

1. $\text{KL}_{n-1}^A(\mathbf{x}) = \text{CKL}_{n-1}^A(\mathbf{y}|\mathbf{z}_o) = \text{CKL}_{n-1}^A(\mathbf{y}_p|\mathbf{z}_o) = 0$ for $A \in \{\text{RF-full}, \text{RF-stand}, \text{LF-full}, \text{LF-auto}\}$
2. $\text{KL}_{m+1}^A(\mathbf{x}) \leq \text{KL}_m^A(\mathbf{x})$, for all A in Table 1
3. $\text{CKL}_{m+1}^A(\mathbf{y}_p|\mathbf{y}_o, \mathbf{z}_o) \leq \text{CKL}_m^A(\mathbf{y}_p|\mathbf{y}_o, \mathbf{z}_o)$ for $A \in \{\text{RF-full}, \text{RF-stand}, \text{RF-ind}, \text{LF-full}, \text{LF-ind}\}$
4. $\text{CKL}_{m+1}^A(\mathbf{y}|\mathbf{z}_o) \leq \text{CKL}_m^A(\mathbf{y}|\mathbf{z}_o)$ for $A \in \{\text{RF-full}, \text{RF-stand}, \text{RF-ind}\}$
5. $\text{CKL}_{m+1}^A(\mathbf{y}_p|\mathbf{z}_o) \leq \text{CKL}_m^A(\mathbf{y}_p|\mathbf{z}_o)$ for $A \in \{\text{RF-stand}, \text{RF-ind}\}$
6. $\text{KL}_m^{\text{RF-full}}(\mathbf{x}) \leq \text{KL}_m^{\text{RF-stand}}(\mathbf{x})$ and $\text{CKL}_m^{\text{RF-full}}(\mathbf{y}|\mathbf{z}_o) \leq \text{CKL}_m^{\text{RF-stand}}(\mathbf{y}|\mathbf{z}_o)$
7. For a GP with exponential covariance function in one dimension: $\text{KL}_1^{\text{LF-auto}}(\mathbf{x}) = 0$

To summarize, the Vecchia approximations tend to become more accurate as the conditioning-set size m increases. For all approaches in Table 1, this increase in accuracy is in terms of the KL divergence for $f(\mathbf{x})$, but for certain methods the same is also guaranteed for distributions such as $f(\mathbf{y}|\mathbf{z}_o)$ or $f(\mathbf{y}_p|\mathbf{z}_o)$ that are more relevant for prediction. Indeed, we have $\hat{f}(\mathbf{y}|\mathbf{z}_o) = f(\mathbf{y}|\mathbf{z}_o)$ in the limit of $m = n - 1$ for most methods. However, for RF-ind and LF-ind, due to the assumption of conditional independence of the entries in \mathbf{y}_p given \mathbf{z}_o and \mathbf{y}_o , this holds only marginally, in the sense that $\hat{f}(\mathbf{y}_i|\mathbf{z}_o) = f(\mathbf{y}_i|\mathbf{z}_o)$ but $\hat{f}(\mathbf{y}|\mathbf{z}_o) \neq f(\mathbf{y}|\mathbf{z}_o)$. This is also shown numerically in Figs. 4 and 5, top row. RF-full can be expected to result in more accurate approximations of $f(\mathbf{y}|\mathbf{z}_o)$ than RF-stand. LF-auto is exact with $m \geq 1$ for a GP with exponential covariance function in one dimension. More generally, for a GP with Matérn covariance with smoothness ν , we obtain nearly exact representations if $m > \nu$ (cf. Katzfuss and Guinness 2019, Fig. 2a).

4.3.3. Computational Complexity

As laid out in Sect. 3.3, the relevant quantities for prediction can be obtained by computing \mathbf{U} , calculating \mathbf{V} from \mathbf{U} , carrying out a selected inversion based on \mathbf{V} , and performing triangular solves in \mathbf{V} . For all methods, the matrix \mathbf{U} has only m off-diagonal nonzero elements per column by construction, and it can be computed in $\mathcal{O}(nm^3)$ time using (9). The cost for each triangular solve in \mathbf{V} is on the order of the number of nonzero entries in \mathbf{V} . The cost of computing \mathbf{V} and the selected inversion is proportional to the sum of squares of the number of nonzero entries per column in \mathbf{V} .

For all response-first methods (i.e., RF-full, RF-stand, and RF-ind), we have shown below (5) that \mathbf{V} is simply a submatrix of \mathbf{U} . For LF-auto and MRA, $\mathbf{V} = \text{rchol}(\mathbf{W})$ must be explicitly computed, but these methods' special conditioning structures ensure that there is no fill-in. Hence, for these methods the number of off-diagonal nonzero entries in each




Figure 3. Average number of nonzero entries per column (ANNZC) in $\mathbf{V} = \text{rchol}(\mathbf{W})$ with $m = 15$ neighbors for increasing $n_O = n_P$ on a unit square using maxmin ordering. **a** Obtained using “brute-force” Cholesky based on reverse row-column ordering of \mathbf{W} , resulting in a dense upper triangle of \mathbf{V} with $\text{ANNZC} = (n_O + 1)/2$ for RF-full and LF-full. **b** Obtained using (5) or (7), with $\text{ANNZC} < m$ for the RF methods, and some reduction in nonzero entries using approximate-minimum-degree (AMD) ordering of \mathbf{W}_{oo} for the LF methods.

column of \mathbf{V} is guaranteed to be at most m . This implies that \mathbf{V} and the selected inverse can hence be computed in $\mathcal{O}(nm^2)$ time, $\mathbf{V}^{-1}\mathbf{H}'$ in $\mathcal{O}(nml)$ time, and $(\mathbf{V}^{-1}\mathbf{H}')'(\mathbf{V}^{-1}\mathbf{H}')$ can be computed in $\mathcal{O}(n\ell^2)$ time. (An additional approximation can be necessary for selected inversion in the case of OP methods—see Appendix D.1 for details.) Thus, the complexity of GP prediction using LF-auto, MRA, and all response-first methods is linear in n .

In contrast, LF-full and LF-ind result in more nonzero entries in \mathbf{W} and hence more potential for fill-in in \mathbf{V} (see Figs. 3 and S1 for illustration). As a result, linear complexity is not guaranteed for these two methods, even when using the block form in (7). For example, for locations on a regular two-dimensional grid ordered according to their coordinates, Katzfuss and Guinness (2019, Prop. 5) proved that the time required for computing $\mathbf{V}_{oo} = \text{rchol}(\mathbf{W}_{oo})$ using reverse ordering grows quadratically with n_O . As shown in Fig. 6, while fill-reducing permutations can reduce computation times somewhat, inference might still be slow or fail due to memory issues.

For the latent NNGP model underlying LF-ind and LF-full, Datta et al. (2016) proposed a sequential Gibbs sampler that samples every element of \mathbf{y} from its full-conditional distribution. As this approach can lead to non-convergence issues, Finley et al. (2019) instead proposed to sample \mathbf{y}_o jointly and then to sample \mathbf{y}_p conditional on \mathbf{y}_o . This avoids numerical nonzero entries similarly to our block-form \mathbf{V} in (7), but it still requires the expensive factorization of \mathbf{W} , and it can result in additional sampling error.

4.3.4. Consistent Framework

Under OP ordering, the likelihood $\hat{f}(\mathbf{z}_o) = \int \hat{f}(\mathbf{x}) d\mathbf{y}$ is unchanged when \mathbf{y}_p is removed from \mathbf{x} , and the Vecchia approximation is applied to the resulting vector $\mathbf{x}_{\bar{o}} = \mathbf{x} \setminus \mathbf{y}_p$:

$$\hat{f}(\mathbf{z}_o) = \int \int \hat{f}(\mathbf{x}_{\bar{o}}) \hat{f}(\mathbf{y}_p | \mathbf{x}_{\bar{o}}) d\mathbf{y}_p d\mathbf{y}_o = \int \hat{f}(\mathbf{x}_{\bar{o}}) \int \hat{f}(\mathbf{y}_p | \mathbf{x}_{\bar{o}}) d\mathbf{y}_p d\mathbf{y}_o = \int \hat{f}(\mathbf{x}_{\bar{o}}) d\mathbf{y}_o. \quad (8)$$

Thus, likelihood inference can first be carried out based on $\mathbf{x}_{\tilde{o}}$ (i.e., only based on \mathbf{y}_o and \mathbf{z}_o) using (10) as described in [Katzfuss and Guinness \(2019\)](#). Then, \mathbf{y}_p can be appended at the end of \mathbf{x} when predictions are desired, without changing the distribution $\hat{f}(\mathbf{z}_o)$. This has the advantage that parameter inference and prediction can be carried out in a consistent framework as described in Appendix C, and the prediction locations do not need to be known when training the model. If \mathbf{U} has already been calculated for $\mathbf{x}_{\tilde{o}}$, it is also possible to reuse this matrix and simply append to it the columns corresponding to \mathbf{y}_p . Similarly, if prediction at additional locations is desired later, these can be ordered last and the existing matrices can be augmented, ensuring that the distribution of existing variables is unchanged.

For the response-first methods, the likelihood reduces to the standard [Vecchia \(1988\)](#) likelihood. However, if only predictions are desired, we can set $g(i) = \emptyset$ for $i = 1, \dots, n_O$ without changing the approximation of $f(\mathbf{y}|\mathbf{z}_o)$, resulting in computational savings. To see this, note that this choice only affects \mathbf{U}_{rr} , which does not appear in any of the quantities in Sect. 3.3, because $\mathbf{V} = \mathbf{U}_{\ell\ell}$ and $\boldsymbol{\mu} = -(\mathbf{U}_{r\ell}\mathbf{U}_{\ell\ell}^{-1})'\mathbf{z}_o$ for response-first.

Strictly speaking, LF-auto does not obey the OP restriction and hence has the undesirable property that changing \mathcal{S} (e.g., by adding prediction locations) might change the joint distribution of other variables in \mathbf{x} . However, as shown in Figs. 2 and 4, the LF-auto approximation is so accurate in one dimension that, even for small m , there is little difference to the exact GP, and so all joint distributions are almost identical to the exact ones.

5. SIMULATION STUDY

We carried out a numerical comparison of the methods in Table 1. This systematic comparison was enabled by our R package `GPvecchia`, which implements all of the methods as special cases of the general Vecchia framework.

We simulated datasets at locations $\mathcal{S} = \mathcal{S}_o \cup \mathcal{S}_p$, consisting of randomly drawn locations \mathcal{S}_o from an independent uniform distribution on \mathbb{D} , combined with an equidistant grid \mathcal{S}_p on \mathbb{D} . We simulated \mathbf{y} at \mathcal{S} from the true distribution $f(\mathbf{y})$ induced by a GP with Matérn covariance function with variance 1 and smoothness parameter ν , and then, we sampled data \mathbf{z}_o by adding independent Gaussian noise with constant variance τ^2 to \mathbf{y}_o . We call $1/\tau^2$ the signal-to-noise ratio (SNR). Effective range is the distance at which the correlation is 0.05.

Our simulation study focuses on the approximation accuracy of summaries of $\hat{f}(\mathbf{y}|\mathbf{z}_o)$, assuming that any potential hyperparameters $\boldsymbol{\theta}$ are fixed and known. This results in more precise statements regarding the distinctions between the different methods and avoids confounding with issues that are not the focus of our study, such as choice of inference algorithms, tuning parameters, or choice of prior distributions for $\boldsymbol{\theta}$.

We computed the KL divergence for the joint distributions $\hat{f}(\mathbf{y}_p|\mathbf{z}_o)$ and $\hat{f}(\mathbf{y}_o|\mathbf{z}_o)$, and averaged over the results from multiple simulations for each method. This approximates the KL divergence for which the expectation is taken with respect to the joint distribution of the observations \mathbf{z}_o (see Sect. 4.3.2) and the observation locations \mathcal{S}_o . We also computed the average marginal KL divergences for $\hat{f}(\mathbf{y}_i|\mathbf{z}_o)$ for $i \in p$ and for $i \in o$.

For ease of presentation, comparisons to the MRA and to an extension of sparse general Vecchia ([Katzfuss and Guinness 2019](#)) are omitted here and shown in Section S3 instead.




Figure 4. For a GP in **one dimension**, KL divergences as a function of conditioning-set size m . Rows correspond to average KL divergence for $\hat{f}(\mathbf{y}_p|\mathbf{z}_o)$ (Joint Pred), $\hat{f}(\mathbf{y}_o|\mathbf{z}_o)$ (Joint Obs), $\hat{f}(\mathbf{y}_i|\mathbf{z}_o)$ for $i \in p$ (Marginal Pred), and $\hat{f}(\mathbf{y}_i|\mathbf{z}_o)$ for $i \in o$ (Marginal Obs), respectively. We used a modified log scale for the y-axes, with values below 10^{-10} treated as zero and indicated by grey bars at the bottom.

5.1. NUMERICAL COMPARISON IN 1-D

First, we considered the unit interval, $\mathbb{D} = [0, 1]$, with $n_O = n_P = 100$, effective range 0.15, and 40 repetitions of the simulation. All methods used coordinate (left-to-right) ordering. To avoid numerical error when computing the KL divergence due to finite machine precision, we constrained the locations in \mathcal{S}_o to be at least 10^{-4} units apart. LF-auto is exact for $v = 0.5$ with any $m \geq 1$. As shown in Fig. 4, the method was much more accurate at the prediction locations than any of the other approaches for $v = 1.5$. RF-ind and LF-ind, which do not condition on prediction locations, could only achieve a certain level of accuracy, with the joint KL divergence for the prediction locations leveling off as m increased. The performance of RF-ind and LF-ind improved on the marginal measures.




Figure 5. For a GP in **two dimensions**, KL divergences (on a log scale) as a function of conditioning-set size m . Rows correspond to average KL divergence for $\hat{f}(y_p|z_o)$ (Joint Pred), $\hat{f}(y_o|z_o)$ (Joint Obs), $\hat{f}(y_i|z_o)$ for $i \in p$ (Marginal Pred), and $\hat{f}(y_i|z_o)$ for $i \in o$ (Marginal Obs), respectively.

5.2. NUMERICAL COMPARISON IN 2-D

On the unit square, $\mathbb{D} = [0, 1]^2$, we used $n_O = n_P = 4,900$ and effective range 0.15, averaging over 20 repetitions. All methods used maxmin ordering. Figure 5 shows the results of the simulations. LF-full is not computationally scalable (see Sect. 5.3) but performed best in terms of accuracy, because it conditions only on latent variables (which contain more information about the process of interest than the noisy response variables). RF-full performed well on both joint and marginal accuracy measures. For approaches using independent conditioning for the prediction locations (RF-ind and LF-ind), the joint KL divergence at the prediction locations did not converge to zero, but these methods were more competitive with the other methods on marginal measures, as expected.

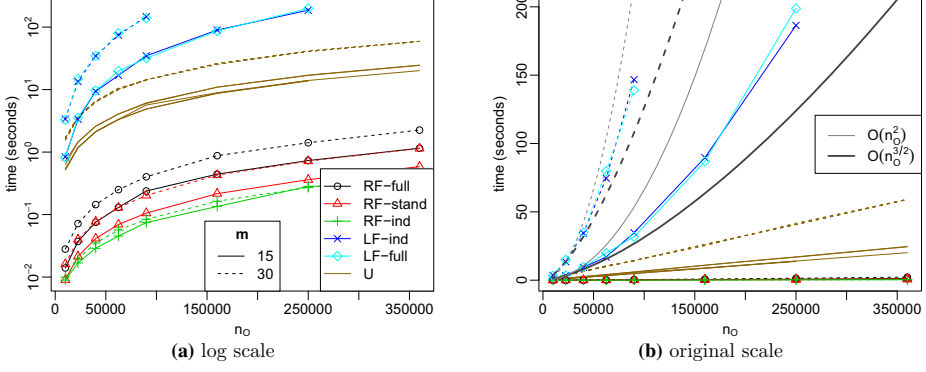


Figure 6. Time for computing \mathbf{U} and \mathbf{V} for $n_O = n_P$ observed and prediction locations on a unit square, as a function of n_O . Time for computing \mathbf{U} was similar for all methods. For RF methods, time for computing \mathbf{V} using (5) was negligible. For LF methods, \mathbf{V} was computed using (7) and using the approximate-minimum-degree fill-reducing permutation for \mathbf{W}_{oo} .

5.3. TIMING COMPARISON

We also carried out a timing study that examined the time for computing \mathbf{U} and \mathbf{V} on a unit square. Figure 6 shows median computation times from five repetitions on a 4-core machine (Intel Core i7-3770) with 3.4GHz and 16GB RAM. Consistent with our theoretical results, the time for computing \mathbf{U} increased roughly linearly with n and was similar for all methods for given n and m . (The time was slightly longer for the RF methods for small m , but this is solely due to an inefficiency in our RF code.) For the response-first methods, the time for computing \mathbf{V} was negligible relative to that for computing \mathbf{U} . For LF-ind and LF-full, the time for computing \mathbf{V} using a fill-reducing permutation increased roughly between $\mathcal{O}(n_O^{3/2})$ and $\mathcal{O}(n_O^2)$, and the computation failed for large n_O due to memory limitations. Computing \mathbf{V}_{oo} based on reverse ordering was even slower (see Figure S2).

5.4. COMPARISON FOR LARGE n

We further compared the scalable response-first methods for large $n_O = n_P$, with smoothness $\nu = 0.5$ and effective range 0.15 on a unit square. Two modifications to previous comparisons were necessary due to the large data size. First, we simulated the GP values on a regular $1,000 \times 1,000$ grid, using a regular subgrid of size n_P as \mathcal{S}_P and subsampling n_O of the remaining grid points as \mathcal{S}_O . Second, as it was impossible to compute the exact KL divergence, we approximated it by subtracting $\log \hat{f}(\mathbf{y}|\mathbf{z}_o)$ for each method from $\log \hat{f}(\mathbf{y}|\mathbf{z}_o)$ as approximated by a “very accurate” RF-full model with $m = 60$, all averaged over ten simulated datasets. As shown in Fig. 7, RF-full was much more accurate than the other methods in all settings.

5.5. HEATON COMPARISON

We compared the response-first methods, all with conditioning-set size $m = 15$, to the methods considered in a recent review and comparison paper (Heaton et al. 2019). We




Figure 7. Approximate KL divergences for the joint distribution $\hat{f}(\mathbf{y}|\mathbf{z}_o)$ in 2D with smoothness $\nu = 0.5$. **a** Fixed $n_O = n_P = 90,000$, varying m . **b** Fixed $m = 15$, increasing $n_O = n_P$ from 30,000 to 250,000 under in-fill (fixed-domain) asymptotics.

Table 2. Comparison using data from Heaton et al. (2019, Sec. 4.1), with lowest (i.e., best) scores in bold.

| | Method | RMSE | CRPS | JLS | Time (min) |
|----------|-------------------|-------------|-------------|--------------|------------|
| $m = 15$ | RF-full | 0.82 | 0.43 | 368.5 | 0.47 |
| | RF-stand | 0.83 | 0.43 | 368.5 | 0.45 |
| | RF-ind | 0.87 | 0.45 | 784.7 | 0.41 |
| Heaton | NNGP ($m = 15$) | 0.88 | 0.46 | | 1.99 |
| | MRA ($m > 500$) | 0.83 | 0.43 | | 13.57 |
| | LatticeKrig | 0.87 | 0.45 | | 25.58 |
| | Partition | 0.86 | 0.47 | | 77.56 |
| | SPDE | 0.86 | 0.59 | | 138.34 |

First three rows: Our methods with conditioning-set size $m = 15$. Bottom five rows: Results for the five best methods taken directly from (Heaton et al. 2019, Tab. 2). CRPS: continuous rank probability score. JLS: Joint log score for test subsets.

used the data from Heaton et al. (2019, Sec. 4.1), which consist of $n_O = 105,569$ training data and $n_P = 44,431$ test data simulated from a GP with exponential covariance. For our methods, we estimated the mean as the sample average of the training data and then estimated the process variance, noise variance, and range parameter (with true values 16.4, 0.05, and 4/3, respectively) by maximizing the likelihood approximated by sparse general Vecchia (Katzfuss and Guinness 2019) on a randomly chosen training subset of size 10,000.

We compared to the results of the top five methods reported in Heaton et al. (Heaton et al. (2019), Tab. 2), based on the RMSE, the continuous rank probability score (CRPS), and the computation time. Timing results in Heaton et al. (2019) were obtained on the Becker computing environment at Brigham Young University; the times for maximizing the likelihood and computing predictions for our methods were obtained on a basic desktop computer (Intel Core i5-3570 CPU @ 3.4GHz), ignoring some setup costs. Only marginal predictions were considered in Heaton et al. (2019); for our methods, we also computed the average log score for ten randomly selected subsets of size 500 of the test data. CRPS and log score are proper scoring rules that evaluate the approximation error in the predictive distribution (e.g., Gneiting and Katzfuss 2014) and simultaneously reward accurate point prediction (i.e., posterior mean) and accurate uncertainty quantification.

The results are shown in Table 2. RF-full and RF-stand had similar scores, because the assumed noise level was negligible. Both approaches outperformed all other methods; only MRA achieved comparable scores, but required much larger m and computation time. (A separate, more thorough comparison to MRA can be found in Section S3.) The NNGP results reported in Heaton et al. (2019) were obtained using a variant called NNGP-conjugate in Finley et al. (2019), which can be viewed as a Bayesian version of RF-ind; indeed, the marginal scores for the two methods were very similar. The joint log score for RF-ind was much worse than for RF-full and RF-stand.

6. APPLICATION TO SATELLITE DATA

We applied the scalable response-first methods from Table 1 to Level-2 bias-corrected solar-induced chlorophyll fluorescence (SIF) retrievals over land from the Orbiting Carbon Observatory 2 (OCO-2) satellite (OCO-2 Science Team et al. 2015). SIF is an important proxy for the amount of biomass produced from photosynthesis (Sun et al. 2017, 2018) and can be used to monitor the productivity of crops (Guan et al. 2016). The OCO-2 satellite has a sun-synchronous orbit with a period of 99 minutes and approximately repeats its spatial coverage every 16 days. Many remote-sensing satellites follow a similar sun-synchronous orbit, an orbital pattern that produces very dense observations along the trajectory of the orbit but wide gaps in space or time between orbits. This pattern is very common, and it presents a challenge for existing nearest neighbor-based prediction methods, because the nearest neighbors to any point in space or time will almost always be a sequence of densely packed points along one of the orbits. As we show below, this choice can be suboptimal and produce unrealistic artifacts in predicted maps.

We analyzed chlorophyll fluorescence data collected between August 1 and August 31, 2018, over the contiguous USA. During this time period, there were a total of 245,236 observations, plotted in Fig. 8. There was little evidence of temporal change during the time period, so we restricted our attention to a purely spatial model. We modeled the data with the spatial Gaussian process

$$z(\mathbf{s}_i) = \beta_0 + \sum_{j=1}^p \beta_j X_j(\mathbf{s}_i) + y(\mathbf{s}_i) + \epsilon_i, \quad \mathbf{s}_1, \dots, \mathbf{s}_n \in \mathbb{D},$$

where X_j were Gaussian basis functions centered at knots chosen as the first $p = 50$ locations in a maxmin ordering of the data locations, which ensured that no two knots were placed close to each other. The basis range was selected to be 637km (10% of Earth radius). The basis functions were included to capture a large amount of unstructured long-range variability that could not be explained by simple linear functions of latitude and longitude. The parameters β_0, \dots, β_p were estimated using least squares.

The residual field (shown in Figure S4) was modeled as $y(\cdot) \sim GP(0, K)$, where K was assumed to be an isotropic Matérn covariance function with three parameters: variance, range, and smoothness. The noise terms ϵ_i were assumed to be independent and identically distributed as $N(0, \tau^2)$. For covariance-parameter estimation on the residuals, we used the sparse general Vecchia likelihood (Katzfuss and Guinness 2019) with maxmin ordering,

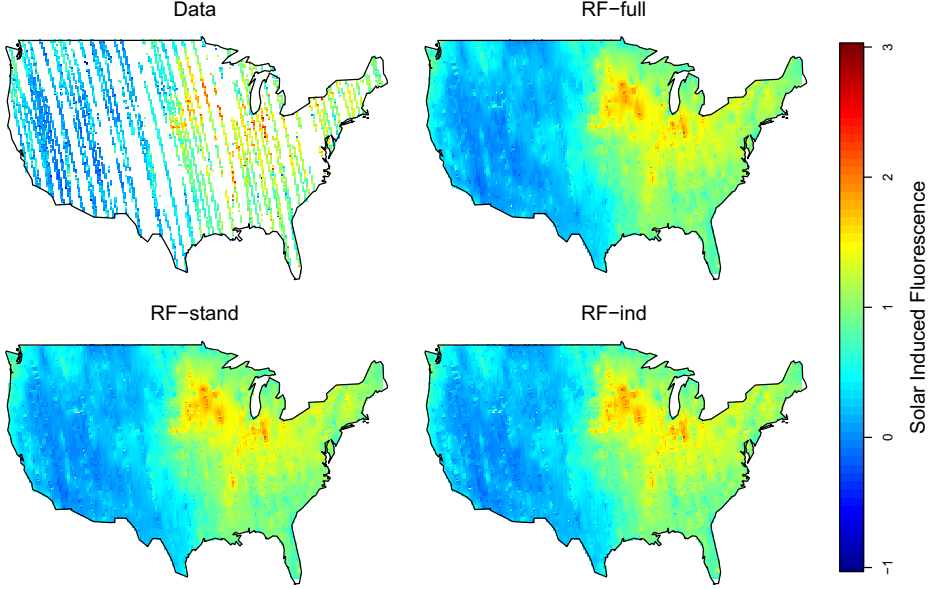


Figure 8. Satellite data, together with Vecchia predictions using $m = 30$ neighbors.

increasing m up to 40, beyond which the estimates did not change significantly. The estimated parameters were: variance = 0.1097, range = 100.8 km, smoothness = 0.0982, noise variance $\hat{\tau}^2 = 0.1869$.

Using the estimated covariance function and noise variance, we computed predictions for the scalable methods RF-full, RF-stand, and RF-ind (local kriging) at a grid of size $n_P = 24,407$ over the contiguous USA. Figure 8 shows predictions (i.e., posterior means of $\beta_0 + \sum_{j=1}^p \beta_j X_j(\mathbf{s}) + y(\mathbf{s})$) for $m = 30$ neighbors. Because the observations cover the study region quite well, the predictions using the various methods look similar, as might also be expected from the second row of Fig. 5. Upon closer inspection, however, the RF-ind predictions appear noisier and exhibit a “streaky” behavior. Figure 9 shows predictions at a higher resolution of $n_P = 18,576$ locations over the state of Texas. We can see that the RF-ind estimates are noisier and have more pronounced discontinuities parallel and perpendicular to the swaths of data. We conjecture that because the data locations are so dense along each swath, for RF-ind, two nearby prediction locations can condition on entirely different sets of observations if the two locations are nearly equidistant from two different swaths. Section S4 contains additional plots showing prediction uncertainties (i.e., posterior standard deviations) for $m = 30$ and predictions with $m = 60$. For $m = 60$, all predicted maps appeared smoother, but the Texas maps still had clearly visible streaks for RF-ind.

We also compared the prediction accuracy of RF-full, RF-stand, and RF-ind in two cross-validation experiments. First, we selected ten separate prediction sets, each consisting of 4000 randomly sampled data locations, to evaluate short-range prediction. Second, we held out 10 of the 67 swaths, one at a time, to evaluate long-range prediction. The average held-out swath size was 3493 locations. For each of the two cross-validation experiments, we computed the root mean squared error (RMSE) and the total log score, obtained by summing




Figure 9. Satellite data, together with Vecchia predictions over Texas using $m = 30$ neighbors.

the negative log predictive densities for each held-out test set. The log scores are reported relative to the lowest achieved log score.

The resulting prediction scores are shown in Fig. 10. For all settings, RF-full performed best. RF-ind (i.e., local kriging) was not competitive in terms of long-range predictions on the swath test sets. For the random test sets, RF-stand and local kriging performed similarly, with both methods roughly requiring $m = 20$ to achieve the same accuracy as RF-full with $m = 10$. These differences can be substantial in terms of computation times, which scale cubically in m . While the absolute differences in the RMSE values were not large, this was at least partially due to all comparisons being carried out relative to the (very noisy) test data \mathbf{z}_p , as the true fluorescence values \mathbf{y}_p are unknown. The “convergence” of RF-full, with similar values for $m = 20$ as for $m = 40$, indicates that even the exact GP without any approximation likely would not achieve significantly lower RMSE. Differences in the log scores were more pronounced, indicating that our new RF-full method can substantially outperform local kriging in terms of uncertainty quantification.

7. CONCLUSIONS

Vecchia approximation of Gaussian processes (GPs) is a powerful computational tool for fast analysis of large spatial datasets. While Vecchia approximations have been very popular for likelihood approximations, their use for the very important task of GP prediction or kriging had not been fully examined. Here, we proposed a general Vecchia framework for GP predictions, which includes as special cases some existing and several novel computational

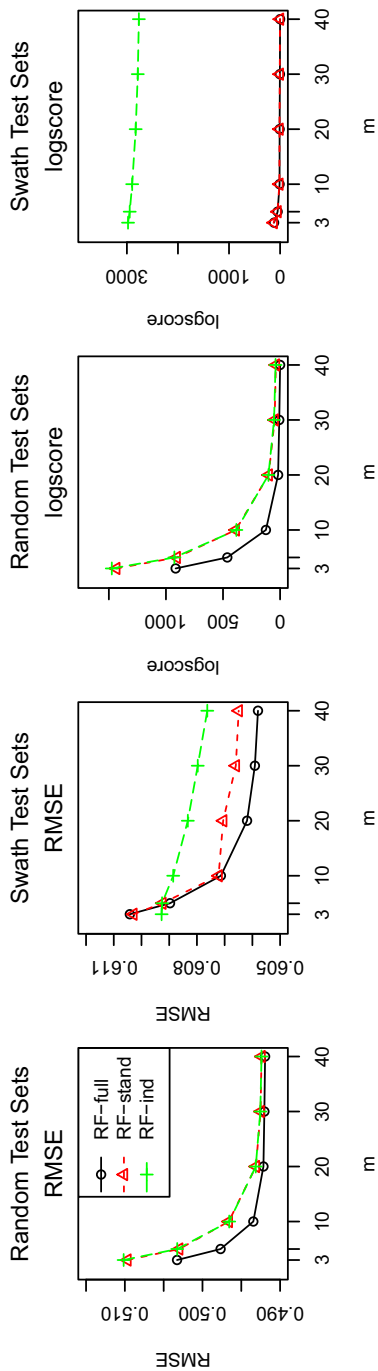


Figure 10. For fluorescence data, comparison of prediction scores as a function of conditioning-set size m .

approaches. We studied the accuracy and computational properties of the prediction methods both theoretically and numerically. In the case of unknown hyperparameters, all methods can be extended straightforwardly as described in Appendix C.

Based on our results, we make the following recommendations, which are also summarized briefly in Table 1. On a one-dimensional domain, LF-auto clearly had the best performance in all of the settings we considered. The autoregressive structure in LF-auto also affords linear computational scaling, and so we recommend LF-auto without any qualifications when the domain is one-dimensional. In two dimensions, we generally recommend RF-full, as it scales linearly and performed well on all accuracy measures. RF-stand is less accurate for noisy data, but has some computational advantages when the number of prediction locations is much smaller than the number of observations. LF-full can be very accurate, but it does not scale linearly in the data size. Local kriging (RF-ind) is fast and can provide accurate marginal predictive distributions, but it ignores dependence in the joint predictive distributions. LF-ind does not scale linearly and its joint predictive distributions at unobserved locations were often less accurate than those from RF-full. These inferential limitations are evident in Fig. 2 and in the top rows of Figs. 4 and 5.

The methods and algorithms proposed here are implemented in the R package `GPvecchia` (Katzfuss et al. 2020b), with default settings reflecting the recommendations in the previous paragraph. In principle, our methods and code are applicable in more than two dimensions, but a thorough investigation of their properties in this context is warranted. For example, a follow-up paper (Katzfuss et al. 2020a) shows that Vecchia-based approximations, using our observed-prediction ordering and appropriate extensions, can be highly accurate for computer-model emulation in higher dimensions. A further follow-up paper (Zilber and Katzfuss 2019) extends our methods to Vecchia–Laplace approximations of generalized GPs for non-Gaussian spatial data.

ACKNOWLEDGEMENTS

Katzfuss’ research was partially supported by National Science Foundation (NSF) Grant DMS–1521676 and NSF CAREER Grant DMS–1654083. Guinness’ research was partially supported by NSF Grant DMS–1613219 and NIH Grant No. R01ES027892. The authors would like to thank Anirban Bhattacharya, David Jones, Jennifer Hoeting, and several anonymous reviewers for helpful comments and suggestions. Jingjie Zhang and Marcin Jurek contributed to the R package `GPvecchia`, and Florian Schäfer provided C code for the exact maxmin ordering.

[Received January 2020. Accepted June 2020.]

A. VECTOR AND INDEXING NOTATION

As an example, define $\mathbf{y} = (\mathbf{y}_1, \mathbf{y}_2, \mathbf{y}_3, \mathbf{y}_4, \mathbf{y}_5)$ as a vector of vectors. Vectors of indices are used for defining subvectors. For example, if $\mathbf{o} = (4, 1, 2)$ is a vector of indices, then $\mathbf{y}_{\mathbf{o}} = (\mathbf{y}_4, \mathbf{y}_1, \mathbf{y}_2)$. Unions of vectors are vectors and are defined when the two vectors have the same type and when the ordering is specified. For example, if $\mathbf{z} = (\mathbf{z}_1, \mathbf{z}_2)$, then $\mathbf{y}_{\mathbf{o}} \cup \mathbf{z} = (\mathbf{y}_4, \mathbf{y}_1, \mathbf{z}_1, \mathbf{y}_2, \mathbf{z}_2)$ defines the union of $\mathbf{y}_{\mathbf{o}}$ and \mathbf{z} . Intersection is defined similarly

and uses the \cap notation. The ordering of the elements of the union or intersection can be defined alternatively via an index function $\#$ taking in an element and a vector and returning the index occupied by the element in the vector. Continuing the example above, $\#(\mathbf{y}_4, \mathbf{y}) = 4$, whereas $\#(\mathbf{y}_4, \mathbf{y}_o \cup \mathbf{z}) = 1$. A full description of the vector and indexing notation can be found in Katzfuss and Guinness (2019, App. A).

B. COMPUTING \mathbf{U}

We recapitulate here the formulas for computing \mathbf{U} from Katzfuss and Guinness (2019). Let $g(i)$ denote the vector of indices of the elements in \mathbf{x} on which x_i conditions. Also define $C(x_i, x_j)$ as the covariance between x_i and x_j implied by the true model in Sect. 2; that is, $C(y_i, y_j) = C(z_i, y_j) = K(\mathbf{s}_i, \mathbf{s}_j)$ and $C(z_i, z_j) = K(\mathbf{s}_i, \mathbf{s}_j) + \mathbb{1}_{i=j}\tau_i^2$. Then, the (j, i) th element of \mathbf{U} can be calculated as

$$\mathbf{U}_{ji} = \begin{cases} d_i^{-1/2}, & i = j, \\ -b_i^{(j)} d_i^{-1/2}, & j \in g(i), \\ 0, & \text{otherwise,} \end{cases} \quad (9)$$

where $\mathbf{b}'_i = C(x_i, \mathbf{x}_{g(i)})C(\mathbf{x}_{g(i)}, \mathbf{x}_{g(i)})^{-1}$, $d_i = C(x_i, x_i) - \mathbf{b}'_i C(\mathbf{x}_{g(i)}, x_i)$, and $b_i^{(j)}$ denotes the k th element of \mathbf{b}_i if j is the k th element in $g(i)$ (i.e., $b_i^{(j)}$ is the element of \mathbf{b}_i corresponding to x_j).

C. PREDICTION WITH UNKNOWN PARAMETERS

In practice, most GP models depend on an unknown parameter vector $\boldsymbol{\theta}$, which we will make explicit here. Likelihood approximation for parameter inference is discussed in detail in Katzfuss and Guinness (2019), but we will review it briefly here. Integration of $\hat{f}(\mathbf{x}|\boldsymbol{\theta})$ in (3) with respect to \mathbf{y} results in the following Vecchia likelihood (Katzfuss and Guinness 2019, Prop. 2):

$$\begin{aligned} -2 \log \hat{f}(\mathbf{z}_o|\boldsymbol{\theta}) &= -2 \sum_i \log \mathbf{U}_{ii} + 2 \sum_i \log \mathbf{V}_{ii} + \tilde{\mathbf{z}}'\tilde{\mathbf{z}} \\ &\quad - (\mathbf{V}^{-1} \mathbf{U}_{\ell, \bullet} \tilde{\mathbf{z}})' (\mathbf{V}^{-1} \mathbf{U}_{\ell, \bullet} \tilde{\mathbf{z}}) + n \log(2\pi), \end{aligned} \quad (10)$$

where \mathbf{U} and \mathbf{V} implicitly depend on $\boldsymbol{\theta}$, and $\tilde{\mathbf{z}} = \mathbf{U}'_{r, \bullet} \mathbf{z}_o$. The computational cost for evaluating this Vecchia likelihood is often low, and Katzfuss and Guinness (2019) provide conditions on the $g(i)$ under which the cost is guaranteed to be linear in n .

This allows for various forms of likelihood-based parameter inference. In a frequentist setting, we can compute $\hat{\boldsymbol{\theta}} = \arg \max_{\boldsymbol{\theta}} \log \hat{f}(\mathbf{z}_o|\boldsymbol{\theta})$ and then compute summaries of the posterior predictive distribution $\hat{f}(\mathbf{y}|\mathbf{z}_o) = \mathcal{N}_n(\boldsymbol{\mu}(\hat{\boldsymbol{\theta}}), \boldsymbol{\Sigma}(\hat{\boldsymbol{\theta}}))$ as described in Sect. 3.3. This often has low computational cost but ignores uncertainty in $\hat{\boldsymbol{\theta}}$. An example is given in Sect. 6.

In a Bayesian setting, given a prior distribution $f(\boldsymbol{\theta})$, a Metropolis–Hastings sampler can be used for parameters whose posterior or full-conditional distribution are not available in

closed form. At the $(l + 1)$ th step of the algorithm, one would propose a new value $\theta^{(P)} \sim q(\theta|\theta^{(l)})$ and accept it with probability $\min(1, h(\theta^{(P)}, \theta^{(l)})/h(\theta^{(l)}, \theta^{(P)}))$, where $h(\theta, \tilde{\theta}) = f(\theta)\hat{f}(\mathbf{z}_o|\theta)q(\tilde{\theta}|\theta)$. After burn-in and thinning, this results in a sample, say, $\theta^{(1)}, \dots, \theta^{(L)}$, leading to a Gaussian-mixture prediction: $\hat{f}(\mathbf{y}|\mathbf{z}_o) = (1/L) \sum_{l=1}^L \mathcal{N}_n(\mu(\theta^{(l)}), \Sigma(\theta^{(l)}))$. See Katzfuss and Guinness (2019, App. E) for an example for the use of RF-full in this setting. For more complicated Bayesian hierarchical models, inference can be carried out using a Gibbs sampler in which \mathbf{y} is sampled from its full-conditional distribution as described in item 4. in Sect. 3.3.

D. NUMERICAL NONZEROS IN \mathbf{V}

For simplicity, we focus here on RF-full, although numerical nonzeros can similarly occur for RF-stand, LF-full, and LF-ind (see Figure S1). For RF-full, the upper triangle of \mathbf{W} is at least as dense as the upper triangle of \mathbf{V} . Specifically, for $j < i$, $\mathbf{V}_{ji} = \mathbf{U}_{\ell_j \ell_i} = 0$ unless $j \in q_y(i)$. From Katzfuss and Guinness (Katzfuss and Guinness (2019), Prop. 3.2), we have that $\mathbf{W}_{ji} = 0$ unless $j \in q_y(i)$ or $\exists k > i$ such that $i, j \in q_y(k)$. Thus, for any pair $j < i$ such that $j \notin q_y(i)$ but $i, j \in q_y(k)$ for some $k > i$, we generally have $\mathbf{V}_{ji} = 0$ and $\mathbf{W}_{ji} \neq 0$.

From the standard Cholesky algorithm, we can derive that the algorithm for $\mathbf{V} = \text{rchol}(\mathbf{W})$ computes \mathbf{V}_{ji} as

$$\mathbf{V}_{ji} = (\mathbf{W}_{ji} - \sum_{k=i+1}^n \mathbf{V}_{ik}\mathbf{V}_{jk})/\mathbf{V}_{ii}.$$

Thus, for any pair $j < i$ such that $\mathbf{V}_{ji} = 0$ but $\mathbf{W}_{ji} \neq 0$, we know that $\mathbf{W}_{ji} = \sum_{k=i+1}^n \mathbf{V}_{ik}\mathbf{V}_{jk}$ theoretically, but due to potential numerical error it is not guaranteed that this equation holds exactly. A numerical nonzero is introduced in \mathbf{V} whenever a rounding error occurs in $\sum_{k=i+1}^n \mathbf{V}_{ik}\mathbf{V}_{jk}$, which relies on (potentially many) previous calculations in the Cholesky algorithm. Such numerical nonzeros are avoided by extracting $\mathbf{V} = \mathbf{U}_{\ell\ell}$ as a submatrix of \mathbf{U} (as proposed in Sect. 4.1), instead of explicitly carrying out the Cholesky factorization $\mathbf{V} = \text{rchol}(\mathbf{W})$.

D.1. IMPLICATIONS FOR SELECTED INVERSE

When \mathbf{V} is computed by copying a submatrix from \mathbf{U} to avoid numerical nonzeros, the selected inverse of this \mathbf{V} is not guaranteed to return the exact posterior variances of \mathbf{y} , unless \mathbf{V} is “padded” with zeros, which results in additional costs. This is because the selected inverse operates on the symbolic nonzero elements; that is, it operates on all elements that have to be computed in the Cholesky, even if they cancel to zero numerically (which is the case for many entries in our case). Denoting by \mathbf{S} the selected inverse of \mathbf{W} based on \mathbf{V} , a close look at the Takahashi recursions reveals that for all j, k with $j, k \in q_y(i)$, we need \mathbf{S}_{ji} and \mathbf{S}_{kj} . The latter element is only calculated if $j \in q_y(k)$. However, if $j \notin q_y(k)$, $\mathbf{S}_{kj} = \text{cov}(y_j, y_k|\mathbf{z})$ will typically be very small (if m is reasonably large), because their corresponding locations will likely be far away from each other and data can be observed in

between. In our experiments, the additional approximation error introduced by the selected inverse was negligible relative to the error introduced by the Vecchia approximation itself. When m becomes large enough for the Vecchia approximation to be accurate, the additional approximation error introduced by SelInv goes to zero as well.

If the exact variances implied by the Vecchia approximation $\widehat{f}(\mathbf{y}|\mathbf{z}_o)$ are desired, they can be computed as $\text{diag}(\text{var}(\mathbf{y}_p|\mathbf{z}_o)) = ((\mathbf{V}^{-1}\mathbf{I}_{\bullet p}) \circ (\mathbf{V}^{-1}\mathbf{I}_{\bullet p}))'\mathbf{1}_n$. For RF-full and RF-stand, this requires $\mathcal{O}(nmn_p)$ time, as described in Sect. 4.3.3, and so the overall computational complexity would not be increased if $n_p = \mathcal{O}(m^2)$.

E. PROOFS

In this section, we provide proofs for the propositions stated throughout the article.

Proof of Proposition 1. Part 1 of this proposition is equivalent to Thm. 1 in Guinness (2018). We prove the statement here again in a different way, which can be easily extended to prove Part 2.

First, consider a generic Vecchia approximation of a vector \mathbf{x} of length n : $\widehat{f}(\mathbf{x}) = \prod_{i=1}^n f(x_i|\mathbf{x}_{g(i)}) = \prod_{i=1}^n \mathcal{N}(x_i|\mu_{i|g(i)}, \sigma_{i|g(i)}^2)$. Then,

$$(-2) \mathbb{E}^{\mathbf{x}} \log \widehat{f}(\mathbf{x}) = \sum_{i=1}^n \log \text{var}(x_i|\mathbf{x}_{g(i)}) + \sum_{i=1}^n \mathbb{E}^{\mathbf{x}} w_i^2 + n \log(2\pi),$$

where $w_i = (x_i - \mu_{i|g(i)})/\sigma_{i|g(i)}$. We have $\mathbb{E}(w_i) = 0$, because $\mathbb{E}(\mu_{i|g(i)}) = \mathbb{E}\mathbb{E}(x_i|\mathbf{x}_{g(i)}) = \mathbb{E}(x_i)$. We also have $\text{var}(w_i) = 1$, because

$$\begin{aligned} \text{var}(x_i - \mu_{i|g(i)}) &= \text{var} \mathbb{E}(x_i - \mathbb{E}(x_i|\mathbf{x}_{g(i)})|\mathbf{x}_{g(i)}) + \text{var}(\mathbb{E}(x_i - \mathbb{E}(x_i|\mathbf{x}_{g(i)})|\mathbf{x}_{g(i)})) \\ &= 0 + \sigma_{i|g(i)}^2. \end{aligned}$$

Hence, $w_i \sim \mathcal{N}(0, 1)$, and so $\mathbb{E}^{\mathbf{x}} w_i^2 = 1$ and

$$(-2) \mathbb{E}^{\mathbf{x}} \log \widehat{f}(\mathbf{x}) = \sum_{i=1}^n \log \text{var}(x_i|\mathbf{x}_{g(i)}) + n + n \log(2\pi).$$

Because the exact density $f(\mathbf{x})$ is a special case of Vecchia with $g(i) = (1, \dots, i-1)$, we have

$$\text{KL}(f(\mathbf{x})\|\widehat{f}(\mathbf{x})) = \mathbb{E}^{\mathbf{x}} \log f(\mathbf{x}) - \mathbb{E}^{\mathbf{x}} \log \widehat{f}(\mathbf{x}) = \frac{1}{2} \sum_{i=1}^n \log \frac{\text{var}(x_i|\mathbf{x}_{g(i)})}{\text{var}(x_i|\mathbf{x}_{1:i-1})}. \quad (11)$$

For $g_1(i) \subset g_2(i)$, we can write, say, $g_2(i) = g_1(i) \cup c(i)$. Using the law of total variance,

$$\text{var}(x_i|\mathbf{x}_{g_1(i)}) = \text{var}(x_i|\mathbf{x}_{g_1(i)}, \mathbf{x}_{c(i)}) + \text{var}(\mathbb{E}(x_i|\mathbf{x}_{g_1(i)}, \mathbf{x}_{c(i)})|\mathbf{x}_{c(i)}) \geq \text{var}(x_i|\mathbf{x}_{g_2(i)}). \quad (12)$$

Now, Part 1 follows by combining (11) and (12) with the assumption that $g_1(i) \subset g_2(i)$ for all i .

For Part 2, we consider $\mathbf{x} = (\mathbf{x}^{(1)}, \mathbf{x}^{(2)})$ with $\mathbf{x}^{(1)} = (x_1, \dots, x_{n_1})$ and $\mathbf{x}^{(2)} = (x_{n_1+1}, \dots, x_{n_1+n_2})$. Then,

$$\begin{aligned} & \text{CKL}(f(\mathbf{x}^{(2)}|\mathbf{x}^{(1)}) \parallel \hat{f}(\mathbf{x}^{(2)}|\mathbf{x}^{(1)})) \\ &= \int f(\mathbf{x}^{(1)}) \int f(\mathbf{x}^{(2)}|\mathbf{x}^{(1)}) \log(f(\mathbf{x}^{(2)}|\mathbf{x}^{(1)})/\hat{f}(\mathbf{x}^{(2)}|\mathbf{x}^{(1)})) d\mathbf{x}^{(2)} d\mathbf{x}^{(1)} \\ &= \mathbb{E}^{\mathbf{x}} \log f(\mathbf{x}^{(2)}|\mathbf{x}^{(1)}) - \mathbb{E}^{\mathbf{x}} \log \hat{f}(\mathbf{x}^{(2)}|\mathbf{x}^{(1)}) \\ &= \frac{1}{2} \sum_{i=n_1+1}^{n_1+n_2} \log \frac{\text{var}(x_i|\mathbf{x}_{g(i)})}{\text{var}(x_i|\mathbf{x}_{1:i-1})}, \end{aligned}$$

where the last equality can be shown almost identically to (11) above, noting that $\hat{f}(\mathbf{x}^{(2)}|\mathbf{x}^{(1)}) = \hat{f}(\mathbf{x})/\hat{f}(\mathbf{x}^{(1)}) = \prod_{i=n_1+1}^{n_1+n_2} f(x_i|\mathbf{x}_{g(i)})$. Part 2 follows by combining this result and (12) with the assumption that $g_1(i) \subset g_2(i)$ for all $i = n_1 + 1, \dots, n_1 + n_2$. \square

Proof of Proposition 2. For all parts of this proposition, note that all variables in the model are conditionally independent of z_j given y_j , and so conditioning on y_j is equivalent to conditioning on y_j and z_j . For Part 1, we can thus verify easily that $g(i)$ is equivalent to $(1, \dots, i-1)$, for all i and all methods under consideration, and so Part 1 follows from (11). Using Proposition 1, the proof for all other parts simply consists of showing that certain conditioning vectors contain certain other conditioning vectors (similar to Katzfuss and Guinness 2019, Prop. 4). For example, for Part 3, all response-first methods are based on the ordering $\mathbf{x} = (\mathbf{z}_o, \mathbf{y}_o, \mathbf{y}_p)$ with nearest neighbor conditioning (under some restrictions for RF-stand and RF-ind), and so it is easy to see that $g_m(i) \subset g_{m+1}(i)$ for all $i \in \ell_p = (n+1, \dots, n+n_p)$. LF-full and LF-ind can equivalently be defined based on the ordering $(\mathbf{y}_o, \mathbf{z}_o, \mathbf{y}_p)$, in which case we also have $g_m(i) \subset g_{m+1}(i)$ for all $i \in \ell_p = (n+1, \dots, n+n_p)$. For Part 5, note that in RF-stand and RF-ind, \mathbf{y}_p does not condition on \mathbf{y}_o . Hence, the distribution $\hat{f}(\mathbf{y}_p|\mathbf{z}_o)$ is equivalent to the distribution obtained under a simplified Vecchia approximation based on $\mathbf{x} = (\mathbf{z}_o, \mathbf{y}_p)$ (i.e., with \mathbf{y}_o removed completely). It is straightforward to show that Part 5 holds for this simplified approximation. For Part 6, note that $g(i)$ is the same for RF-full and RF-stand for all $i = 1, \dots, n$. For $i \in p$, letting $a(i) = q_y^{\text{RF-full}}(i) \cap o$, we have $\mathbf{x}_{g^{\text{RF-full}}(i)} = (\mathbf{y}_{q_y^{\text{RF-stand}}}, \mathbf{y}_{a(i)}) = (\mathbf{y}_{q_y^{\text{RF-stand}}}, \mathbf{y}_{a(i)}, \mathbf{z}_{a(i)})$ and $\mathbf{x}_{g^{\text{RF-stand}}(i)} = (\mathbf{y}_{q_y^{\text{RF-stand}}}, \mathbf{z}_{a(i)})$, and so $g^{\text{RF-stand}}(i) \subset g^{\text{RF-full}}(i)$ for all $i \in p$. For Part 7, note that a GP with exponential covariance in 1-D is a Markov process, and so in (6) with $q_y(i) = i-1$ and left-to-right ordering, we have $f(y_i|\mathbf{y}_{q_y(i)}) = f(y_i|y_1, \dots, y_{i-1})$ and hence $\hat{f}(\mathbf{x}) = f(\mathbf{x})$. \square

REFERENCES

- Banerjee, S., Carlin, B. P., and Gelfand, A. E. (2004). *Hierarchical Modeling and Analysis for Spatial Data*. Chapman & Hall, Boca Raton.
- Banerjee, S., Gelfand, A. E., Finley, A. O., and Sang, H. (2008). Gaussian predictive process models for large spatial data sets. *Journal of the Royal Statistical Society, Series B*, 70(4):825–848.
- Cressie, N. and Johannesson, G. (2008). Fixed rank kriging for very large spatial data sets. *Journal of the Royal Statistical Society, Series B*, 70(1):209–226.

- Cressie, N. and Wikle, C. K. (2011). *Statistics for Spatio-Temporal Data*. Wiley, Hoboken, NJ.
- Curriero, F. C. and Lele, S. (1999). A composite likelihood approach to semivariogram estimation. *Journal of Agricultural, Biological, and Environmental Statistics*, 4(1):9–28.
- Datta, A., Banerjee, S., Finley, A. O., and Gelfand, A. E. (2016). Hierarchical nearest-neighbor Gaussian process models for large geostatistical datasets. *Journal of the American Statistical Association*, 111(514):800–812.
- Du, J., Zhang, H., and Mandrekar, V. S. (2009). Fixed-domain asymptotic properties of tapered maximum likelihood estimators. *The Annals of Statistics*, 37:3330–3361.
- Eidsvik, J., Shaby, B. A., Reich, B. J., Wheeler, M., and Niemi, J. (2014). Estimation and prediction in spatial models with block composite likelihoods using parallel computing. *Journal of Computational and Graphical Statistics*, 23(2):295–315.
- Erismann, A. M. and Tinney, W. F. (1975). On computing certain elements of the inverse of a sparse matrix. *Communications of the ACM*, 18(3):177–179.
- Finley, A. O., Datta, A., Cook, B. C., Morton, D. C., Andersen, H. E., and Banerjee, S. (2019). Efficient algorithms for Bayesian nearest neighbor Gaussian processes. *Journal of Computational and Graphical Statistics*, 28(2):401–414.
- Finley, A. O., Sang, H., Banerjee, S., and Gelfand, A. E. (2009). Improving the performance of predictive process modeling for large datasets. *Computational Statistics & Data Analysis*, 53(8):2873–2884.
- Foreman-Mackey, D., Agol, E., Ambikasaran, S., and Angus, R. (2017). Fast and scalable Gaussian process modeling with applications to astronomical time series. *The Astronomical Journal*, 154:220.
- Furrer, R., Genton, M. G., and Nychka, D. (2006). Covariance tapering for interpolation of large spatial datasets. *Journal of Computational and Graphical Statistics*, 15(3):502–523.
- Gneiting, T. and Katzfuss, M. (2014). Probabilistic forecasting. *Annual Review of Statistics and Its Application*, 1(1):125–151.
- Gramacy, R. B. and Apley, D. W. (2015). Local Gaussian process approximation for large computer experiments. *Journal of Computational and Graphical Statistics*, 24(2):561–578.
- Guan, K., Berry, J. A., Zhang, Y., Joiner, J., Guanter, L., Badgley, G., and Lobell, D. B. (2016). Improving the monitoring of crop productivity using spaceborne solar-induced fluorescence. *Global Change Biology*, 22(2):716–726.
- Guinness, J. (2018). Permutation methods for sharpening Gaussian process approximations. *Technometrics*, 60(4):415–429.
- Guinness, J. (2019). Spectral density estimation for random fields via periodic embeddings. *Biometrika*, 106(2):267–286.
- Heaton, M. J., Datta, A., Finley, A. O., Furrer, R., Guinness, J., Guhaniyogi, R., Gerber, F., Gramacy, R. B., Hammerling, D., Katzfuss, M., Lindgren, F., Nychka, D. W., Sun, F., and Zammit-Mangion, A. (2019). A case study competition among methods for analyzing large spatial data. *Journal of Agricultural, Biological, and Environmental Statistics*, 24(3):398–425.
- Higdon, D. (1998). A process-convolution approach to modelling temperatures in the North Atlantic Ocean. *Environmental and Ecological Statistics*, 5(2):173–190.
- Jones, D. R., Schonlau, M., and W. J. Welch (1998). Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13:455–492.
- Jurek, M. and Katzfuss, M. (2018). Multi-resolution filters for massive spatio-temporal data. [arXiv:1810.04200](https://arxiv.org/abs/1810.04200).
- Katzfuss, M. (2017). A multi-resolution approximation for massive spatial datasets. *Journal of the American Statistical Association*, 112(517):201–214.
- Katzfuss, M. and Cressie, N. (2011). Spatio-temporal smoothing and EM estimation for massive remote-sensing data sets. *Journal of Time Series Analysis*, 32(4):430–446.
- Katzfuss, M. and Gong, W. (2019). A class of multi-resolution approximations for large spatial datasets. *Statistica Sinica*. <https://doi.org/10.5705/ss.202018.0285>.
- Katzfuss, M. and Guinness, J. (2019). A general framework for Vecchia approximations of Gaussian processes. *Statistical Science*, accepted.

- Katzfuss, M., Guinness, J., and Lawrence, E. (2020a). Scaled Vecchia approximation for fast computer-model emulation. [arXiv:2005.00386](https://arxiv.org/abs/2005.00386).
- Katzfuss, M., Jurek, M., Zilber, D., Gong, W., Guinness, J., Zhang, J., and Schäfer, F. (2020b). *GPvecchia: Fast Gaussian-process inference using Vecchia approximations*. R package version 0.1.3.
- Kaufman, C. G., Schervish, M. J., and Nychka, D. W. (2008). Covariance tapering for likelihood-based estimation in large spatial data sets. *Journal of the American Statistical Association*, 103(484):1545–1555.
- Kelly, B. C., Becker, A. C., Sobolewska, M., Siemiginowska, A., and Uttley, P. (2014). Flexible and scalable methods for quantifying stochastic variability in the era of massive time-domain astronomical data sets. *Astrophysical Journal*, 788(1):33.
- Kennedy, M. C. and O’Hagan, A. (2001). Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B*, 63(3):425–464.
- Le, N. D. and Zidek, J. V. (2006). *Statistical Analysis of Environmental Space-Time Processes*. Springer, Berlin.
- Li, S., Ahmed, S., Klimeck, G., and Darve, E. (2008). Computing entries of the inverse of a sparse matrix using the FIND algorithm. *Journal of Computational Physics*, 227(22):9408–9427.
- Lin, L., Yang, C., Meza, J., Lu, J., Ying, L., and Weinan, E. (2011). SelInv—An algorithm for selected inversion of a sparse symmetric matrix. *ACM Transactions on Mathematical Software*, 37(4):40.
- Lindgren, F., Rue, H., and Lindström, J. (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B*, 73(4):423–498.
- Nychka, D. W., Bandyopadhyay, S., Hammerling, D., Lindgren, F., and Sain, S. R. (2015). A multi-resolution Gaussian process model for the analysis of large spatial data sets. *Journal of Computational and Graphical Statistics*, 24(2):579–599.
- OCO-2 Science Team, Gunson, M., and Eldering, A. (2015). OCO-2 Level 2 bias-corrected solar-induced fluorescence and other select fields from the IMAP-DOAS algorithm aggregated as daily files, retrospective processing V7r. https://disc.gsfc.nasa.gov/datacollection/OCO2_L2_Lite_SIF_7r.html
- Quiñonero-Candela, J. and Rasmussen, C. E. (2005). A unifying view of sparse approximate Gaussian process regression. *Journal of Machine Learning Research*, 6:1939–1959.
- Rasmussen, C. E. and Williams, C. K. I. (2006). *Gaussian Processes for Machine Learning*. MIT Press, Cambridge.
- Rue, H. and Held, L. (2005). *Gaussian Markov Random Fields: Theory and Applications*. CRC Press, Boca Raton.
- Sang, H., Jun, M., and Huang, J. Z. (2011). Covariance approximation for large multivariate spatial datasets with an application to multiple climate model errors. *Annals of Applied Statistics*, 5(4):2519–2548.
- Schäfer, F., Katzfuss, M., and Owhadi, H. (2020). Sparse Cholesky factorization by Kullback–Leibler minimization. [arXiv:2004.14455](https://arxiv.org/abs/2004.14455).
- Schäfer, F., Sullivan, T. J., and Owhadi, H. (2017). Compression, inversion, and approximate PCA of dense kernel matrices at near-linear computational complexity. [arXiv:1706.02205](https://arxiv.org/abs/1706.02205).
- Snelson, E. and Ghahramani, Z. (2007). Local and global sparse Gaussian process approximations. In *Artificial Intelligence and Statistics 11 (AISTATS)*.
- Snook, J., Larochelle, H., and Adams, R. P. (2012). Practical Bayesian optimization of machine learning algorithms. In *Neural Information Processing Systems*.
- Stein, M. L., Chi, Z., and Welty, L. (2004). Approximating likelihoods for large spatial data sets. *Journal of the Royal Statistical Society: Series B*, 66(2):275–296.
- Stroud, J. R., Stein, M. L., and Lysen, S. (2017). Bayesian and maximum likelihood estimation for Gaussian processes on an incomplete lattice. *Journal of Computational and Graphical Statistics*, 26(1):108–120.
- Sun, Y., Frankenberg, C., Jung, M., Joiner, J., Guanter, L., Köhler, P., and Magney, T. (2018). Overview of solar-induced chlorophyll fluorescence (sif) from the orbiting carbon observatory-2: Retrieval, cross-mission comparison, and global monitoring for gpp. *Remote Sensing of Environment*, 209:808–823.
- Sun, Y., Frankenberg, C., Wood, J. D., Schimel, D. S., Jung, M., Guanter, L., Drewry, D., Verma, M., Porcar-Castell, A., Griffiths, T. J., et al. (2017). Oco-2 advances photosynthesis observation from space via solar-induced chlorophyll fluorescence. *Science*, 358(6360):eaam5747.

- Sun, Y. and Stein, M. L. (2016). Statistically and computationally efficient estimating equations for large spatial datasets. *Journal of Computational and Graphical Statistics*, 25(1):187–208.
- Tzeng, S. and Huang, H.-C. (2018). Resolution adaptive fixed rank kriging. *Technometrics*, 60(2):198–208.
- Vecchia, A. (1988). Estimation and model identification for continuous spatial processes. *Journal of the Royal Statistical Society, Series B*, 50(2):297–312.
- Vecchia, A. (1992). A new method of prediction for spatial regression models with correlated errors. *Journal of the Royal Statistical Society, Series B*, 54(3):813–830.
- Wang, Y., Khardon, R., and Protopapas, P. (2012). Nonparametric Bayesian estimation of periodic light curves. *Astrophysical Journal*, 756(1):67.
- Wikle, C. K. and Cressie, N. (1999). A dimension-reduced approach to space-time Kalman filtering. *Biometrika*, 86(4):815–829.
- Zhang, B., Sang, H., and Huang, J. Z. (2019). Smoothed full-scale approximation of Gaussian process models for computation of large spatial datasets. *Statistica Sinica*, 29:1711–1737.
- Zilber, D. and Katzfuss, M. (2019). Vecchia-Laplace approximations of generalized Gaussian processes for big non-Gaussian spatial data. [arXiv:1906.07828](https://arxiv.org/abs/1906.07828).

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.