



Vecchia–Laplace approximations of generalized Gaussian processes for big non-Gaussian spatial data

Daniel Zilber, Matthias Katzfuss^{*}

Department of Statistics, Texas A&M University, USA

ARTICLE INFO

Article history:

Received 7 January 2020

Received in revised form 12 August 2020

Accepted 23 August 2020

Available online 28 August 2020

Keywords:

Exponential family

Geostatistics

Kriging

Nearest Neighbor

Sparse inverse Cholesky

Spatial Generalized Linear Mixed Model

ABSTRACT

Generalized Gaussian processes (GGPs) are highly flexible models that combine latent GPs with potentially non-Gaussian likelihoods from the exponential family. GGPs can be used in a variety of settings, including GP classification, nonparametric count regression, modeling non-Gaussian spatial data, and analyzing point patterns. However, inference for GGPs can be analytically intractable, and large datasets pose computational challenges due to the inversion of the GP covariance matrix. A Vecchia–Laplace approximation for GGPs is proposed, which combines a Laplace approximation to the non-Gaussian likelihood with a computationally efficient Vecchia approximation to the GP, resulting in simple, general, scalable, and accurate methodology. Numerical studies and comparisons on simulated and real spatial data are provided. The methods are implemented in a freely available R package.

© 2020 Elsevier B.V. All rights reserved.

1. Introduction

Dependent non-Gaussian data are ubiquitous in time series, geospatial applications, and more generally in nonparametric regression and classification. A popular model for such data is obtained by combining a latent Gaussian process (GP) with conditionally independent, potentially non-Gaussian likelihoods from the exponential family. This is traditionally referred to as a spatial generalized linear mixed model (SGLMM) in the spatial statistics literature (Diggle et al., 1998), but the same model has more recently also been referred to as a generalized GP (GGP) (e.g., Chan and Dong, 2011); we will use the latter, more concise term throughout. GGPs are highly flexible, interpretable, and allow for natural, probabilistic uncertainty quantification. However, inference for GGPs can be analytically intractable, and large datasets pose additional computational challenges due to the inversion of the GP covariance matrix.

Popular techniques to numerically perform the intractable marginalization necessary for inference are, in order of increasing speed: Markov chain Monte Carlo (MCMC), expectation propagation, variational methods, and Laplace approximations. See Shang and Chan (2013) for a recent review of deterministic techniques and Filippone and Girolami (2014) for a comparison of MCMC and expectation propagation. Rue et al. (2009) argue that variational methods and expectation propagation suffer from underestimated and overestimated posterior variances, respectively. Here, we consider the Laplace approximation (e.g., Tierney and Kadane, 1986; Williams and Barber, 1998), a classic technique for integral evaluation based on second-order Taylor expansion. Bonat and Ribeiro (2016) show numerically that the Laplace approximation can be a practical and accurate method for GGP inference.

It has long been recognized that the computational cost for GPs is cubic in the data size. Much work has been done on GP approximations that address this problem in the context of Gaussian noise (as recently reviewed in Heaton et al.,

^{*} Correspondence to: 3143 TAMU, College Station, TX 77843, USA.

E-mail addresses: dzilber@tamu.com (D. Zilber), katzfuss@gmail.com (M. Katzfuss).

2019). Low-rank approaches (e.g., Higdon, 1998; Quiñero-Candela and Rasmussen, 2005; Banerjee et al., 2008; Cressie and Johannesson, 2008; Katzfuss and Cressie, 2011) have limitations in the presence of fine-scale structure (Stein, 2014), but they have proved popular due to their simplicity. Approximations relying on sparsity in covariance matrices (Furrer et al., 2006; Kaufman et al., 2008) by definition can only capture local, short-range dependence and cannot guarantee low computation cost. Approaches that take advantage of Toeplitz or Kronecker structure (e.g., Dietrich and Newsam, 1997; Flaxman et al., 2015; Guinness and Fuentes, 2017) can be extremely efficient but are not as generally applicable. Recently, methods relying on sparsity in precision matrices (Rue and Held, 2005; Lindgren et al., 2011; Nychka et al., 2015) have gained popularity due to their accuracy and flexibility. In particular, a class of highly promising GP approximations (Vecchia, 1988; Stein et al., 2004; Datta et al., 2016; Guinness, 2018; Katzfuss and Guinness, 2019; Katzfuss et al., 2020a) rely on ordered conditional independence that can guarantee linear scaling in the data size while resolving dependence at all scales.

There are also a number of existing methods for large non-Gaussian datasets modeled using GPs. A popular approach is to combine a low-rank GP with an approximation of the non-Gaussian likelihood, as the dimension reduction inherent in the low-rank approximation carries through even to the non-Gaussian case. Sengupta and Cressie (2013) estimate parameters using an expectation–maximization algorithm with low-rank and Laplace approximations. Sheth et al. (2015) use variational inference to obtain the posterior and select a set of conditioning points for their low-rank approximation. Some methods of dimension reduction, including random sketching (e.g., Yang et al., 2017) and projection, offer theoretical guarantees and can be combined with MCMC methods for the analysis of non-Gaussian data (e.g., Hughes and Haran, 2013; Guan and Haran, 2018), but are still subject to the limitations of low-rank methods. Nickisch et al. (2018) develop state-space models for one-dimensional non-Gaussian time series, which can perform inference in linear time and memory using a set of knots in time, a form of low-rank approximation. Alternate priors (e.g., log gamma priors for count data, Bradley et al., 2018) are an interesting but specialized approach to completely avoid the intractability issues with GPs.

Similar to what we shall propose, some authors have combined a sparse-precision approach with a non-Gaussian approximation. A prominent example is Lindgren et al. (2011), in which an integrated nested Laplace approximation (INLA) is combined with a sparse-precision approximation of the GP using its representation as the solution to a stochastic partial differential equation. Datta et al. (2016) proposed to apply the GP approximation of Vecchia (1988) to a latent GP, but did not provide an explicit algorithm for large non-Gaussian data. While both Lindgren et al. (2011) and Datta et al. (2016) limit the number of nonzero entries per row or column in the precision matrix to a small constant, the computational complexity for decomposing this sparse $n \times n$ matrix is not linear in n , but rather $\mathcal{O}(n^{3/2})$ in two dimensions (Lipton et al., 1979, Thm. 6), and at least $\mathcal{O}(n^2)$ in higher dimensions. In the Gaussian setting, this scaling problem can be overcome by applying a Vecchia approximation to the observed data (Vecchia, 1988) or to the joint distribution of the observed data and the latent GP (Katzfuss and Guinness, 2019).

To handle both scaling and intractability, we propose a Vecchia–Laplace (VL) approximation for GPs. The posterior mode necessary for the Laplace approximation is found using the Newton–Raphson algorithm, which can be viewed as iterative GP inference based on Gaussian pseudo-data. At each iteration of our VL algorithm, the joint Gaussian distribution of the pseudo-data and the latent GP realizations is approximated using the general Vecchia framework (Katzfuss and Guinness, 2019; Katzfuss et al., 2020a). By modeling the joint distribution of pseudo-data and GP realizations at each iteration, our VL approach can ensure sparsity and guarantee linear scaling in n for any dimension, overcoming the scaling issues of the sparse-matrix approaches mentioned above.

To our knowledge, we provide the first explicit algorithm extending and applying the highly promising class of general Vecchia GP approximations to large non-Gaussian data. We believe it to be a useful addition to the literature due to its speed, simplicity, guaranteed numerical performance, and wide applicability (e.g., binary, count, right-skewed, and point-pattern data). In addition, as shown in Katzfuss and Guinness (2019), the general Vecchia approximation includes many popular GP approximations (e.g., Vecchia, 1988; Snelson and Ghahramani, 2007; Finley et al., 2009; Sang et al., 2011; Datta et al., 2016; Katzfuss, 2017; Katzfuss and Gong, 2019) as special cases, and so our VL methodology also directly provides extensions of these GP approximations to non-Gaussian data.

The remainder of this document is organized as follows. In Section 2, we review the Laplace approximation and general Vecchia. In Section 3, we introduce and examine our VL method, including parameter inference and predictions at unobserved locations. In Sections 4 and 5, we study and compare the performance of VL on simulated and real data, respectively. Some details are left to the appendix. A separate Supplementary Material document contains Sections S1–S6 with additional derivations, simulations, and discussion. The methods and algorithms proposed here are implemented in the R package *GPvecchia* (Katzfuss et al., 2020) with sensible default settings, so that a wide audience of practitioners can immediately use the code with little background knowledge. Our results and figures can be reproduced using the code and data at <https://github.com/katzfuss-group/GPvecchia-Laplace>.

2. Review of existing results

2.1. Generalized Gaussian processes

Let $y(\cdot) \sim GP(\mu, K)$ be a latent Gaussian process with mean function μ and kernel or covariance function K on a domain $\mathcal{D} \subset \mathbb{R}^d$, $d \in \mathbb{N}^+$. Observations $\mathbf{z} = (z_1, \dots, z_n)'$ at locations $\mathbf{s}_i \in \mathcal{D}$ are assumed to be conditionally independent,

Table 1

Examples of popular likelihoods, together with the Gaussian pseudo-data and pseudo-variances implied by the Laplace approximation. The non-canonical logarithmic link function is used for the Gamma likelihood to ensure that the second parameter, ae^{-y} , is positive.

Distribution	Likelihood $g(z y)$	Pseudo-data t_y	Pseudo-variance $d(y)$
Gaussian	$\mathcal{N}(y, \tau^2)$	z	τ^2
Bernoulli	$\mathcal{B}(\text{logit}^{-1}(y))$	$y + \frac{(1+e^y)^2}{e^y} (z - \frac{e^y}{1+e^y})$	$(1 + e^{-y})(1 + e^y)$
Poisson	$\mathcal{P}(e^y)$	$y + e^{-y}(z - e^y)$	e^{-y}
Gamma	$\mathcal{G}(a, ae^{-y})$	$y + (1 - z^{-1})e^y$	$e^y/(az)$

$z_i|\mathbf{y} \stackrel{\text{ind.}}{\sim} g_i(z_i|y_i)$, where $\mathbf{y} = (y_1, \dots, y_n)'$ and $y_i = y(\mathbf{s}_i)$. We assume that the observation densities or likelihoods g_i are from the exponential family. Parameters θ in μ , K , or the g_i will be assumed fixed and known for now; for example, θ may contain regression coefficients determining the mean function μ , or variance, smoothness, and range parameters determining a Matérn covariance K .

Our goal is to obtain an approximation of the posterior of \mathbf{y} , which takes the form

$$p(\mathbf{y}|\mathbf{z}) = \frac{\mathcal{N}_n(\mathbf{y}|\boldsymbol{\mu}, \mathbf{K}) \prod_{i=1}^n g_i(z_i|y_i)}{p(\mathbf{z})}, \quad (1)$$

where $\boldsymbol{\mu} = (\mu(\mathbf{s}_1), \dots, \mu(\mathbf{s}_n))'$, and \mathbf{K} is an $n \times n$ covariance matrix with (i, j) entry $(\mathbf{K})_{ij} = K(\mathbf{s}_i, \mathbf{s}_j)$. Once an approximation of the posterior (1) has been obtained, it is conceptually straightforward to extend this result to other quantities of interest, such as the integrated likelihood for inference on parameters θ (see Section 3.2), and prediction of $y(\cdot)$ at unobserved locations (see Section 3.3).

2.2. Review of the Laplace approximation

The normalizing constant $p(\mathbf{z})$ in (1) is not available in closed form for non-Gaussian likelihoods. A popular approach to this issue is the Laplace approximation (e.g., Williams and Barber, 1998; Rasmussen and Williams, 2006, Section 3.4), which approximates $p(\mathbf{z}) = \int \exp(\log p(\mathbf{z}|\mathbf{y}))p(\mathbf{y})d\mathbf{y}$ via a second-order Taylor expansion of $\log p(\mathbf{z}|\mathbf{y})$ at the mode of the posterior density $p(\mathbf{y}|\mathbf{z})$. As this results in an exponentiated quadratic form in \mathbf{y} , it is equivalent to a Gaussian approximation of the likelihood. The mode of $\log p(\mathbf{y}|\mathbf{z})$ does not depend on the normalizing constant, and so it can be obtained using standard optimization procedures such as the Newton–Raphson algorithm. The crucial observation for our later developments is that each Newton–Raphson update in the GGP setting is equivalent to computing the posterior mean of \mathbf{y} given Gaussian pseudo-data (e.g., Rasmussen and Williams, 2006, Sect. 3.4.1). Upon convergence of the algorithm, we have a Laplace approximation for the normalizing constant and a Gaussian approximation for the likelihood, which gives us a Gaussian posterior.

We now go into the details of this approximation. Based on the first and second derivatives of $\log g_i$, we define

$$u_i(y_i) = \frac{\partial}{\partial y_i} \log g_i(z_i|y_i) \quad \text{and} \quad d_i(y_i) = -\left(\frac{\partial^2}{\partial y_i^2} \log g_i(z_i|y_i)\right)^{-1}, \quad i = 1, \dots, n.$$

Stacking these quantities as $\mathbf{u}_y = (u_1(y_1), \dots, u_n(y_n))'$ and $\mathbf{D}_y = \text{diag}(d_1(y_1), \dots, d_n(y_n))$, it is easy to see that $\frac{\partial}{\partial \mathbf{y}} \log p(\mathbf{y}|\mathbf{z}) = -\mathbf{K}^{-1}(\mathbf{y} - \boldsymbol{\mu}) + \mathbf{u}_y$ and $-\frac{\partial^2}{\partial \mathbf{y} \partial \mathbf{y}'} \log p(\mathbf{y}|\mathbf{z}) = \mathbf{K}^{-1} + \mathbf{D}_y^{-1} =: \mathbf{W}_y$. When given the posterior mode $\boldsymbol{\alpha} = \arg\max_{\mathbf{y} \in \mathbb{R}^n} \log p(\mathbf{y}|\mathbf{z})$, the combined Gaussian/Laplace approximation of the posterior is

$$\hat{p}_L(\mathbf{y}|\mathbf{z}) = \mathcal{N}_n(\mathbf{y}|\boldsymbol{\alpha}, \mathbf{W}_\alpha^{-1}).$$

The subscript $\boldsymbol{\alpha}$ in \mathbf{W}_α^{-1} implies evaluation of \mathbf{W}_y at the mode $\boldsymbol{\alpha}$, rather than at an arbitrary \mathbf{y} . To obtain the mode $\boldsymbol{\alpha}$ with the Newton–Raphson algorithm, we start with an initial value $\mathbf{y}^{(0)}$, and update the current guess for $\ell = 0, 1, 2, \dots$ until convergence as $\mathbf{y}^{(\ell+1)} = \mathbf{h}(\mathbf{y}^{(\ell)})$, where

$$\mathbf{h}(\mathbf{y}) = \mathbf{y} - \left(\frac{\partial^2}{\partial \mathbf{y} \partial \mathbf{y}'} \log p(\mathbf{y}|\mathbf{z})\right)^{-1} \left(\frac{\partial}{\partial \mathbf{y}} \log p(\mathbf{y}|\mathbf{z})\right). \quad (2)$$

This Newton–Raphson update is equivalent to computing the posterior mean of \mathbf{y} given Gaussian pseudo-data $\mathbf{t}_y = \mathbf{y} + \mathbf{D}_y \mathbf{u}_y$ with noise covariance matrix \mathbf{D}_y . Specifically, we can write the Newton–Raphson update in (2) as:

$$\mathbf{h}(\mathbf{y}) = \boldsymbol{\mu} + \mathbf{W}_y^{-1} \mathbf{D}_y^{-1} (\mathbf{t}_y - \boldsymbol{\mu}) = \mathbb{E}(\mathbf{y}|\mathbf{t}_y), \quad (3)$$

which is the conditional mean of \mathbf{y} given Gaussian pseudo-data $\mathbf{t}_y|\mathbf{y} \sim \mathcal{N}_n(\mathbf{y}, \mathbf{D}_y)$. The derivation of (3) is straightforward and included in Appendix A for completeness. This means we can obtain the mode $\boldsymbol{\alpha}$ by iterating between (a) computing pseudo-data $\mathbf{t}_{y^{(\ell)}}$ with i th entry $y_i^{(\ell)} + d_i(y_i^{(\ell)})u_i(y_i^{(\ell)})$, and (b) obtaining the posterior mean $\mathbf{y}^{(\ell+1)}$ of \mathbf{y} given $\mathbf{t}_{y^{(\ell)}}$ assuming independent Gaussian noise with variances $d_1(y_1^{(\ell)}), \dots, d_n(y_n^{(\ell)})$.

Some examples of popular likelihoods and the corresponding pseudo-data and pseudo-variances are summarized in Table 1. The Bernoulli and Poisson cases are also illustrated in Fig. 1.

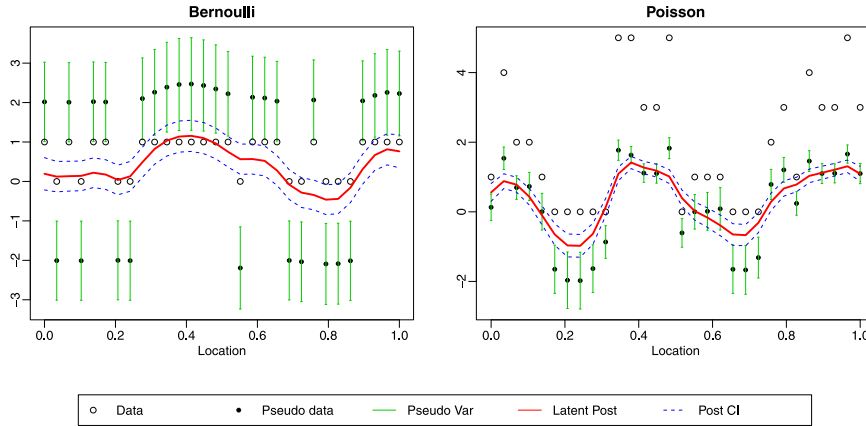


Fig. 1. Pseudo-data \mathbf{t}_α plus or minus half the standard deviation of the pseudo-noise for simulated data \mathbf{z} in one spatial dimension, along with the latent posterior mode α plus or minus half the posterior standard deviation. Note that the data exhibit a different scale than the pseudo-data due to the link function.

Once the algorithm has converged (i.e., $\alpha := \mathbf{y}^{(\ell+1)} = \mathbf{y}^{(\ell)}$), we can use the second-order expansion of the loglikelihood at the mode as a Gaussian approximation of the likelihood based on pseudo-data,

$$\hat{p}_L(\mathbf{z}|\mathbf{y}) = p(\mathbf{t}_\alpha|\mathbf{y}) = \mathcal{N}_n(\mathbf{t}_\alpha|\mathbf{y}, \mathbf{D}_\alpha), \quad (4)$$

or combine it with the Laplace approximation to get a Gaussian approximation of the posterior conditional on pseudo-data,

$$\hat{p}_L(\mathbf{y}|\mathbf{z}) = p(\mathbf{y}|\mathbf{t}_\alpha) = \mathcal{N}_n(\mathbf{y}|\alpha, \mathbf{W}_\alpha^{-1}). \quad (5)$$

For conciseness, we henceforth refer to (5) as the “Laplace approximation”, rather than the more precise “combined Gaussian and Laplace approximation”.

2.3. Review of the general Vecchia approximation

The Laplace approximation described in Section 2.2 allows us to deal with non-Gaussian likelihoods, but it still requires decomposing the $n \times n$ matrix \mathbf{K} and thus scales as $\mathcal{O}(n^3)$. To achieve computational feasibility even for data sizes n in the tens of thousands or more, we also apply a general Vecchia approximation (Katzfuss and Guinness, 2019), which we will briefly review here.

Assume that $\mathbf{y} \sim \mathcal{N}_n(\boldsymbol{\mu}, \mathbf{K})$ is a vector of GP realizations and $\mathbf{t}|\mathbf{y} \sim \mathcal{N}_n(\mathbf{y}, \mathbf{D})$ a vector of noisy data, where \mathbf{D} is diagonal. Then, consider a vector $\mathbf{x} = \mathbf{y} \cup \mathbf{t}$ consisting of the $2n$ elements of \mathbf{y} and \mathbf{t} in some ordering (more details below). It is well known that the density function, $p(\mathbf{x})$, can be factored into a product of univariate conditional densities, $p(\mathbf{x}) = \prod_{i=1}^{2n} p(x_i|\mathbf{x}_{1:i-1})$. The general Vecchia framework extends the approximation of Vecchia (1988) to the vector \mathbf{x} consisting of latent GP realizations and noisy data, resulting in the approximate density

$$\hat{p}(\mathbf{x}) = \prod_{i=1}^{2n} p(x_i|\mathbf{x}_{c(i)}), \quad (6)$$

where $c(i) \subset \{1, \dots, i-1\}$ is a conditioning index set of size m (or of size $i-1$ for $i \leq m$). A small m can lead to enormous computational savings and good approximations; Schäfer et al. (2020) show that under some settings, the approximation error can be bounded when m increases only polylogarithmically with n . While Vecchia is related to composite likelihood (e.g., Varin et al., 2011), most variants of the latter assume some form of marginal or unordered conditional independence, which may reduce approximation accuracy; for more details and numerical comparisons, see Katzfuss and Guinness (2019, Sect. 3.8 and App. D).

As $y_i = y(\mathbf{s}_i)$ is indexed by location and t_i is the corresponding noisy observation, the ordering within \mathbf{y} and within \mathbf{t} is determined by an ordering of the observed locations, $\mathbf{s}_1, \dots, \mathbf{s}_n$. We will use a coordinate-based (left-to-right) ordering in one spatial dimension. In higher-dimensional spaces, we recommend a maxmin ordering (Guinness, 2018; Schäfer et al., 2017), which sequentially chooses each location in the ordering to maximize the minimum distance to previous locations in the ordering.

By straightforward extension of the proof of Prop. 1 in Katzfuss and Guinness (2019) to the case $\boldsymbol{\mu} \neq \mathbf{0}$, it can be shown that the approximation in (6) implies a multivariate normal joint distribution, $\hat{p}(\mathbf{x}) = \mathcal{N}(\boldsymbol{\mu}_x, \mathbf{Q}^{-1})$, where $\mu_{x,i} = \mu(\mathbf{s}_j)$ if $x_i = y_j$ or $x_i = t_j$, $\mathbf{Q} = \mathbf{U}\mathbf{U}'$, and \mathbf{U} is the sparse upper triangular Cholesky factor based on a reverse row-column ordering

of \mathbf{Q} . We write this as $\mathbf{U} = \text{rchol}(\mathbf{Q}) := \text{rev}(\text{chol}(\text{rev}(\mathbf{Q})))$, where $\text{rev}(\cdot)$ reverse-orders the rows and columns of its matrix argument. The nonzero entries of \mathbf{U} are computed directly based on the covariance function K as described in [Appendix B](#).

Let \mathbf{U}_y and \mathbf{U}_t be the submatrices of \mathbf{U} consisting of the rows of \mathbf{U} corresponding to \mathbf{y} and \mathbf{t} , respectively. Then, the sparse matrix $\mathbf{W} = \mathbf{U}_y \mathbf{U}_y'$ is the general Vecchia approximation to the posterior precision matrix of \mathbf{y} given \mathbf{t} . Defining $\mathbf{V} := \text{rchol}(\mathbf{W})$, we can obtain the posterior mean of \mathbf{y} as $E(\mathbf{y}|\mathbf{t}) = \boldsymbol{\mu} - (\mathbf{V}')^{-1} \mathbf{V}^{-1} \mathbf{U}_y \mathbf{U}_t' (\mathbf{t} - \boldsymbol{\mu})$.

2.4. Ordering in Vecchia approximations

We now describe two specific approximations within the general Vecchia framework, which are based on how the elements of \mathbf{y} and \mathbf{t} are ordered in the vector \mathbf{x} in (6): Interweaved (IW) ordering and response-first (RF) ordering. While other ordering and conditioning schemes can also be used in the Vecchia–Laplace methodology to be introduced in Section 3, we recommend these specific schemes to achieve high accuracy while ensuring linear complexity.

2.4.1. Interweaved (IW) ordering

Vecchia-Interweaved (IW) is the sparse general Vecchia approach proposed for likelihood inference in [Katzfuss and Guinness \(2019\)](#), reviewed briefly here. It is a special case of general Vecchia in (6), in which $\mathbf{x} = (y_1, t_1, y_2, t_2, \dots, y_n, t_n)'$ is specified using an interweaved ordering of latent \mathbf{y} and responses \mathbf{t} . We consider the following specific expression for (6):

$$\hat{p}_{IW}(\mathbf{x}) = \prod_{i=1}^n p(t_i | y_i) p(y_i | \mathbf{y}_{q_y(i)}, \mathbf{t}_{q_t(i)}). \quad (7)$$

If $x_j = t_i$, we only condition on y_i , because \mathbf{D} is diagonal and therefore t_i is conditionally independent of all other variables in \mathbf{y} and \mathbf{t} given y_i . If $x_j = y_i$, we condition on $\mathbf{y}_{q_y(i)}$ and $\mathbf{t}_{q_t(i)}$, where $q(i) = q_y(i) \cup q_t(i)$ is the conditioning index vector consisting of the indices of the nearest m locations previous to i in the ordering. For splitting $q(i)$ into $q_y(i)$ and $q_t(i)$, we attempt to maximize $q_y(i)$ while ensuring linear complexity ([Katzfuss and Guinness, 2019](#)). Specifically, for $i = 1, \dots, n$, we set $q_y(i) = (k_i) \cup (q_y(k_i) \cap q(i))$, where $k_i \in q(i)$ is the index whose latent-conditioning set has the most overlap with $q(i)$: $k_i = \arg \max_{j \in q(i)} |q_y(j) \cap q(i)|$, choosing the closest k_i in space to \mathbf{s}_i in case of a tie. In one-dimensional space with coordinate ordering, this results in $q_y(i) = q(i) = (\max(1, i - m), \dots, i - 1)$ and $q_z(i) = \emptyset$. In higher-dimensional space, we may not be able to condition entirely on \mathbf{y} , so the remaining conditioning indices are assigned to $q_t(i) = q(i) \setminus q_y(i)$. These conditioning rules guarantee that \mathbf{U} and \mathbf{V} are both highly sparse with at most m nonzero off-diagonal elements per column. [Katzfuss and Guinness \(2019\)](#) showed that these matrices, and the resulting posterior mean and precision matrix, can be obtained in $\mathcal{O}(nm^3)$ time.

2.4.2. Response-first (RF) ordering

For approximating predictions at observed locations in Algorithm 1 in more than one dimension, we recommend the new RF-full method described in [Katzfuss et al. \(2020a\)](#), reviewed briefly here. RF-full orders first all response variables, then all latent variables: $\mathbf{x} = (\mathbf{t}', \mathbf{y}')' = (t_1, \dots, t_n, y_1, \dots, y_n)'$. We consider the following specific expression for (6):

$$\hat{p}_{RF}(\mathbf{x}) = \prod_{i=1}^n p(t_i) p(y_i | \mathbf{y}_{q_y(i)}, \mathbf{t}_{q_t(i)}).$$

The responses t_i do not condition on anything and are considered independent; this implies a poor approximation to $p(\mathbf{t})$, but it does not affect the posterior distribution $p(\mathbf{y}|\mathbf{t})$, which is the relevant quantity for our purposes. We now assume $q(i) = q_y(i) \cup q_t(i)$ to be set of indices corresponding to the m locations closest to \mathbf{s}_i (including \mathbf{s}_i), not considering the ordering. For any $j \in q(i)$, we then let y_i condition on y_j if it is ordered previously in \mathbf{x} ; otherwise, we condition on t_j . More precisely, we set $q_y(i) = \{j \in q(i) : j < i\}$ and $q_t(i) = \{j \in q(i) : j \geq i\}$. Similar to IW, RF-full inference can be carried out in $\mathcal{O}(nm^3)$ time ([Katzfuss et al., 2020a](#)).

3. Vecchia–Laplace methods

We now introduce our Vecchia–Laplace (VL) approximation, which allows fast inference for large datasets modeled using GGP, by combining the Laplace and general Vecchia approximations reviewed in Section 2.

3.1. The VL algorithm

To apply a Laplace approximation, it is first necessary to find the mode of the posterior density of \mathbf{y} . Rapid convergence to the mode can be achieved using a Newton–Raphson algorithm, which can be viewed as iteratively computing a new value $\mathbf{y}^{(l+1)}$ as the posterior mean of the latent GP realization \mathbf{y} based on Gaussian pseudo-data $\mathbf{t} = \mathbf{t}_{y^{(l)}}$, as discussed in Section 2.2. Our VL algorithm applies a general Vecchia approximation $\hat{p}(\mathbf{x})$ to the joint distribution of $\mathbf{x} = \mathbf{y} \cup \mathbf{t}$ at each iteration l , and computes the posterior mean of \mathbf{y} given \mathbf{t} under this approximate distribution. We recommend IW ordering

(Section 2.4.1) in one spatial dimension, and RF ordering (Section 2.4.2) when working in more than one dimension. The resulting VL algorithm is presented as Algorithm 1. After convergence, we obtain the approximation

$$\hat{p}_{VL}(\mathbf{y}|\mathbf{z}) = \mathcal{N}_n(\mathbf{y}|\alpha_V, \mathbf{W}_V^{-1}). \quad (8)$$

Algorithm 1 Vecchia–Laplace (VL)

```

1: procedure VECCHIA-SPECIFY( $\mathcal{S}, m$ )                                ▷ Define Vecchia Structure
2:   Order locations  $\mathcal{S}$  using coordinate (in 1D) or maxmin ordering (in 2D or higher)
3:   For VL-IW, determine variable ordering and conditioning as in Section 2.4.1
4:   For VL-RF, determine variable ordering and conditioning as in Section 2.4.2
5:   return ordering and conditioning info in Vecchia Approximation Object VAO
6: end procedure

7: procedure VL-INFERENCE( $\mathbf{z}, \text{VAO}, g_i, \boldsymbol{\mu}, K$ )                      ▷ Maximize GP Posterior
8:   Derive  $u_i(\cdot) = \frac{\partial}{\partial \mathbf{y}} \log g_i|_{(\cdot)}$  and  $d_i(\cdot) = -(\frac{\partial^2}{\partial \mathbf{y}^2} \log g_i)^{-1}|_{(\cdot)}$ 
9:   Initialize  $\mathbf{y}^{(0)} = \boldsymbol{\mu}$ 
10:  for  $l=0,1,\dots$  do
11:    Compute  $\mathbf{u} = (u_1(y_1^{(l)}), \dots, u_n(y_n^{(l)}))'$  and  $\mathbf{D} = \text{diag}(d_1(y_1^{(l)}), \dots, d_n(y_n^{(l)}))$ 
12:    Update pseudo-data  $\mathbf{t} = \mathbf{y}^{(l)} + \mathbf{D}\mathbf{u}$ 
13:    Compute  $\mathbf{U}$  (see Appendix B) based on  $\mathbf{D}, K$ , and VAO
14:    Extract submatrices  $\mathbf{U}_y$  and  $\mathbf{U}_t$ 
15:    Compute  $\mathbf{W} = \mathbf{U}_y\mathbf{U}_y'$  and  $\mathbf{V} = \text{rchol}(\mathbf{W})$ 
16:    Compute the new posterior mean:  $\mathbf{y}^{(l+1)} = \boldsymbol{\mu} - (\mathbf{V}')^{-1}\mathbf{V}^{-1}\mathbf{U}_y\mathbf{U}_t'(\mathbf{t} - \boldsymbol{\mu})$ 
17:    if  $\|\mathbf{y}^{(l+1)} - \mathbf{y}^{(l)}\| < \epsilon$  then
18:      return  $\alpha_V = \mathbf{y}^{(l+1)}$  and  $\mathbf{W}_V = \mathbf{W}$                                 ▷ Posterior Mode Estimate
19:    end if
20:  end for
21: end procedure

```

Once the algorithm has converged and the posterior mean α_V and precision \mathbf{W}_V have been obtained, the posterior distribution in (8) can be used for estimation of the integrated likelihood (Section 3.2) and for prediction at unobserved locations (Section 3.3). As we will see in our simulation studies later, even for moderate m , the VL procedure in Algorithm 1 essentially finds the exact mode of the posterior.

3.2. Integrated likelihood for parameter inference

In the case of unknown parameters $\boldsymbol{\theta}$ in μ , K , or in the g_i , we would like to carry out parameter inference based on the integrated likelihood,

$$\mathcal{L}(\boldsymbol{\theta}) = p(\mathbf{z}|\boldsymbol{\theta}) = \int p(\mathbf{z}|\mathbf{y}, \boldsymbol{\theta})p(\mathbf{y}|\boldsymbol{\theta})d\mathbf{y}.$$

However, this quantity is exactly the unknown normalizing constant in the denominator of (1), and the integral can generally not be carried out analytically. Instead, we will base parameter inference on the integrated likelihood implied by our VL approximation. In the following, we will again suppress dependence on $\boldsymbol{\theta}$ for ease of notation.

First, rearranging terms in (1), we have $p(\mathbf{z}) = p(\mathbf{z}|\mathbf{y})p(\mathbf{y})/p(\mathbf{y}|\mathbf{z})$. The Laplace approximation approximates the posterior in the denominator as $\hat{p}_L(\mathbf{y}|\mathbf{z}) = p(\mathbf{y}|\mathbf{t}_\alpha)$ (see (5)). Noting that rearranging the definition of a conditional density gives $p(\mathbf{y}) = p(\mathbf{y}, \mathbf{t})/p(\mathbf{t}|\mathbf{y})$, we obtain the Laplace approximation of the integrated likelihood:

$$\mathcal{L}_L(\boldsymbol{\theta}) = \hat{p}_L(\mathbf{z}) = \frac{p(\mathbf{y}, \mathbf{t})}{p(\mathbf{y}|\mathbf{t})} \cdot \frac{p(\mathbf{z}|\mathbf{y})}{p(\mathbf{t}|\mathbf{y})} = p(\mathbf{t}) \cdot \frac{p(\mathbf{z}|\mathbf{y})}{p(\mathbf{t}|\mathbf{y})}, \quad (9)$$

where the terms are evaluated at $\mathbf{y} = \boldsymbol{\alpha}$ and $\mathbf{t} = \mathbf{t}_\alpha$. In this form, the approximation of the integrated likelihood of the data \mathbf{z} can be interpreted as a product of the integrated likelihood of the Gaussian pseudo-data $p(\mathbf{t})$, times a correction term given by the ratio of the true likelihood to the Gaussian likelihood of the pseudo-data: $p(\mathbf{z}|\mathbf{y})/p(\mathbf{t}|\mathbf{y}) = \prod_{i=1}^n g_i(z_i|y_i)/\mathcal{N}(t_i|y_i, d_i)$.

To achieve scalability, we approximate the density $p(\mathbf{t}) = p(\mathbf{x})/p(\mathbf{y}|\mathbf{t})$ as implied by the IW approximation $\hat{p}_{IW}(\mathbf{x})$ in (7). The resulting expression for $\hat{p}_{IW}(\mathbf{t})$ is derived in Katzfuss and Guinness (2019) for the case of $\boldsymbol{\mu} = \mathbf{0}$. We show in Section S1 that the approximate density essentially has the same form if the prior mean is not zero:

$$-2 \log \hat{p}_{IW}(\mathbf{t}) = -2 \sum_i \log \mathbf{U}_{ii} + 2 \sum_i \log \mathbf{V}_{ii} + \tilde{\mathbf{t}}'\tilde{\mathbf{t}} - \tilde{\mathbf{t}}'\tilde{\mathbf{t}} + n \log(2\pi),$$

where $\tilde{\mathbf{t}} = \mathbf{U}_t'(\mathbf{t} - \boldsymbol{\mu})$ and $\tilde{\mathbf{t}} = \mathbf{V}^{-1}\mathbf{U}_y\tilde{\mathbf{t}}$.

Thus, for a specific parameter value θ , we run Algorithm 1 based on θ to obtain α_V , set $\mathbf{y} = \alpha_V$, $\mathbf{t} = \mathbf{t}_{\alpha_V}$, and $d_i = (\mathbf{D}_{\alpha_V})_{ii}$, and then evaluate the VL integrated likelihood as

$$\mathcal{L}_{VL}(\theta) = \hat{p}_{VL}(\mathbf{z}|\theta) = \hat{p}_{IW}(\mathbf{t}) \prod_{i=1}^n \frac{g_i(z_i|y_i)}{\mathcal{N}(t_i|y_i, d_i)}. \quad (10)$$

We can plug $\mathcal{L}_{VL}(\theta)$ into any numerical likelihood-based inference procedure, such as numerical optimization for finding the maximum likelihood estimator of θ , or sampling-based algorithms for finding the posterior of θ . In an iterative inference procedure, we recommend initializing $\mathbf{y}^{(0)}$ in Algorithm 1 at the mode α_V obtained for the previous parameter value. Our integrated likelihood can also be used directly to evaluate the posterior of θ over a grid of high-probability points (Rue et al., 2009, Sect. 3.1). An extension to the integrated nested Laplace approximation (INLA) that improves the accuracy of the marginal posteriors of the y_i (Rue et al., 2009, Sect. 3.2) is straightforward.

3.3. Predictions at unobserved locations

We now consider making predictions at n^* unobserved locations, $\mathcal{S}^* = \{\mathbf{s}_1^*, \dots, \mathbf{s}_{n^*}^*\}$, by obtaining the posterior distribution of $\mathbf{y}^* = (y_1^*, \dots, y_{n^*}^*)'$ with $y_i^* = y(\mathbf{s}_i^*)$. Using the Laplace approximation as expressed in (4), GGP predictions are approximated as GP predictions given Gaussian pseudo-data \mathbf{t}_α with noise covariance matrix \mathbf{D}_α .

Hence, to obtain scalable predictions at unobserved locations, we use the recommended prediction methods in Katzfuss et al. (2020a) that apply Vecchia approximations to the multivariate normal vector $\tilde{\mathbf{x}} = \mathbf{t} \cup \mathbf{y} \cup \mathbf{y}^*$. For one-dimensional space, we use an extension of IW called LF-auto in Katzfuss et al. (2020a), and for higher-dimensional space we use the RF-full method of Katzfuss et al. (2020a). In both cases, the pseudo-data $\mathbf{t} = \mathbf{t}_{\alpha_V}$ and the noise variances $\mathbf{D} = \mathbf{D}_{\alpha_V}$ are evaluated at the approximate mode α_V obtained using Algorithm 1. Based on this approximation, we can compute the implied posterior distribution of $\tilde{\mathbf{y}} = \mathbf{y} \cup \mathbf{y}^*$ as described in Section 2.3: $\hat{p}(\tilde{\mathbf{y}}|\mathbf{t}) \sim \mathcal{N}(\tilde{\boldsymbol{\mu}}, (\tilde{\mathbf{V}}\tilde{\mathbf{V}}')^{-1})$. Katzfuss et al. (2020a) describe how to efficiently extract quantities of interest from this distribution, including the posterior mean and variances at unobserved locations. Finally, summaries or samples from the posterior of $\tilde{\mathbf{y}}$ can be transformed to the data scale using the likelihood function $g(z|y)$, if desired. Sometimes it is difficult to compute certain predictive summaries at the data scale analytically, but it is always possible to approximate them via sampling.

Algorithm 2 in Section S2 provides pseudo-code for maximum-likelihood estimation of parameters and for prediction.

3.4. Properties

3.4.1. Complexity

Inference for GPs with independent Gaussian noise using the Vecchia approximations considered here requires $\mathcal{O}(nm^3)$ time, where m is the maximum size of the conditioning sets $q(i)$, and can be easily parallelized (Katzfuss and Guinness, 2019; Katzfuss et al., 2020a). Our VL Algorithm 1 iteratively computes the Vecchia approximation multiple times until convergence, only adding $\mathcal{O}(n)$ cost at each iteration for computing the pseudo-data $\mathbf{t}_{y^{(l)}}$. Hence, the VL algorithm requires $\mathcal{O}(knm^3)$ time, where k , the number of iterations required until convergence, can be very small (often, $k < 10$).

Once α_V has been determined using Algorithm 1, evaluating the integrated likelihood (10) for parameter inference requires $\mathcal{O}(nm^3)$ time (Katzfuss and Guinness, 2019), and prediction at n^* unobserved locations requires $\mathcal{O}((n + n^*)m^3)$ time (Katzfuss et al., 2020a). Thus, all computational costs are linear in n for fixed m .

3.4.2. Approximation errors

Our VL approximation $\hat{p}_{VL}(\mathbf{y}|\mathbf{z}) = \mathcal{N}_n(\mathbf{y}|\alpha_V, \mathbf{W}_V^{-1})$ in (8) has two sources of error relative to the true posterior $p(\mathbf{y}|\mathbf{z})$: the Vecchia approximation and the Laplace approximation. Both errors are difficult to quantify in general, but our numerical experiments in Section 4 show that our approximation can be very accurate. The error due to the Vecchia approximation can always be reduced by increasing m (e.g., Katzfuss et al., 2020a).

The error of the Laplace approximation is known to depend on the likelihood being approximated. Laplace is exact for Gaussian likelihoods, in which case the VL approximation reverts to the general Vecchia approximation. For non-Gaussian spatial data, theoretical error bounds are difficult to obtain Rue et al. (e.g., 2009, Sect. 4.1). From an empirical point of view, Fong et al. (2010) affirm the non-spatial results of Breslow and Clayton (1993), showing that INLA, an extension of the Laplace approximation, generally performs well for GGPs, with the exception of some types of binomial data. Bonat and Ribeiro (2016) provide a thorough simulation study comparing Laplace to MCMC methods for parameter estimation in the case of binomial, Poisson, and negative-binomial spatial data; they conclude that the Laplace approximation is “a safe option” that is computationally practical.

3.4.3. Convergence

For GGPs as described in Section 2.1, the log-posterior in (1) is concave under appropriate parameterizations. Existing results show that the Newton–Raphson algorithm used in the Laplace approximation is then theoretically guaranteed to converge to its mode (e.g., Boyd and Vandenberghe, 2004, Section 9.5.2). In our VL Algorithm 1, the distribution $\hat{p}(\mathbf{y})$ implied by the general Vecchia approximation changes at each iteration, which makes it difficult to theoretically guarantee




Fig. 2. RRMSE versus time (on a log scale) for Bernoulli data of size $n = 625$ on the unit square. Laplace is run once until convergence. For VL-RF, we considered $m \in \{1, 5, 10, 20, 40\}$. The number of HMC iterations varies from 10,100 to 1,000,000 in increments of 100, with the first 10,000 considered burn-in.

convergence, except in special cases. Fortunately, empirical testing of Algorithm 1 under different parameter and data settings showed that convergence can always be expected when machine precision is not an issue.

4. Simulations and comparisons

We compared our VL approaches to other methods using simulated data. Throughout Section 4, unless specified otherwise, we simulated realizations \mathbf{y} on a grid of size $\sqrt{n} \times \sqrt{n}$ on the unit square from a GP with mean zero and a Matérn covariance function with variance 1, smoothness ν , and range parameter $\lambda = 0.05$. Gridded locations allow us to carry out simulations for large n using Fourier methods. The data were then generated conditional on \mathbf{y} using the four likelihoods in Table 1, with $a = 2$ in the Gamma case.

As low-rank approximations are very popular for large spatial data, we also considered a fully independent conditional or modified-predictive-process approximation to Laplace with m knots (abbreviated as LowRank here), which is equivalent to VL-IW except that each conditioning set $q_y(i) = (1, \dots, m)$ simply consists of the first m latent variables in maxmin ordering. This equivalence allowed us to run VL and LowRank using the same code base, thus avoiding differences solely due to programming.

Criteria used for comparison are the run time (on a 2017 MacBook Pro), the relative root mean square error (RRMSE) and the difference in log scores (dLS). Results are averaged over 100 simulated datasets, unless noted otherwise. The RRMSE is the root mean square error of the posterior mean of \mathbf{y} obtained by one of the approximation methods relative to the true simulated \mathbf{y} , divided by the RMSE of the Laplace approximation. The log score is computed as the negative logarithm of the approximated posterior density of \mathbf{y} evaluated at the true \mathbf{y} , with low values corresponding to well calibrated and sharp posterior distributions (e.g., Gneiting and Katzfuss, 2014, Sect. 3). The dLS is the log score of an approximation method minus the log score for the Laplace approximation. When averaged over a sufficient number of simulated data, the dLS can be shown to approximate the difference between the Kullback–Leibler (KL) divergence of the exact posterior distribution and the considered approximation, minus the KL divergence between the exact distribution and the Laplace approximation.

4.1. Comparison to MCMC

Non-Gaussian spatial models are often fitted using Markov chain Monte Carlo (MCMC), which under mild regularity conditions is “exact approximate”, converging to the true posterior as the number of iterations approaches infinity. For finite computation time and large n , however, MCMC results can be very poor relative to the Laplace approximation. We demonstrate this with a single simulated dataset consisting of $n = 625$ Bernoulli observations based on a GP with smoothness parameter $\nu = 0.5$ on the unit square. We compared Laplace and VL-RF to Hamiltonian Monte Carlo (HMC) (Neal et al., 2011), a MCMC method well suited to sampling correlated variables. As shown in Fig. 2, VL quickly achieved the same accuracy as Laplace as m increased, but at a fraction of the computing time. In contrast, HMC took orders of magnitude longer to achieve similar accuracy. Even with 1 million iterations, the RMSE for HMC was slightly higher than for VL; this is in line with existing simulation studies suggesting that the Laplace approximation error may be negligible in many GGP settings (see Section 3.4.2). We expect the relative performance of HMC to degrade further as n increases. More details and results can be found in Section S3.

4.2. Computational scaling of Laplace approximations

While the Laplace approximation is very useful for moderate data sizes n , we now briefly illustrate the computational infeasibility for large n due to its cubic scaling. In Fig. 3, we show the average computation time for observations with

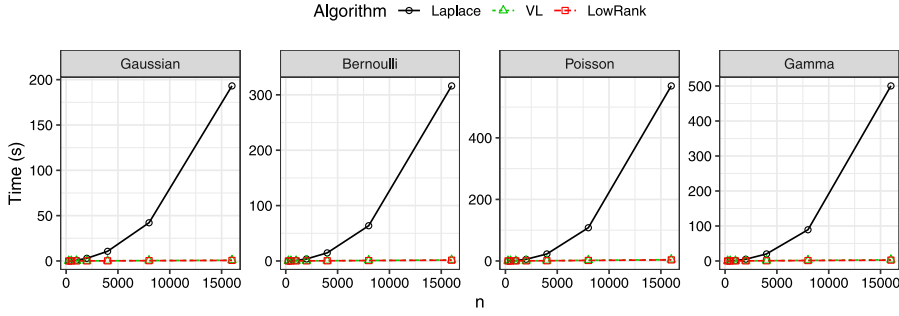


Fig. 3. For sample size n between 250 and 16,000, computing time for the Laplace approximation based on Newton–Raphson, compared to VL and LowRank using Algorithm 1 with $m = 10$.

smoothness $\nu = 0.5$ in the setting described later in Section 4.4. Clearly, Laplace using Newton–Raphson quickly became infeasibly slow as n increased, while VL and LowRank were much faster.

4.3. VL accuracy in one-dimensional space

We now compare the accuracy of the VL and LowRank approximations. Both approaches scale linearly in n for fixed m , and both approaches converge to the Laplace approximation as m increases, with equivalence guaranteed for $m = n - 1$.

Fig. 4 shows the average results for 100 simulated datasets of size $n = 2500$ each, on the unit interval. For the Gaussian likelihood, the noise variance was $\tau^2 = 0.1^2$. Clearly, VL-IW was extremely accurate and delivered essentially equivalent results to the Laplace approximation, even for very small m . For exponential covariance (i.e., Matérn with smoothness $\nu = 0.5$), an exact screening effect holds in one-dimensional space, and so VL-IW is exactly equal to Laplace for any $m \geq 1$. LowRank required much larger m to achieve equivalence to Laplace.

4.4. VL accuracy in two-dimensional space

Fig. 5 shows results for the same simulation study as in Section 4.3, except that the data were simulated on the two-dimensional unit square, with noise variance $\tau^2 = 0.1$ for the Gaussian likelihood. While all methods are again equivalent to Laplace for $m = n - 1$, the two-dimensional problem is considerably more difficult, and higher values of m were required for accurate approximations. As we can see, the recommended VL-RF had roughly equivalent performance to Laplace once m reached 20, and it was more accurate than VL-IW for $m > 10$. LowRank performed considerably worse than the VL methods, and further simulations (not shown) showed that in some cases LowRank approached the accuracy of Laplace only when m was almost as large as n . A simulation with larger range parameter $\lambda = 0.2$ is shown in Section S4.1 of the supplement; while VL-RF was still more accurate than LowRank for all settings, the larger range reduced the amount of fine-scale variation, thus reducing the advantage of VL over LowRank relative to Fig. 5, especially for logistic regression models. The relative performance of the methods was similar in higher dimensions; plots for 3 and 4 dimensions are shown in Section S4.2.

For larger n , the differences between LowRank and VL became even more pronounced. Fig. 6 shows the RMSE for simulations with increasing sample size n but fixed m . VL-RF improved in accuracy under this asymptotic in-fill scenario almost as fast as Laplace, while LowRank failed to improve.

4.5. Simulations for log-Gaussian Cox processes

Point patterns are sets of points or locations $\mathbf{s}_1, \dots, \mathbf{s}_N$ in a domain \mathcal{D} . A popular model for point patterns is the log-Gaussian Cox process (LGCP), a doubly stochastic Poisson process whose intensity function $\lambda(\cdot)$ is modeled as random, $\log \lambda(\cdot) = y(\cdot) \sim GP(\mu, C)$. Inference for LGCPs is difficult due to stochastic integrals.

A natural approximation (e.g., Diggle et al., 2013) for LGCPs relies on partitioning the domain \mathcal{D} into n grid cells A_1, \dots, A_n with center points $\mathbf{a}_1, \dots, \mathbf{a}_n$, respectively. The number of observed points falling into A_i is treated as the data, $z_i = z(A_i) = \sum_{j=1}^N 1_{\mathbf{s}_j \in A_i}$. These gridded data conditionally follow a Poisson distribution, $z_1, \dots, z_n \mid y(\cdot) \stackrel{\text{ind.}}{\sim} \mathcal{P}(\mu(A_i))$, where

$$\mu(A_i) = \int_{A_i} \lambda(\mathbf{s}) d\mathbf{s} \approx |A_i| \lambda(\mathbf{a}_i) = |A_i| e^{y(\mathbf{a}_i)}.$$

This model falls under the GGP framework, so we can apply our VL methods to obtain fast inference for point patterns.

Fig. 7 shows a LGCP whose log-intensity is modeled as a GP with Matérn covariance with range parameter 2.5 on a spatial domain $\mathcal{D} = [0, 50]^2$, discretized into $n = 2500 = 50 \times 50$ unit-square grid cells. This is equivalent to the simulation in Section 4.4, as the domain can be scaled to a unit-square domain with range 0.05, but on the original scale

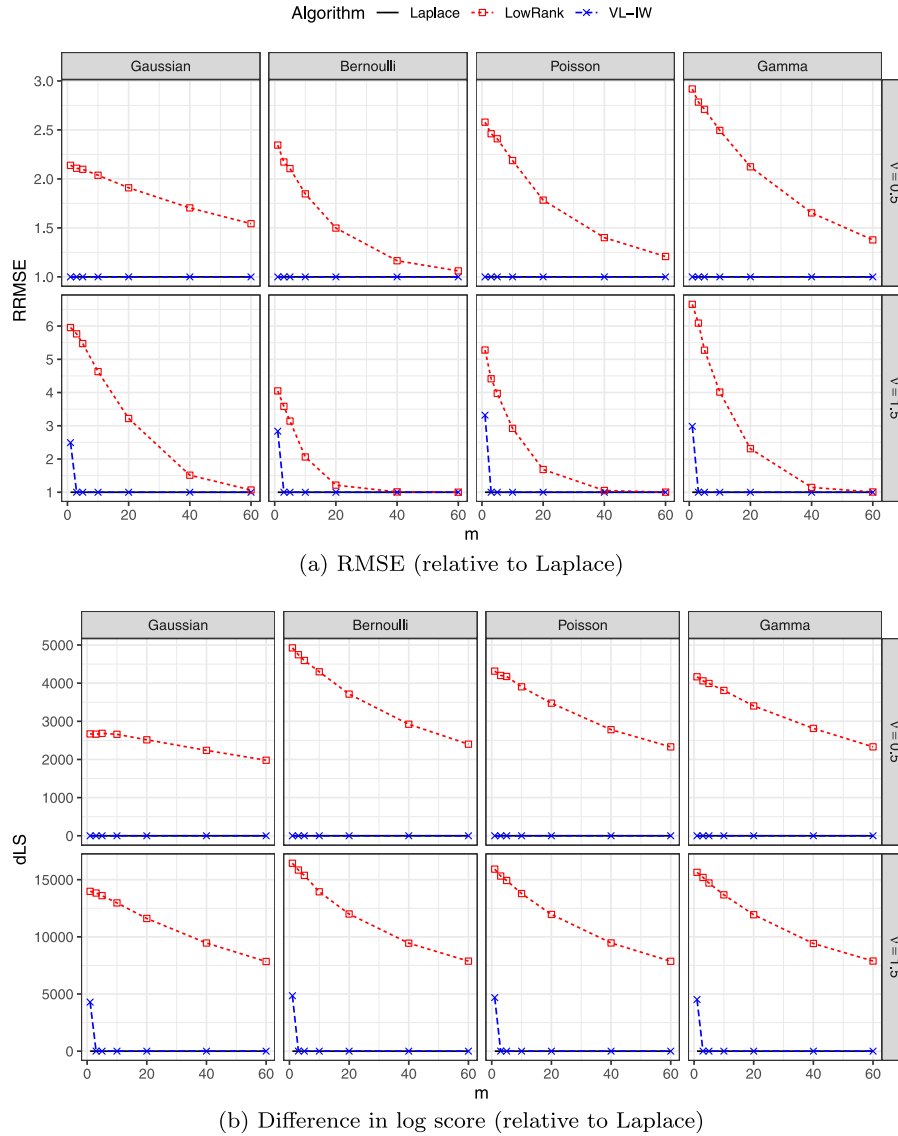


Fig. 4. Simulation results for $n = 2500$ observations on a **one-dimensional** spatial domain.

the grid induces areal regions with unit area, $|A_i| = 1$, with intensity function $\mu(A_i) = \exp(y(\mathbf{a}_i))$. Thus, the averaged results for fitting repeatedly simulated datasets from this LGCP are equivalent to the Poisson results shown in the third column of Fig. 5, indicating that VL can be used to obtain virtually equivalent inference to that using a Laplace algorithm, albeit at much lower computational cost for large n .

4.6. Parameter estimation

We also explored parameter estimation based on each method's integrated-likelihood approximation. Specifically, we considered Poisson data at $n = 625$ locations in the unit square, based on a GP with true smoothness $\nu = 0.5$ and range $\lambda = 0.05$.

First, we simulated a single realization of the spatial data. Holding the variance fixed at the true value of one, we sequentially evaluated the integrated likelihood on a grid of values for the range and smoothness parameters, using the Laplace approximation in (9), and the VL-RF approximation with $m = 20$ in (10). The exact integrated likelihood is intractable. As shown in Fig. 8, the integrated likelihoods as approximated by Laplace and by VL were almost identical, while the LowRank approximation was quite poor. These likelihood approximations are equivalent to approximations to the posterior distribution $p(\boldsymbol{\theta}|\mathbf{z})$ assuming flat priors for $\boldsymbol{\theta}$. This indicates that Bayesian inference for GGPs can be carried out quickly and accurately using the VL approximation.

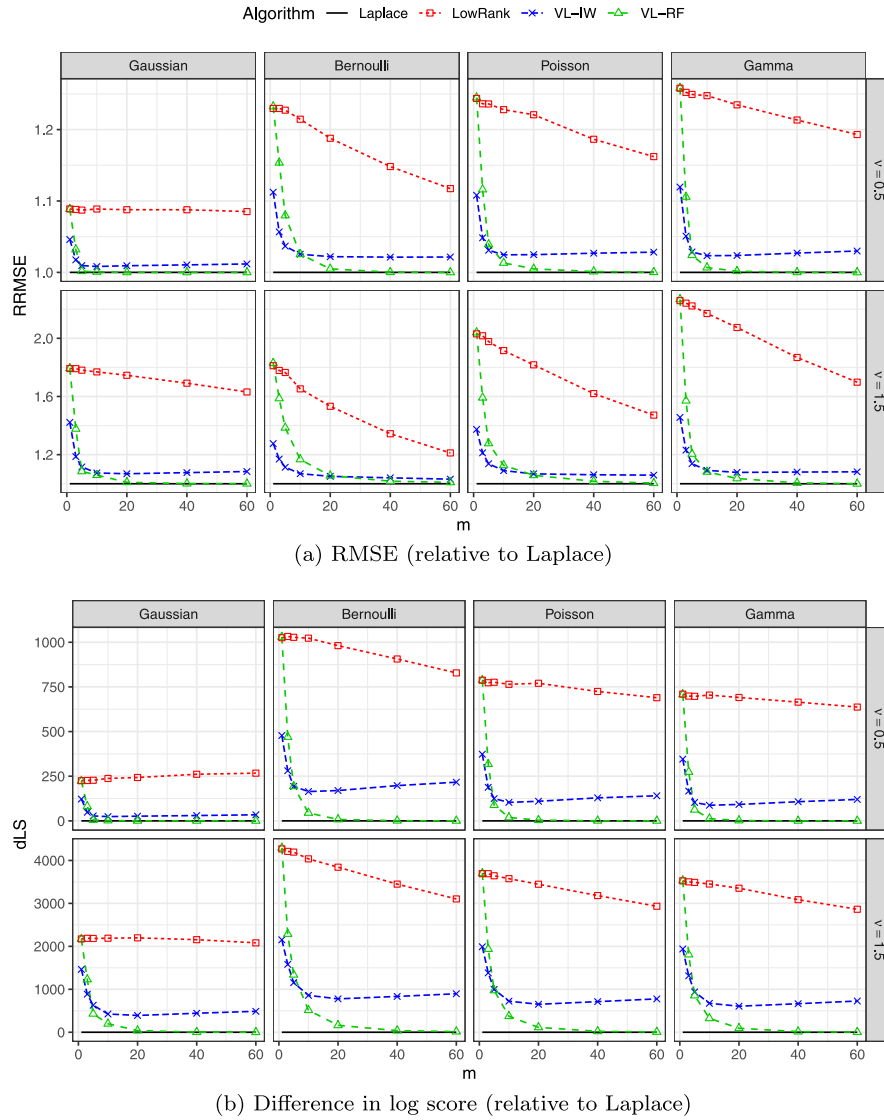


Fig. 5. Simulation results for $n = 2500$ observations on a **two-dimensional** spatial domain.




Fig. 6. On a two-dimensional domain with $v = 0.5$ and fixed $m = 10$, RMSE between true \mathbf{y} and posterior mode α_v for increasing sample size n (on a log scale) up to 300,000. Laplace without further approximation becomes prohibitively expensive for large n , so we only computed it up to $n = 16,000$.






Fig. 7. Gridding a simulated LGCP point pattern: The latent log-intensity $y(\cdot)$ (left), a corresponding simulated point pattern (center), and the down-sampled Poisson count data used for analysis on a $n = 50 \times 50 = 2500$ grid (right).



(a) For a single realization of Poisson data, contour lines of the integrated likelihood at $(-650, -900, -2000, -5000, -10000, -15000)$



(b) Parameter estimates obtained by optimizing the integrated likelihood for 100 sample realizations; also see Table 2

Fig. 8. For Poisson data at $n = 625$ locations in the unit square, comparison of different approximations to the integrated likelihood, using conditioning sets of size $m = 20$ for VL and LowRank. Red dots show the true parameter values.

We then simulated 100 different realizations of the spatial Poisson data and examined the parameter estimates obtained by maximizing the different approximations to the integrated likelihood. The scatter plot in Fig. 8 shows the parameter estimates, using $m = 20$ conditioning points for VL and LowRank. While the estimates using Laplace and VL were similar, LowRank had significant outliers that increased the RMSE of the parameter estimates (see Table 2). The LowRank parameter estimation frequently diverged due to the rough likelihood surface, and for those cases we repeated the optimization with bounds $(0.001, 20)$ for both range and smoothness, but LowRank still failed repeatedly.

4.7. Interpretation of simulation results

In our simulations, VL provided similar accuracy as Laplace with a considerably smaller number m of conditioning points compared to LowRank. The time required per iteration for VL approaches is $\mathcal{O}(nm^3)$. At the expense of fully parallel computation, LowRank can be carried out in $\mathcal{O}(nm^2)$ time by computing the decomposition of the covariance of the conditioning set once at the beginning of the procedure. However, as VL with any given m , say $m = \tilde{m}$, was substantially more accurate than LowRank with $m = \tilde{m}^{3/2}$, we conclude that VL is more computationally efficient than LowRank for a given approximation accuracy, except for very smooth posteriors. The improvement in accuracy for VL relative to LowRank became even more pronounced as we increased the sample size under in-fill asymptotics.

Table 2

For 100 simulated Poisson datasets at $n = 625$ locations in the unit square, RMSE for parameter estimates based on different approximations to the integrated likelihood. Both range and smoothness parameters were bounded to the interval $[0.001, 20]$, but LowRank estimation still failed repeatedly.

	Range			Smoothness		
	$m = 10$	$m = 20$	$m = 40$	$m = 10$	$m = 20$	$m = 40$
LowRank	0.107	0.407	0.098	9.17	8.40	12.60
VL	0.293	0.040	0.023	1.01	0.78	0.47
Laplace			0.023			0.51

Table 3

For the MODIS data, comparison of prediction scores (lower is better) between VL and LowRank.

Method	MSE	CRPS
VL	0.0149	0.144
LowRank	0.0528	0.170
Ratio	$3.54 \times$	$1.18 \times$

5. Application to satellite data

We applied our methodology to a large, spatially correlated, non-Gaussian dataset of column water vapor. These data were collected by NASA's Moderate Resolution Imaging Spectroradiometer (MODIS), which is mounted on the NASA Aqua satellite (Borbás et al., 2017). We considered a Level-2 near-infrared dataset of total precipitable water at a $1354 \times 2030 = 2,746,820$ grid of 1km pixels. We used up to 500,000 of these data points for our demonstration. Our dataset was observed between 13:45 and 13:50 on March 28, 2019 over a rectangular region off the coast of west Africa with west, north, east, and south bounding coordinates -42.707 , 67.476 , 4.443 , and 45.126 , respectively and was found on the NASA Earthdata website, <https://earthdata.nasa.gov>.

Precipitable water amounts are continuous and strictly positive, with values near 0 corresponding to clear skies and larger values implying more water. Exploratory plots showed a right-skewed density, so we assumed that the data can be modeled using a spatial generalized GP with a Gamma likelihood:

$$z(\mathbf{s}_i)|y(\mathbf{s}_i) \stackrel{\text{ind.}}{\sim} \mathcal{G}(a, ae^{-y(\mathbf{s}_i)}), \quad y(\cdot) \sim \mathcal{N}(\mu, K),$$

where $E(z(\mathbf{s})|y(\mathbf{s})) = \exp(y(\mathbf{s}))$, $\mu(\mathbf{s}) = \beta_1 + \beta_2 \text{lat}(\mathbf{s})$ is a linear trend consisting of an intercept and a latitudinal gradient, and K is an isotropic Matérn covariance function with variance σ^2 , smoothness ν , and range parameter ρ . We estimated the parameter values $\beta_1 = -1.515$, $\beta_2 = 0.000766$, $a = 0.89$, $\sigma^2 = .25$, $\rho = 31$ km, and $\nu = 3$ as described in Section S6.

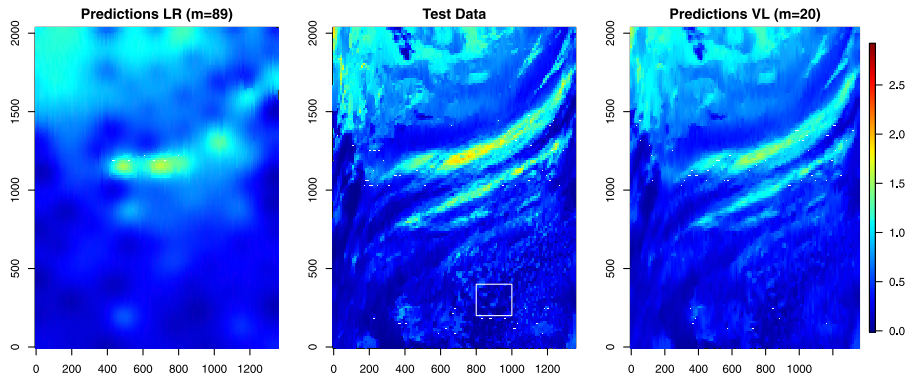
We again compared our VL approach to a LowRank method. We randomly sampled $n = 250,000$ observations \mathbf{z} of the full dataset as training data, and 250,000 of the remaining observations as test data \mathbf{z}^* at locations \mathcal{S}^* . For VL, we set $m = 20$ following our recommendations in Section 4.4 and further justified in Section S6.. For LowRank, we used $m = 89 \approx (20)^{3/2}$ for a computationally fair comparison. On an Intel Xeon E5-2690 CPU with 64 GB RAM, Algorithm 1 for VL required 10 iterations with a total run time of about 18 min (1.8 min per iteration). Taking advantage of an implementation that achieves the $\mathcal{O}(nm^2)$ scaling, each iteration for LowRank required 1.3 min on average across 6 iterations. Note that, based on our numerical experiments, we estimate that Laplace without further approximation would take months of computing time, while HMC-based approaches would take years to achieve the same accuracy as VL.

Fig. 9(a) shows prediction maps of the posterior mean $E(\mathbf{z}^*|\mathbf{z}) = E(\exp(\mathbf{y}^*)|\mathbf{z})$ with i th entry $\exp(E(y_i^*|\mathbf{z}) + \text{var}(y_i^*|\mathbf{z})/2)$. Clearly, much of the fine-scale structure was lost when using LowRank. To further illustrate this issue, we made predictions on a 200×200 grid over a small subregion. As shown in Fig. 9(b), the LowRank predictions were virtually useless at this scale, while VL was able to recover much of the important spatial structure from the noisy and incomplete training data.

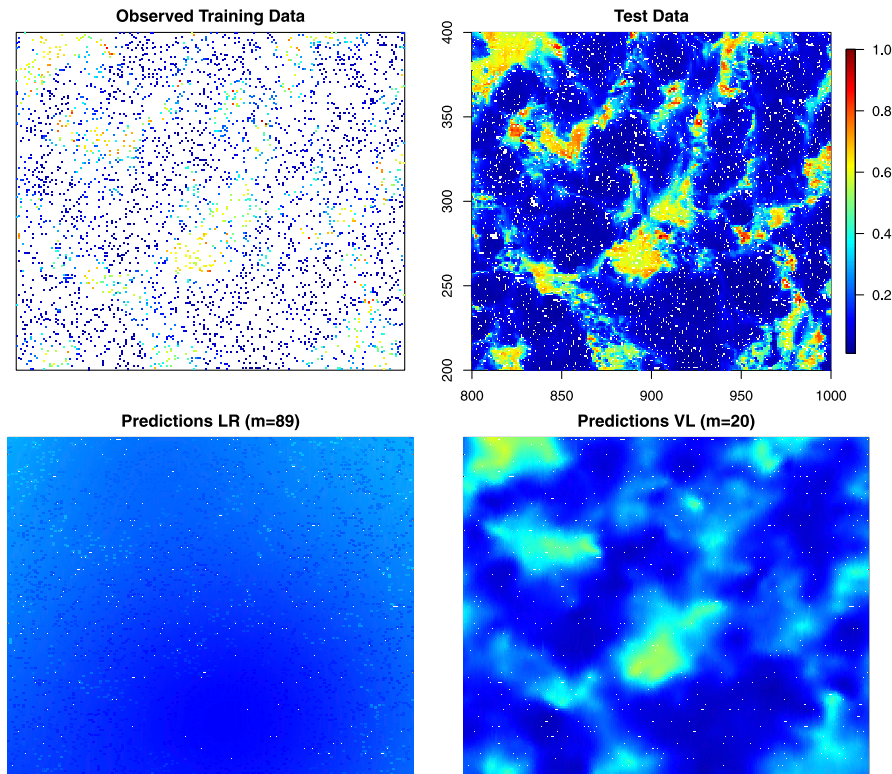
Table 3 quantifies the improvement in predictions using VL over LowRank. We computed the MSE based on the posterior mean $E(\mathbf{z}^*|\mathbf{z})$. To compare the accuracy of the uncertainty quantification, we also computed the continuous ranked probability score (CRPS) (e.g., Gneiting and Katzfuss, 2014), which encourages well calibrated and sharp predictive distributions. Table 3 shows that VL strongly outperformed LowRank for comparable computational complexity.

6. Conclusions and future work

In this work, we presented a novel combination of techniques that allow for efficient analysis of large, spatially correlated, non-Gaussian datasets or point patterns. The key idea is to apply a Vecchia approximation to the Gaussian (and hence tractable) joint distribution of GP realizations and pseudo-data at each iteration of a Newton–Raphson algorithm, leading to a Gaussian Laplace approximation. This allows us to carry out inference for non-Gaussian data by iteratively



(a) Entire spatial domain



(b) Zooming into the white square shown in Panel (a)

Fig. 9. Prediction maps for MODIS data using VL and LowRank (LR).

applying existing Vecchia approximations for Gaussian pseudo-data, which are updated at each iteration. Our Vecchia-Laplace (VL) techniques guarantee linear complexity in the data size while capturing spatial dependence at all scales. Compared to alternative methods such as low-rank approximations or sampling-based approaches, our VL approximations can achieve higher accuracy at a fraction of the computation time.

Vecchia approximations require specification of an ordering of the model variables and of a conditioning set for each variable, and these two issues also play a critical role in the performance of our VL approaches. Through simulation studies, we showed that, in one-dimensional space, interleaving the GP realizations and the pseudo-data (Katzfuss and Guinness, 2019) can provide results that are virtually indistinguishable from Laplace, even for very small conditioning sets. For two-dimensional space, we recommend the response-first Vecchia approximation (Katzfuss et al., 2020a). Due to the computational efficiency of our approach, it is also possible to use a VL approximation of the integrated likelihood for parameter inference, for which we recommend the interleaved ordering in any dimension.

The methods and algorithms proposed here are implemented in the R package *GPvecchia* (Katzfuss et al., 2020). The default settings of the package functions reflect the recommendations in the previous paragraph. The tuning parameter m , which controls a trade-off between accuracy and computation cost, can be set by the user. In practice, a useful strategy is to start with a relatively small value of m and gradually increase it until the inference converges or the computational resources are exhausted.

Our methods and code are applicable in more than two dimensions, but a thorough investigation of their properties in this context will be carried out in future work. For example, Katzfuss et al. (2020b) show that Vecchia-based approximations with appropriate extensions can be highly accurate for computer-model emulation in up to ten dimensions; a combination with our VL methods could allow emulation of non-Gaussian computer-model output. Other potential future work includes extending the Laplace approximation in our methods to an integrated nested Laplace approximation (INLA) that improves the accuracy of the marginal posteriors of the latent variables (Rue et al., 2009, Sect. 3.2); the use of conjugate-gradient (Zhang et al., 2019) or incomplete-Cholesky (Schäfer et al., 2020) methods that allow the computation of the latent posterior mean in linear time even for completely latent Vecchia approximations; or extensions to spatio-temporal filtering using Vecchia approximations based on domain partitioning (Jurek and Katzfuss, 2018, 2020).

Acknowledgments

Katzfuss' research was partially supported by National Science Foundation (NSF) grant DMS-1521676 and NSF CAREER grant DMS-1654083. The authors would like to thank Joe Guinness, Jianhua Huang, Leonid Koralov, Boris Vainberg, and several anonymous reviewers for helpful comments and suggestions. Wenlong Gong, Joe Guinness, Marcin Jurek, and Jingjie Zhang contributed to the R package *GPvecchia*, and Florian Schäfer provided C code for the exact maxmin ordering.

Appendix A. Newton-Raphson update using pseudo-data

The desired Newton-Raphson update has the form

$$\mathbf{h}(\mathbf{y}) = \mathbf{y} - \left(\frac{\partial^2}{\partial \mathbf{y} \partial \mathbf{y}'} \log p(\mathbf{y}|\mathbf{z}) \right)^{-1} \left(\frac{\partial}{\partial \mathbf{y}} \log p(\mathbf{y}|\mathbf{z}) \right). \quad (\text{A.1})$$

As shown in Section 2.2, we have $\frac{\partial}{\partial \mathbf{y}} \log p(\mathbf{y}|\mathbf{z}) = \mathbf{K}^{-1}(\boldsymbol{\mu} - \mathbf{y}) + \mathbf{u}_y$ and $-\frac{\partial^2}{\partial \mathbf{y} \partial \mathbf{y}'} \log p(\mathbf{y}|\mathbf{z}) = \mathbf{K}^{-1} + \mathbf{D}_y^{-1} = \mathbf{W}_y$. Using an idea similar to iterative weighted least squares (Section 2.5, McCullagh and Nelder, 1989), we can premultiply the variable \mathbf{y} by the Hessian to combine terms, and then rearrange and pull out the prior mean. Dropping the iteration subscript of \mathbf{y} for ease of notation, we can write (A.1) as

$$\begin{aligned} \mathbf{h}(\mathbf{y}) &= \mathbf{y} + \mathbf{W}^{-1}(\mathbf{K}^{-1}(\boldsymbol{\mu} - \mathbf{y}) + \mathbf{u}) \\ &= \mathbf{W}^{-1}((\mathbf{K}^{-1} + \mathbf{D}^{-1})\mathbf{y} - \mathbf{K}^{-1}\mathbf{y} + (\mathbf{K}^{-1}\boldsymbol{\mu} + \mathbf{D}^{-1}\boldsymbol{\mu}) - \mathbf{D}^{-1}\boldsymbol{\mu} + \mathbf{D}^{-1}\mathbf{D}\mathbf{u}) \\ &= \boldsymbol{\mu} + \mathbf{W}^{-1}(\mathbf{D}^{-1}(\mathbf{y} + \mathbf{D}\mathbf{u} - \boldsymbol{\mu})) \\ &= \boldsymbol{\mu} + \mathbf{W}^{-1}\mathbf{D}^{-1}(\mathbf{t} - \boldsymbol{\mu}), \end{aligned}$$

where $\mathbf{t} = \mathbf{y} + \mathbf{D}\mathbf{u}$.

Now consider the posterior mean in the case of a Gaussian likelihood $\mathbf{t}|\mathbf{y} \sim \mathcal{N}_n(\mathbf{y}, \mathbf{D})$ with a conjugate Gaussian prior, $\mathbf{y} \sim \mathcal{N}_n(\boldsymbol{\mu}, \mathbf{K})$. Employing a well-known formula, we have

$$\mathbf{E}(\mathbf{y}|\mathbf{t}) = (\mathbf{K}^{-1} + \mathbf{D}^{-1})^{-1}(\mathbf{K}^{-1}\boldsymbol{\mu} + \mathbf{D}^{-1}\mathbf{t}) = \boldsymbol{\mu} + \mathbf{W}^{-1}\mathbf{D}^{-1}(\mathbf{t} - \boldsymbol{\mu}).$$

Thus, we have $\mathbf{h}(\mathbf{y}) = \mathbf{E}(\mathbf{y}|\mathbf{t})$, the posterior mean under the assumption of Gaussian pseudo-data \mathbf{t} .

Appendix B. Computing U

Consider a general Vecchia approximation of the form (6). To obtain \mathbf{U} , define $C(x_i, x_j)$ as the covariance between x_i and x_j implied by the true model; that is, $C(y_i, y_j) = C(t_i, t_j) = K(\mathbf{s}_i, \mathbf{s}_j)$ and $C(t_i, t_j) = K(\mathbf{s}_i, \mathbf{s}_j) + \mathbb{1}_{i=j}d_i$. Then, the (j, i) th element of \mathbf{U} can be calculated as

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

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

Appendix C. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.csda.2020.107081>.

References

- Banerjee, S., Gelfand, A.E., Finley, A.O., Sang, H., 2008. Gaussian predictive process models for large spatial data sets. *J. R. Stat. Soc. Ser. B Stat. Methodol.* (ISSN: 1369-7412) 70 (4), 825–848. <http://dx.doi.org/10.1111/j.1467-9868.2008.00663.x>.
- Bonat, W.H., Ribeiro, Jr., P.J., 2016. Practical likelihood analysis for spatial generalized linear mixed models. *Environmetrics* 27 (2), 83–89.
- Borbas, E., Menzel, P., Gao, B.C., 2017. MODIS atmosphere L2 water vapor product. In: NASA MODIS Adaptive Processing System. Goddard Space Flight Center, http://dx.doi.org/10.5067/MODIS/MYD05_L2.061.
- Boyd, S., Vandenberghe, L., 2004. *Convex Optimization*. Cambridge University Press.
- Bradley, J.R., Holan, S.H., Wikle, C.K., 2018. Computationally efficient multivariate spatio-temporal models for high-dimensional count-valued data (with discussion). *Bayesian Anal.* 13 (1), 253–310.
- Breslow, N.E., Clayton, D.G., 1993. Approximate inference in generalized linear mixed models. *J. Amer. Statist. Assoc.* 88 (421), 9–25.
- Chan, A.B., Dong, D., 2011. Generalized Gaussian process models. In: *CVPR*, pp. 2681–2688.
- Cressie, N., Johannesson, G., 2008. Fixed rank kriging for very large spatial data sets. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 70 (1), 209–226.
- Datta, A., Banerjee, S., Finley, A.O., Gelfand, A.E., 2016. Hierarchical nearest-neighbor Gaussian process models for large geostatistical datasets. *J. Amer. Statist. Assoc.* (ISSN: 0162-1459) 111 (514), 800–812. <http://dx.doi.org/10.1080/01621459.2015.1044091>.
- Dietrich, C.R., Newsam, G.N., 1997. Fast and exact simulation of stationary Gaussian processes through circulant embedding of the covariance matrix. *SIAM J. Sci. Comput.* 18 (4), 1088–1107.
- Diggle, P.J., Moraga, P., Rowlingson, B., Taylor, B.M., 2013. Spatial and spatio-temporal log-Gaussian Cox processes: Extending the geostatistical paradigm. *Statist. Sci.* 28 (4), 542–563.
- Diggle, P., Tawn, J., Moyeed, R., 1998. Model-based geostatistics. *J. R. Stat. Soc. Ser. C. Appl. Stat.* 47 (3), 299–350.
- Filippone, M., Girolami, M., 2014. Pseudo-marginal Bayesian inference for Gaussian processes. *IEEE Trans. Pattern Anal. Mach. Intell.* (ISSN: 0162-8828) 36 (11), 2214–2226. <http://dx.doi.org/10.1109/TPAMI.2014.2316530>.
- Finley, A.O., Sang, H., Banerjee, S., Gelfand, A.E., 2009. Improving the performance of predictive process modeling for large datasets. *Comput. Stat. Data Anal.* (ISSN: 0167-9473) 53 (8), 2873–2884. <http://dx.doi.org/10.1016/j.csda.2008.09.008>.
- Flaxman, S., Wilson, A., Neill, D., Nickisch, H., Smola, A., 2015. Fast Kronecker inference in Gaussian processes with non-Gaussian likelihoods. In: *International Conference on Machine Learning*, pp. 607–616.
- Fong, Y., Rue, H., Wakefield, J., 2010. Bayesian inference for generalized linear mixed models. *Biostatistics* 11 (3), 397–412.
- Furrer, R., Genton, M.G., Nychka, D., 2006. Covariance tapering for interpolation of large spatial datasets. *J. Comput. Graph. Statist.* 15 (3), 502–523.
- Gneiting, T., Katzfuss, M., 2014. Probabilistic forecasting. *Annu. Rev. Stat. Appl.* 1, 125–151.
- Guan, Y., Haran, M., 2018. A computationally efficient projection-based approach for spatial generalized linear mixed models. *J. Comput. Graph. Statist.* 27 (4), 701–714.
- Guinness, J., 2018. Permutation methods for sharpening Gaussian process approximations. *Technometrics* 60 (4), 415–429. <http://dx.doi.org/10.1080/00401706.2018.1437476>.
- Guinness, J., Fuentes, M., 2017. Circulant embedding of approximate covariances for inference from Gaussian data on large lattices. *J. Comput. Graph. Statist.* 26 (1), 88–97.
- 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., Zammit-Mangion, A., 2019. A case study competition among methods for analyzing large spatial data. *J. Agric. Biol. Environ. Stat.* 24 (3), 398–425.
- Higdon, D., 1998. A process-convolution approach to modelling temperatures in the North Atlantic Ocean. *Environ. Ecol. Stat.* (ISSN: 1352-8505) 5 (2), 173–190.
- Hughes, J., Haran, M., 2013. Dimension reduction and alleviation of confounding for spatial generalized linear mixed models. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 75 (1), 139–159.
- Jurek, M., Katzfuss, M., 2018. Multi-resolution filters for massive spatio-temporal data. [arXiv:1810.04200](https://arxiv.org/abs/1810.04200).
- Jurek, M., Katzfuss, M., 2020. Hierarchical sparse Cholesky decomposition with applications to high-dimensional spatio-temporal filtering. [arXiv:2006.16901](https://arxiv.org/abs/2006.16901).
- Katzfuss, M., 2017. A multi-resolution approximation for massive spatial datasets. *J. Amer. Statist. Assoc.* 112 (517), 201–214.
- Katzfuss, M., Cressie, N., 2011. Spatio-temporal smoothing and EM estimation for massive remote-sensing data sets. *J. Time Series Anal.* 32 (4), 430–446. <http://dx.doi.org/10.1111/j.1467-9892.2011.00732.x>.
- Katzfuss, M., Gong, W., 2019. A class of multi-resolution approximations for large spatial datasets. *Statist. Sinica* (in press).
- Katzfuss, M., Guinness, J., 2019. A general framework for Vecchia approximations of Gaussian processes. *Statist. Sci.* (in press).
- Katzfuss, M., Guinness, J., Gong, W., Zilber, D., 2020a. Vecchia approximations of Gaussian-process predictions. *J. Agric. Biol. Environ. Stat.* 25 (3), 383–414.
- Katzfuss, M., Guinness, J., Lawrence, E., 2020b. 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., Schäfer, F., 2020. GPvecchia: Fast Gaussian-process inference using Vecchia approximations.
- Kaufman, C.G., Schervish, M.J., Nychka, D.W., 2008. Covariance tapering for likelihood-based estimation in large spatial data sets. *J. Amer. Statist. Assoc.* 103 (484), 1545–1555.
- Lindgren, F., Rue, H., Lindström, J., 2011. An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 73 (4), 423–498.
- Lipton, R.J., Rose, D.J., Tarjan, R.E., 1979. Generalized nested dissection. *SIAM J. Numer. Anal.* (ISSN: 0036-1429) 16 (2), 346–358. <http://dx.doi.org/10.1137/0716027>.
- McCullagh, P., Nelder, J.A., 1989. *Generalized Linear Models*. CRC Press.
- Neal, R.M., et al., 2011. MCMC using Hamiltonian dynamics. In: *Handbook of Markov Chain Monte Carlo*, Vol. 2, No. 11. CRC Press New York, NY.
- Nickisch, H., Solin, A., Grigorievskiy, A., 2018. State-space Gaussian processes with non-Gaussian likelihood. [arXiv preprint arXiv:1802.04846](https://arxiv.org/abs/1802.04846).
- Nychka, D.W., Bandyopadhyay, S., Hammerling, D., Lindgren, F., Sain, S.R., 2015. A multi-resolution Gaussian process model for the analysis of large spatial data sets. *J. Comput. Graph. Statist.* 24 (2), 579–599.
- Quiñero-Candela, J., Rasmussen, C.E., 2005. A unifying view of sparse approximate Gaussian process regression. *J. Mach. Learn. Res.* 6, 1939–1959.
- Rasmussen, C.E., Williams, C.K.I., 2006. *Gaussian Processes for Machine Learning*. MIT Press, ISBN: 026218253X, <http://dx.doi.org/10.1142/S0129065704001899>.
- Rue, H., Held, L., 2005. *Gaussian Markov Random Fields: Theory and Applications*. CRC press.
- Rue, H., Martino, S., Chopin, N., 2009. Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *J. R. Stat. Soc. Ser. B Stat. Methodol.* (ISSN: 13697412) 71 (2), 319–392. <http://dx.doi.org/10.1111/j.1467-9868.2008.00700.x>.
- Sang, H., Jun, M., Huang, J.Z., 2011. Covariance approximation for large multivariate spatial datasets with an application to multiple climate model errors. *Ann. Appl. Stat.* 5 (4), 2519–2548.
- Schäfer, F., Katzfuss, M., Owahdi, H., 2020. Sparse Cholesky factorization by Kullback–Leibler minimization. [arXiv:2004.14455](https://arxiv.org/abs/2004.14455).

- Schäfer, F., Sullivan, T.J., 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).
- Sengupta, A., Cressie, N., 2013. Hierarchical statistical modeling of big spatial datasets using the exponential family of distributions. *Spat. Stat.* (ISSN: 22116753) 4, 14–44. <http://dx.doi.org/10.1016/j.spasta.2013.02.002>.
- Shang, L., Chan, A.B., 2013. On approximate inference for generalized Gaussian process models. [arXiv:1311.6371](https://arxiv.org/abs/1311.6371).
- Sheth, R., Wang, Y., Khardon, R., 2015. Sparse variational inference for generalized Gaussian Process models. In: *Proceedings of the 32nd International Conference on Machine Learning*, Vol. 37.
- Snelson, E., Ghahramani, Z., 2007. Local and global sparse Gaussian process approximations. In: *Artificial Intelligence and Statistics*, Vol. 11. AISTATS.
- Stein, M.L., 2014. Limitations on low rank approximations for covariance matrices of spatial data. *Spat. Stat.* (ISSN: 22116753) 8, 1–19. <http://dx.doi.org/10.1016/j.spasta.2013.06.003>.
- Stein, M.L., Chi, Z., Welty, L., 2004. Approximating likelihoods for large spatial data sets. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 66 (2), 275–296.
- Tierney, L., Kadane, J.B., 1986. Accurate approximations for posterior moments and marginal densities. *J. Amer. Statist. Assoc.* 81 (393), 82–86.
- Varin, C., Reid, N., Firth, D., 2011. An overview of composite likelihood methods. *Statist. Sinica* 5–42.
- Vecchia, A., 1988. Estimation and model identification for continuous spatial processes. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 50 (2), 297–312.
- Williams, C.K., Barber, D., 1998. Bayesian classification with Gaussian processes. *IEEE Trans. Pattern Anal. Mach. Intell.* 20 (12), 1342–1351.
- Yang, Y., Pilanci, M., Wainwright, M.J., et al., 2017. Randomized sketches for kernels: Fast and optimal nonparametric regression. *Ann. Statist.* 45 (3), 991–1023.
- Zhang, L., Datta, A., Banerjee, S., 2019. Practical Bayesian modeling and inference for massive spatial data sets on modest computing environments. *Stat. Anal. Data Min.: ASA Data Sci. J.* 12 (3), 197–209.