



Correlation-based sparse inverse Cholesky factorization for fast Gaussian-process inference

Myeongjong Kang¹ · Matthias Katzfuss¹

Received: 23 January 2022 / Accepted: 2 March 2023

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2023

Abstract

Gaussian processes are widely used as priors for unknown functions in statistics and machine learning. To achieve computationally feasible inference for large datasets, a popular approach is the Vecchia approximation, which is an ordered conditional approximation of the data vector that implies a sparse Cholesky factor of the precision matrix. The ordering and sparsity pattern are typically determined based on Euclidean distance of the inputs or locations corresponding to the data points. Here, we propose instead to use a correlation-based distance metric, which implicitly applies the Vecchia approximation in a suitable transformed input space. The correlation-based algorithm can be carried out in quasilinear time in the size of the dataset, and so it can be applied even for iterative inference on unknown parameters in the correlation structure. The correlation-based approach has two advantages for complex settings: It can result in more accurate approximations, and it offers a simple, automatic strategy that can be applied to any covariance, even when Euclidean distance is not applicable. We demonstrate these advantages in several settings, including anisotropic, nonstationary, multivariate, and spatio-temporal processes. We also illustrate our method on multivariate spatio-temporal temperature fields produced by a regional climate model.

Keywords Covariance approximation · Maximum–minimum-distance ordering · Nearest neighbors · Spatial statistics · Vecchia approximation

1 Introduction

Gaussian processes (GPs) are used for modeling functions in a variety of settings, including geostatistics (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., Gu and Wang 2018; Kennedy and O'Hagan 2001; Sacks et al. 1989), and optimization (Jones et al. 1998). GPs can also be used to represent wide neural networks (Yang 2019). However, direct application of GPs requires working with and decomposing the data covariance matrix at a cost that is cubic in the data size, which is often too expensive for today's large datasets.

Many approaches have been proposed to scale GP inference to large numbers of observations (see Heaton et al. 2019; Liu et al. 2020, for recent reviews). Among these,

probably the most promising class of approximations in spatial statistics consists of the Vecchia approximation (Vecchia 1988) and its extensions (e.g., Stein et al. 2004; Datta et al. 2016a; Guinness 2018; Sun and Stein 2016; Katzfuss and Guinness 2021; Katzfuss et al. 2020; Schäfer et al. 2021). As detailed in Katzfuss and Guinness (2021), the class also contains many other popular GP approximations as special cases (e.g., Snelson and Ghahramani 2007; Finley et al. 2009; Sang et al. 2011; Eidsvik et al. 2012; Datta et al. 2016a; Katzfuss and Gong 2020; Katzfuss 2017) and it is closely related to composite-likelihood methods (e.g., Varin 2008; Eidsvik et al. 2014). Vecchia approximations obtain a sparse Cholesky factor of the precision matrix via an ordered conditional approximation, based on removing conditioning variables in a factorization of the joint density of the GP observations into a product of conditional distributions.

The performance of a Vecchia approximation depends heavily on the choice of ordering of the variables and the choice of conditioning sets (which determines the Cholesky sparsity pattern). So far, Vecchia approximations have been mostly applied in geospatial applications featuring isotropic GPs in low-dimensional input spaces, for which the ordering

✉ Matthias Katzfuss
katzfuss@gmail.com

¹ Department of Statistics, Texas A&M University, College Station, TX, USA

and conditioning can be carried out based on the inputs or locations. Specifically, the observations can be ordered using a maximum–minimum-distance algorithm, and the sparsity is determined by nearest-neighbor conditioning (Guinness 2018). Both ordering and conditioning are typically carried out based on Euclidean distance of the corresponding inputs. We call this existing approach Euclidean-based Vecchia (EVecchia). EVecchia has also been used for nonisotropic settings, including for nonstationary (Konomi et al. 2019; Risser and Turek 2020), multivariate (Zhang et al. 2021), space-time (White and Porcu 2019), and periodic GPs (Datta et al. 2016a, Supplement A.9).

Here, we propose Vecchia approximations whose ordering and conditioning employ a correlation-based distance metric; we refer to this approach as correlation-based Vecchia or CVecchia. Correlation-based conditioning (but not ordering) was already mentioned in the early Vecchia papers (Vecchia 1988; Jones and Zhang 1997; Stein et al. 2004), but it was dismissed and not thoroughly explored, mainly due to concerns about high computational cost and instability. In contrast, we argue that CVecchia can improve approximation accuracy, and it can be carried out efficiently even in the presence of unknown parameters, allowing both frequentist and Bayesian parameter inference. Yu et al. (2017) proposed a related correlation-based idea in the context of hierarchical low-rank compression (but not factorization) of a positive-definite matrix. So far, all previous Vecchia approaches have based the ordering on spatial or temporal locations, without considering the covariance function to be approximated. Conditioning sets have also been selected based on the locations; one exception is the dynamic spatio-temporal nearest-neighbor GP (Datta et al. 2016b), whose adaptive neighbor-selection scheme defines a space-time distance as a function of the spatio-temporal covariance.

EVecchia and CVecchia are equivalent for strictly decreasing isotropic correlation functions (Jones and Zhang 1997; Stein et al. 2004), but CVecchia has two advantages for more complex situations, such as anisotropic, nonstationary, multivariate, and spatio-temporal processes: It can provide much higher accuracy, and it offers a simple, automatic strategy even when Euclidean distance is not applicable. Thus, CVecchia greatly expands the applicability of the Vecchia approach; in fact, CVecchia can be applied to any covariance matrix whose individual entries can be obtained or computed quickly, as the approximation only relies on evaluating or accessing a near-linear number of entries. CVecchia implicitly applies a Vecchia approximation in a suitable transformed input domain, in which the GP of interest is isotropic and Euclidean distance is meaningful.

The remainder of this document is organized as follows. In Sect. 2, we review Vecchia approximations from a perspective that enables our extensions. In Sect. 3, we introduce correlation-based Vecchia and discuss its properties.

Section 4 provides numerical comparisons. In Sect. 5, we illustrate the performance of our method using output from a regional climate model. Section 6 concludes and discusses future work. Appendix A contains proofs. The code for running our method and reproducing figures can be found at <https://github.com/katzfuss-group/correlationVecchia>.

2 Review of Euclidean-based Vecchia

2.1 The Vecchia approximation

Consider a centered Gaussian random vector $\mathbf{y} = (y_1, y_2, \dots, y_n)^\top \sim \mathcal{N}_n(\mathbf{0}, \mathbf{K})$, where \mathbf{K} is an $n \times n$ positive-definite covariance matrix. For example, \mathbf{y} may be a vector of observations of a GP. Evaluating the Gaussian density $p(\mathbf{y})$, which typically relies on Cholesky decomposition of \mathbf{K} , generally requires $\mathcal{O}(n^3)$ computing time and $\mathcal{O}(n^2)$ memory; this is often too expensive for large $n \gg 10^3$.

A promising approach to reduce the computational effort is the Vecchia approximation. Motivated by the exact factorization $p(\mathbf{y}) = \prod_{i=1}^n p(y_i | \mathbf{y}_{1:i-1})$ with $\mathbf{y}_{1:0} := \emptyset$, the Vecchia approximation is given by

$$\hat{p}(\mathbf{y}) = \prod_{i=1}^n p(y_i | \mathbf{y}_{c(i)}) = \mathcal{N}_n(\mathbf{0}, \hat{\mathbf{K}}), \quad (1)$$

where $c(1) = \emptyset$ and $c(i) \subset \{1, \dots, i-1\}$ for $i = 2, \dots, n$. We assume that all conditioning sets are at most of size m , $|c(i)| = \min(m, i-1)$, for some integer $m \ll n$. The approximate covariance matrix $\hat{\mathbf{K}}$ has a sparse inverse Cholesky factor: $\hat{\mathbf{K}}^{-1} = \mathbf{U}\mathbf{U}^\top$, where \mathbf{U} is a sparse upper triangular matrix with at most m off-diagonal nonzeros per column, given by $\mathbf{U}_{\tilde{c}(i),i} = (\mathbf{K}_{\tilde{c}(i),\tilde{c}(i)})^{-1}\mathbf{e}_1 / (\mathbf{e}_1^\top (\mathbf{K}_{\tilde{c}(i),\tilde{c}(i)})^{-1}\mathbf{e}_1)^{1/2}$, where $\tilde{c}(i) = \{i\} \cup c(i)$ and \mathbf{e}_1 is a vector of length $m+1$ with the first entry equal to one and all other entries equal to zero (Schäfer et al. 2021). Each of the n columns of \mathbf{U} can be computed in $\mathcal{O}(m^3)$ time, completely in parallel. Further, the \mathbf{U} implied by the Vecchia approximation is the optimal sparse inverse Cholesky factor of \mathbf{K} in terms of Kullback–Leibler (KL) divergence between $\mathcal{N}(\mathbf{0}, \mathbf{K})$ and $\mathcal{N}(\mathbf{0}, (\mathbf{U}\mathbf{U}^\top)^{-1})$ for the sparsity pattern for \mathbf{U} implied by the $c(i)$ as above (Schäfer et al. 2021).

The size of conditioning sets, m , acts as a tuning parameter that trades off sparsity and computational speed against approximation accuracy. In particular, if $m = 0$, then the Vecchia approximation assumes diagonal $\hat{\mathbf{K}}$ and yields independent y_1, \dots, y_n . If $c(i) = \{1, \dots, i-1\}$ and hence $m = n-1$, then the Vecchia approximation is exact. In general, adding indices to the conditioning sets is guaranteed to result in lower or equal KL divergence (Guinness 2018). In many settings, high accuracy can be achieved even using rel-

atively small m . In practice, often $m < 100$ is chosen with respect to available computational resources (see, e.g., the guidelines and discussion in Katzfuss and Guinness 2021).

2.2 Ordering and conditioning

For given m , the accuracy of a Vecchia approximation depends on the choice of ordering of the variables y_1, \dots, y_n in \mathbf{y} , and on the choice of conditioning sets $c(m+2), \dots, c(n)$. Arguably the preferred approach in this setting is to combine a maximum–minimum-distance ordering (MM; Guinness 2018) and nearest-neighbor conditioning (NN), as illustrated in Fig. 1.

Specifically, for MM ordering, the first index i_1 can be selected arbitrarily (e.g., $i_1 = 1$), and then the subsequent indices are selected for $k = 2, \dots, n$ as

$$i_k = \arg \max_{i \in \mathcal{I} \setminus \mathcal{I}_{1:k-1}} \min_{j \in \mathcal{I}_{1:k-1}} \tau(i, j), \quad (2)$$

where $\mathcal{I} = \{1, \dots, n\}$ and $\mathcal{I}_{1:k-1} = \{i_1, \dots, i_{k-1}\}$, using a predefined distance measure τ between the entries of \mathbf{y} . For simplicity of notation, assume henceforth and in (1) that $\mathbf{y} = (y_1, \dots, y_n)$ follows MM ordering (i.e., $y_k = y_{i_k}$).

For NN conditioning, y_i conditions on the $\min(m, i-1)$ previously ordered variables $y_{c(i)}$ that are nearest to y_i in terms of τ . Specifically, for $1 < i \leq m+1$, we have $c(i) = \{1, \dots, i-1\}$; for $i > m+1$, we have

$$c(i) \subset \{1, \dots, i-1\} \text{ of size } |c(i)| = m, \text{ s.t. } \tau(i, j) \leq \tau(i, k) \forall j \in c(i), k \in \{1, \dots, i-1\} \setminus c(i). \quad (3)$$

We also employ an algorithm that groups similar conditioning sets (Guinness 2018) to lessen overall computational cost of Vecchia approximation. Although we only consider conditioning sets consisting of the m nearest neighbors here, our framework also allows the use of other neighbor-selection strategies. For instance, Schäfer et al. (2021) uses conditioning sets consisting of all variables within a ball of a certain radius, which decreases systematically with the MM-ordering index i ; however, we carried out exploratory numerical studies, in which this radius-based approach was often significantly less accurate than NN conditioning, especially for irregularly spaced inputs.

As we can see, specifying a Vecchia approximation requires a choice of distance $\tau(i, j)$ between pairs (y_i, y_j) to determine MM and NN. So far, the Vecchia approximation has been applied in the setting where \mathbf{y} is a realization of a GP $y(\cdot) \sim \mathcal{GP}(0, K)$ at inputs $\mathbf{x}_1, \dots, \mathbf{x}_n$, so that $y_i = y(\mathbf{x}_i)$ and $\mathbf{K}_{ij} = K(\mathbf{x}_i, \mathbf{x}_j)$. Then, the ordering and conditioning for y_1, \dots, y_n are typically based on the Euclidean distance between corresponding inputs:

$$\tau_E(i, j) = \|\mathbf{x}_i - \mathbf{x}_j\|,$$

which we call Euclidean-based maximum–minimum-distance ordering (E-MM) and Euclidean-based nearest neighbor conditioning (E-NN), respectively. E-MM and E-NN are illustrated in Fig. 1. We refer to a Vecchia approximation based on this approach as EVecchia (which is then only a function of m). EVecchia has been shown to outperform Vecchia approximations based on other ordering and conditioning schemes for GPs in low-dimensional input spaces (e.g., Guinness 2018; Katzfuss and Guinness 2021; Schäfer et al. 2021).

3 Correlation-based Vecchia approximation

3.1 Definition and overview

We propose a correlation-based Vecchia (CVecchia) approximation of $\mathbf{y} \sim \mathcal{N}_n(\mathbf{0}, \mathbf{K})$. CVecchia consists of a Vecchia approximation (1) for which the MM ordering (2) and NN conditioning (3) are carried out using a correlation-based distance,

$$\begin{aligned} \tau_C(i, j) &= (1 - |\rho_{ij}|)^{1/2}, \quad \text{where} \\ \rho_{ij} &= \mathbf{K}_{ij}/(\mathbf{K}_{ii}\mathbf{K}_{jj})^{1/2}, \quad i, j \in \mathcal{I} = \{1, \dots, n\} \end{aligned} \quad (4)$$

which we will call C-MM and C-NN, respectively.

As we will explore in more detail below, CVecchia is equivalent to EVecchia for many popular isotropic kernels; CVecchia can be more accurate than EVecchia for nonisotropic kernels (e.g., anisotropic, nonstationary, spatio-temporal); and CVecchia is applicable even when EVecchia is not (e.g., multivariate GPs, GPs based on discrete or non-Euclidean inputs such as in text analysis or natural language processing).

Provided that \mathbf{K} is positive-definite, $\tau_C : \mathcal{I} \times \mathcal{I} \rightarrow [0, 1]$ in (4) is a proper distance metric (Van Dongen and Enright 2012) and, in particular, satisfies the triangle inequality. This allows us to rapidly compute C-MM and C-NN with an adaptation of the algorithm in Schäfer et al. (2021) in quasilinear time in n , assuming that each entry of \mathbf{K} can be computed in $\mathcal{O}(1)$ time; in practice, this computational cost is often small relative to that of the core Vecchia approximation in (1), and so the computational complexity of CVecchia can still be thought of as $\mathcal{O}(nm^3)$, same as for EVecchia. Among other things, this means that CVecchia is useful even when \mathbf{K} depends on unknown parameters that must be inferred.

As MM and NN only depend on the ranking of distances (not the distances themselves), other correlation-based distance metrics that are ordinally equivalent (e.g., Van Dongen and Enright 2012) to τ_C in (4) will result in equivalent CVecchia approximations.

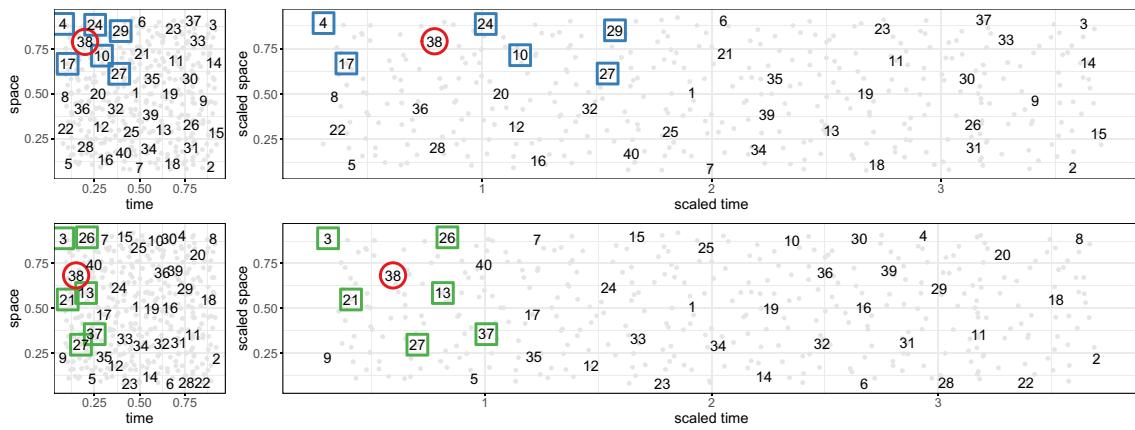


Fig. 1 Euclidean (top) and correlation-based (bottom) maximum-minimum-distance ordering (MM) and nearest-neighbor conditioning (NN) for $n = 400$ spatio-temporal inputs (grey points), assuming a spatio-temporal covariance (8) with the ratio of temporal to spatial range, $r_t/r_s = 4$. The first 40 ordered inputs are numbered, and boxes denote the nearest $m = 6$ previously ordered neighbors $c(i)$ of

the $i = 38$ th input (red circle), in the unit-square input space $[0, 1]^2$ (left panels) and the transformed input space $[0, 4] \times [0, 1]$ (right). Correlation-based MM and NN can be thought of as Euclidean MM and NN in the transformed input space (bottom right). This figure is inspired by Figure 1 in Katzfuss et al. (2022)

By definition of the correlation-based distance in (4), C-MM and C-NN ignore the marginal variances of the variables y_1, \dots, y_n . Thus, one may ask whether, for example, a better conditioning set $c(i)$ could be obtained based on a distance metric that takes into account highly varying marginal variances. However, this is not the case. To see this, note that we have $\text{KL}(p(\mathbf{y})|\hat{p}(\mathbf{y})) = \sum_i \log(\text{var}(y_i|\mathbf{y}_{c(i)})) / \text{var}(y_i|\mathbf{y}_{1:i-1})/2$ (e.g., Guinness 2018), where $\text{var}(y_i|a\mathbf{y})$ is the same for any $a \neq 0$ and so $\text{var}(y_i|\mathbf{y}_{c(i)})$ does not depend on the marginal variances of the conditioning variables.

3.2 Properties of CVecchia in the special case of reducible GPs

CVecchia is equivalent to EVecchia if \mathbf{K} is the covariance matrix of a realization of an isotropic GP with strictly decreasing positive covariance function: $\mathbf{K}_{ij} = K(\mathbf{x}_i, \mathbf{x}_j) = \sigma^2 \rho(\tau_E(i, j))$, where $\rho : \mathbb{R}_0^+ \rightarrow [0, 1]$ is strictly decreasing; examples include Matérn and power-exponential covariance functions. Taken one step further, this finding suggests that CVecchia can be interpreted as EVecchia on a transformed input space in the special case of *reducible* GPs, which we define as follows:

Definition 1 (q -reducibility) A zero-mean Gaussian process $y(\cdot)$ on \mathbb{R}^d with $d \geq 1$ is q -reducible if there exists a $\psi : \mathbb{R}^d \rightarrow \mathbb{R}^q$ such that $y(\psi^{-1}(\cdot))$ is a Gaussian process with a strictly decreasing isotropic covariance function. In particular, y is bijectively reducible if $q = d$.

Definition 1 is broad enough to include many GPs of interest. For some covariance functions, the deformation

function ψ can be easily identified, including (geometrically) anisotropic GPs, automatic relevance determination, and latent-dimension (i.e., dimension-expansion) approaches to multivariate and spatio-temporal GPs. Also, some popular nonstationary GPs are explicitly constructed in the way we define the reducibility (e.g., Perrin and Monestiez 1999; Schmidt and O'Hagan 2003; Vu et al. 2020).

A major advantage of CVecchia is that it is not required to identify the deformation ψ explicitly, but that it automatically carries out the approximation in a transformed space in which Euclidean distance is meaningful:

Proposition 1 Assume that a zero-mean Gaussian process $y(\cdot)$ is q -reducible with respect to ψ . If the first index is chosen to be the same for both C-MM and E-MM, then CVecchia of $y(\cdot)$ at inputs $\mathbf{x}_1, \dots, \mathbf{x}_n$ is identical to EVecchia of $y(\psi^{-1}(\cdot))$ at the transformed inputs $\psi(\mathbf{x}_1), \dots, \psi(\mathbf{x}_n)$.

The dimension q in Proposition 1 is important, in that EVecchia approximations become more challenging as the input dimension increases. There have been studies on necessary and sufficient conditions for reducibility and how large q must be (e.g., Perrin and Senoussi 2000; Curriero 2006), and sufficient conditions for related concepts have been identified (e.g., Porcu et al. 2010; Perrin and Meiring 2003; Perrin and Schlather 2007). In some settings, theoretical guarantees depending on q on the performance of CVecchia for reducible GPs can be provided using recent results for isotropic GPs (Schäfer et al. 2021). For example, if a process is q -reducible to an isotropic GP whose kernel is the Green's function of an elliptic PDE (which is equivalent to a Matérn covariance up to edge effects), then CVecchia can provide an ϵ -accurate approximation in $\mathcal{O}(n \log^{3q}(n/\epsilon))$ time.

While Proposition 1 provides an explanation for why CVecchia can produce adaptive approximations to some popular nonisotropic GPs, this property deserves further investigation for its relationship to Euclidean embeddings (Matousek 2013; Maehara 2013; Witsenhausen 1986). It is well-known that an exact representation of a given metric space into Euclidean space is not easy to find and that is why approximate embeddings have been studied. For instance, the Johnson-Lindenstrauss flattening lemma (Johnson and Lindenstrauss 1984) states the existence of low-distortion (no more than a factor of $1 \pm \epsilon$) Euclidean embedding of a given finite metric space to q -dimensional Euclidean space where $q \geq \mathcal{O}(\log(n)/\epsilon)$. This may provide another way to carry out performance evaluation of CVecchia approximations.

3.3 Estimation of parameters

So far, we have assumed a fixed \mathbf{K} and $p(\mathbf{y}) = \mathcal{N}_n(\mathbf{y}|\mathbf{0}, \mathbf{K})$, but in practice $\mathbf{K} = \mathbf{K}_\theta$ and hence p_θ often depend on unknown parameters θ . Our CVecchia approximation $\hat{p}(\mathbf{y}) = \prod_{i=1}^n p(y_i|\mathbf{y}_{c(i)})$ depends on θ both via p_θ and via the correlation distance τ_C^θ in (4) used to determine the MM ordering of y_1, \dots, y_n and the NNs in the $c(i)$. To emphasize this, we will sometimes use $\hat{p}_{\theta_1}^{\theta_2}(\mathbf{y})$ to denote a CVecchia approximation of p_{θ_1} based on $\tau_C^{\theta_2}$.

For frequentist inference, Guinness (2021) proposed to find the maximum likelihood estimator of θ by optimizing the Vecchia loglikelihood via Fisher scoring. Given that

$$\begin{aligned} \log \hat{p}(\mathbf{y}) &= \sum_{i=1}^n \log p(y_i|\mathbf{y}_{c(i)}) \\ &= \sum_{i=1}^n (\log p(\mathbf{y}_{\tilde{c}(i)}) - \log p(\mathbf{y}_{c(i)})), \end{aligned} \quad (5)$$

the score $\mathbf{g}^{(k)}$ and the Fisher information $\mathbf{M}^{(k)}$ of $\hat{p}(\mathbf{y})$ at the k th iteration of the Fisher-scoring algorithm can be computed by addition and subtraction of the score and Fisher information of each of the $2n$ normal distributions of dimension at most $m + 1$ on the right-hand side of (5). The parameter vector is then updated as $\theta^{(k+1)} = \theta^{(k)} + (\mathbf{M}^{(k)})^{-1} \mathbf{g}^{(k)}$.

For CVecchia, we propose to use a modified Fisher-scoring algorithm, where we now compute $\mathbf{g}^{(k)} = \frac{\partial}{\partial \theta} \log \hat{p}_{\theta^{(k)}}(\mathbf{y})|_{\theta=\theta^{(k)}}$ and $\mathbf{M}^{(k)} = -\mathbb{E} \frac{\partial^2}{\partial \theta^2} \log \hat{p}_{\theta^{(k)}}(\mathbf{y})|_{\theta=\theta^{(k)}}$ with fixed ordering and conditioning based on $\tau_C^{\theta^{(k)}}$. In other words, when computing derivatives of the CVecchia log-likelihood for the Fisher-scoring updates, the dependence of the ordering and conditioning on θ is ignored. Instead, the ordering and conditioning are updated based on $\tilde{\theta}^{(k)} = \theta^{(k)}$ after certain iterations $k \in \mathcal{G}$, and $\tilde{\theta}^{(k)} = \tilde{\theta}^{(k-1)}$ otherwise. For simplicity, we can update the ordering and conditioning

at the end of each iteration, $\mathcal{G} = \{1, 2, 3, 4, \dots\}$. Alternatively, the computational cost can be reduced by setting $\mathcal{G} = \{1, 2, 4, 8, \dots\}$ and thus skipping this update for exponentially increasing numbers of iterations, exploiting that the parameter values tend to change less and less with increasing iteration numbers. In either case, repeatedly updating the ordering and conditioning over the course of the Fisher-scoring algorithm did not introduce convergence problems in our numerical experiments.

As the Vecchia approximation implies a valid density $\hat{p}(\mathbf{y}) = \mathcal{N}_n(\mathbf{y}|\mathbf{0}, \hat{\mathbf{K}})$, it is also possible to carry out Bayesian inference on θ , assuming a prior $p(\theta)$ has been specified. However, the dependence of C-MM and C-NN on θ again presents a challenge. In the context of a spatio-temporal covariance, Datta et al. (2016b) essentially proposed to approximate the posterior as $\hat{p}(\theta|\mathbf{y}) \propto p(\theta) \hat{p}_\theta^\theta(\mathbf{y})$ based on C-NN, meaning that the conditioning sets $c(i)$ are recomputed for every θ at which the posterior is evaluated. However, in our exploratory studies, we found this approach to lead to unstable and sinuous approximate posteriors. Instead, we propose to first obtain a maximum likelihood or maximum a posteriori estimate $\hat{\theta}$ using Fisher scoring, as above. Then, we approximate the posterior as $\hat{p}(\theta|\mathbf{y}) \propto p(\theta) \hat{p}_{\hat{\theta}}^{\hat{\theta}}(\mathbf{y})$, with fixed correlation distance $\tau_C^{\hat{\theta}}$ and hence fixed C-MM and C-NN based on $\theta = \hat{\theta}$. This approach leads to smooth posteriors, as illustrated in Sect. 4.7.

3.4 Prediction

Our method can be used for accurate and efficient prediction of an unobserved vector $\mathbf{y}^* = (y_1^*, \dots, y_{n^*}^*)$ with $(\mathbf{y}^\top, \mathbf{y}^{*\top})^\top \sim \mathcal{N}_{n+n^*}(\mathbf{0}, \mathbf{K}^{all})$. For prediction and uncertainty analysis, the goal is to obtain the joint posterior predictive distribution $p(\mathbf{y}^*|\mathbf{y})$. Following Katzfuss et al. (2020), we apply a Vecchia approximation to $(\mathbf{y}^\top, \mathbf{y}^{*\top})^\top$ with the entries of \mathbf{y}^* ordered after those of \mathbf{y} to obtain a CVecchia approximation of the posterior predictive distribution,

$$\hat{p}(\mathbf{y}^*|\mathbf{y}) = \prod_{i=1}^{n^*} p(y_i^*|\mathbf{y}_{c_o(i)}, \mathbf{y}_{c_u(i)}^*) = \mathcal{N}_{n^*}(\boldsymbol{\mu}^*, \hat{\mathbf{K}}^*), \quad (6)$$

where $y_1^*, \dots, y_{n^*}^*$ are assumed to follow a restricted C-MM ordering, which is obtained from a C-MM ordering of all (observed and unobserved) variables under the restriction of having the observed variables be ordered first (in which case the ordering of the unobserved variables takes the distances to observed variables into account). As recommended in Katzfuss et al. (2020), we allow the unobserved variables to condition on both observed and (previously ordered) unobserved variables. Specifically, y_i^* conditions on the nearest (in terms of correlation-based distance with respect to \mathbf{K}^{all})

m variables among $y_1, \dots, y_n, y_1^*, \dots, y_{i-1}^*$. For notational convenience, in (6) the resulting conditioning set is split into indices $c_o(i)$ corresponding to observed variables and $c_u(i)$ corresponding to unobserved variables; either $c_o(i)$ or $c_u(i)$ can be an empty set for any i . Each of the conditional distributions in the product in (6) can be computed in $\mathcal{O}(m^3)$ time, resulting in fast prediction or joint simulation even for large n and n^* .

In practice, \mathbf{K}^{all} will typically depend on unknown parameters θ . Predictions can then be based on a frequentist estimate of θ or based on samples from the Bayesian posterior of θ , which can be obtained using the observed data \mathbf{y} as described in Sect. 3.3. Then, given a frequentist estimate $\hat{\theta}$, the posterior predictive distribution is obtained as $\hat{p}_{\hat{\theta}}^*(\mathbf{y}^*|\mathbf{y})$ using similar notation as in Sect. 3.3. Given samples $\theta^{(1)}, \dots, \theta^{(L)}$ from the posterior, we can account for posterior uncertainty in θ and obtain an averaged posterior predictive distribution $\hat{p}(\mathbf{y}^*|\mathbf{y}) = (1/L) \sum_{l=1}^L \hat{p}_{\theta^{(l)}}^*(\mathbf{y}^*|\mathbf{y})$, where $\hat{\theta}$ is again a maximum likelihood or maximum a posteriori estimate.

3.5 Noise

The methods discussed so far are most appropriate if \mathbf{y} is observed without noise. However, data in many application areas are typically modeled as a GP with additive noise. Suppose now that we observe $\mathbf{z} = \mathbf{y} + \boldsymbol{\epsilon}$ with $\boldsymbol{\epsilon} = (\epsilon_1, \dots, \epsilon_n)^\top \sim \mathcal{N}_n(\mathbf{0}, \mathbf{D})$, where \mathbf{D} is diagonal.

A straightforward way of extending our methods to this noisy setting is to apply the same CVecchia approach to the covariance matrix of \mathbf{z} , which is $\Sigma = \mathbf{K} + \mathbf{D}$. However, in this approach the noise terms weaken the screening effect and hence an accurate CVecchia approximation will often require a larger m than in the noise-free case. Interestingly, if the signal and noise variances are both constant (i.e., $\mathbf{K}_{ii} = \mathbf{K}_{jj}$ and $\mathbf{D}_{ii} = \mathbf{D}_{jj}$ for all i, j), then C-MM and C-NN do not depend on the noise variance (even if it is zero). This can be seen by noting that $\tau_C(i, j) \leq \tau_C(i, k)$ if and only if $\tau_C^{+D}(i, j) \leq \tau_C^{+D}(i, k)$, where $\tau_C^{+D}(i, j) = (1 - |\rho_{ij}^{+D}|)^{1/2}$ with $\rho_{ij}^{+D} = \Sigma_{ij}/(\Sigma_{ii} \Sigma_{jj})^{1/2}$ for $i, j \in \mathcal{I}$. For varying noise variances, high-noise observations move farther away from other observations in terms of correlation distance, and so they are less likely to be included in conditioning sets; this makes intuitive sense, in that their high noise means that they contain less information about \mathbf{y} .

An alternative way of extending our methods to the noisy setting is to apply CVecchia to the (now latent) noise-free variables \mathbf{y} as before and then add noise. In other words, we set $\hat{\Sigma} = \hat{\mathbf{K}} + \mathbf{D}$, where $\hat{\mathbf{K}}$ is obtained using CVecchia as in previous sections. While this is conceptually simple, inference then requires obtaining the Cholesky factor of the posterior precision matrix $\text{var}(\mathbf{y}|\mathbf{z})^{-1} = \hat{\mathbf{K}}^{-1} + \mathbf{D}^{-1}$, which can be

very expensive due to fill-in. Fortunately, the computational speed of CVecchia can be maintained without introducing meaningful additional approximation error by approximating the Cholesky factor using an incomplete Cholesky factorization (IC), as proposed for EVecchia in Schäfer et al. (2021). This approach is useful both for parameter inference based on evaluating the CVecchia likelihood and for making predictions. We demonstrate numerically in Sect. 4.7 that this IC-based approach can be highly accurate in the context of CVecchia as well.

4 Examples and numerical comparisons

We conducted simulation experiments to demonstrate that CVecchia is widely applicable and highly accurate. Specifically, we considered anisotropic, nonstationary, multivariate, and spatio-temporal GPs, and an example without any explicit inputs. We begin by assuming that the covariance matrices are known; then, we demonstrate parameter estimation and prediction using our methods. Throughout, our proposed CVecchia approach is denoted by C-MM + C-NN. We compared to existing or other reasonable competing Vecchia approximations, which necessarily differ between simulation scenarios, because none of them are meaningfully applicable across all the scenarios. We compared the different Vecchia methods in terms of the KL-divergence between the exact distribution $\mathcal{N}(\mathbf{0}, \mathbf{K})$ and the approximate distribution $\mathcal{N}(\mathbf{0}, \hat{\mathbf{K}})$, averaged over 10 simulations in settings with known covariance structure and over 200 simulations in parameter-inference or prediction settings. Comparisons are carried out as a function of m , as all considered Vecchia methods become more accurate and more computationally expensive as m increases, with a time complexity of $\mathcal{O}(nm^3)$.

4.1 Anisotropic and nonstationary GPs

We considered nonstationary GPs at $n = 30^2 = 900$ inputs selected uniformly at random on the unit square, $\mathcal{X} = [0, 1]^2$. We compared various combinations of ordering (E-MM, C-MM, X-ord, Y-ord) and conditioning (E-NN, C-NN) schemes, where X-ord and Y-ord denote ordering by the first or second coordinate of the input space, respectively. EVecchia corresponds to E-MM + E-NN. Vecchia (1988)'s original approach is given by Y-ord + E-NN.

We used a nonstationary Matérn covariance function (Paciorek and Schervish 2006; Stein 2005):

$$K(\mathbf{x}, \mathbf{x}') = \sigma^2 \frac{|\mathbf{A}(\mathbf{x})|^{1/4} |\mathbf{A}(\mathbf{x}')|^{1/4}}{|\tilde{\mathbf{A}}(\mathbf{x}, \mathbf{x}')|^{1/2}} M_{\frac{v(\mathbf{x})+v(\mathbf{x}')}{2}} \left(((\mathbf{x} - \mathbf{x}')^\top \tilde{\mathbf{A}}(\mathbf{x}, \mathbf{x}')^{-1} (\mathbf{x} - \mathbf{x}'))^{1/2} \right), \quad \mathbf{x}, \mathbf{x}' \in \mathcal{X}, \quad (7)$$

where $M_\nu(0) = 1$, $M_\nu(x) = x^\nu B_\nu(x)$ for $x > 0$, B_ν is a modified Bessel function of order ν , $\nu : \mathcal{X} \rightarrow \mathbb{R}^+$ is the smoothness, $\mathbf{A} : \mathcal{X} \rightarrow \mathbb{R}^{d \times d}$ is a (positive definite) anisotropy matrix, and $\tilde{\mathbf{A}}(\mathbf{x}, \mathbf{x}') = (\mathbf{A}(\mathbf{x}) + \mathbf{A}(\mathbf{x}'))/2$. For simplicity, we assumed $\sigma = 1$.

We considered the following settings as special cases of (7):

Anisotropic: $\nu(\mathbf{x}) \equiv 0.5$, $\mathbf{A}(\mathbf{x}) \equiv 10^{-2} \text{diag}(a^{-2}, 1)$, where a is the degree of anisotropy.

Varying smoothness: $\nu(\mathbf{x}) = 0.2 + 1.3x_1$ (i.e., varying as a function of the first coordinate), $\mathbf{A} \equiv 10^{-2} \text{diag}(1, 1)$.

Varying rotation: $\nu(\mathbf{x}) \equiv 0.5$,

$$\mathbf{A}(\mathbf{x}) = \begin{pmatrix} \cos \eta(\mathbf{x}) & \sin \eta(\mathbf{x}) \\ -\sin \eta(\mathbf{x}) & \cos \eta(\mathbf{x}) \end{pmatrix}^\top$$

$$\text{diag}(10^{-4}, 10^{-2}) \begin{pmatrix} \cos \eta(\mathbf{x}) & \sin \eta(\mathbf{x}) \\ -\sin \eta(\mathbf{x}) & \cos \eta(\mathbf{x}) \end{pmatrix}$$

is a rotation matrix with spatially varying angle $\eta(\mathbf{x}) = \frac{\pi x_1}{2}$.

In the anisotropic setting, the correlation-based distance $\tau_C(i, j)$ is a strictly increasing function of $\|\tilde{\mathbf{x}}_i - \tilde{\mathbf{x}}_j\|$, where $\tilde{\mathbf{x}}_i = \mathbf{A}^{-1/2}\mathbf{x}_i$, because $B_\nu(\cdot)$ is strictly decreasing; thus, CVecchia is equivalent to EVecchia applied to the transformed inputs $\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_n$. For varying smoothness and rotation, the transformed space is not easily identified. However, as shown in Fig. 2, C-NN was always more accurate than E-NN. In addition, using C-MM instead of E-MM led to further improvements for the anisotropic and varying-rotation setting. The improvement of CVecchia over existing methods was especially pronounced for strong anisotropy (i.e., large a) and for varying rotation.

4.2 Multivariate GP

We considered a p -variate GP, $\mathbf{y}(\cdot) = (y^{(1)}(\cdot), \dots, y^{(p)}(\cdot))^\top \sim \mathcal{GP}(0, K)$, with a cross-covariance function based on a latent dimension separating the processes (Apanasovich and Genton 2010),

$$K_{i,j}(\mathbf{x}, \mathbf{x}') = \text{cov}(y^{(i)}(\mathbf{x}), y^{(j)}(\mathbf{x}'))$$

$$= \sigma^2 \exp(-\|\tilde{\mathbf{x}}_i - \tilde{\mathbf{x}}'_j\|/r),$$

$$\mathbf{x}, \mathbf{x}' \in \mathcal{X}, i, j \in \{1, \dots, p\},$$

where $\tilde{\mathbf{x}}_i = (\mathbf{x}^\top, \nu_i)^\top \in \mathbb{R}^{2+1}$, and ν_i represents the location of the i -th component of the multivariate GP in the latent dimension. Thus, the dependence between $y^{(i)}(\cdot)$ and $y^{(j)}(\cdot)$ decreases with their latent distance $|\nu_i - \nu_j|$. We assumed $\sigma^2 = 1$, $r = 0.1$, and $\nu_1 = 0$. We considered a total of n observations stacked into a vector $\mathbf{y} = (y^{(1)\top}, \dots, y^{(p)\top})^\top$,

where $\mathbf{y}^{(j)} = (y_1^{(j)}, \dots, y_{n_j}^{(j)})^\top$ with $y_i^{(j)} = y^{(j)}(\mathbf{x}_i^{(j)})$, and $n = \sum n_j$.

Here, $\tau_C(i, j)$ is a strictly increasing function of $\|\tilde{\mathbf{x}}_i - \tilde{\mathbf{x}}_j\|$, and so CVecchia is equivalent to EVecchia applied to the transformed inputs $\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_n$ in the expanded $(2 + 1)$ -dimensional input space. The competing methods considered in Sect. 4.1 are not directly applicable in this multivariate setting, and so we considered the following alternative approaches. S-E-MM separately orders the entries of each $\mathbf{y}^{(j)}$ according to an MM ordering of the corresponding inputs $\mathbf{x}_1^{(j)}, \dots, \mathbf{x}_{n_j}^{(j)}$, and then orders $\mathbf{y}^{(1)}$, then $\mathbf{y}^{(2)}$, and so forth, in \mathbf{y} . To construct conditioning sets of size m , J-E-NN considers the nearest m inputs in \mathcal{X} among all previously ordered variables in the joint vector \mathbf{y} , while S-E-NN carries out nearest-neighbor conditioning separately for each $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(p)}$. D-E-NN divides m by p and finds the m/p nearest previously ordered neighbors among each of the components $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(p)}$ (according to their inputs in \mathcal{X}).

We compared these various Vecchia approaches for bivariate ($p = 2$) and trivariate ($p = 3$) GPs, with each process observed at $n_j = 400$ randomly sampled locations in \mathcal{X} . In both cases, we assumed that the processes were observed in a misaligned manner (i.e., $\mathbf{x}_i^{(j)} \neq \mathbf{x}_i^{(k)}$ for $j \neq k$). As shown in Fig. 3, C-NN outperformed other conditioning approaches; C-MM provided additional improvements in some settings over S-E-MM. We also considered the setting of identical observation locations for the different processes (i.e., $\mathbf{x}_i^{(j)} = \mathbf{x}_i^{(k)}$), but the results were very similar to the misaligned case and are hence not shown.

4.3 Spatio-temporal GP

We considered a spatio-temporal GP indexed by a space-time input coordinate $\mathbf{x} = (\mathbf{s}^\top, t)^\top$, where we assumed that space is scaled to the unit square, $\mathbf{s} \in [0, 1]^2$, and time is scaled to the unit interval, $t \in [0, 1]$. We considered a space-time covariance function of the form

$$K(\mathbf{x}, \mathbf{x}') = \sigma^2 \exp(-\|(\mathbf{s} - \mathbf{s}')\|/r_s - |t - t'|\|/r_t)$$

$$= \sigma^2 \exp(-\|\mathbf{A}^{-1}(\mathbf{x} - \mathbf{x}')\|), \quad (8)$$

where r_s and r_t are the spatial and temporal range parameters, and $\mathbf{A} = \text{diag}(r_s, r_s, r_t)$. We assumed that $\sigma^2 = 1$, $r_s = 0.1$, and $r_t = 1.0$.

Here, $\tau_C(i, j)$ is a strictly increasing function of $\|\tilde{\mathbf{x}}_i - \tilde{\mathbf{x}}_j\|$, where $\tilde{\mathbf{x}}_i = \mathbf{A}^{-1/2}\mathbf{x}$, and so CVecchia is equivalent to EVecchia applied to the transformed inputs $\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_n$. As space and time are not commensurable, the previous competing methods are again not meaningful. We considered ordering by time (T-ord), and conditioning on the NN in time (T-NN). Note that, when inputs are taken at the same time point, T-ord orders the inputs according to the values of the second spa-

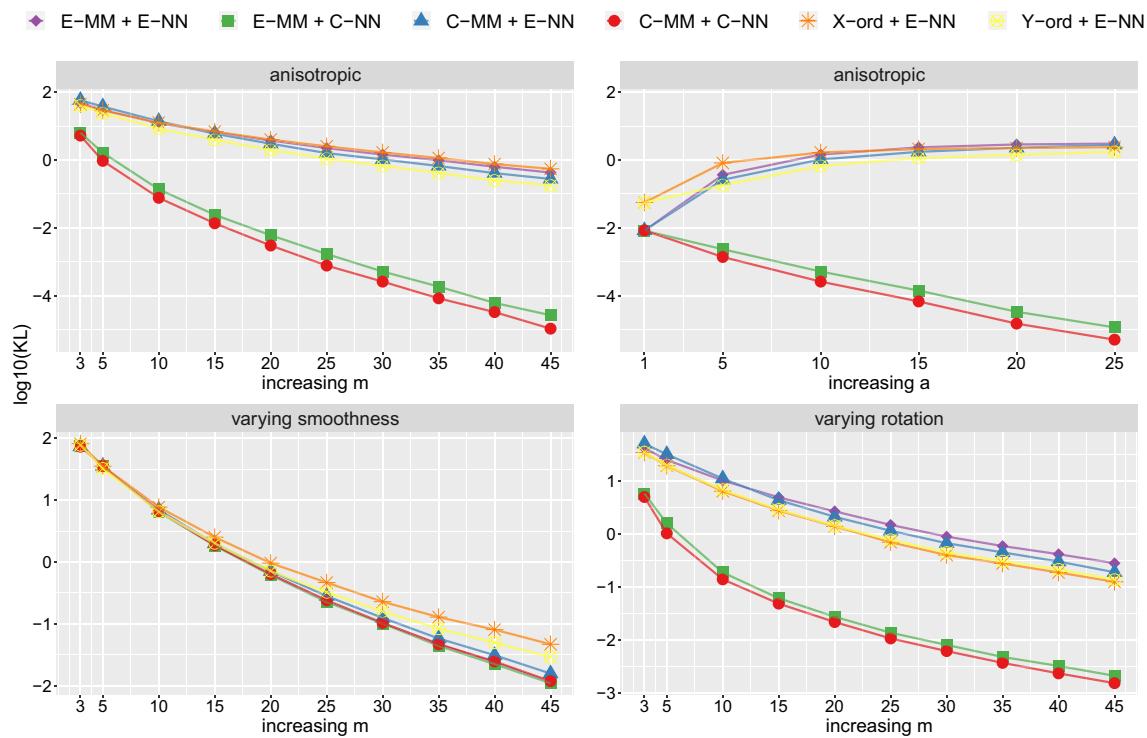


Fig. 2 Log10-scale KL divergences between true and approximate likelihoods of GPs with the degree of anisotropy $a = 10$ for increasing size of conditioning sets m (top left), with $m = 30$ for increasing a (top right), with varying smoothness parameter $\nu = 0.2 + 1.3x_1$ (a function of the first coordinate) for increasing m (bottom left), and with varying rotation angle $\eta = \frac{\pi x_1}{2}$ for increasing m (bottom right)

tial coordinate. If these values are again the same, it uses the values of the first coordinate. Further, we considered E-NN based on the distance of the (unit-scaled) space-time coordinates, $\|\mathbf{x} - \mathbf{x}'\|$. To our understanding, the correlation-based conditioning approach proposed in Datta et al. (2016b) corresponds to T-ord + C-NN.

As illustrated in Fig. 4, we simulated $n = 900$ space-time observations on the unit cube according to four different simulation scenarios, the latter three of which were chosen to mimic common observation patterns for environmental data:

- Random** Space-time coordinates are selected uniformly at random, and so they are irregular in space and time.
- Station** Observations are obtained at 9 regular time points at 100 irregularly spaced “monitoring stations.”
- Gridded** Observations are obtained at 9 regular time points on a regular grid of size $10 \times 10 = 100$ in space (e.g., mimicking output from climate models).
- Satellite** Similar to data from polar-orbiting satellites, at 900 regularly spaced time points, we have 90 observations along each of 5 one-dimensional tracks at two repeat cycles.

As shown in Fig. 5, CVecchia outperformed the competing methods.

Note that we repeated the experiments from Sects. 4.1–4.3 for larger $n = 3600$, but we found out that the shapes of the KL curves were very similar to those in Figs. 2, 3 and 5.

4.4 Gaussian hierarchical model

We have so far considered only cases in which covariance structure is computed based on inputs. In this subsection, we offer an example that has no inputs, so that CVecchia is applicable but EVecchia is not. Motivated by hierarchical models which are widely used for combining information and describing heterogeneity between sub-populations, we assumed that $\mu \sim \mathcal{N}(0, \sigma_0^2)$ and

$$\begin{aligned} \mu_{i_1, \dots, i_j} \mid \mu_{i_1, \dots, i_{j-1}} &\stackrel{i.i.d.}{\sim} \mathcal{N}(\mu_{i_1, \dots, i_{j-1}}, \sigma_j^2) \\ i_k = 1, 2, \quad k = 1, \dots, j, \quad j = 1, \dots, J, \end{aligned}$$

where $\sigma_0^2 = \sigma_1^2 = \dots = \sigma_k^2 = 1$. We observe $\mathbf{y} = \{y_{i_1, \dots, i_J} : i_k = 1, 2, k = 1, \dots, J\}$ with $y_{i_1, \dots, i_J} = \mu_{i_1, \dots, i_J}$ at the finest level. This hierarchical model is illustrated for depth $J = 3$ in the left panel of Fig. 6. We have $\text{cov}(y_{i_1, \dots, i_J}, y_{l_1, \dots, l_J}) = \sum_{r=0}^{\alpha} \sigma_r^2$, where α is the level up to which y_{i_1, \dots, i_J} and y_{l_1, \dots, l_J} have a common ancestor.

We conducted a numerical comparison using depth $J = 12$, and so $n = 2^{12} = 4,096$. Because the competing methods

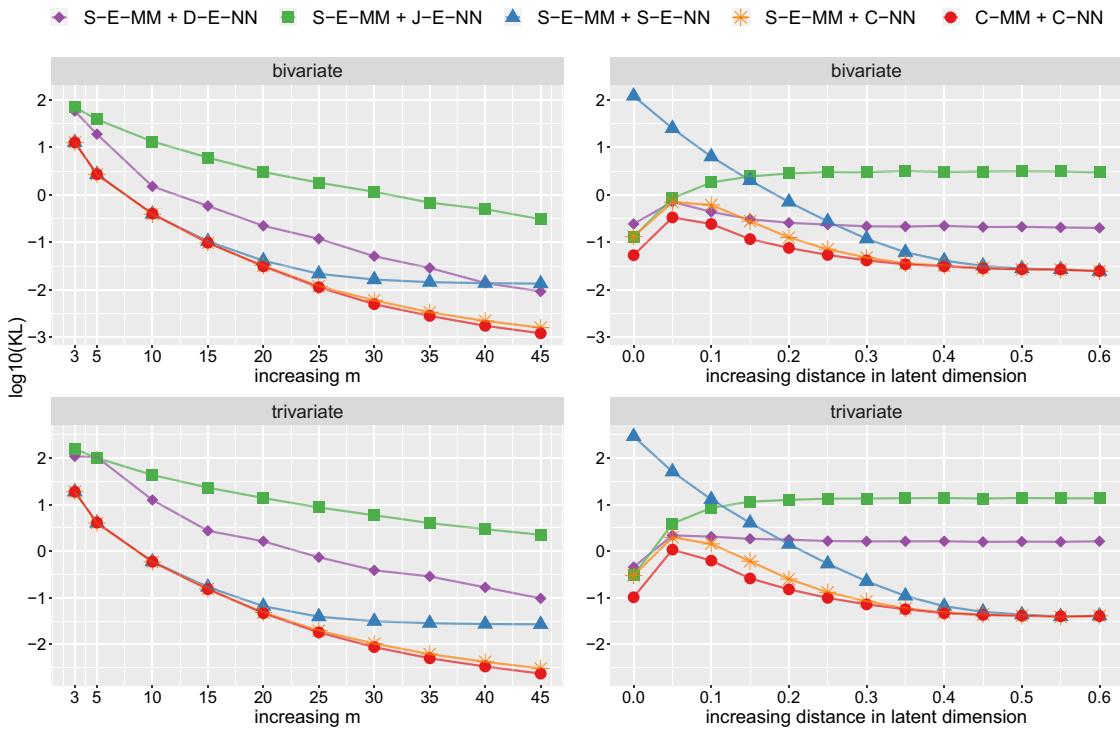


Fig. 3 Log10-scale KL divergences between true and approximate likelihoods of distinctly observed bivariate GPs with distance in latent dimension $\Delta = |\nu_1 - \nu_2| = 0.4$ for increasing size of conditioning sets m (top left) and with $m = 20$ for increasing Δ (top right), and distinctly

observed trivariate GPs with distance in latent dimension $\Delta = 0.4$, where $\nu_j = (j - 1)\Delta$, for increasing m (bottom left) and with $m = 20$ for increasing Δ (bottom right)

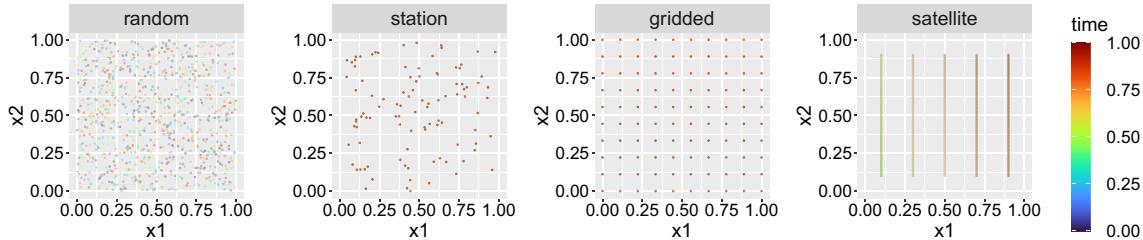


Fig. 4 Inputs for spatio-temporal simulation scenarios plotted against spatial coordinates x_1 and x_2 . The inputs are color-coded by time, although later time points exactly cover earlier time points in the Station and Gridded case, and also to some degree in the Satellite scenario due to its two repeat cycles

used in other experiments are again not directly applicable, we compared three variants of the Vecchia approximation in the right panel of Fig. 6, where L-ord denotes lexicographic (or simply left-to-right) ordering, R-ord denotes random ordering, and R-N conditions on randomly selected previously ordered entries. We repeated R-ord + R-N 200 times, but interestingly the resulting KL divergences appear quite similar when plotted on the log scale. CVecchia strongly outperformed the other two methods.

4.5 Parameter estimation

We examined the performance of frequentist parameter estimation using Fisher scoring (Sect. 3.3) in the Station and Satellite space-time scenarios of Sect. 4.3. The task was to estimate the range parameters r_s and r_t . We updated the ordering and conditioning at every Fisher-scoring iteration ($\mathcal{G} = \{1, 2, 3, \dots\}$).

We compared the different approximation methods described in Sect. 4.3. For reference, we also considered “optimal” parameter estimation using the exact GP without Vecchia approximation (or, equivalently, a Vecchia approximation with $m = n - 1$). The methods were compared in

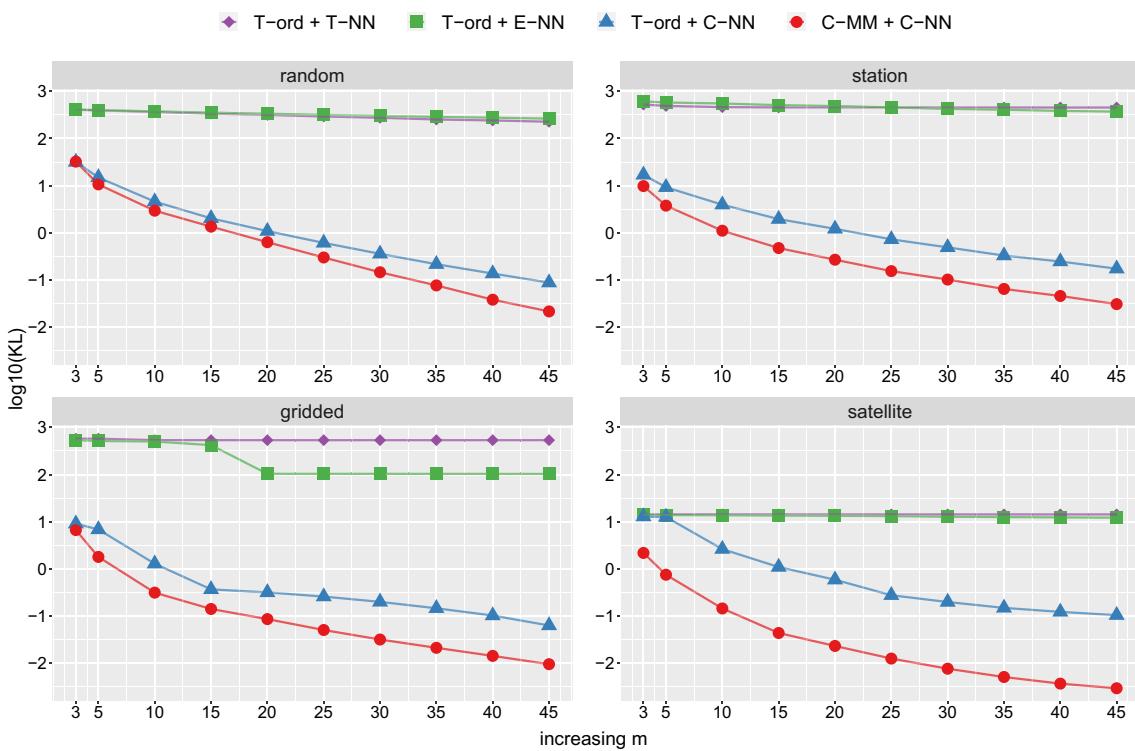


Fig. 5 Log10-scale KL divergences between true and approximate likelihoods of spatio-temporal GPs (Fig. 4) for increasing size of conditioning sets m

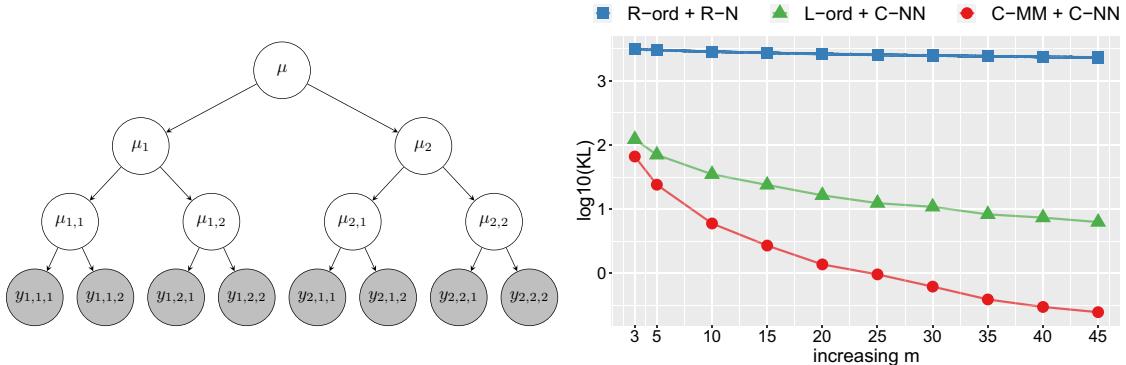


Fig. 6 A graphical representation of the hierarchical normal model with $J = 3$ (left) and log10-scale KL divergences between true and approximate likelihoods with $J = 12$ for increasing m (right)

terms of the average KL divergence between the true distribution (using the true parameter values) and the approximate distribution (using each method's estimated parameters). We also computed the root mean squared difference (RMSD) between the values of r_s and r_t as estimated by the exact GP and as estimated by the different Vecchia approximations.

As shown in Fig. 7, CVecchia produced by far the most accurate estimated distributions, which were similar to those based on the exact GP for $m \geq 25$. While the RSMDs were quite noisy, despite averaging over 200 simulated datasets, CVecchia also generally performed best in terms of RMSD.

4.6 Prediction

To illustrate prediction performance, we again considered the Random, Station, and Satellite space-time scenarios from Sect. 4.3. Of the 900 space-time observations, 100 were randomly selected as test data, and so the training data consisted of the remaining $n = 800$ observations. To lessen the computational cost of our many comparisons, we assumed that the covariance parameters were known.

Figure 8 shows the prediction performance for the 100 test data, as measured by the logarithmic score (see Gneiting and Katzfuss 2014, for details) averaged over 200 simulation

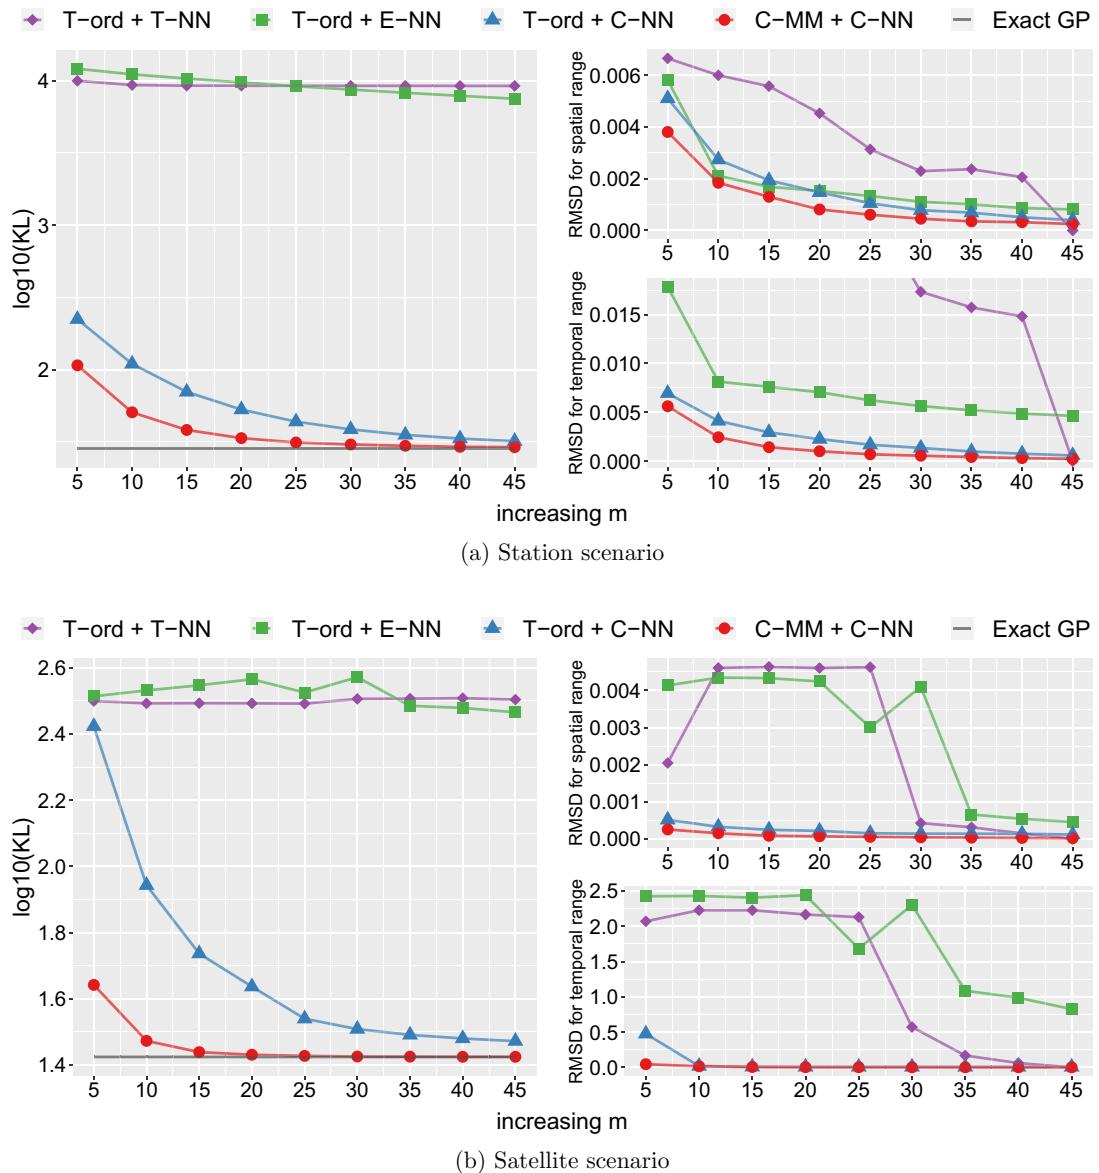


Fig. 7 Performance in parameter estimation using Fisher scoring under two spatio-temporal scenarios (Station and Satellite) in Fig. 4. Each subfigure contains three plots: Log10-scale KL divergences between

true and estimated likelihoods (left) and root mean squared difference (RMSD) for spatial range parameter (top right) and for temporal range parameter (bottom right), for increasing size of conditioning sets m

runs. In the Random and Station scenarios, CVecchia and T-ord + C-NN both performed well. In the Satellite scenario, CVecchia performed best.

4.7 Bayesian inference for noisy data

We considered Bayesian inference with CVecchia for noisy data under the Random, Station, and Satellite space-time scenarios of Sect. 4.3. The task was to calculate posterior densities of the range parameters r_s and r_t . We assumed that the priors were $\log(r_s) \sim \mathcal{N}(\log(0.1), 0.6^2)$ and $\log(r_t) \sim \mathcal{N}(\log(1.0), 0.6^2)$, with constant noise variances,

$\mathbf{D} = (0.4)\mathbf{I}_n$. Figure 9 presents two different approaches described in Sect. 3.5: One is the naive approach that directly uses the covariance matrix of the noisy observations, and the other is the IC-based approach that applies CVecchia to the noise-free variables and then adds the noise. As claimed in Sect. 3.5, Fig. 9 shows that, while CVecchia provided reliable approximate posteriors compared to the other methods, the IC-based approach provided further improvements. C-MM and C-NN were fixed based on the true values of θ ; we also tried updating C-MM for each evaluated θ value, but this resulted in unstable posteriors.

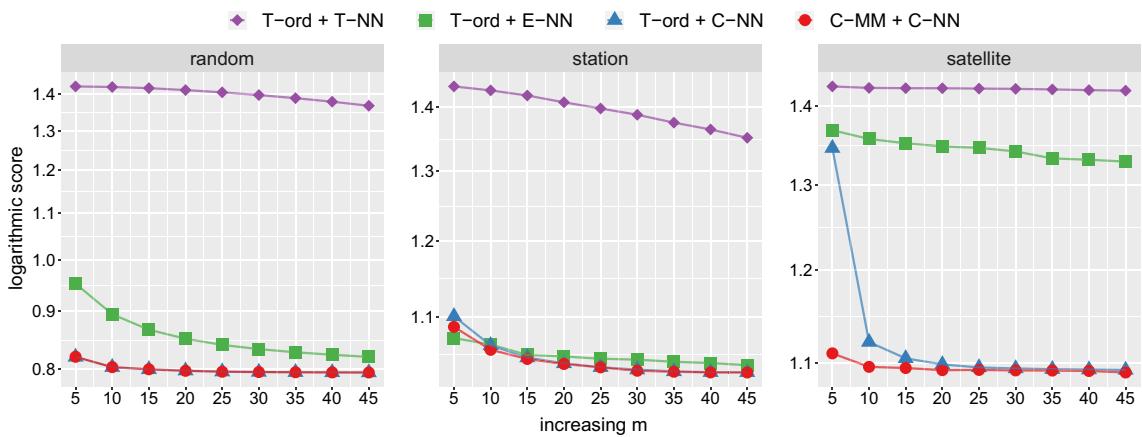
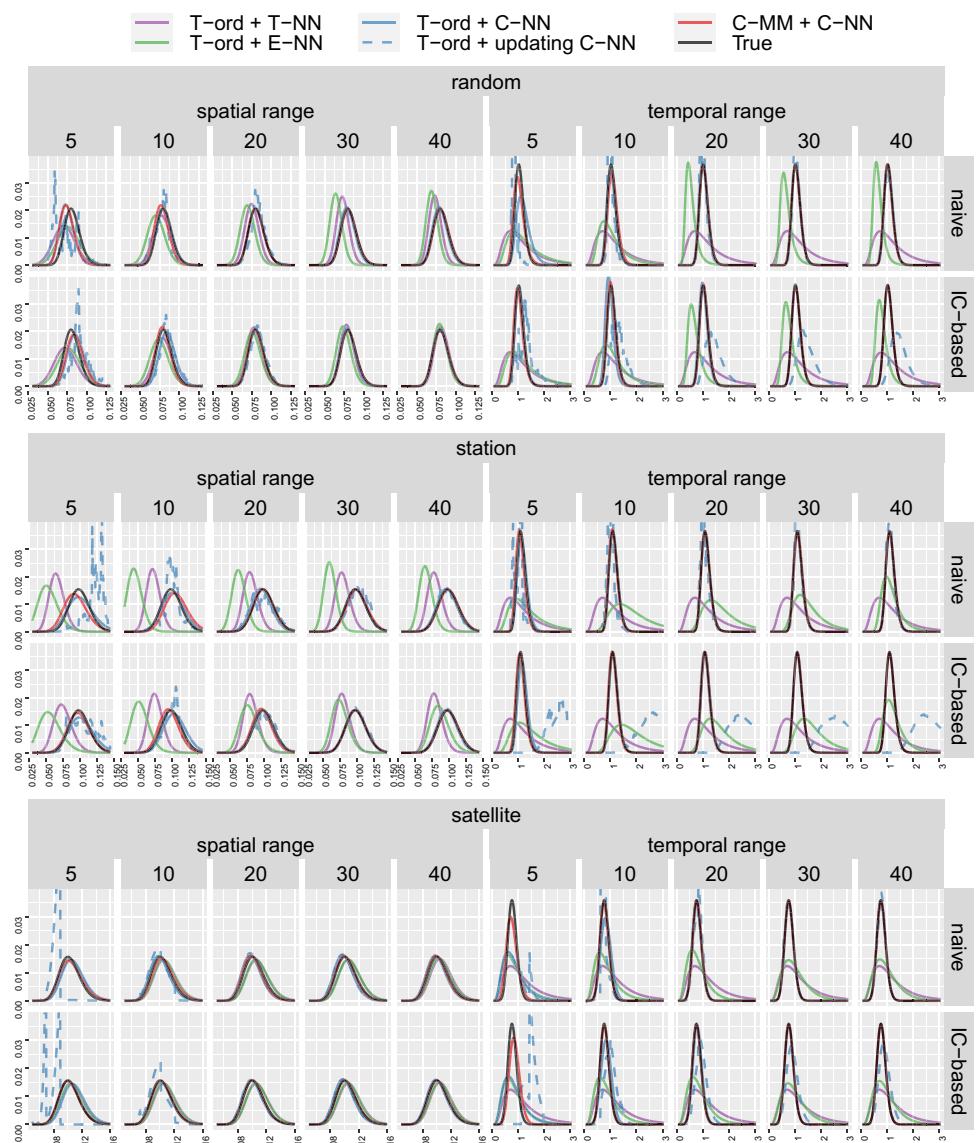


Fig. 8 Logarithmic score for predictions under three space-time scenarios: Random (left), Station (center), and Satellite (right). Note that the y-axes of the panels are on a log scale

Fig. 9 True and approximate posterior densities of spatial and temporal ranges under three space-time scenarios from Fig. 4: Random (top), Station (center), and Satellite (bottom). For each scenario, the left (right) five columns are posteriors of the spatial (temporal) range parameter. For each range parameter, the first (second) row presents posteriors with the naive (IC-based) approach for size of conditioning sets $m = 5, 10, 20, 30, 40$. For $m \geq 10$, some lines are not visible because they are covered by the (exact) black lines



5 Application to real data

We assessed the use and efficacy of CVecchia to fit and predict regional climate model (RCM) outputs. The North American Regional Climate Change Assessment Program (NARCCAP; Mearns et al. 2009) is a research program designed to (1) provide high-resolution projections of climate change, (2) investigate uncertainties in regional climate change simulations based on different atmosphere-ocean general circulation models (AOGCMs), and (3) evaluate RCM performance over North America (Mearns et al. 2012). While the program ran 50-km spatial resolution simulation based on multiple RCMs driven by multiple AOGCMs, we only considered the Canadian regional climate model (CRCM) using the NCEP-DOE Reanalysis II (NCEP) as boundary conditions. The details on RCMs and AOGCMs in the NARCCAP are available from <https://www.narccap.ucar.edu/>.

In particular, we studied a bivariate spatio-temporal dataset given by maximum and minimum daily surface air temperature (tasmax and tasmin) fields for June–August 2001 (92 days) in the South region (Arkansas, Kansas, Louisiana, Mississippi, Oklahoma, and Texas; see Karl and Koss 1984). Figure 10 shows tasmax and tasmin fields in the South region on selected days. The cartographic boundary files of the south region are available from <https://www.census.gov/>. The total sample size is $n_{total} = 78,384 \times 2 = 156,768$. We split the dataset into training ($n_{train} = 114,298$) and test ($n_{test} = 42,470$) sets in the following manner: (1) randomly select 12 locations for each time slice; (2) assign observations (for both variables) corresponding to space-time locations on the $5^2 \times 3$ space-time cube centered at the selected locations to the test set; and (3) assign the remaining space-time locations to the training set.

We fit a joint model of tasmin and tasmax using the training set and then carried out predictions on the test set. Let \mathbf{y}_{tasmin} and \mathbf{y}_{tasmax} be training vectors of tasmin and tasmax, respectively. We modeled them as

$$\begin{bmatrix} \mathbf{y}_{tasmin} \\ \mathbf{y}_{tasmax} \end{bmatrix} \sim \mathcal{N}_{n+n} \left(\begin{bmatrix} \mathbf{1} \mathbf{0} \\ \mathbf{1} \mathbf{1} \end{bmatrix} \begin{bmatrix} \beta_0 \\ \beta_1 \end{bmatrix}, \mathbf{K} \right)$$

using a Matérn covariance function with a different range parameter for each dimension (latitude, longitude, time, and latent dimension); that is,

$$\begin{aligned} \mathbf{K}_{i,j} &= K(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j) \\ &= \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \|\mathbf{A}^{-1}(\tilde{\mathbf{x}}_i - \tilde{\mathbf{x}}_j)\|^\nu B_\nu(\|\mathbf{A}^{-1}(\tilde{\mathbf{x}}_i - \tilde{\mathbf{x}}_j)\|), \end{aligned}$$

$\tilde{\mathbf{x}} = (\mathbf{x}^\top, \xi)^\top$, \mathbf{x} is a space-time coordinate, ξ is an indicator variable that indicates whether $\tilde{\mathbf{x}}$ corresponds to tasmin, Γ is the gamma function, B_ν is the modified Bessel function of

the second kind, and $\mathbf{A} = diag(r_{lat}, r_{lon}, r_t, r_l)$. Assuming that $\nu = 0.75$ (based on preliminary analyses) and a nugget of zero, we estimated the unknown parameters $\beta_0, \beta_1, r_{lat}, r_{lon}, r_t, r_l$ using the Fisher scoring approach described in Sect. 3.3; the result is given in Table 1.

Figure 11 shows the prediction performance for the test set, as measured by the root mean square prediction error (RMSPE), compared to five other Vecchia variants. S-E-MM + S-E-NN and S-E-MM + J-E-NN are from Sect. 4.2 and based on Euclidean distance between unit-scaled space-time coordinates. T-ord + T-NN is from Sect. 4.3. Note that T-ord separately orders observations of each temperature field by time and then joins them. S-C-NN carries out C-NN conditioning separately for each temperature field, while J-C-NN searches C-NN in the joint vector. We applied a grouping algorithm (Guinness 2018) for improving computational efficiency to all methods except T-ord + T-NN, because interestingly it resulted in a doubling of the computational cost for that method.

CVecchia (C-MM + C-NN) provided the lowest RMSPE for any m considered. The improvement was substantial, with CVecchia's accuracy with $m = 10$ surpassing that of S-E-MM + J-E-NN with $m = 50$, whose computational cost is roughly two orders of magnitude higher due to the cubic scaling in m . Moreover, as shown in the right panel of Fig. 11, CVecchia offered a better trade-off between run time and prediction accuracy. The run-time analysis was performed on a 64-bit workstation with 16 GB RAM and an Intel Core i7-8700K CPU running at 3.70 GHz. We also carried out a comparison in terms of the logarithmic score, but the resulting curves looked almost identical to the RMSPE curves in Fig. 11.

6 Conclusions

We have introduced CVecchia, a covariance approximation that results in a sparse inverse Cholesky factor, whose ordering and sparsity pattern are based on the correlation structure. For reducible GPs, CVecchia implicitly applies a Euclidean-based Vecchia approximation in a transformed input domain in which the GP is isotropic. CVecchia is applicable to any covariance matrix, and it even allows for likelihood-based inference on unknown covariance parameters. We numerically demonstrated the applicability of CVecchia to a variety of covariance structures, some of which had no applicable existing Vecchia approximations. In settings with suitable existing approximations, CVecchia strongly outperformed them.

Special cases of our general CVecchia idea have already been successfully employed in several applications (which were started later but completed earlier than the present paper): Katzfuss et al. (2022) used the idea to approxi-

Table 1 For the NARCCAP data, parameter estimates for the six methods ($m = 50$) using a Matérn covariance function with a different range parameter for each dimension, smoothness $\nu = 0.75$, and zero nugget

	$\hat{\beta}_0$	$\hat{\beta}_1$	$\hat{\sigma}^2$	\hat{r}_{lat}	\hat{r}_{lon}	\hat{r}_t	\hat{r}_l
S-E-MM + S-E-NN	278.513	12.084	61.983	0.819	0.742	0.036	2.000
S-E-MM + J-E-NN	278.200	14.033	62.906	0.828	0.750	0.036	2.681
T-ord + T-NN	275.143	13.613	70.134	0.890	0.792	0.001	2.000
T-ord + S-C-NN	268.826	11.451	51.585	0.719	0.656	0.040	2.000
T-ord + J-C-NN	266.677	12.198	51.331	0.716	0.654	0.040	2.498
C-MM + C-NN	276.754	13.054	38.859	0.593	0.543	0.026	1.668

Temperatures are in Kelvin, the spatial region is scaled to fit into the unit square (without changing its shape), and the time period is scaled to the unit interval.

Fig. 10 Minimum (top) and maximum (bottom) surface air temperature fields (in degrees Kelvin) in the South region (Arkansas, Kansas, Louisiana, Mississippi, Oklahoma and Texas) from NARCCAP

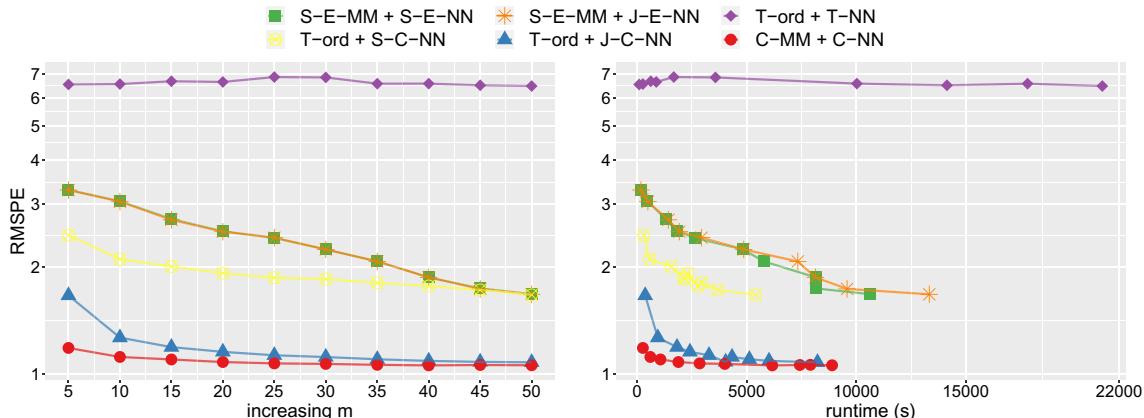
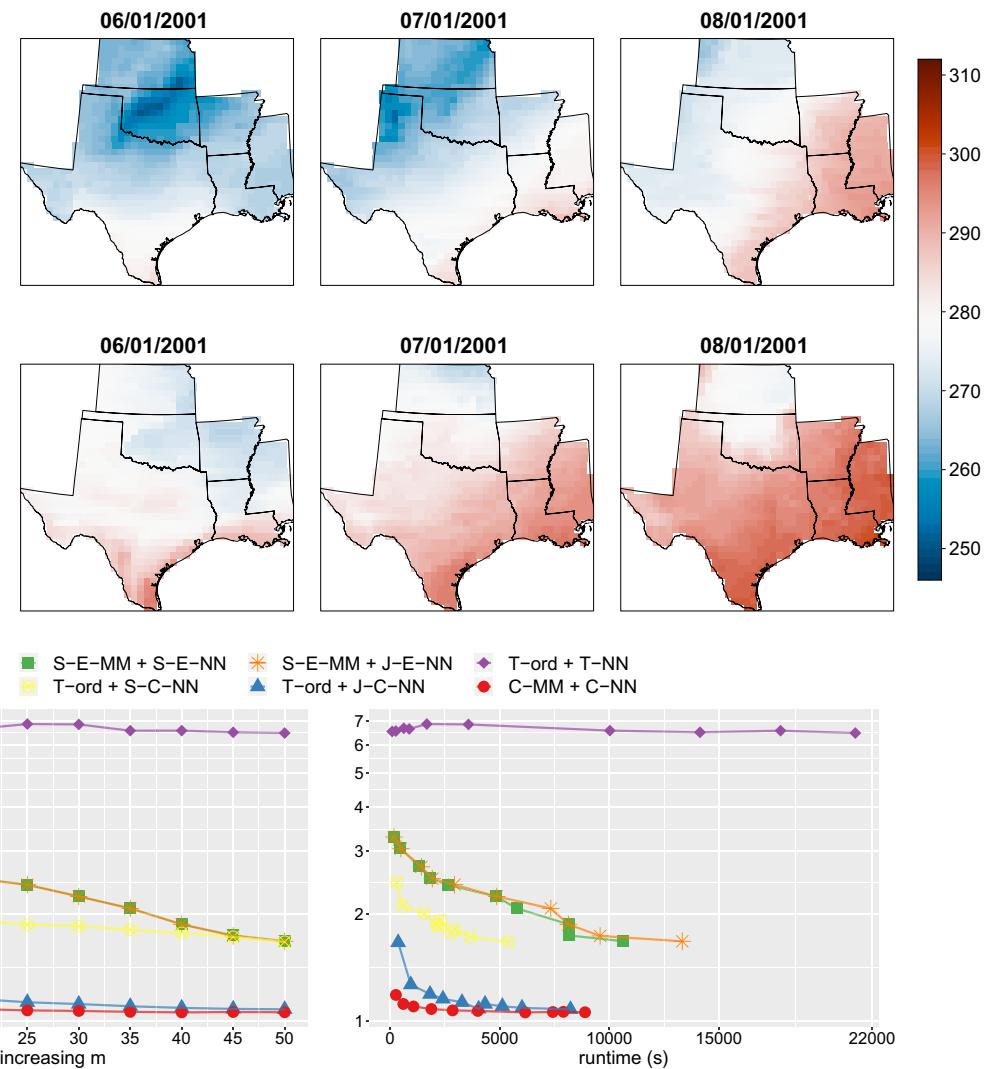


Fig. 11 For the NARCCAP data, root mean squared prediction error (on a log scale) at held-out test points as a function of m (left) and as a function of training time of the Fisher-scoring algorithm (right)

mate anisotropic GPs for computer-model emulation in high input dimension; Messier and Katzfuss (2021) approximated spatio-temporal land-use regression for ground-level nitrogen dioxide; and in the context of nonparametric inference (Kidd and Katzfuss 2021), ideas related to CVecchia were used with sample correlations instead of parametric correlations.

While we have largely focused on geospatial settings here, CVecchia can also be applied to large-scale machine learning settings where input domains are not Euclidean and there is no explicit expression of covariance function. Examples include: multi-task learning (Groot et al. 2011; Williams et al. 2007), where multiple observations are collected from multiple related tasks and joint modeling utilizes intra- and inter-task relatedness; natural language processing (NLP; see Min et al. 2021, for recent review), where words are represented in a latent vector space using word embedding methods (e.g., Beck et al. 2014; Beck 2017; Deriu et al. 2017) and CVecchia can be applied efficiently based on only geometric relations between word vectors; and modeling of multiple interacting latent chemical species in biochemical interaction networks (Gao et al. 2008), where the covariance function of gene expression is an integral equation of inputs.

Acknowledgements Both authors were partially supported by National Science Foundation (NSF) Grant DMS-1953005. Katzfuss' research was also partially supported by NSF Grants DMS-1654083 and CCF-1934904. We would like to thank Florian Schäfer and Joseph Guinness for helpful comments and discussions.

A Proofs

Before proving the main result (Proposition 1 in Sect. 3.2), we shall need the following simple lemma.

Lemma 1 Suppose that $y(\cdot) \sim \mathcal{GP}(0, K)$ on \mathbb{R}^d is q -reducible with respect to ψ . Define $y_0(\cdot) = y(\psi^{-1}(\cdot)) \sim \mathcal{GP}(0, K_0)$. Then,

$$K(\mathbf{x}, \mathbf{x}') = K_0(\|\psi(\mathbf{x}) - \psi(\mathbf{x}')\|), \quad \mathbf{x}, \mathbf{x}' \in \mathbb{R}^d.$$

Proof of Lemma 1 Note that the isotropic covariance function K_0 is only a function of Euclidean distance between inputs. For any inputs $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^d$,

$$\begin{aligned} K(\mathbf{x}, \mathbf{x}') &= \text{cov}(y(\mathbf{x}), y(\mathbf{x}')) \\ &= \text{cov}(y(\psi^{-1}(\psi(\mathbf{x}))), y(\psi^{-1}(\psi(\mathbf{x}')))) \\ &= K_0(\|\psi(\mathbf{x}) - \psi(\mathbf{x}')\|), \end{aligned}$$

where $\psi(\mathbf{x}), \psi(\mathbf{x}') \in \mathbb{R}^q$. \square

Proof of Proposition 1 From Lemma 1,

$$\begin{aligned} \tau_C(i, j) &= \left(1 - \frac{K(\mathbf{x}_i, \mathbf{x}_j)}{\sqrt{K(\mathbf{x}_i, \mathbf{x}_i)}\sqrt{K(\mathbf{x}_j, \mathbf{x}_j)}}\right)^{1/2} \\ &= \left(1 - \frac{K_0(\|\psi(\mathbf{x}_i) - \psi(\mathbf{x}_j)\|)}{\sqrt{K_0(\|\psi(\mathbf{x}_i) - \psi(\mathbf{x}_i)\|)}\sqrt{K_0(\|\psi(\mathbf{x}_j) - \psi(\mathbf{x}_j)\|)}}\right)^{1/2} \\ &= \left(1 - \frac{K_0(\|\psi(\mathbf{x}_i) - \psi(\mathbf{x}_j)\|)}{K_0(0)}\right)^{1/2}, \end{aligned}$$

which is strictly increasing in $\tau_E^\psi(i, j) = \|\psi(\mathbf{x}_i) - \psi(\mathbf{x}_j)\|$, the Euclidean distances between the corresponding transformations. Then, since each step of the MM ordering only depends on the ranking of distances between inputs, for each k ,

$$\begin{aligned} &\arg \max_{i \in \mathcal{I} \setminus \mathcal{I}_{1:k-1}} \min_{j \in \mathcal{I}_{1:k-1}} \tau_C(i, j) \\ &= \arg \max_{i \in \mathcal{I} \setminus \mathcal{I}_{1:k-1}} \min_{j \in \mathcal{I}_{1:k-1}} \tau_E^\psi(i, j), \end{aligned}$$

and so C-MM of the inputs is identical to E-MM of their transformations. For the same reason, C-NN of the inputs is identical to E-NN of their transformations. Therefore, if the first index is chosen to be same for both C-MM and E-MM orderings, CVecchia of $y(\cdot)$ at the inputs $\mathbf{x}_1, \dots, \mathbf{x}_n$ is equivalent to EVecchia of $y(\psi^{-1}(\cdot))$ at the transformed inputs $\psi(\mathbf{x}_1), \dots, \psi(\mathbf{x}_n)$. \square

References

- Apanasovich, T.V., Genton, M.G.: Cross-covariance functions for multivariate random fields based on latent dimensions. *Biometrika* **97**(1), 15–30 (2010)
- Banerjee, S., Carlin, B.P., Gelfand, A.E.: *Hierarchical Modeling and Analysis for Spatial Data*. Chapman & Hall (2004)
- Beck, D.: Modelling representation noise in emotion analysis using Gaussian processes. In: *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pp. 140–145 (2017)
- Beck, D., Cohn, T., Specia, L.: Joint emotion analysis via multi-task Gaussian processes. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1798–1803. ACL (2014)
- Cressie, N., Wikle, C.K.: *Statistics for Spatio-Temporal Data*. Wiley, Hoboken (2011)
- Curriero, F.C.: On the use of non-Euclidean distance measures in geostatistics. *Math. Geol.* **38**(8), 907–926 (2006)
- Datta, A., Banerjee, S., Finley, A.O., Gelfand, A.E.: Hierarchical nearest-neighbor Gaussian process models for large geostatistical datasets. *J. Am. Stat. Assoc.* **111**(514), 800–812 (2016)
- Datta, A., Banerjee, S., Finley, A.O., Hamm, N.A.S., Schaap, M.: Non-separable dynamic nearest-neighbor Gaussian process models for large spatio-temporal data with an application to particulate matter analysis. *Ann. Appl. Stat.* **10**(3), 1286–1316 (2016)
- Deriu, J., Lucchi, A., De Luca, V., Severyn, A., Müller, S., Cieliebak, M., Hofmann, T., Jaggi, M.: Leveraging large amounts of weakly

supervised data for multi-language sentiment classification. In: Proceedings of the 26th International Conference on World Wide Web, pp. 1045–1052 (2017)

Eidsvik, J., Finley, A.O., Banerjee, S., Rue, H.: Approximate Bayesian inference for large spatial datasets using predictive process models. *Comput. Stat. Data Anal.* **56**(6), 1362–1380 (2012)

Eidsvik, J., Shaby, B.A., Reich, B.J., Wheeler, M., Niemi, J.: Estimation and prediction in spatial models with block composite likelihoods using parallel computing. *J. Comput. Graph. Stat.* **23**(2), 295–315 (2014)

Finley, A.O., Sang, H., Banerjee, S., Gelfand, A.E.: Improving the performance of predictive process modeling for large datasets. *Comput. Stat. Data Anal.* **53**(8), 2873–2884 (2009)

Gao, P., Honkela, A., Ratnay, M., Lawrence, N.D.: Gaussian process modelling of latent chemical species: applications to inferring transcription factor activities. *Bioinformatics* **24**(16), i70–i75 (2008)

Gneiting, T., Katzfuss, M.: Probabilistic forecasting. *Annu. Rev. Stat. Appl.* **1**(1), 125–151 (2014)

Groot, P., Burlutiu, A., Heskes, T.: Learning from multiple annotators with Gaussian processes. In: International Conference on Artificial Neural Networks, pp. 159–164. Springer (2011)

Gu, M., Wang, L.: Scaled Gaussian stochastic process for computer model calibration and prediction. *SIAM/ASA J. Uncertain. Quant.* **6**(4), 1555–1583 (2018)

Guinness, J.: Permutation and grouping methods for sharpening Gaussian process approximations. *Technometrics* **60**(4), 415–429 (2018)

Guinness, J.: Gaussian process learning via Fisher scoring of Vecchia's approximation. *Stat. Comput.* **31**(25) (2021)

Heaton, M.J., Datta, A., Finley, A.O., Furrer, R., Guinness, J., Guhaniyogi, R., Gerber, F., Gramacy, R.B., Hammerling, D.M., Katzfuss, M., Lindgren, F., Nychka, D.W., Sun, F., Zammit-Mangion, A.: A case study competition among methods for analyzing large spatial data. *J. Agric. Biol. Environ. Stat.* **24**(3), 398–425 (2019)

Johnson, W.B., Lindenstrauss, J.: Extensions of Lipschitz mappings into a Hilbert space. *Contemp. Math.* **26**(189–206), 1 (1984)

Jones, D.R., Schonlau, M., Welch, W.J.: Efficient global optimization of expensive black-box functions. *J. Global Optim.* **13**, 455–492 (1998)

Jones, R.H., Zhang, Y.: Models for continuous stationary space-time processes. In: Gregoire, T.G., Brillinger, D.R., Diggle, P.J., Russek-Cohen, E., Warren, W.G., Wolfinger, R.D. (eds.) *Modelling Longitudinal and Spatially Correlated Data*, pp. 289–298. Springer, New York (1997)

Karl, T., Koss, W.J.: Regional and national monthly, seasonal, and annual temperature weighted by area, 1895–1983. *Historical Climatology Series 4–3*, p. 38 (1984)

Katzfuss, M.: A multi-resolution approximation for massive spatial datasets. *J. Am. Stat. Assoc.* **112**(517), 201–214 (2017)

Katzfuss, M., Gong, W.: A class of multi-resolution approximations for large spatial datasets. *Stat. Sin.* **30**(4), 2203–2226 (2020)

Katzfuss, M., Guinness, J.: A general framework for Vecchia approximations of Gaussian processes. *Stat. Sci.* **36**(1), 124–141 (2021)

Katzfuss, M., Guinness, J., Gong, W., Zilber, D.: Vecchia approximations of Gaussian-process predictions. *J. Agric. Biol. Environ. Stat.* **25**(3), 383–414 (2020)

Katzfuss, M., Guinness, J., Lawrence, E.: Scaled Vecchia approximation for fast computer-model emulation. *SIAM/ASA Journal on Uncertainty Quantification*, accepted (2022)

Kennedy, M.C., O'Hagan, A.: Bayesian calibration of computer models. *J. R. Stat. Soc. B* **63**(3), 425–464 (2001)

Kidd, B., Katzfuss, M.: Bayesian nonstationary and nonparametric covariance estimation for large spatial data. *Bayesian Analysis*, accepted (2021)

Konomi, B.A., Hanandeh, A.A., Ma, P., Kang, E.L.: Computationally efficient nonstationary nearest-neighbor Gaussian process models using data-driven techniques. *Environmetrics* **30**(8), e2571 (2019)

Liu, H., Ong, Y.-S., Shen, X., Cai, J.: When Gaussian process meets big data: a review of scalable GPs. *IEEE Trans. Neural Netw. Learn. Syst.* (2020)

Maehara, H.: Euclidean embeddings of finite metric spaces. *Discrete Math.* **313**(23), 2848–2856 (2013)

Matousek, J.: *Lectures on Discrete Geometry*, vol. 212. Springer, Berlin (2013)

Mearns, L.O., Arritt, R., Biner, S., Bukovsky, M.S., McGinnis, S., Sain, S., Caya, D., Correia, J., Flory, D., Gutowski, W., et al.: The North American regional climate change assessment program: overview of phase I results. *Bull. Am. Meteor. Soc.* **93**(9), 1337–1362 (2012)

Mearns, L.O., Gutowski, W., Jones, R., Leung, R., McGinnis, S., Nunes, A., Qian, Y.: A regional climate change assessment program for North America. *EOS Trans. Am. Geophys. Union* **90**(36), 311–311 (2009)

Messier, K.P., Katzfuss, M.: Scalable penalized spatiotemporal land-use regression for ground-level nitrogen dioxide. *Ann. Appl. Stat.* **15**(2), 688–710 (2021)

Min, B., Ross, H., Sulem, E., Veyseh, A.P.B., Nguyen, T.H., Sainz, O., Agirre, E., Heinz, I., Roth, D.: Recent advances in natural language processing via large pre-trained language models: a survey. [arXiv:2111.01243](https://arxiv.org/abs/2111.01243) (2021)

Paciorek, C., Schervish, M.: Spatial modelling using a new class of nonstationary covariance functions. *Environmetrics* **17**(5), 483–506 (2006)

Perrin, O., Meiring, W.: Nonstationarity in \mathbb{R}^n is second-order stationarity in \mathbb{R}^{2n} . *J. Appl. Probab.* **40**(3), 815–820 (2003)

Perrin, O., Monestiez, P.: Modelling of non-stationary spatial structure using parametric radial basis deformations. In: *geoENV II-Geostatistics for Environmental Applications*, pp. 175–186. Springer (1999)

Perrin, O., Schlather, M.: Can any multivariate Gaussian vector be interpreted as a sample from a stationary random process? *Stat. Probab. Lett.* **77**(9), 881–884 (2007)

Perrin, O., Senoussi, R.: Reducing non-stationary random fields to stationarity and isotropy using a space deformation. *Stat. Probab. Lett.* **48**(1), 23–32 (2000)

Porcu, E., Matkowski, J., Mateu, J.: On the non-reducibility of non-stationary correlation functions to stationary ones under a class of mean-operator transformations. *Stoch. Environ. Res. Risk Assess.* **24**(5), 599–610 (2010)

Rasmussen, C.E., Williams, C.K.I.: *Gaussian Processes for Machine Learning*. MIT Press, Cambridge (2006)

Risser, M.D., Turek, D.: Bayesian inference for high-dimensional non-stationary Gaussian processes. *J. Stat. Comput. Simul.* (2020)

Sacks, J., Welch, W., Mitchell, T., Wynn, H.: Design and analysis of computer experiments. *Stat. Sci.* **4**(4), 409–435 (1989)

Sang, H., Jun, M., Huang, J.Z.: Covariance approximation for large multivariate spatial datasets with an application to multiple climate model errors. *Ann. Appl. Stat.* **5**(4), 2519–2548 (2011)

Schäfer, F., Katzfuss, M., Owhadi, H.: Sparse Cholesky factorization by Kullback–Leibler minimization. *SIAM J. Sci. Comput.* **43**(3), A2019–A2046 (2021)

Schmidt, A.M., O'Hagan, A.: Bayesian inference for non-stationary spatial covariance structure via spatial deformations. *J. R. Stat. Soc. B* **65**(3), 743–758 (2003)

Snelson, E., Ghahramani, Z.: Local and global sparse Gaussian process approximations. In: *Artificial Intelligence and Statistics 11 (AISTATS)* (2007)

Stein, M.L.: Nonstationary spatial covariance functions. Technical Report No. 21, University of Chicago (2005)

Stein, M.L., Chi, Z., Welty, L.: Approximating likelihoods for large spatial data sets. *J. R. Stat. Soc. B* **66**(2), 275–296 (2004)

Sun, Y., Stein, M.L.: Statistically and computationally efficient estimating equations for large spatial datasets. *J. Comput. Graph. Stat.* **25**(1), 187–208 (2016)

Van Dongen, S., Enright, A.J.: Metric distances derived from cosine similarity and Pearson and Spearman correlations. [arXiv:1208.3145](https://arxiv.org/abs/1208.3145) (2012)

Varin, C.: On composite marginal likelihoods. *AStA Adv. Stat. Anal.* **92**(1), 1–28 (2008)

Vecchia, A.: Estimation and model identification for continuous spatial processes. *J. R. Stat. Soc. B* **50**(2), 297–312 (1988)

Vu, Q., Zammit-Mangion, A., Cressie, N.: Modeling nonstationary and asymmetric multivariate spatial covariances via deformations. [arXiv:2004.08724](https://arxiv.org/abs/2004.08724) (2020)

White, P., Porcu, E.: Nonseparable covariance models on circles cross time: a study of Mexico City ozone. *Environmetrics*, p. e2558 (2019)

Williams, C., Bonilla, E.V., Chai, K.M.: Multi-task Gaussian process prediction. *Advances in Neural Information Processing Systems*, pp. 153–160 (2007)

Witsenhausen, H.S.: Minimum dimension embedding of finite metric spaces. *J. Comb. Theory Ser. A* **42**(2), 184–199 (1986)

Yang, G.: Wide feedforward or recurrent neural networks of any architecture are Gaussian processes. In: *Advances in Neural Information Processing Systems*, pp. 9951–9960 (2019)

Yu, C.D., Levitt, J., Reiz, S., Biros, G.: Geometry-oblivious FMM for compressing dense SPD matrices. In: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, p. 53. ACM (2017)

Zhang, L., Banerjee, S., Finley, A.O.: High-dimensional multivariate geostatistics: a Bayesian matrix-normal approach. *Environmetrics* **32**(4), e2675 (2021)

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

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.