

UCLA

Research Reports

Title

Highly Scalable Bayesian Geostatistical Modeling via Meshed Gaussian Processes on Partitioned Domains

Permalink

<https://escholarship.org/uc/item/8pm5v0f8>

Authors

Peruzzi, Michele
Banerjee, Sudipto
Finley, Andrew O

Publication Date

2020-11-24

DOI

10.1080/01621459.2020.1833889

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed





Highly Scalable Bayesian Geostatistical Modeling via Meshed Gaussian Processes on Partitioned Domains

Michele Peruzzi , Sudipto Banerjee & Andrew O. Finley


To cite this article: Michele Peruzzi , Sudipto Banerjee & Andrew O. Finley (2020): Highly Scalable Bayesian Geostatistical Modeling via Meshed Gaussian Processes on Partitioned Domains, Journal of the American Statistical Association, DOI: [10.1080/01621459.2020.1833889](https://doi.org/10.1080/01621459.2020.1833889)

To link to this article: <https://doi.org/10.1080/01621459.2020.1833889>

 View supplementary material [↗](#)

 Published online: 24 Nov 2020.

 Submit your article to this journal [↗](#)

 Article views: 178

 View related articles [↗](#)

 View Crossmark data [↗](#)



Highly Scalable Bayesian Geostatistical Modeling via Meshed Gaussian Processes on Partitioned Domains

Michele Peruzzi^{a,b}, Sudipto Banerjee^c, and Andrew O. Finley^a

^aDepartment of Forestry, Michigan State University, East Lansing, MI; ^bDepartment of Statistical Science, Duke University, Durham, NC; ^cDepartment of Biostatistics, UCLA Fielding School of Public Health, Los Angeles, CA

ABSTRACT

We introduce a class of scalable Bayesian hierarchical models for the analysis of massive geostatistical datasets. The underlying idea combines ideas on high-dimensional geostatistics by partitioning the spatial domain and modeling the regions in the partition using a sparsity-inducing directed acyclic graph (DAG). We extend the model over the DAG to a well-defined spatial process, which we call the meshed Gaussian process (MGP). A major contribution is the development of an MGPs on tessellated domains, accompanied by a Gibbs sampler for the efficient recovery of spatial random effects. In particular, the cubic MGP (Q-MGP) can harness high-performance computing resources by executing all large-scale operations in parallel within the Gibbs sampler, improving mixing and computing time compared to sequential updating schemes. Unlike some existing models for large spatial data, a Q-MGP facilitates massive caching of expensive matrix operations, making it particularly apt in dealing with spatiotemporal remote-sensing data. We compare Q-MGPs with large synthetic and real world data against state-of-the-art methods. We also illustrate using Normalized Difference Vegetation Index data from the Serengeti park region to recover latent multivariate spatiotemporal random effects at millions of locations. The source code is available at github.com/mkln/meshgp. Supplementary materials for this article are available online.

ARTICLE HISTORY

Received March 2020
Accepted October 2020

KEYWORDS

Bayesian; Domain partitioning; Graphical models; Large n ; Sparsity; Spatial



1. Introduction


Collecting large quantities of spatial and spatiotemporal data is now commonplace in many fields. In ecology and forestry, massive datasets are collected using satellite imaging and other remote sensing instruments such as LiDAR that periodically record high-resolution images. Unfortunately, clouds frequently obstruct the view resulting in large regions with missing information. Figure 1 shows this phenomenon in Normalized Difference Vegetation Index (NDVI) data from the Serengeti region. Filling such gaps in the data is an important goal as is quantifying uncertainty in predictions. This goal is achieved through stochastic modeling of the underlying phenomenon, which involves the specification of a spatial or spatiotemporal process characterizing dependence from a finite realization. Gaussian processes (GPs) are a customary choice to characterize spatial dependence, but their implementation is notoriously burdened by their $O(n^3)$ computational complexity. Consequently, intense research has been devoted in recent years to developing scalable models for large spatial datasets—see detailed reviews by Sun, Li, and Genton (2011) and Banerjee (2017).

Computational complexity can be reduced by considering low-rank models; among these, knot-based methods motivated by “kriging” ideas enjoy some optimality properties but oversmooth the estimates of spatial random effects unless the number of knots is large, and require corrections to avoid


overestimation of the nugget (Banerjee et al. 2008; Cressie and Johannesson 2008; Banerjee et al. 2010; Guhaniyogi et al. 2011; Finley, Banerjee, and Gelfand 2012). Other methods reduce the computational burden by introducing sparsity in the covariance matrix; strategies include tapering (Furrer, Genton, and Nychka 2006; Kaufman, Schervish, and Nychka 2008) or partitioning of the spatial domain into regions with a typical assumption of independence across regions (Sang and Huang 2012; Stein 2014). These can be improved by considering a recursive partitioning scheme, resulting in a multi-resolution approximation (MRA; Katzfuss 2017). Other assumptions on conditional independence assumptions also have a good track record in terms of scalability to large spatial datasets: Gaussian random Markov fields (GMRF; Rue and Held 2005), composite likelihood methods (Eidsvik et al. 2014), and neighbor-based likelihood approximations (Vecchia 1988) belong to this family.

The recent literature has witnessed substantial activity surrounding the so called Vecchia approximation (Vecchia 1988). This approximation can be regarded as a special case of the GMRF approximations with a simplified neighborhood structure motivated from a directed acyclic graphical (DAG) representation of a GP likelihood. Extensions leading to well-defined spatial processes to accommodate inference at arbitrary locations by extending the DAG representation to the entire domain include nearest neighbor Gaussian processes (NNGPs; Datta,

CONTACT Sudipto Banerjee  sudipto@ucla.edu  Department of Biostatistics, UCLA Fielding School of Public Health, 650 Charles E. Young Drive South, Los Angeles, CA 90095.

 Supplementary materials for this article are available online. Please go to www.tandfonline.com/r/JASA.

NDVI



Elevation




Figure 1. Left: NDVI in the Serengeti region on 2016-12-17. White areas correspond to missing data due to cloud cover. Right: Elevation data for the same region.

Banerjee, Finley, Gelfand, et al. 2016; Datta, Banerjee, Finley, Hamm, et al. 2016) and further generalizations by constructing DAGs over the augmented space of outcomes and spatial effects (Katzfuss and Guinness 2017). These approaches render computational scalability by introducing sparsity in the precision matrix. The DAG relies upon a specific topological ordering of the locations, which also determine the construction of neighborhood sets, and certain orderings tend to deliver improved performance of such models (Katzfuss and Guinness 2017; Guinness 2018).

When inference on the latent process is sought, Bayesian inference has the benefits of providing direct probability statements based upon the posterior distribution of the process. Inference based on asymptotic approximations are avoided, but there remain challenges in computing the posterior distribution given that inference is sought on a very high-dimensional parameter space (including the realizations of the latent process). One possibility, available for Gaussian first-stage likelihoods, is to work with a collapsed or marginalized likelihood by integrating out the spatial random effects. However, Gibbs samplers and other MCMC algorithms for the collapsed models can be inexorably slow and are impractical when data are in the millions. A sequential Gibbs sampler that updates the latent spatial effects (Datta, Banerjee, Finley, Gelfand, et al. 2016) is faster in updating the parameters but suffers from high autocorrelation and slow mixing. Another possibility emerges when interest lies in prediction or imputation of the outcome variable only and not the latent process. Here, a so called “response” model that models the outcome itself using an NNGP can be constructed. This model is much faster and enjoys superior convergence properties, but we lose inference on the latent process and its predictive performance tends to be inferior to the latent process model. Furthermore, these options are unavailable in non-Gaussian first-stage hierarchical models or when the focus is not uniquely on prediction. A detailed comparison of different approaches for computing Bayesian NNGP models is presented in Finley et al. (2019).

Our current contribution introduces a class of *meshed Gaussian process* (MGP) models for Bayesian hierarchical modeling

of large spatial datasets. This class builds upon the aforementioned works that build upon Vecchia (1988) and other DAG based models. The inferential focus remains within the context of massive spatial datasets over very large domains. We exploit the demonstrated benefits of the DAG based models, but we now adapt them to partitioned domains. We describe dependence across regions of a partitioned domain using a small, patterned DAG which we refer to as a *mesh*. Within each region, some locations are selected as *reference* and collectively mapped to a single node in the DAG. Relationships among nodes are governed by kriging ideas. In the resulting MGP, regions in the spatial domain depend on each other through the reference locations. Realizations at all other locations are assumed independent, conditional upon the reference locations. This construction leads to a valid standalone spatial process.

As a particular subclass of MGPs, we propose a novel partitioning and graph design based on domain tessellations. Unlike methods that build sparse DAGs by limiting dependence to m nearest neighbors, our approach shapes the underlying DAG with a known, repeating pattern corresponding to the chosen tessellation geometry. The underlying sparse DAG enables scaling computations to large data settings and its known pattern guarantees the availability of block-parallel sampling schemes; furthermore, large computational savings can be achieved at no additional approximation cost if data are collected on patterned lattices. Finally, extensions to spatiotemporal and/or multivariate data are straightforward once a suitable covariance function has been defined. We use axis-parallel domain partitioning and the corresponding cubic DAG—resulting in cubic MGPs or Q-MGPs—to show substantial improvements in computational time and inferential performance relative to other models with data sizes ranging from the thousands to the several millions, for both spatial and spatiotemporal data and using multivariate spatial processes.

The present work may appear to share similarities with the block-NNGP model of Quiroz, Prates, and Dey (2019), who advocate building sparse DAGs on grouped locations based on their ordering and subsequent identification of m “past” neighbors. Unlike block-NNGPs, our tessellated GPs

consider the domain tessellation as generating the DAG; the number of parents of any node is thus fixed and depends on the geometry of the chosen tessellation rather than on a user-defined parameter. Inclusion of more distant locations in the parent set of any location will, therefore, not proceed by increasing the number of neighbors m , but rather by increasing the regions' size and/or modifying their shape. Central to tessellated GPs is the idea of forcing a DAG with known *coloring* on the data, resulting in guaranteed efficiencies when recovering the latent spatial effects. This strategy is analogous in spirit to multi-resolution approximations (Gramacy and Lee 2008; Katzfuss 2017), which also force a DAG on the data, resulting in conditional independence patterns that are known in advance and that can be used to improve computations. However, while multi-resolution approximations are defined by branching graphs associated to recursive domain partitioning, tessellated GPs use a single domain partition, with each region connected in the DAG only to its immediate neighbors. Compared to treed graphs, tessellated GPs are associated to DAGs with fewer conditionally independent groups and whose repeated patterns facilitate the identification of redundant matrix operations arising when one or more coordinate margins are gridded. We also note that while the idea of partitioning domains to create approximations is not new, construction of the DAG-based approximation over partitioned domains has received considerably less attention. Finally, our focus here is in developing tessellated GPs as a methodology that enables the efficient recovery of the latent spatial random effects and the Bayesian estimation of covariance parameters via MCMC; we are thus not focusing on alternative computational algorithms (see, e.g., Finley et al. 2019), which have been developed for NNGPs but can nonetheless all be adapted to general MGP models.

The balance of this article proceeds as follows. Section 2 introduces our general framework for hierarchical Bayesian modeling of spatial processes using networks of grouped spatial locations. The MGP is outlined in Section 3, where we provide a general, scalable computing algorithm in Section 3.1. Tessellation-based schemes and the specific case of Q-MGPs are outlined in Section 4, which highlights their properties and computational advantages. We illustrate the performance of our proposed approach in Section 5 using simulation experiments and an application on a massive dataset with millions of spatiotemporal locations. We conclude the article with a discussion and pointers to further research. Supplementary materials accompanying this article as an appendix are available online and contain further comparisons of Q-MGPs with several state-of-the-art methods for spatial data.

2. Spatial Processes on Partitioned Domains

A $q \times 1$ spatial process assigns a probability law on $\{\mathbf{w}(\ell) : \ell \in \mathcal{D}\}$, where $\mathbf{w}(\ell)$ is a $q \times 1$ random vector with elements $w_i(\ell)$ for $i = 1, 2, \dots, q$. In the following general discussion we will not distinguish between spatial ($\mathcal{D} \subset \mathbb{R}^d$) and spatiotemporal domains ($\mathcal{D} \subset \mathbb{R}^{d+1}$), and denote spatial or spatiotemporal locations as ℓ , s , or \mathbf{u} .

For any finite set of spatial locations $\{\ell_1, \ell_2, \dots, \ell_{n_{\mathcal{L}}}\} = \mathcal{L} \subset \mathcal{D}$ of size $n_{\mathcal{L}}$, let $P(\cdot)$ denote the probability law of the $n_{\mathcal{L}}q \times 1$ random vector $\mathbf{w}_{\mathcal{L}} = (\mathbf{w}(\ell_1)^\top, \mathbf{w}(\ell_2)^\top, \dots, \mathbf{w}(\ell_{n_{\mathcal{L}}})^\top)^\top$ with probability density $p(\cdot)$. The joint density of $\mathbf{w}_{\mathcal{L}}$ can be expressed as a DAG (or a Bayesian network model) with respect to the ordered set of locations \mathcal{L} as

$$p(\mathbf{w}_{\mathcal{L}}) = \prod_{i=1}^{n_{\mathcal{L}}} p(\mathbf{w}(\ell_i) | \mathbf{w}(\ell_1), \dots, \mathbf{w}(\ell_{i-1})), \quad (1)$$

where the conditional set for each $\mathbf{w}(\ell_i)$ can be interpreted as the set of its parents in a large, dense Bayesian network. Defining a simplified valid joint density on \mathcal{L} by reducing the size of the conditioning sets is a popular strategy for fast likelihood approximations in the context of large spatial datasets. One typically limits dependence to “past” neighboring locations with respect to the ordering in (1) (Vecchia 1988; Stein, Chi, and Welty 2004; Gramacy and Apley 2015; Datta, Banerjee, Finley, Gelfand, et al. 2016; Katzfuss and Guinness 2017). The neighbors are defined and fixed and model performance may benefit from the addition of some distant locations (Stein, Chi, and Welty 2004). The ordering in \mathcal{L} is also fixed and inferential performance may benefit from the use of some fixed permutations (Guinness 2018). The result of shrinking the conditional sets to a smaller set of neighbors from the past yields a sparse DAG or Bayesian network, which yields potentially massive computational gains.

We proceed in a similar manner, but instead of defining a sparse DAG at the level of each individual location, we map entire groups of locations to nodes in a much smaller graph; the same graph will be used to model the dependence between any location in the spatial domain and, therefore, to define a spatial process. Let $\mathcal{P} = \{\mathcal{D}_1, \dots, \mathcal{D}_M\}$ be a partition of \mathcal{D} into M mutually exclusive subsets so that $\mathcal{D} = \cup_{i=1}^M \mathcal{D}_i$ and $\mathcal{D}_i \cap \mathcal{D}_j = \emptyset$ whenever $i \neq j$. Similar to the nomenclature in the NNGP, we fix a *reference set* $\mathcal{S} = \{s_1, \dots, s_{n_{\mathcal{S}}}\} \subset \mathcal{D}$, which itself is partitioned using \mathcal{P} by letting $\mathcal{S}_j = \mathcal{D}_j \cap \mathcal{S}$. The set of nonreference locations is similarly partitioned with $\mathcal{U}_j = \mathcal{D}_j \setminus \mathcal{S}_j$ so that $\mathcal{D}_j = \mathcal{S}_j \cup \mathcal{U}_j$ for each $j = 1, 2, \dots, M$. We now construct a DAG to model dependence within and between \mathcal{S} and \mathcal{U} . Let $\mathcal{G} = \{\mathbf{V}, \mathbf{E}\}$ be a graph with nodes $\mathbf{V} = \mathbf{A} \cup \mathbf{B}$, where we refer to $\mathbf{A} = \{\mathbf{a}_1, \dots, \mathbf{a}_M\}$ as the *reference nodes* and to $\mathbf{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_M\}$ as the *nonreference*, or simply “other”, nodes. Let $\mathbf{A} \cap \mathbf{B} = \emptyset$. We introduce a map $\eta : \mathcal{D} \rightarrow \mathbf{V}$ such that

$$\eta(\ell) = \begin{cases} \mathbf{a}_j \in \mathbf{A} & \text{if } \ell \in \mathcal{S}_j, \\ \mathbf{b}_j \in \mathbf{B} & \text{if } \ell \in \mathcal{U}_j. \end{cases} \quad (2)$$

This surjective many-to-one map links each location in \mathcal{S}_j and \mathcal{U}_j to a node in \mathcal{G} . The edges connecting nodes in \mathcal{G} are $\mathbf{E} = \{\text{Pa}[\mathbf{v}_1], \dots, \text{Pa}[\mathbf{v}_{2M}]\}$ where $\text{Pa}[\mathbf{v}] \subset \mathbf{V}$ denotes the set of parents of any $\mathbf{v} \in \mathbf{V}$ and, hence, identifies the directed edges pointing to \mathbf{v} . We let \mathcal{G} be acyclic, that is, there is no chain $\{\mathbf{v}_{i_1} \rightarrow \mathbf{v}_{i_2} \rightarrow \dots \rightarrow \mathbf{v}_{i_t}\}$ of elements of \mathbf{V} such that $\mathbf{v}_{i_j} \in \text{Pa}[\mathbf{v}_{i_{j+1}}]$ and $\mathbf{v}_{i_{j+1}} \in \text{Pa}[\mathbf{v}_{i_1}]$. Crucially, we assume that $\text{Pa}[\mathbf{v}] \subset \mathbf{A}$ for all $\mathbf{v} \in \mathbf{V}$, that is, that only reference nodes have children, to distinguish the reference nodes \mathbf{A} from the other nodes \mathbf{B} . Apart from the assumption that $\mathbf{a}_j \in \text{Pa}[\mathbf{b}_j]$, we refrain from defining the parents of a node, thereby retaining flexibility. In general, however, all locations in \mathcal{U}_j will share the same parent set. In Section 4, we will consider meshes associated to domain tessellations.

Consider the enumeration $\mathcal{S}_i = \{s_{i_1}, \dots, s_{i_{n_i}}\}$, where $\{i_1, i_2, \dots, i_{n_i}\} \subset \{1, 2, \dots, n_{\mathcal{S}}\}$, and let $\mathbf{w}_i = (\mathbf{w}(s_{i_1})^\top, \mathbf{w}(s_{i_2})^\top, \dots, \mathbf{w}(s_{i_{n_i}})^\top)^\top$ be the $n_i q \times 1$ random vector listing elements of $\mathbf{w}(s)$ for each $s \in \mathcal{S}_i$. We now rewrite (1) as a product of M conditional densities

$$p(\mathbf{w}_{\mathcal{S}}) = p(\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_M) = \prod_{i=1}^M p(\mathbf{w}_i | \mathbf{w}_1, \dots, \mathbf{w}_{i-1}). \quad (3)$$

The conditioning sets are then reduced based on the graph \mathcal{G} :

$$\tilde{p}(\mathbf{w}_{\mathcal{S}}) = \prod_{i=1}^M p(\mathbf{w}_i | \mathbf{w}_{[i]}), \quad (4)$$

where we denote $\mathbf{w}_{[i]} = \{\mathbf{w}_j : \mathbf{a}_j \in \text{Pa}[\mathbf{a}_i]\}$, and $\text{Pa}[\mathbf{a}_i] \subset \{\mathbf{a}_1, \dots, \mathbf{a}_{i-1}\} \subset \mathcal{A}$. This is a proper multivariate joint density since the graph is acyclic (Lauritzen 1996). It is instructive to note how the above approximation behaves when the size of the parent set shrinks, for a given domain partitioning scheme. To this end, we adapt a result in Banerjee (2020) and show that sparser DAGs correspond to a larger Kullback–Leibler (KL) divergence from the base density p . This result has been proved earlier for Gaussian likelihoods by Guinness (2018), but the argument given below is free of distributional assumptions and is linked to the submodularity of entropy and the “information never hurts” principle (see, e.g., Cover and Thomas 1991).

Consider random vector \mathbf{w} and some partition of the domain \mathcal{P} corresponding to nodes $\mathbf{V} = \{v_1, \dots, v_M\}$ via map η . Let the base process correspond to graph $\mathcal{G}_0 = \{\mathbf{V}, \mathbf{E}_0\}$ where $\mathbf{E}_0 = \{\text{Pa}_0[v_1], \dots, \text{Pa}_0[v_M]\}$. Then, let $\mathcal{G}_1 = \{\mathbf{V}, \mathbf{E}_1\}$ where $\mathbf{E}_1 = \{\text{Pa}_1[v_1], \dots, \text{Pa}_1[v_M]\}$ and $\text{Pa}_1[v_i] \subseteq \text{Pa}_0[v_i]$ for all $i \in \{1, \dots, M\}$. Finally construct $\mathcal{G}_2 = \{\mathbf{V}, \mathbf{E}_2\}$ by letting $\text{Pa}_2[v_i^*] = \text{Pa}_1[v_i^*] \setminus \{v_i^*\}$ for some $v_i^* \in \text{Pa}_1[v_i^*]$. In other words, graph \mathcal{G}_2 is obtained by removing the directed edge $v_i^* \rightarrow v_i^*$ from \mathcal{G}_1 . We approximate p using densities p_1 and p_2 based on \mathcal{G}_1 and \mathcal{G}_2 , respectively, obtaining

$$\frac{p_1(\mathbf{w})}{p_2(\mathbf{w})} = \prod_{i=1}^M \frac{p(\mathbf{w}_i | \mathbf{w}_{[i]_1})}{p(\mathbf{w}_i | \mathbf{w}_{[i]_2})} = \frac{p(\mathbf{w}_{i^*} | \mathbf{w}_{[i^*]_1})}{p(\mathbf{w}_{i^*} | \mathbf{w}_{[i^*]_2})}. \quad (5)$$

Considering the KL divergence of each density from p , and denoting $\mathbf{V}^* = \mathbf{V} \setminus \{i^*\} \cup \text{Pa}_1[i^*]$, we find

$$\begin{aligned} & \text{KL}(p_2 \| p) - \text{KL}(p_1 \| p) \\ &= \int \left\{ \log \left(\frac{p(\mathbf{w})}{p_2(\mathbf{w})} \right) - \log \left(\frac{p(\mathbf{w})}{p_1(\mathbf{w})} \right) \right\} p(\mathbf{w}) d\mathbf{w} \\ &= \int \log \left(\frac{p_1(\mathbf{w})}{p_2(\mathbf{w})} \right) p(\mathbf{w}) d\mathbf{w} \\ &= \int \log \left(\frac{p(\mathbf{w}_{i^*} | \mathbf{w}_{[i^*]_1})}{p(\mathbf{w}_{i^*} | \mathbf{w}_{[i^*]_2})} \right) p(\mathbf{w}) d\mathbf{w} \\ &= \int \log \left(\frac{p(\mathbf{w}_{i^*} | \mathbf{w}_{[i^*]_1})}{p(\mathbf{w}_{i^*} | \mathbf{w}_{[i^*]_2})} \right) p(\mathbf{w}_{i^*}, \mathbf{w}_{[i^*]_1}) d\mathbf{w}_{i^*} d\mathbf{w}_{[i^*]_1} \\ &= \int \left\{ \int \log \left(\frac{p(\mathbf{w}_{i^*} | \mathbf{w}_{[i^*]_1})}{p(\mathbf{w}_{i^*} | \mathbf{w}_{[i^*]_2})} \right) p(\mathbf{w}_{i^*} | \mathbf{w}_{[i^*]_1}) d\mathbf{w}_{i^*} \right\} \\ & \quad \times p(\mathbf{w}_{[i^*]_1}) d\mathbf{w}_{[i^*]_1} \geq 0, \end{aligned} \quad (6)$$

where we use (5), the fact that \mathbf{V}^* and $\{i^*\} \cup \text{Pa}_1[i^*]$ are disjoint, and Jensen’s inequality. This result implies that larger parent sets

are preferable as they correspond to better approximations to the full model; the choice of sparser graphs will be driven by computational considerations—see Section 3.2.

We construct the spatial process over arbitrary locations by enumerating other locations as $\mathcal{U} = \{\mathbf{u}_1, \dots, \mathbf{u}_{n_{\mathcal{U}}}\} \subset \mathcal{D} \setminus \mathcal{S}$ and extending (4) to the nonreference locations. Given the partition of \mathcal{U} defined earlier with components \mathcal{U}_j for $j = 1, 2, \dots, M$, for each $\mathbf{u} \in \mathcal{U}_j$ we set $\eta(\mathbf{u}) = \mathbf{b}_j$ and recall that $\text{Pa}[\mathbf{b}_j] \subset \mathcal{A}$ by construction. For each $i = 1, \dots, n_{\mathcal{U}}$, we denote $\mathbf{w}_{[u_i]} = \{\mathbf{w}_j : \mathbf{a}_j \in \text{Pa}[\eta(\mathbf{u}_i)]\} \subset \mathbf{w}_{\mathcal{S}}$ and define the conditional density of $\mathbf{w}_{\mathcal{U}}$ given $\mathbf{w}_{\mathcal{S}}$ as

$$\tilde{p}(\mathbf{w}_{\mathcal{U}} | \mathbf{w}_{\mathcal{S}}) = \prod_{\mathbf{u}_i \in \mathcal{U}} p(\mathbf{w}(\mathbf{u}_i) | \mathbf{w}_{[u_i]}) = \prod_{j=1}^M p(\mathbf{w}_{\mathcal{U}_j} | \mathbf{w}_{[b_j]}). \quad (7)$$

Therefore, for any finite subset of spatial locations $\mathcal{L} \subset \mathcal{D}$ we can let $\mathcal{U} = \mathcal{L} \setminus \mathcal{S}$ and obtain

$$\tilde{p}(\mathbf{w}_{\mathcal{L}}) = \int \tilde{p}(\mathbf{w}_{\mathcal{U}} | \mathbf{w}_{\mathcal{S}}) \tilde{p}(\mathbf{w}_{\mathcal{S}}) \prod_{s_i \in \mathcal{S} \setminus \mathcal{L}} d(\mathbf{w}(s_i)).$$

We show (see Appendix A, available online) that this is a well-defined process by verifying the Kolmogorov consistency conditions. This new process can be built starting from a base process, a fixed reference set, domain partition \mathcal{P} and a graph \mathcal{G} . Next, we elucidate with GPs.

3. Meshed Gaussian Processes

Let $\{\mathbf{w}(\boldsymbol{\ell}) : \boldsymbol{\ell} \in \mathcal{D}\}$ be a q -variate multivariate GP, denoted as $\mathbf{w}(\boldsymbol{\ell}) \sim \text{GP}(\mathbf{0}, \mathbf{C}(\cdot, \cdot | \boldsymbol{\theta}))$. The cross-covariance $\mathbf{C}(\cdot, \cdot | \boldsymbol{\theta})$ indexed by parameters $\boldsymbol{\theta}$ is a function $\mathbf{C} : \mathcal{D} \times \mathcal{D} \rightarrow \mathcal{M}_{q \times q}$, where $\mathcal{M}_{q \times q}$ is a subset of $\mathfrak{R}^{q \times q}$ (the space of all $q \times q$ real matrices) such that the (i, j) th entry of $\mathbf{C}(\boldsymbol{\ell}, \boldsymbol{\ell}' | \boldsymbol{\theta})$ evaluates the covariance between the i th and j th elements of $\mathbf{w}(\boldsymbol{\ell})$ at $\boldsymbol{\ell}$ and $\boldsymbol{\ell}'$, respectively, that is, $\text{cov}(w_i(\boldsymbol{\ell}), w_j(\boldsymbol{\ell}'))$. We omit dependence on $\boldsymbol{\theta}$ to simplify notation. The cross-covariance function itself needs to be neither symmetric nor positive-definite, but must satisfy the following two properties: (i) $\mathbf{C}(\boldsymbol{\ell}, \boldsymbol{\ell}') = \mathbf{C}(\boldsymbol{\ell}', \boldsymbol{\ell})^\top$; and (ii) $\sum_{i=1}^n \sum_{j=1}^n z_i^\top \mathbf{C}(\boldsymbol{\ell}_i, \boldsymbol{\ell}_j) z_j > 0$ for any integer n and any finite collection of points $\{\boldsymbol{\ell}_1, \boldsymbol{\ell}_2, \dots, \boldsymbol{\ell}_n\}$ and for all $z_i \in \mathfrak{R}^q \setminus \{\mathbf{0}\}$. See Genton and Kleiber (2015) for a review of cross-covariance functions for multivariate processes. The (partial) realization of the multivariate process over any finite set \mathcal{L} has a multivariate normal distribution $\mathbf{w}_{\mathcal{L}} \sim N(\mathbf{0}, \mathbf{C}_{\mathcal{L}})$ where $\mathbf{w}_{\mathcal{L}}$ is the $qn_{\mathcal{L}} \times 1$ column vector and $\mathbf{C}_{\mathcal{L}}$ is the $qn_{\mathcal{L}} \times qn_{\mathcal{L}}$ block matrix with the $q \times q$ matrix $\mathbf{C}(\boldsymbol{\ell}_i, \boldsymbol{\ell}_j)$ as its (i, j) block for $i, j = 1, \dots, n_{\mathcal{L}}$.

We construct the MGP from a base, or *parent*, multivariate GP for $\mathbf{w}(\boldsymbol{\ell})$ and then, using the graph \mathcal{G} defined in Section 2, represent the joint density at the reference set \mathcal{S} as

$$\tilde{p}(\mathbf{w}_{\mathcal{S}}) = \prod_{j=1}^M N(\mathbf{w}_j | \mathbf{H}_j \mathbf{w}_{[j]}, \mathbf{R}_j), \quad (8)$$

where $\mathbf{H}_1 = \mathbf{O}_{n_1 \times 1}$, $\mathbf{R}_1 = \mathbf{C}_{\mathcal{S}_j}$ and for $j > 1$, $\mathbf{H}_j = \mathbf{C}_{\mathcal{S}_j, \mathcal{S}_{[j]}} \mathbf{C}_{\mathcal{S}_{[j]}}^{-1}$ and $\mathbf{R}_j = \mathbf{C}_{\mathcal{S}_j} - \mathbf{C}_{\mathcal{S}_j, \mathcal{S}_{[j]}} \mathbf{C}_{\mathcal{S}_{[j]}}^{-1} \mathbf{C}_{\mathcal{S}_{[j]}, \mathcal{S}_j}$. The resulting joint density $\tilde{p}(\mathbf{w}_{\mathcal{S}})$ is multivariate normal with covariance $\tilde{\mathbf{C}}_{\mathcal{S}}$ and a precision matrix $\tilde{\mathbf{C}}_{\mathcal{S}}^{-1}$. The precision matrix for Gaussian

graphical models is easily derived using customary linear model representations for each conditional regression. Consider the DAG in (4). Each \mathbf{w}_i is $n_i q \times 1$ and let $J_i = |\text{Pa}[\mathbf{a}_i]|$ be the number of parents for \mathbf{a}_i in the graph \mathcal{G} . Furthermore, let \mathbf{C}_{ij} be the $n_i q \times n_j q$ covariance matrix between \mathbf{w}_i and \mathbf{w}_j , $\mathbf{C}_{i,[i]}$ be the $n_i q \times J_i q$ covariance matrix between \mathbf{w}_i and $\mathbf{w}_{[i]}$, and $\mathbf{C}_{[i],[i]}$ be the $J_i q \times J_i q$ covariance matrix between $\mathbf{w}_{[i]}$ and itself. Representing each conditional density in (4) as a linear regression on \mathbf{w}_i , we get

$$\begin{aligned} \mathbf{w}_1 &= \boldsymbol{\omega}_1 \sim N(\mathbf{0}, \mathbf{R}_1); & \mathbf{w}_i &= \sum_{\{j: \mathbf{a}_j \in \text{Pa}[\mathbf{a}_i]\}} \mathbf{H}_{ij} \mathbf{w}_j + \boldsymbol{\omega}_i, \\ & & i &= 2, 3, \dots, M, \end{aligned} \quad (9)$$

where each \mathbf{H}_{ij} is an $n_i q \times n_j q$ is a coefficient matrix representing the multivariate regression of \mathbf{w}_j given $\mathbf{w}_{[i]}$, $\boldsymbol{\omega}_i \stackrel{\text{ind}}{\sim} N(\mathbf{0}, \mathbf{R}_i)$ for $i = 1, 2, \dots, M$, and each \mathbf{R}_i is an $n_i q \times n_i q$ residual covariance matrix. We set $\mathbf{H}_{ii} = \mathbf{O}$ and $H_{ij} = \mathbf{O}$, where \mathbf{O} is the matrix of zeros, whenever $j \in \{j : \mathbf{a}_j \notin \text{Pa}[\mathbf{a}_i]\}$. For $j \in \{j : \mathbf{a}_j \in \text{Pa}[\mathbf{a}_i]\}$, let $\{j_1, j_2, \dots, j_{j_i}\}$ be the indices in $\text{Pa}[\mathbf{a}_i]$ and let $\mathbf{H}_{i,[i]} = [\mathbf{H}_{i,j_1}, \mathbf{H}_{i,j_2}, \dots, \mathbf{H}_{i,j_{j_i}}]$ be the $n_i q \times (\sum_{k=1}^{j_i} n_{j_k} q)$ block matrix formed by stacking \mathbf{H}_{i,j_k} side by side for each $\mathbf{a}_{j_k} \in \text{Pa}[\mathbf{a}_i]$. Since $\mathbb{E}[\mathbf{w}_i | \mathbf{w}_{[i]}] = \mathbf{H}_{i,[i]} \mathbf{w}_{[i]} = \mathbf{C}_{i,[i]} \mathbf{C}_{[i],[i]}^{-1} \mathbf{w}_{[i]}$, we obtain $\mathbf{H}_{i,[i]} = \mathbf{C}_{i,[i]} \mathbf{C}_{[i],[i]}^{-1}$ and each \mathbf{H}_{ij} can be obtained from the respective submatrix of $\mathbf{H}_{i,[i]}$. We also obtain $\mathbf{R}_i = \text{var}\{\mathbf{w}_i | \mathbf{w}_{[i]}\} = \mathbf{C}_{i,i} - \mathbf{C}_{i,[i]} \mathbf{C}_{[i],[i]}^{-1} \mathbf{C}_{[i],i}$. Therefore, all the \mathbf{H}_{ij} 's and \mathbf{R}_i 's can be computed from the base cross-covariance function.

The distribution of $\mathbf{w} = [\mathbf{w}_1^\top, \mathbf{w}_2^\top, \dots, \mathbf{w}_M^\top]^\top$ can be obtained by noting that $\mathbf{w} = \mathbf{H}\mathbf{w} + \boldsymbol{\omega}$, where $\mathbf{H} = \{\mathbf{H}_{ij}\}$ is the $(\sum_{i=1}^M n_i q) \times (\sum_{i=1}^M n_i q)$ block matrix with $\{\mathbf{H}_{ij}\}$ as (i, j) th block. Therefore, $\tilde{\mathbf{C}}_{\mathcal{S}} = \text{var}(\mathbf{w}) = (\mathbf{I} - \mathbf{H})^{-1} \mathbf{R} (\mathbf{I} - \mathbf{H})^{-\top}$, where \mathbf{R} is block-diagonal with \mathbf{R}_i as the (i, i) th block. Note that $\mathbf{I} - \mathbf{H}$ is block lower-triangular with 1's on the diagonal, hence nonsingular. Also, the precision matrix $\tilde{\mathbf{C}}_{\mathcal{S}}^{-1} = (\mathbf{I} - \mathbf{H})^\top \mathbf{R}^{-1} (\mathbf{I} - \mathbf{H})$ is sparse because of $\mathbf{H}_{ij} = \mathbf{O}$ whenever $\mathbf{a}_j \notin \text{Pa}[\mathbf{a}_i]$. Block-sparsity of $\tilde{\mathbf{C}}_{\mathcal{S}}^{-1}$ can be induced by building \mathcal{G} with few, carefully placed directed edges among nodes in \mathcal{A} ; Appendix B, available online, contains a more in-depth treatment. We extend (8) to the collection of nonreference locations $\mathcal{U} \subset \mathcal{D} \setminus \mathcal{S}$:

$$\begin{aligned} \tilde{p}(\mathbf{w}_{\mathcal{U}} | \mathbf{w}_{\mathcal{S}}) &= \prod_{j=1}^M N(\mathbf{w}_{\mathcal{U}_j} | \mathbf{H}_{\mathcal{U}_j} \mathbf{w}_{\mathcal{B}_j}, \mathbf{R}_{\mathcal{U}_j}) \\ &= N(\mathbf{w}_{\mathcal{U}} | \mathbf{H}_{\mathcal{U}} \mathbf{w}_{\mathcal{S}}, \mathbf{R}_{\mathcal{U}}), \end{aligned} \quad (10)$$

where $\mathbf{H}_{\mathcal{U}_j} = \mathbf{C}_{\mathcal{U}_j, \mathcal{S}_{[\mathcal{B}_j]}} \mathbf{C}_{\mathcal{S}_{[\mathcal{B}_j]}}^{-1}$ and $\mathbf{R}_{\mathcal{U}_j} = \mathbf{C}_{\mathcal{U}_j} - \mathbf{C}_{\mathcal{U}_j, \mathcal{S}_{[\mathcal{B}_j]}} \mathbf{C}_{\mathcal{S}_{[\mathcal{B}_j]}}^{-1} \mathbf{C}_{\mathcal{S}_{[\mathcal{B}_j]}, \mathcal{U}_j}$, analogously to (8), while $\mathbf{H}_{\mathcal{U}}$ and $\mathbf{R}_{\mathcal{U}}$ are analogous to $\mathbf{H}_{\mathcal{S}}$ and $\mathbf{R}_{\mathcal{S}}$. Clearly, given that all the \tilde{p} densities are Gaussian, all finite dimensional distributions will also be Gaussian. We have constructed a GP with the following cross-covariance function for any two locations $\ell_1, \ell_2 \in \mathcal{D}$

$$\begin{aligned} \text{cov}_{\tilde{p}}(\mathbf{w}(\ell_1), \mathbf{w}(\ell_2)) &= \begin{cases} \tilde{\mathbf{C}}_{s_i, s_j} & \text{if } \ell_1 = s_i, \ell_2 = s_j \text{ and } s_i, s_j \in \mathcal{S}, \\ \mathbf{H}_{\ell_1} \tilde{\mathbf{C}}_{\mathcal{S}_{\{\ell_1\}}, s_j} & \text{if } \ell_1 \in \mathcal{D} \setminus \mathcal{S}, \ell_2 = s_j \text{ and } s_j \in \mathcal{S}, \\ \delta_{(\ell_1 = \ell_2)} \mathbf{R}_{\ell_1} & \\ + \mathbf{H}_{\ell_1} \tilde{\mathbf{C}}_{\mathcal{S}_{\{\ell_1\}}, \mathcal{S}_{\{\ell_2\}}} \mathbf{H}_{\ell_2}^\top & \text{otherwise.} \end{cases} \end{aligned}$$

For a given base Gaussian covariance function \mathbf{C} , domain partitioning \mathcal{P} , mesh \mathcal{G} , and reference set \mathcal{S} , we denote the corresponding MGP as $\text{MGP}(\mathcal{G}, \mathcal{P}, \mathcal{S}, \mathbf{C})$.

3.1. Bayesian Hierarchical Model and Gibbs Sampler

Meshed GPs produce block-sparse precision matrices that are constructed cheaply from their block-sparse Cholesky factors by solving small linear systems. General purpose sparse-Cholesky algorithms (Davis 2006; Chen et al. 2008) can then be used to obtain collapsed samplers as in Finley et al. (2019). Unfortunately, these algorithms can only be used on Gaussian first stage models and are computationally impracticable for data in the millions. Hence, we develop a more general scalable Gibbs sampler for the recovery of spatial random effects in hierarchical MGP models that entirely circumvents large matrix computations.

Consider a multivariate spatiotemporally varying regression model at $\ell \in \mathcal{D} \subset \mathfrak{N}^{d+1}$,

$$\mathbf{y}(\ell) = \mathbf{X}(\ell)^\top \boldsymbol{\beta} + \mathbf{Z}(\ell)^\top \mathbf{w}(\ell) + \boldsymbol{\varepsilon}(\ell), \quad (11)$$

where $\mathbf{y}(\ell) \in \mathfrak{N}^l$ is the multivariate point-referenced outcome, $\mathbf{X}(\ell)^\top = \text{blockdiag}\{\mathbf{x}_i(\ell)^\top\}_{i=1}^l$ is a $l \times p = l \times \sum p_i$ matrix of spatially referenced predictors linked to constant coefficients $\boldsymbol{\beta}$, $\mathbf{w}(\ell)$ is the spatial process, $\mathbf{Z}(\ell)$ is a $l \times q$ design matrix, $\boldsymbol{\varepsilon}(\ell)$ is measurement error such that $\boldsymbol{\varepsilon}(\ell) \stackrel{\text{iid}}{\sim} N(\mathbf{0}, \mathbf{D})$ and $\mathbf{D} = \text{diag}(\tau_1^2, \dots, \tau_l^2)$. A simple univariate regression model with a spatially varying intercept can be obtained with $l = 1$, $\mathbf{Z}(\ell) = 1$. For observed locations $\mathcal{T} = \{\ell_1, \dots, \ell_n\}$, we write the above model compactly $\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{w} + \boldsymbol{\varepsilon}$, where $\mathbf{y} = (\mathbf{y}(\ell_1)^\top, \dots, \mathbf{y}(\ell_n)^\top)^\top$, \mathbf{w} and $\boldsymbol{\varepsilon}$ are similarly defined, $\mathbf{X} = [\mathbf{X}(\ell_1) : \dots : \mathbf{X}(\ell_n)]^\top$, $\mathbf{Z} = \text{blockdiag}(\{\mathbf{Z}(\ell_i)^\top\}_{i=1}^n)$, and $\mathbf{D}_n = \text{blockdiag}(\{\mathbf{D}_i\}_{i=1}^n)$.

For subsets $\{\ell_1, \dots, \ell_{n_A}\} = \mathcal{A} \subset \mathcal{T}$, let $\mathbf{y}(\mathcal{A}) = (\mathbf{y}(\ell_1)^\top, \dots, \mathbf{y}(\ell_{n_A})^\top)^\top$, with analogous definitions for $\mathbf{w}(\mathcal{A})$ and $\boldsymbol{\varepsilon}(\mathcal{A})$, $\mathbf{X}(\mathcal{A}) = [\mathbf{X}(\ell_1) : \dots : \mathbf{X}(\ell_{n_A})]^\top$, $\mathbf{Z}_{\mathcal{A}} = \text{blockdiag}(\{\mathbf{Z}(\ell_i)^\top\}_{i=1}^{n_A})$ and $\mathbf{D}_{\mathcal{A}} = \text{blockdiag}(\{\mathbf{D}_i\}_{i=1}^{n_A})$. After fixing a reference set \mathcal{S} , we obtain $\mathcal{S}^* = \mathcal{T} \cap \mathcal{S}$ and $\mathcal{U} = \mathcal{T} \setminus \mathcal{S}$. We partition the domain as above to obtain $\mathcal{S}_j, \mathcal{S}_j^*, \mathcal{U}_j$ for $j = 1, \dots, M$ and model $\mathbf{w}(\ell)$ using the MGP which yields $\mathbf{w} \sim N(\mathbf{0}, \tilde{\mathbf{C}}_{\mathcal{S}}^{-1})$. We complete the model specification by assigning $\boldsymbol{\beta} \sim N(\boldsymbol{\beta} | \boldsymbol{\mu}_{\boldsymbol{\beta}}, \boldsymbol{\Sigma}_{\boldsymbol{\beta}})$, $\tau_j^2 \sim \text{Inv.Gamma}(\tau_j^2 | a_{\tau_j}, b_{\tau_j})$, $\boldsymbol{\theta} \sim p(\boldsymbol{\theta})$.

The resulting full conditional distribution for $\boldsymbol{\beta}$ is $N(\boldsymbol{\Sigma}_{\boldsymbol{\beta}}^* \boldsymbol{\mu}_{\boldsymbol{\beta}}^*, \boldsymbol{\Sigma}_{\boldsymbol{\beta}}^*)$, where $\boldsymbol{\Sigma}_{\boldsymbol{\beta}}^* = (\boldsymbol{\Sigma}_{\boldsymbol{\beta}}^{-1} + \mathbf{X}^\top \mathbf{D}_n^{-1} \mathbf{X})^{-1}$, $\boldsymbol{\mu}_{\boldsymbol{\beta}}^* = \boldsymbol{\Sigma}_{\boldsymbol{\beta}}^{-1} \boldsymbol{\mu}_{\boldsymbol{\beta}} + \mathbf{X}^\top \mathbf{D}_n^{-1} (\mathbf{y} - \mathbf{Z}\mathbf{w})$. For τ_r^2 , $r = 1, \dots, q$, the full conditional is Inverse-Gamma with parameters $a_{\tau_r} + n/2$ and $b_{\tau_r} + \frac{1}{2} \mathbf{E}_r^\top \mathbf{E}_r$ where $\mathbf{E}_r = \mathbf{y}_r - \mathbf{X}_r \boldsymbol{\beta} - \mathbf{Z}_r \mathbf{w}$ and $\mathbf{y}_r, \mathbf{X}_r, \mathbf{Z}_r$ are the subsets of $\mathbf{y}, \mathbf{X}, \mathbf{Z}$ corresponding to outcome r (out of q).

The Gibbs update of the $\mathbf{w}_{\mathcal{U}}$ components can proceed simultaneously as all blocks in \mathcal{U} have no children and their parents are in \mathcal{S} . The full conditional for $\mathbf{w}_{\mathcal{U}_j}$ for $j = 1, \dots, M$ is thus $N(\boldsymbol{\Sigma}_{\mathcal{U}_j}^* \boldsymbol{\mu}_{\mathcal{U}_j}^*, \boldsymbol{\Sigma}_{\mathcal{U}_j}^*)$ where $\boldsymbol{\Sigma}_{\mathcal{U}_j}^* = (\mathbf{Z}(\mathcal{U}_j) \mathbf{D}^{-1} \mathbf{Z}(\mathcal{U}_j)^\top + \mathbf{R}_{\mathcal{U}_j}^{-1})^{-1}$ and $\boldsymbol{\mu}_{\mathcal{U}_j}^* = \mathbf{Z}(\mathcal{U}_j) \mathbf{D}^{-1} (\mathbf{y}(\mathcal{U}_j) - \mathbf{X}(\mathcal{U}_j)^\top \boldsymbol{\beta}) + \mathbf{R}_{\mathcal{U}_j}^{-1} \mathbf{H}_{\mathcal{U}_j} \mathbf{w}_{\mathcal{B}_j}$, where $\mathbf{w}_{\mathcal{B}_j}$ is the spatial process at locations corresponding to the parents of $\mathcal{B}_j \in \mathcal{B} \subset \mathcal{V}$.

We update $\mathbf{w}_{S_j} = \mathbf{w}_j$ for $j = 1, \dots, M$ via its full conditional $N(\boldsymbol{\Sigma}_j^* \boldsymbol{\mu}_j^*, \boldsymbol{\Sigma}_j^*)$. Let $\mathbf{1}_j = (\text{In}(s_1 \in S_j^*), \dots, \text{In}(s_{n_j} \in S_j^*))^\top$ be the vector of indicators that identify locations with nonmissing outputs, and let $\mathbf{a}_j \in \mathcal{V}$ be the node in \mathcal{G} corresponding to S_j . Then,

$$\begin{aligned} \boldsymbol{\Sigma}_j^{*-1} &= \mathbf{Z}_j^\top \tilde{\mathbf{D}}_{n_j}^{-1} \mathbf{Z}_j + \mathbf{R}_j^{-1} + \sum_{i=1}^{|\text{Ch}[\mathbf{a}_j]|} \mathbf{H}_i^{[j]\top} \mathbf{R}_i^{[j]-1} \mathbf{H}_i^{[j]}, \\ \boldsymbol{\mu}_j^* &= \mathbf{R}_j^{-1} \mathbf{H}_j \mathbf{w}_{[j]} + \mathbf{Z}_j^\top \tilde{\mathbf{D}}_{n_j}^{-1} \tilde{\mathbf{y}}_j + \sum_{i=1}^{|\text{Ch}[\mathbf{a}_j]|} \mathbf{H}_i^{[j]\top} \mathbf{R}_i^{[j]-1} \mathbf{w}_i^{[j]}, \end{aligned} \quad (12)$$

where $\tilde{\mathbf{D}}_{n_j}^{-1} = \mathbf{I}_j \odot \mathbf{D}_{n_j}^{-1}$ with $\mathbf{I}_j = \mathbf{1}_j \mathbf{1}_j^\top$, and $\tilde{\mathbf{y}}_j = \mathbf{1}_j \odot (\mathbf{y}_j - \mathbf{X}_j \boldsymbol{\beta})$ and \odot denotes the Hadamard or Schur (element-by-element) product. Finally, $\boldsymbol{\theta}$ is updated via a Metropolis step with target density $p(\boldsymbol{\theta})N(\mathbf{w}_S \mid \mathbf{0}, \tilde{\mathbf{C}}_S)N(\mathbf{w}_U \mid \mathbf{H}_U \mathbf{w}_S, \mathbf{R}_U)$ using (8) and (10). The Gibbs sampling algorithm will iterate across the above steps and, upon convergence, will produce samples from $p(\boldsymbol{\beta}, \{\tau_j^2\}_{j=1}^q, \mathbf{w} \mid \mathbf{y})$.

We obtain posterior predictive inference at arbitrary $\boldsymbol{\ell} \in \mathcal{D}$ by evaluating $p(\mathbf{y}(\boldsymbol{\ell}) \mid \mathbf{y})$. If $\boldsymbol{\ell} \in \mathcal{S} \cup \mathcal{U}$, then we draw one sample of $\mathbf{y}(\boldsymbol{\ell}) \sim N(\mathbf{X}(\boldsymbol{\ell})^\top \boldsymbol{\beta} + \mathbf{Z}(\boldsymbol{\ell})^\top \mathbf{w}(\boldsymbol{\ell}), \mathbf{D})$ for each draw of the parameters from $p(\boldsymbol{\beta}, \{\tau_j^2\}_{j=1}^q, \mathbf{w} \mid \mathbf{y})$. Otherwise, considering that $\boldsymbol{\ell} \in \mathcal{D}_j$ for some j and thus $\boldsymbol{\eta}(\boldsymbol{\ell}) = \mathbf{b}_j$, with parent nodes $\text{Pa}[\mathbf{b}_j]$ and children $\text{Ch}[\mathbf{b}_j] = \emptyset$, we sample $\mathbf{w}(\boldsymbol{\ell})$ from the full conditional $N(\boldsymbol{\Sigma}_\ell^* \boldsymbol{\mu}_\ell^*, \boldsymbol{\Sigma}_\ell^*)$, where $\boldsymbol{\Sigma}_\ell^* = (\mathbf{Z}(\boldsymbol{\ell}) \mathbf{D}^{-1} \mathbf{Z}(\boldsymbol{\ell})^\top + \mathbf{R}_\ell^{-1})^{-1}$ and $\boldsymbol{\mu}_\ell^* = \mathbf{Z}(\boldsymbol{\ell}) \mathbf{D}^{-1} (\mathbf{y}(\boldsymbol{\ell}) - \mathbf{X}(\boldsymbol{\ell})^\top \boldsymbol{\beta}) + \mathbf{R}_\ell^{-1} \mathbf{H}_\ell \mathbf{w}_{[\mathbf{b}_j]}$, then draw $\mathbf{y}(\boldsymbol{\ell}) \sim N(\mathbf{X}(\boldsymbol{\ell})^\top \boldsymbol{\beta} + \mathbf{Z}(\boldsymbol{\ell})^\top \mathbf{w}(\boldsymbol{\ell}), \mathbf{D})$.

3.2. Nonseparable Multivariate Spatiotemporal Covariances

We provide an account of the computational cost of general MGPs as a starting point to motivate the introduction of more efficient tessellated MGPs, and specifically Q-MGPs, in Section 4. We consider (11) and take $l = 1$ to simplify our exposition. In the resulting model, $\boldsymbol{\beta}$ is the regression coefficient on the p point-referenced regressors with a static effect on the outcome, whereas the q -variate spatiotemporal process $\mathbf{w}(\cdot)$ captures the dynamic effect of the \mathbf{Z} regressors. Typically in geostatistical modeling p and q are small, hence sampling $\boldsymbol{\beta}$ and τ^2 carries a negligible computational cost. The cost of each Gibbs iteration is dominated by updates of $\boldsymbol{\theta}$ and \mathbf{w} . Let us assume, solely for expository purposes, that each of the M blocks comprise the same number of locations, that is, $|\mathcal{S}_j| = |\mathcal{U}_j| = m$, for all $j = 1, \dots, M$. Thus, $m = \frac{n}{2M}$ and the graph nodes have J or fewer parents and L or fewer children.

The evaluation of $N(\mathbf{w}_S \mid \mathbf{0}, \tilde{\mathbf{C}}_S) = \prod_{j=1}^M N(\mathbf{w}_j \mid \mathbf{H}_j \mathbf{w}_{[j]}, \mathbf{R}_j)$ and $N(\mathbf{w}_U \mid \mathbf{H}_U \mathbf{w}_S, \mathbf{R}_U) = \prod_{j=1}^M N(\mathbf{w}_{U_j} \mid \mathbf{H}_{U_j} \mathbf{w}_{[\mathbf{b}_j]}, \mathbf{R}_{U_j})$ dominates the computation. Each term in the product entails \mathbf{R}_j^{-1} and $\mathbf{R}_{U_j}^{-1}$, both of size $qm \times qm$, and their determinants. These require $\mathbf{C}_{[j]}^{-1}$ of size $Jqm \times Jqm$ or less, resulting in $O(2M(q^3 m^3 + J^3 q^3 m^3)) = O(2Mq^3 m^3 (J^3 + 1)) \approx O(2Mq^3 m^3 J^3) = O(\frac{n^3 q^3 J^3}{M^2})$ flops via Cholesky decomposition. Reasonably, J and m are fixed so M may grow linearly with sample size and the cost is $O(nq^3 J^3)$

considering $M \propto n$. The total computing time is $\sim O(\frac{nq^3 J^3}{K})$ with K processors for computing the $2M$ densities. Sampling \mathbf{w}_S and \mathbf{w}_U from their full conditional distributions requires $O(2Mq^3 m^3 + MLq^2 m^2 + Mq^2 m^2)$ flops, assuming \mathbf{R}_j^{-1} and $\mathbf{R}_{U_j}^{-1}$ are stored in the previous step. The first term in the complexity order is due to the Cholesky decomposition of covariance matrices, the second is due to sampling the reference nodes, and the third comes from sampling other nodes. Without further assumptions, parallelization reduces complexity to $O(\frac{2Mq^3 m^3}{K} + \frac{Mq^2 m^2}{K} + MLq^2 m^2)$, since the covariances can be computed beforehand and the M components of \mathbf{w}_U are independent given \mathbf{w}_S . With fixed block size m , the overall complexity for a Gibbs iteration is $O(\frac{2}{K} Mq^3 m^3 (J^3 + 1) + \frac{1}{K} 2Mq^3 m^3 + \frac{1}{K} Mq^2 m^2 + MLq^2 m^2) \approx O(\frac{1}{K} J^3 q^3 n + q^2 n) \approx O(n)$, linear in the sample size and cubic in J , highlighting the computational speedup of sparse graphs (J small), the negative impact of large q , and the serial sampling of \mathbf{w}_S .

In terms of storage, \mathbf{H}_j and \mathbf{R}_j correspond to a storage requirement of $O(4Mq^2 m^2) = O(q^2 n)$. The matrix \mathbf{Z} of size $qn \times qn$ can be represented as a list of $2M$ block-diagonal (hence sparse) \mathbf{Z}_j matrices. Furthermore, computing $\mathbf{Z} \mathbf{w}$ (dimension $n \times 1$) can be vectorized as the row-wise sum of $\mathbf{Z}^* \odot \mathbf{w}^*$ where \mathbf{Z}^* and \mathbf{w}^* are $n \times q$ matrices with j th column representing the j th space-time varying predictor. The cost of storing \mathbf{Z} is thus $O(2qn)$.

Complexity is further reduced by considering a graph with small J or a finer partition resulting in large M and small m , whereas the overall time can be reduced by distributing computations on K processors. Possible choices for \mathcal{G} include nearest-neighbor graphs and multiresolution trees. In settings with large q , adjusting J and M may be insufficient to reduce the computational burden. Covariance functions that are separable in the variables (but perhaps nonseparable in space and time) bring the cost of Cholesky factorizations of $Jqm \times Jqm$ matrices from $O(J^3 q^3 m^3)$ to $O(J^3 m^3 + q^3)$ because $\mathbf{C}^{-1} = (\mathbf{C}_{h,u} \otimes \mathbf{C}_v)^{-1} = \mathbf{C}_{h,u}^{-1} \otimes \mathbf{C}_v^{-1}$, where $\mathbf{C}_{h,u}$ is the $Jm \times Jm$ space-time component of the cross-covariance, and \mathbf{C}_v the $q \times q$ variable component. Savings accrue when evaluating the likelihood and in sampling from the full-conditionals at the cost of realism in describing the spatial process.

The next section develops a novel MGP design based on domain tessellations or tiling—that is, partitions of the domain forming repeated patterns—to which we associate similarly patterned meshes. If observations are also located in patterns, the bulk of the largest linear solvers will be redundant, resulting in a significant reduction in computational time. In either scenario, sampling \mathbf{w}_S will also proceed in parallel with improved mixing.

4. MGPs Based on Domain Tessellation or Tiling

We construct MGPs based on a tessellation or tiling of the domain. For spatial domains ($d = 2$, Figure 2), regular tiling results in triangular, quadratic, or hexagonal tessellations; mixed designs are also possible. These partition schemes can be linked to a DAG \mathcal{G} by drawing directed outgoing edges starting from an originating node/tile. The same fixed pattern can be repeated over a surface of any size. In dimensions $d > 2$, which may




Figure 2. Regular tessellation base units and corresponding MGP graphs for spatial domains.

include time, space-filling tessellations or honeycombs can be constructed analogously, along with their corresponding meshes. Constructions of MGPs based on these ideas simply requires partitioning the locations \mathcal{S} into subsets based on the chosen tessellation.

This subclass of MGP models corresponds by design to graphs with known *coloring*, with each color linked to a subgraph conditionally independent of all nodes of other colors, regardless of the dimension of the domain. This feature enables large-scale parallel sampling of $w_{\mathcal{S}}$ and improves mixing without the need to implement heuristic graph-coloring algorithms. Furthermore, regions in a tessellated domain are typically translations and/or rotations of a single geometric shape. Carefully choosing \mathcal{S} , it will be possible to avoid computing the bulk of linear solvers, resulting in substantial computational gains. Subsequently, we focus on axis-parallel partitioning (quadratic or cubic tessellation) and cubic meshes, but analogous constructions and the same properties hold with other tessellation schemes.

A cubic MGP (Q-MGP) is constructed by partitioning each coordinate axis into intervals. In $d + 1$ dimensions, splitting each axis into L intervals results in L^{d+1} regions. Consider a spatiotemporal domain $\mathcal{D} = \times_{r=1}^{d+1} \mathcal{D}^{(r)}$, where $\mathcal{D}^{(d+1)}$ is the time dimension. We partition each coordinate axis into L_r disjoint sets: $\mathcal{D}^{(r)} = \mathcal{I}_{r,1} \cup \dots \cup \mathcal{I}_{r,L_r}$, where $\mathcal{I}_{r,j} \cap \mathcal{I}_{r,k} = \emptyset$ if $j \neq k$ and $\mathcal{I}_{r,s}$ denotes the s th interval in the r th coordinate axis. Solely for exposition, and without loss of generality, assume that $\mathcal{D}^{(r)} = \mathcal{I} = [0, 1]$ and $L_r = L$ for $r = 1, \dots, d + 1$. Any location $\ell = (\ell_1, \dots, \ell_{d+1}) \in \mathcal{D}$ will be such that $\ell \in \mathcal{I}_{1,i_1} \times \dots \times \mathcal{I}_{d+1,i_{d+1}} = \mathcal{D}_j$ for some i_1, \dots, i_{d+1} and with $j = 1, \dots, M$, where $M = L^{d+1}$. We refer to this axis-parallel partition scheme as a cubic tessellation and denote it by $\mathbf{T} = \{\mathcal{I}_{r,s}\}_{r=1, \dots, d+1}^{s=1, \dots, L}$. We use \mathbf{T} to partition the reference set \mathcal{S} as $\mathcal{S}_j = \mathcal{D}_j \cap \mathcal{S}$ for $j = 1, \dots, L^{d+1}$.

Next, we define $\eta(\ell) = (\eta_1(\ell), \dots, \eta_L(\ell)) \in \{1, \dots, L\}^{d+1}$, where $\eta_j = \eta_j(\ell) = r$ if $\ell_j \in \mathcal{I}_{j,r}$. Then, let $\mathcal{Q} = (\mathbf{V}, \mathbf{E})$ be a DAG with $\mathbf{V} = \mathbf{A} \cup \mathbf{B}$ and reference nodes $\mathbf{A} = \{\mathbf{a}_1, \dots, \mathbf{a}_{L^{d+1}}\}$. Therefore, for any $j = 1, \dots, L^{d+1}$ if $\mathbf{s} \in \mathcal{S}_j$ then $\eta(\mathbf{s}) = \mathbf{a}_j \in \mathbf{A} \subset \mathbf{V}$. We write each node $\mathbf{v} \in \mathbf{V}$ as $\mathbf{v} = (v_{\eta_1}, \dots, v_{\eta_L}) \in \{1, \dots, L\}^{d+1}$. The directed edges are constructed using a “line-of-sight” strategy. Suppose $\text{Pa}[\mathbf{v}] = \{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(d+1)}\}$. The h th parent of \mathbf{v} is defined as $\mathbf{x}^{(h)} = (a_{\eta_1}, \dots, a_{\eta_h} - k, \dots, a_{\eta_L}) \cap \{1, \dots, d + 1\}^{d+1}$, where $k \geq 1$ is the smallest integer such that $\mathbf{x}^{(h)} \in \mathbf{A}$. Consequently $\mathbf{x}^{(h)} = \emptyset$ if $a_{\eta_h} = 1$. Thus, the parents

of node $\mathbf{v} = \eta(\ell)$ are the ones that precede it along each of the $d + 1$ coordinates. If $\ell \in \mathcal{D}_j \setminus \mathcal{S}_j$, then $\eta(\ell) = \mathbf{b}_j \in \mathbf{B}$ and $\text{Pa}[\mathbf{b}_j] = \{\mathbf{a}_j\} \cup \text{Pa}[\mathbf{a}_j]$ where $\mathbf{a}_j \in \mathbf{A}$ is a reference node. To avoid $\text{Pa}[\mathbf{b}_j] = \emptyset$ we set $\text{Pa}[\mathbf{b}_j] = \{\mathbf{x}_1^{(1)}, \mathbf{x}_2^{(1)}, \dots, \mathbf{x}_1^{(d+1)}, \mathbf{x}_2^{(d+1)}\}$.

The two parents along the h th dimension are $\mathbf{x}_1^{(h)} = a_{\eta_h} + k_1$, $\mathbf{x}_2^{(h)} = a_{\eta_h} - k_2$ where k_i is the smallest positive integer such that $\mathbf{x}_i^{(h)} \in \mathbf{A}$, $i = 1, 2$. In this setting $J = 2(d + 1)$. The construction is finalized by fixing the cross-covariance function $\mathbf{C}(\ell, \ell')$; Figure 3 shows that the same basic structure can be immediately extended to higher dimensions, including time.

4.1. Caching Redundant Expensive Matrix Operations

The key computational bottleneck for the Gibbs sampler in Section 3.2 is calculating, for $j = 1, \dots, 2M$, of (i) $\mathbf{C}_{[j]}^{-1}$ ($2Mj^3q^3m^3$ flops) and (ii) $\mathbf{R}_j^{-1}, \boldsymbol{\Sigma}_j^{*-1}$ ($4Mq^3m^3$ flops). The former is costlier than the latter by a factor of $J^3/2$. Q-MGPs are designed to greatly reduce this cost. We start with an axis-parallel tessellation of the domain in equally sized regions $\mathcal{D}_1, \dots, \mathcal{D}_M$, storing observed locations in \mathcal{U} to create $\mathcal{U}_1, \dots, \mathcal{U}_M$, which we assume, for simplicity, to be no larger than m in size. Taking a stationary base-covariance function \mathbf{C} , implies that $\mathbf{C}(\mathcal{L}_1, \mathcal{L}_2) = \mathbf{C}(\mathcal{L}_1 + \mathbf{h}, \mathcal{L}_2 + \mathbf{h})$, where $\mathbf{h} \in \mathfrak{R}^{d+1}$ is used to shift all locations in the sets. Recall that the reference set \mathcal{S} of MGPs can include unobserved locations. Hence, we can build \mathcal{S} on a lattice of regularly spaced locations. Since domain partitions have the same size, we have $\mathcal{S}_j = \mathcal{S}^* + \mathbf{h}_j$ for $j = 1, \dots, M$, where \mathcal{S}^* is a single “prototype set” using which one can locate all other reference subsets. Also, since $\text{Pa}[\mathbf{a}_j] \subset \text{Pa}[\mathbf{b}_j]$, there will be $4(d + 1)$ prototype sets for parents, that is, $\mathcal{S}_{\text{Pa}[\mathbf{v}_j]} = \mathcal{S}_r^* + \mathbf{h}_j$ for some $r \in \{1, \dots, 4(d + 1)\}$ and $j = 1, \dots, 2M$. Then, we can build maps $\xi_{\mathcal{S}} : \{1, \dots, M\} \rightarrow \{1, \dots, 4(d + 1)\}$ and $\xi_{\mathcal{U}} : \{1, \dots, M\} \rightarrow \{1, \dots, 4(d + 1)\}$ linking each of \mathcal{S}_j and \mathcal{U}_j to a parent prototype. This ensures that $\mathbf{C}_{[j]}^{-1} = \mathbf{C}_{\mathcal{S}_r^*}^{-1}$ for each $j = 1, \dots, 2M$. One only needs the maps $\xi_{\mathcal{S}}$ and $\xi_{\mathcal{U}}$, cache the r unique inverses, and reuse them. The same method applies to cache $\mathbf{R}_{\mathcal{S}_j}^{-1} = \mathbf{R}_{\mathcal{S}_r^*}^{-1}$ on reference sets, but not on other locations since no redundancy arises in $\mathbf{C}_{\mathcal{U}_j}$ for $j = 1, \dots, M$. See Figure 4 for an illustration. Compared to general MGPs (see Table 1), the number of large linear system solvers is now constant with sample size and $(d + 1) \ll M$ significantly reduces computational cost.

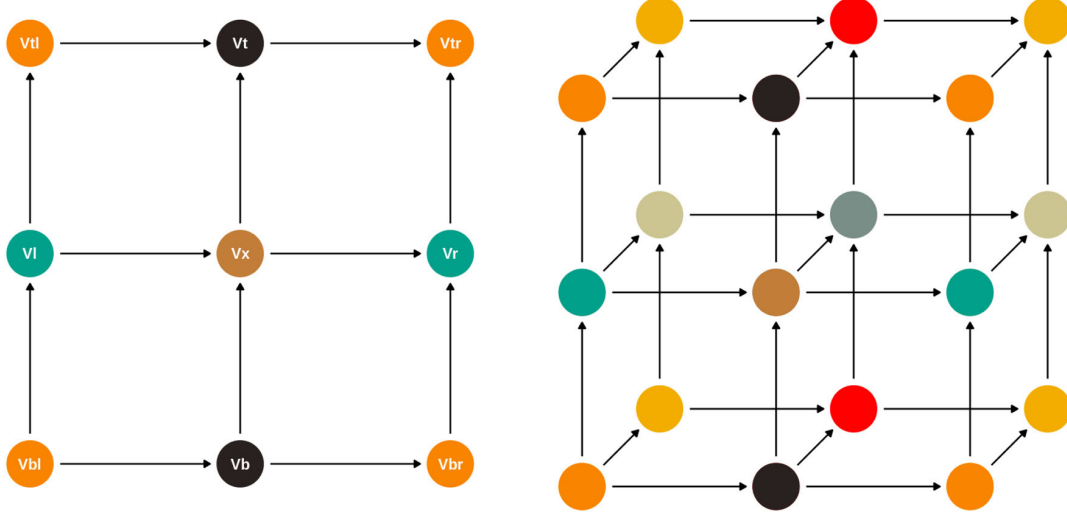


Figure 3. Q-MGP meshes used for spatial data on $d = 2$ (left) can be extended for use on spatiotemporal data $d = 3$ (right). Node colors correspond to Gibbs sampler blocks.




Figure 4. Visualizing redundancies: a spatial domain is partitioned in $M = 25$ regions and linked to a quadratic mesh. The reference set \mathcal{S} is fixed on a regular grid, with $m = 9$. Parent locations of the orange (resp. purple) are in green (resp. blue). Using a stationary covariance, $C_{\text{blue,blue}} = C_{\text{green,green}}$. Therefore, only one inversion is necessary; this can be replicated at no cost across 9 of the 16 regions.

Furthermore, Q-MGPs automatically adjust to settings where observed locations \mathcal{T} are on partly regular lattices, that is, they are located at patterns repeating in space or time which emerge after initial inspections of the data. Appendix G, available online, outlines a simple algorithm to identify such patterns and create maps $\xi_{\mathcal{S}}$ and $\xi_{\mathcal{U}}$. In such cases, we fix $\mathcal{S} \supseteq \mathcal{T}$ and $\mathcal{U} = \emptyset$. In addition to the above mentioned savings, we now do not have to compute $R_{\mathcal{U}_j}^{-1}$ and $\Sigma_{\mathcal{U}_j}^{*-1}$. If \mathcal{T} is not a regular lattice over the whole domain, $4(d+1)$ is a lower bound and in general there are $M^* \ll M$ inverses to compute. If \mathcal{T} is a fully observed regular lattice and if $\mathbf{Z}(\ell) = I$ (a varying intercept model), then we save in computing the full conditional covariances as well, since all $D_j = I$. See Appendix C, available online, for details on choosing \mathcal{S} and \mathcal{U} .

4.2. Improved Mixing via Parallel Sampling

With caching, a much larger proportion of time is spent on sampling; parallelization may in general be achieved via appropriate node coloring algorithms (see, e.g., Molloy and Reed 2002; Gonzalez et al. 2011; Lewis 2016), but this step is unnecessary in Q-MGPs as the colors in \mathcal{Q} are set in advance independently of the data and result in efficient parallel sampling of the latent effects. Reference nodes \mathcal{A} of \mathcal{Q} are colored to achieve independence conditional on realizations of nodes of all other colors. For example, we partition spatial domains ($d = 2$) into $M_1 \times M_2$ regions and link each region to a reference node in a quadratic

Table 1. Summary of computational cost of general MGPs and Q-MGPs.

	$C_{\mathcal{U},\mathcal{U}}^{-1}$	$R_{\mathcal{S}_j}^{-1}$	$R_{\mathcal{U}_j}^{-1}$	$\Sigma_{\mathcal{S}_j}^{*-1}$	$\Sigma_{\mathcal{U}_j}^{*-1}$	Sampling $w_{\mathcal{S}}, w_{\mathcal{U}}$
MGPs (all cases)	$2M^3q^3m^3$	Mq^3m^3	Mq^3m^3	Mq^3m^3	Mq^3m^3	$MLq^2m^2 + Mq^2m^2$
Q-MGPs						
Irregular locations	$4(d+1)^3q^3m^3$	$4(d+1)q^3m^3$	Mq^3m^3	Mq^3m^3	Mq^3m^3	$MLq^2m^2 + Mq^2m^2$
Pattern lattice w/missing	$2M^*j^3q^3m^3$	$2M^*q^3m^3$		Mq^3m^3		MLq^2m^2
Lattice w/ missing	$4(d+1)^3q^3m^3$	$4(d+1)q^3m^3$		Mq^3m^3		MLq^2m^2
Full lattice and $\mathbf{Z}(\ell) = I_q$	$4(d+1)^3q^3m^3$	$4(d+1)q^3m^3$		$2^{(d+2)}(d+1)q^3m^3$		MLq^2m^2

NOTE: Rows are sorted from most expensive (top) to least expensive (bottom).

mesh. A “central” reference node \mathbf{v}_+ will have two parents and two children, that is, $\text{Pa}[\mathbf{v}_+] = \{\mathbf{v}_l, \mathbf{v}_b\}$ and $\text{Ch}[\mathbf{v}_+] = \{\mathbf{v}_r, \mathbf{v}_t\}$, with l, b, r, t , respectively, denoting *left, bottom, right, top*—refer to Figure 3 (left). We have $\text{Pa}[\mathbf{v}_l] = \{\mathbf{v}_+, \mathbf{v}_{tl}\}$ and $\text{Pa}[\mathbf{v}_r] = \{\mathbf{v}_+, \mathbf{v}_{br}\}$. The Markov blanket of \mathbf{v}_+ , denoted as $\text{mb}(\mathbf{v}_+)$, is the set of neighbors of \mathbf{v}_+ in the undirected “moral” graph $\mathcal{Q}^{\mathcal{M}}$, hence $\text{mb}(\mathbf{v}_+) = \text{Pa}[\mathbf{v}_+] \cup \text{Ch}[\mathbf{v}_+] \cup \{\mathbf{v}_{tl}, \mathbf{v}_{br}\}$. The corresponding spatial process is such that $p(\mathbf{w}_+ | \mathbf{w} \setminus \mathbf{w}_+) = p(\mathbf{w}_+ | \mathbf{w}_{\text{mb}(\mathbf{v}_+)})$. Denoting $\mathbf{v}_{bl} = \text{Pa}[\mathbf{v}_l] \cap \text{Pa}[\mathbf{v}_b]$ and $\mathbf{v}_{tr} = \text{Ch}[\mathbf{v}_r] \cap \text{Ch}[\mathbf{v}_t]$, we note that $\{\mathbf{v}_{bl}, \mathbf{v}_{tr}\} \cap \text{mb}(\mathbf{v}_+) = \emptyset$. We partition reference nodes \mathbf{A} into four groups $\{\mathbf{A}^{(1)}, \mathbf{A}^{(2)}, \mathbf{A}^{(3)}, \mathbf{A}^{(4)}\}$, such that $\{\mathbf{v}_+\} \subset \mathbf{A}^{(1)}$, $\{\mathbf{v}_b, \mathbf{v}_t\} \subset \mathbf{A}^{(2)}$, $\{\mathbf{v}_l, \mathbf{v}_r\} \subset \mathbf{A}^{(3)}$, and $\{\mathbf{v}_{tl}, \mathbf{v}_{tr}, \mathbf{v}_{bl}, \mathbf{v}_{br}\} \subset \mathbf{A}^{(4)}$. This 3×3 pattern is repeated over the whole graph. Then, if $\mathbf{v} \in \mathbf{A}^{(j)}$, $\text{mb}(\mathbf{v}) \cap \mathbf{A}^{(j)} = \emptyset$. Denoting by \mathcal{D} the other variables in the Gibbs sampler, we get:

$$p(\mathbf{w}_j | \mathbf{w}_{-j}, \mathcal{D}) = p(\mathbf{w}_j | \mathbf{w}_{\text{mb}(\mathbf{v}_j)}, \mathcal{D}) = \prod_{\mathbf{v}_i \in \mathbf{A}^{(j)}} p(\mathbf{w}_i | \mathbf{w}_{\mathbf{A}^{(-j)}}, \mathcal{D}).$$

Since parallelization is possible within each of the groups, only be four serial steps are needed; time savings are due to $M/4$ typically being orders of magnitude larger than the number of available processors. Extensions to other tessellation schemes and higher dimensional domains and the associated graphs follow analogously.

5. Data Analysis

Satellite imaging and remote sensing data are nowadays frequently collected in large quantities and processed to be used in geology, ecology, forestry, and other fields, but clouds and atmospheric conditions obstruct aerial views and corrupt the data creating gaps. Recovery of the underlying signal and quantification of the associated uncertainty are thus the major goals to enable practitioners in the natural sciences to fully exploit these data sources. Several scalable geostatistical models based on GPs have been implemented on tens or hundreds of thousands of data points, with few exceptions. In considering larger data sizes, one must either have a large time budget—usually several days—or reduce model flexibility and richness. Scalability concerns become the single most important issue in multivariate spatiotemporal settings. In fact, repeated collection of aerial images and multiple spatially referenced predictors modeled to have a variable effect on the outcome have a multiplicative effect on data size. With no separability assumptions, the dimension of the latent spatial random effects that one may wish to recover will be huge even when individual images would be manageable when considered individually.

The lack of software to implement scalable models for spatiotemporal data makes it difficult to compare our proposed approach with others in these settings. On the other hand, a recent article (Heaton et al. 2019) pins many state-of-the-art models against each other in a spatial ($d = 2$) prediction contest. On the same data, we show in Appendix E, available online, that Q-MGPs can outperform all competitors in terms of predictive performance and coverage while using a similar computational budget.

5.1. Nonseparable Multivariate Spatiotemporal Base Covariance

In our analyses, we choose a class of multivariate space-time cross-covariances that models the covariance between variables i and j at the $(\mathbf{h}, u) \in \mathfrak{R}^{d+1}$ space-time lags as

$$\mathbf{C}_{ij}(\mathbf{h}, u) = \frac{\sigma^2}{\left(\psi_1\left(\frac{|u|^2}{\psi_2(\delta_{ij}^2)}\right)\right)^{d/2} \left(\psi_2(\delta_{ij}^2)\right)^{1/2}} \phi_1 \times \left(\frac{\|\mathbf{h}\|^2}{\psi_1\left(\frac{|u|^2}{\psi_2(\delta_{ij}^2)}\right)}\right), \quad (13)$$

where $\delta_{ij} > 0$ (and with $\delta_{ij} = \delta_{ji}$) is the latent dissimilarity between variables i and j . In the resulting cross-covariance function $\mathbf{C}(\mathbf{h}, u, \mathbf{v})$ in \mathfrak{R}^{d+1+k} , each component of the q -variate spatial process is represented by a point in a k -dimensional latent space, $k \leq q$. Refer to Apanasovich and Genton (2010) for a more in-depth discussion. We set $\phi_1(x) = \exp(-cx)$ and $\psi_j(x) = (a_j x^{\alpha_j} + 1)^{\beta_j}$, $j = 1, 2$; see Gneiting (2002) for alternatives. We also fix $\alpha_1 = \alpha_2 = \frac{1}{2}$, and seek to estimate $\boldsymbol{\theta} = (\sigma^2, c, a_1, \beta_1, a_2, \beta_2, \{\delta_{ij}\}_{i < j = 1, \dots, q})$ a posteriori. The usual exponential covariance arises in univariate spatial settings.

5.2. Synthetic Data

We mimick real world satellite imaging data analyzed later in Section 5.3 at a much smaller scale by generating 81 datasets from the model $\mathbf{y}(\boldsymbol{\ell}) = \mathbf{Z}(\boldsymbol{\ell})^\top \mathbf{w}(\boldsymbol{\ell}) + \boldsymbol{\varepsilon}(\boldsymbol{\ell})$, where $\boldsymbol{\varepsilon}(\boldsymbol{\ell}) \sim N(0, \tau^2)$ with $\boldsymbol{\ell} \in \mathcal{T}$ and \mathcal{T} is a regular grid of size $40 \times 40 \times 10$, resulting in $n_{\text{all}} = 16,000$ total locations. We take $\mathbf{w} \sim \text{GP}(\mathbf{0}, \mathbf{C})$ where \mathbf{C} is as in (13), $\psi_2 \equiv 1$ and $\sigma^2 = 1$. We generate one dataset for each combination of $\tau^2 \in \{1/1000, 1/20, 1/10\}$, temporal range $\alpha \in \{5, 50, 500\}$, space-time separability $\beta \in \{1/20, 1/2, 1 - \frac{1}{20}\}$, and spatial range $c \in \{1, 5, 25\}$.

We compare Q-MGPs with the similarly targeted Gapfill method of Gerber et al. (2018) as implemented in the R package `gapfill`. We create “synthetic clouds” of radius $\sqrt{0.1}$ and with center $(c_{1,t}, c_{2,t}) \in [0, 1/20]^2$ where $c_{1,t}, c_{2,t} \stackrel{\text{iid}}{\sim} U[0, 1]$ to cover the outcomes at six randomly selected times for each of the 81 datasets. Outcomes at two of the remaining four time periods were then randomly selected to be completely unobserved at all but 10 locations to avoid errors from `gapfill`. Refer to Figure 5 for an illustration.

A Q-MGP model with $M = 500$ was fit by partitioning each spatial axis into 10 intervals and the time axis into 5 intervals. The priors were $\tau^2 \sim \text{Inv.G.}(2, 1)$, $\sigma^2 \sim \text{Inv.G.}(2, 1)$, $\beta \sim U(0, 1)$, $\alpha \sim U(0, 10^4)$, $c \sim U(0, 10^4)$; 7000 iterations of Gibbs sampling were run, of which 5000 used for burn-in and thinning the remaining 2000 to obtain a posterior sample of size 1000. For each of the 81 datasets we calculate the mean absolute prediction error (MAE) and the root mean squared prediction error (RMSE). Figure 6 compares Gapfill’s 90% intervals with 90% posterior equal-tailed credible intervals for the Q-MGP predictions obtained from 1000 posterior samples. In terms of




Figure 5. Artificial cloud covering in synthetic data.

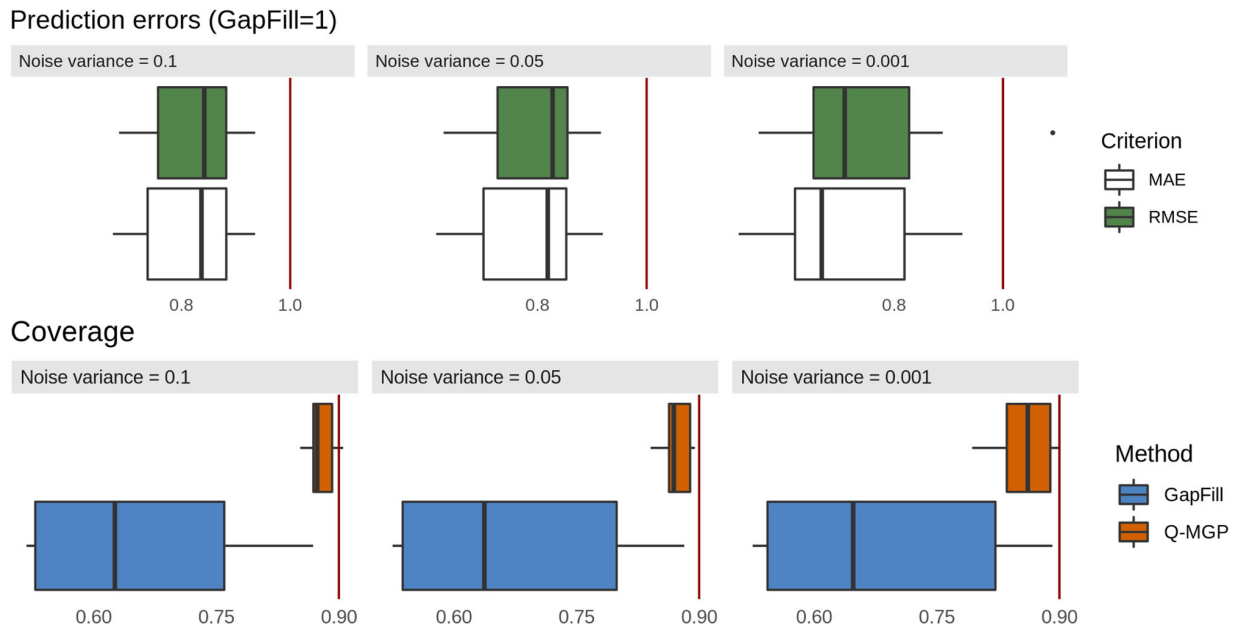


Figure 6. Performance of Q-MGP and Gapfill in out-of-sample predictions in 81 spatiotemporal datasets, at the three tested levels of noise variance τ^2 .

MAE, the Q-MGP model outperformed Gapfill in all datasets; in terms of RMSE, it outperformed Gapfill in all but one dataset. The average MAE of Q-MGP was 0.4094 against Gapfill's 0.5366; the average RMSE was 0.5308 against Gapfill's 0.6820. The Q-MGP also yielded improved coverage of the prediction intervals, although some under-coverage was observed possibly due to the large M . This comparison may favor Q-MGPs as the data were generated from a GP. Appendix K, available online, confirms similar findings on non-Gaussian data (a GIF image).

5.3. NDVI Data From the Serengeti Ecosystem

Time series of NDVI derived from satellite imagery are used to understand spatial patterns in vegetation phenology. For such studies, image pixel-level NDVI values are observed over time to assess seasonal trends in vegetation green-up, growing season length and peak, and senescence. These analyses typically require NDVI values for all pixels over the region and time period of interest. As noted in the beginning of this

section, atmospheric conditions, for example, cloud cover, and sensor malfunction cause missing NDVI pixel values and hence predicting such values, that is, gap-filling, is of key interest to the remote sensing community. Here, we consider NDVI data derived from the LandSat 8 sensor (which provides a $\sim 30 \times 30$ m spatial resolution pixel) taken over Serengeti National Park, Tanzania, Africa. These data were part of a larger study conducted by Desanker, Dahlin, and Finley (2020) that looked at environmental drivers in vegetation phenology change. The data cover an area of $30 \text{ km} \times 30 \text{ km}$ and 34 months, and correspond to 64 images of size 1000×1000 collected at 16-day intervals. Data on NDVI are complemented with elevation and soil moisture data, for a total of three spatially referenced predictors. We are thus interested in understanding their varying effect in space and time, in addition to predicting NDVI output at missing locations. We achieve both these goals by implementing model (11), where $\mathbf{Z}(\ell) = \mathbf{X}(\ell)$ includes the intercept and three predictors; their varying effect will be represented by $\mathbf{w}(\ell)$, which we recover by implementing Q-MGP models. Storing posterior samples of the multivariate spatially varying coefficients for the

full data with $q = 4$ is impossible using our computing resources as each sample would be of size $1000 \times 1000 \times 64 \times 4 = 2.56 \text{E}+8$. For this reason, we consider two feasible setups. Denote by n_{all} the number of observed and missing locations. In model (1), we subsample each image to obtain 64 frames of size 500×500 , and fit a regression model with $\mathbf{Z}(\ell) = 1$ resulting in a spatially varying intercept model on $n_{\text{obs}} = 12,582,484$ observed locations, a total of $n_{\text{all}} = 16,000,000$ locations for prediction, and a latent spatial random effect w of the same size. The Q-MGP was fit using $M = 328,125$ space-time regions of size ~ 48 .

The base covariance of (13) becomes a univariate nonseparable spatiotemporal covariance as in Gneiting (2002). In model (2), we aim to estimate the varying effect of elevation on NDVI. We subsample each image to obtain 64 frames of size 278×278 , each covering an area of $25 \text{ km} \times 25 \text{ km}$, and take $\mathbf{Z}(\ell) = (1 \ \mathbf{X}_{\text{elev}}(\ell))$ resulting in $q = 2$ and targeting the recovery of latent effects of size 9,892,352. Considering the additional computational burden of the multivariate latent effects, in this case we used $M = 156,800$, corresponding to smaller space-time regions of average size ~ 31 . In this model, there is a single unknown δ_{ij} in (13) which corresponds to the latent dissimilarity between the intercept and elevation. We thus consider $\psi_2 = (a_2 \delta_{ij} + 1)^{\beta_2}$ as the unknown parameter. We assign priors $\beta_r \sim N(0, 100)$ for $r = 1, \dots, q$, $\sigma^2 \sim \text{Inv.G.}(2, 1)$, $\tau^2 \sim \text{Inv.G.}(2, 1)$, and uniform priors to other covariance parameters (their support is reported in Table 2).


In both cases, approximate posterior samples of the latent random effects and the other unknown parameters were obtained by running the proposed Gibbs sampler for a total of 25,000 iterations. A posterior sample of size 500 was obtained by using the first 22,000 iterations as burn-in, and thinning the remaining 3000 by a factor of 6. Additional computational details are at Appendix F, available online. Posterior summaries

Table 2. Posterior summaries of Q-MGP models implemented on the Serengeti data.

	Q-MGP model (1)	Q-MGP model (2)
n_{all}	16,000,000	4,946,176
n_{obs}	12,755,856	3,961,715
M	328,125	156,800
q	1	2
$\beta_{\text{elevation}}$	0.0017 _(0.0014,0.0021)	0.0415 _(0.0398,0.0432)
β_{topindex}	$5.54\text{e-}4$ _($4.72\text{e-}4,6.30\text{e-}4$)	-0.0011 _($-0.0012,-0.0008$)
β_{accum}	$-4.84\text{e-}4$ _($-5.66\text{e-}4,-4.02\text{e-}4$)	$7.88\text{e-}4$ _($6.94\text{e-}4,9.06\text{e-}4$)
σ^2	0.0585 _(0.0583,0.0587)	0.0728 _(0.0711,0.0749)
τ^2	$1.05\text{e-}4$ _($1.05\text{e-}4,1.05\text{e-}4$)	$1.27\text{e-}4$ _($1.21\text{e-}4,1.32\text{e-}4$)
$c \sim U(0, 1\text{e}+6)$	7.0331 _(7.0146,7.0519)	3.0710 _(3.0562,3.0846)
$a_1 \sim U(0, 1\text{e}+6)$	433.98 _(429.67,439.50)	3857.6 _(3492.6,4154.7)
$\beta_1 \sim U(0, 1)$	0.0694 _(0.0690,0.0697)	0.1058 _(0.1043,0.1080)
$\psi_2 \sim U(0, 1\text{e}+6)$	–	221.36 _(198.09,240.57)
95% coverage	94.96	95.66
RMSE	0.0175	0.0253
Time/it. (sec)	6.18	4.53
Time (hr)	42.9	31.5

of the unknown parameters for these models are reported in Table 2, along with RMSE in predicting NDVI at 10,000 left-out locations, 95% posterior coverage at those locations, and run times. Both models achieved similar out-of-sample predictive performance and coverage. Figure 7 shows the NDVI predictions of model (2) at one of the 64 time points. This reveals that the varying effect of elevation on NDVI output (see, e.g., Figure 8) is credibly different from zero at 42.54% of the space-time locations (95% CI). In particular, it highlights the extent to which higher elevation reduces vegetation. The spatial range is approximately 4 km; the time range is about 8 days. The large estimated ψ_2 indicates that the correlation between the two covariates of the latent random process is very small at all spatial and temporal lags. The predicted NDVI and latent spatiotemporal effects are supplied as animated GIF images in the supplementary materials.

Observed NDVI



Predicted NDVI






Figure 7. NDVI predictions from Q-MGP model (2) at time 60 (2016-12-17).


Elevation: effect on NDVI (2016-12-01 to 2017-02-19)




Residual spatial effect



Elevation: effect on NDVI



Covariance with estim. parameters



Elevation: nonzero effect locations (95% cred. int.)




Figure 8. Top: the effect of elevation on NDVI output, evolving over five time periods. Middle left: effect on NDVI not explained by elevation; right: effect on NDVI attributable to elevation. Bottom left: Estimated covariance at different space-time lags; right, in blue: locations with credibly nonzero effect of elevation on NDVI output.

6. Discussion

We have developed a class of Bayesian hierarchical models for large spatial and spatiotemporal datasets based on linking domain partitions to DAGs. These models can be tailored for specific algorithmic needs, and we have demonstrated the advantages of using a cubic tessellation scheme (Q-MGP) when targeting the efficient recovery of spatial random effects in Bayesian hierarchical models using Gibbs samplers.

When considering alternative computational strategies, the proposed Q-MGP may not be optimal. For example, Gaussian first stage models enable marginalization of the latent spatial effects; posterior sampling of unknown covariance parameters via MCMC is typically associated by better mixing. Future research may thus focus on identifying “optimal” DAGs for collapsed samplers. Furthermore, the blocked conditional independence structure of Q-MGPs may be suboptimal as it corresponds to possibly restrictive conditional independence assumptions in neighboring locations. While we have not focused on the effect of different tessellations or partitioning choices in this article, alternative tessellation schemes (e.g., hexagonal) may be associated to less stringent assumptions and possibly better performance, while retaining all the desirable features of Q-MGP.

Other natural extensions to high-dimensional spatiotemporal statistics include settings where there are a very large number of spatiotemporal outcomes in addition to a large number of spatial and temporal locations. Here there are a few different avenues. One approach is in the same spirit of joint modeling pursued here, but instead of modeling the cross-covariance functions explicitly, as has been done here, we pursue dimension reduction using factor models (see, e.g., Christensen and Amemiya 2003; Lopes, Salazar, and Gamerman 2008; Ren and Banerjee 2013; Taylor-Rodriguez et al. 2019). The aforementioned references have attempted to model the factor models using spatial processes some of which have used scalable low-rank predictive processes or the NNGP. We believe that modeling latent factors using spatiotemporal MGPs will impart some of the computational and inferential benefits seen here. However, this will need further development especially with regard to identifiability of loading matrices (Lopes, Salazar, and Gamerman 2008; Ren and Banerjee 2013) and process parameters.

A different approach to multivariate spatial modeling has relied upon conditional or hierarchical specifications. This has been well articulated in the text by Cressie and Wikle (2011); see also Royle and Berliner (1999) and the recent developments in Cressie and Zammit-Mangion (2016). An advantage of the hierarchical approach is that the multivariate processes are valid stochastic processes, essentially by construction and without requiring spectral representations, and can also impart considerable computational benefits. It will be interesting to extend the ideas in Cressie and Zammit-Mangion (2016) to augmented spaces of DAGs to further introduce conditional independence, and therefore sparsity, in MGP models with high-dimensional outcomes.

Finally, it is worth pointing out that alternate computational algorithms, particularly tuned for high-dimensional Bayesian models, should also be explored. Recent developments on

algorithms based upon classes of piecewise deterministic Markov processes (see, e.g., Fearnhead et al. 2018; Bierkens, Fearnhead, and Roberts 2019, and references therein) that avoid Gibbs samplers and even reversible MCMC algorithms are being shown to be increasingly effective for high-dimensional Bayesian inference. Adapting such algorithms to MGP and Q-MGP for scalable Bayesian spatial process models will constitute natural extensions of our current offering.

Supplementary Materials

The online supplement includes additional theoretical and computational details on Meshed Gaussian Processes, along with discussions on tessellation designs; the choice of the reference set and partition sizes; an application to multivariate outcomes; and a comparison with other state-of-the-art scalable methods for large spatial data.

Funding

Banerjee was supported by the NSF grants DMS-1513654, IIS-1562303, and DMS-1916349; and by the National Institute of Health grants NIEHS-R01ES027027 and NIEHS-R01ES030210. Finley and Peruzzi were supported by National Science Foundation (NSF) EF-1253225 and DMS-1916395, and National Aeronautics and Space Administration's Carbon Monitoring System project. Peruzzi was supported in part by R01ES028804 of the National Institute of Environmental Health Sciences of the National Institutes of Health and European Union project 856506.

References

- Apanasovich, T. V., and Genton, M. G. (2010), “Cross-Covariance Functions for Multivariate Random Fields Based on Latent Dimensions,” *Biometrika*, 97, 15–30, DOI: 10.1093/biomet/asp078. [9]
- Banerjee, S. (2017), “High-Dimensional Bayesian Geostatistics,” *Bayesian Analysis*, 12, 583–614, DOI: 10.1214/17-BA1056R. [1]
- (2020), “Modeling Massive Spatial Datasets Using a Conjugate Bayesian Linear Modeling Framework,” *Spatial Statistics* (in press), DOI: 10.1016/j.spasta.2020.100417. [4]
- Banerjee, S., Finley, A. O., Waldmann, P., and Ericsson, T. (2010), “Hierarchical Spatial Process Models for Multiple Traits in Large Genetic Trials,” *Journal of American Statistical Association*, 105, 506–521, DOI: 10.1198/jasa.2009.ap09068. [1]
- Banerjee, S., Gelfand, A. E., Finley, A. O., and Sang, H. (2008), “Gaussian Predictive Process Models for Large Spatial Data Sets,” *Journal of the Royal Statistical Society, Series B*, 70, 825–848, DOI: 10.1111/j.1467-9868.2008.00663.x. [1]
- Bierkens, J., Fearnhead, P., and Roberts, G. (2019), “The Zig-Zag Process and Super-Efficient Sampling for Bayesian Analysis of Big Data,” *The Annals of Statistics*, 47, 1288–1320, DOI: 10.1214/18-AOS1715. [13]
- Chen, Y., Davis, T. A., Hager, W. W., and Rajamanickam, S. (2008), “Algorithm 887: CHOLMOD, Supernodal Sparse Cholesky Factorization and Update/Downdate,” *ACM Transactions on Mathematical Software*, 35, 1–14, DOI: 10.1145/1391989.1391995. [5]
- Christensen, W. F., and Amemiya, Y. (2003), “Modeling and Prediction for Multivariate Spatial Factor Analysis,” *Journal of Statistical Planning and Inference*, 115, 543–564, DOI: 10.1016/S0378-3758(02)00173-8. [13]
- Cover, T. M., and Thomas, J. A. (1991), *Elements of Information Theory*, Wiley Series in Telecommunications and Signal Processing, New York: Wiley-Interscience. [4]
- Cressie, N. A., and Johannesson, G. (2008), “Fixed Rank Kriging for Very Large Spatial Data Sets,” *Journal of the Royal Statistical Society, Series B*, 70, 209–226, DOI: 10.1111/j.1467-9868.2007.00633.x. [1]
- Cressie, N. A., and Wikle, C. K. (2011), *Statistics for Spatio-Temporal Data*, Wiley Series in Probability and Statistics, Hoboken, NJ: Wiley. [13]
- Cressie, N. A., and Zammit-Mangion, A. (2016), “Multivariate Spatial Covariance Models: A Conditional Approach,” *Biometrika*, 103, 915–935, DOI: 10.1093/biomet/asw045. [13]

- Datta, A., Banerjee, S., Finley, A. O., and Gelfand, A. E. (2016), “Hierarchical Nearest-Neighbor Gaussian Process Models for Large Geostatistical Datasets,” *Journal of the American Statistical Association*, 111, 800–812, DOI: 10.1080/01621459.2015.1044091. [2,3]
- Datta, A., Banerjee, S., Finley, A. O., Hamm, N. A. S., and Schaap, M. (2016), “Nonseparable Dynamic Nearest Neighbor Gaussian Process Models for Large Spatio-Temporal Data With an Application to Particulate Matter Analysis,” *The Annals of Applied Statistics*, 10, 1286–1316, DOI: 10.1214/16-AOAS931. [2]
- Davis, T. A. (2006), *Direct Methods for Sparse Linear Systems*, Philadelphia, PA: SIAM. [5]
- Desanker, G., Dahlin, K. M., and Finley, A. O. (2020), “Environmental Controls on Landsat-Derived Phenoregions Across an East African Megatransect,” *Ecosphere*, 11, e03143. [10]
- Eidsvik, J., Shaby, B. A., Reich, B. J., Wheeler, M., and Niemi, J. (2014), “Estimation and Prediction in Spatial Models With Block Composite Likelihoods,” *Journal of Computational and Graphical Statistics*, 23, 295–315, DOI: 10.1080/10618600.2012.760460. [1]
- Fearnhead, P., Bierkens, J., Pollock, M., and Roberts, G. O. (2018), “Piecewise Deterministic Markov Processes for Continuous-Time Monte Carlo,” *Statistical Science*, 33, 386–412, DOI: 10.1214/18-STS648. [13]
- Finley, A. O., Banerjee, S., and Gelfand, A. E. (2012), “Bayesian Dynamic Modeling for Large Space-Time Datasets Using Gaussian Predictive Processes,” *Journal of Geographical Systems*, 14, 29–47, DOI: 10.1007/s10109-011-0154-8. [1]
- Finley, A. O., Datta, A., Cook, B. D., Morton, D. C., Andersen, H. E., and Banerjee, S. (2019), “Efficient Algorithms for Bayesian Nearest Neighbor Gaussian Processes,” *Journal of Computational and Graphical Statistics*, 28, 401–414, DOI: 10.1080/10618600.2018.1537924. [2,3,5]
- Furrer, R., Genton, M. G., and Nychka, D. (2006), “Covariance Tapering for Interpolation of Large Spatial Datasets,” *Journal of Computational and Graphical Statistics*, 15, 502–523, DOI: 10.1198/106186006X132178. [1]
- Genton, M. G., and Kleiber, W. (2015), “Cross-Covariance Functions for Multivariate Geostatistics,” *Statistical Science*, 30, 147–163, DOI: 10.1214/14-STS487. [4]
- Gerber, F., Furrer, R., Schaepman-Strub, G., de Jong, R., and Schaepman, M. E. (2018), “Predicting Missing Values in Spatio-Temporal Remote Sensing Data,” *IEEE Transactions on Geoscience and Remote Sensing*, 56, 2841–2853, DOI: 10.1109/TGRS.2017.2785240. [9]
- Gneiting, T. (2002), “Nonseparable, Stationary Covariance Functions for Space-Time Data,” *Journal of the American Statistical Association*, 97, 590–600, DOI: 10.1198/016214502760047113. [9,11]
- Gonzalez, J., Low, Y., Gretton, A., and Guestrin, C. (2011), “Parallel Gibbs Sampling: From Colored Fields to Thin Junction Trees,” in *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, Proceedings of Machine Learning Research (Vol. 15), eds. G. Gordon, D. Dunson, and M. Dudík, PMLR, Fort Lauderdale, FL, USA, pp. 324–332. [8]
- Gramacy, R. B., and Apley, D. W. (2015), “Local Gaussian Process Approximation for Large Computer Experiments,” *Journal of Computational and Graphical Statistics*, 24, 561–578, DOI: 10.1080/10618600.2014.914442. [3]
- Gramacy, R. B., and Lee, H. K. H. (2008), “Bayesian Treed Gaussian Process Models With an Application to Computer Modeling,” *Journal of the American Statistical Association*, 103, 1119–1130, DOI: 10.1198/016214508000000689. [3]
- Guhaniyogi, R., Finley, A. O., Banerjee, S., and Gelfand, A. E. (2011), “Adaptive Gaussian Predictive Process Models for Large Spatial Datasets,” *Environmetrics*, 22, 997–1007, DOI: 10.1002/env.1131. [1]
- Guinness, J. (2018), “Permutation and Grouping Methods for Sharpening Gaussian Process Approximations,” *Technometrics*, 60, 415–429, DOI: 10.1080/00401706.2018.1437476. [2,3,4]
- Heaton, M. J., Datta, A., Finley, A. O., Furrer, R., Guinness, J., Guhaniyogi, R., Gerber, F., Gramacy, R. B., Hammerling, D., Katzfuss, M., Lindgren, F., Nychka, D. W., Sun, F., and Zammit-Mangion, A. (2019), “A Case Study Competition Among Methods for Analyzing Large Spatial Data,” *Journal of Agricultural, Biological and Environmental Statistics*, 24, 398–425, DOI: 10.1007/s13253-018-00348-w. [9]
- Katzfuss, M. (2017), “A Multi-Resolution Approximation for Massive Spatial Datasets,” *Journal of the American Statistical Association*, 112, 201–214, DOI: 10.1080/01621459.2015.1123632. [1,3]
- Katzfuss, M., and Guinness, J. (2017), “A General Framework for Vecchia Approximations of Gaussian Processes,” arXiv no. 1708.06302. [2,3]
- Kaufman, C. G., Schervish, M. J., and Nychka, D. W. (2008), “Covariance Tapering for Likelihood-Based Estimation in Large Spatial Data Sets,” *Journal of the American Statistical Association*, 103, 1545–1555, DOI: 10.1198/016214508000000959. [1]
- Lauritzen, S., L. (1996), *Graphical Models*, Oxford: Clarendon Press. [4]
- Lewis, R. (2016), *A Guide to Graph Colouring*, Cham: Springer. [8]
- Lopes, H. F., Salazar, E., and Gamerman, D. (2008), “Spatial Dynamic Factor Analysis,” *Bayesian Analysis*, 3, 759–792, DOI: 10.1214/08-BA329. [13]
- Molloy, M., and Reed, B. (2002), *Graph Colouring and the Probabilistic Method*, Berlin, Heidelberg: Springer-Verlag. [8]
- Quiroz, Z. C., Prates, M. O., and Dey, D. K. (2019), “Block Nearest Neighbor Gaussian Processes for Large Datasets,” arXiv no. 1908.06437. [2]
- Ren, Q., and Banerjee, S. (2013), “Hierarchical Factor Models for Large Spatially Misaligned Data: A Low-Rank Predictive Process Approach,” *Biometrics*, 69, 19–30, DOI: 10.1111/j.1541-0420.2012.01832.x. [13]
- Royle, J. A., and Berliner, L. M. (1999), “A Hierarchical Approach to Multivariate Spatial Modeling and Prediction,” *Journal of Agricultural, Biological, and Environmental Statistics*, 1, 29–56, DOI: 10.2307/1400420. [13]
- Rue, H., and Held, L. (2005), *Gaussian Markov Random Fields: Theory and Applications*, Boca Raton, FL: Chapman & Hall/CRC. [1]
- Sang, H., and Huang, J. Z. (2012), “A Full Scale Approximation of Covariance Functions for Large Spatial Data Sets,” *Journal of the Royal Statistical Society, Series B*, 74, 111–132, DOI: 10.1111/j.1467-9868.2011.01007.x. [1]
- Stein, M. L. (2014), “Limitations on Low Rank Approximations for Covariance Matrices of Spatial Data,” *Spatial Statistics*, 8, 1–19, DOI: 10.1016/j.spasta.2013.06.003. [1]
- Stein, M. L., Chi, Z., and Welty, L. J. (2004), “Approximating Likelihoods for Large Spatial Data Sets,” *Journal of the Royal Statistical Society, Series B*, 66, 275–296, DOI: 10.1046/j.1369-7412.2003.05512.x. [3]
- Sun, Y., Li, B., and Genton, M. (2011), “Geostatistics for Large Datasets,” in *Advances and Challenges in Space-Time Modelling of Natural Events*, eds. J. Montero, E. Porcu, and M. Schlather, Berlin, Heidelberg: Springer-Verlag, pp. 55–77. [1]
- Taylor-Rodriguez, D., Finley, A. O., Datta, A., Babcock, C., Andersen, H. E., Cook, B. D., Morton, D. C., and Banerjee, S. (2019), “Spatial Factor Models for High-Dimensional and Large Spatial Data: An Application in Forest Variable Mapping,” *Statistica Sinica*, 29, 1155–1180, DOI: 10.5705/ss.202018.0005. [13]
- Vecchia, A. V. (1988), “Estimation and Model Identification for Continuous Spatial Processes,” *Journal of the Royal Statistical Society, Series B*, 50, 297–312, DOI: 10.1111/j.2517-6161.1988.tb01729.x. [1,2,3]

Highly Scalable Bayesian Geostatistical Modeling via Meshed Gaussian Processes on Partitioned Domains

Appendix

Michele Peruzzi, Sudipto Banerjee and Andrew O. Finley

A Spatial Meshed Process

Let $\mathbf{w}(\mathbf{s})$, $\mathbf{s} \in \mathcal{D}$ be the base process, \mathcal{S} the fixed reference set, $\mathcal{L} = \{\ell_1, \dots, \ell_n\} \subset \mathcal{D}$ and $\mathcal{U} = \mathcal{L} \setminus \mathcal{S}$. Then the joint density $\tilde{p}(\mathbf{w}_{\mathcal{L}})$ is proper. In fact, using the definitions of $\tilde{p}(\mathbf{w}_{\mathcal{S}})$ and $\tilde{p}(\mathbf{w}_{\mathcal{U}} | \mathbf{w}_{\mathcal{S}})$ we get

$$\begin{aligned} \int \tilde{p}(\mathbf{w}_{\mathcal{L}}) \prod_{\ell_i \in \mathcal{L}} d(\mathbf{w}(\ell_i)) &= \int \int \tilde{p}(\mathbf{w}_{\mathcal{U}} | \mathbf{w}_{\mathcal{S}}) \tilde{p}(\mathbf{w}_{\mathcal{S}}) \prod_{\mathbf{s}_i \in \mathcal{S} \setminus \mathcal{L}} d(\mathbf{w}(\mathbf{s}_i)) \prod_{\ell_i \in \mathcal{L}} d(\mathbf{w}(\ell_i)) \\ &= \int \tilde{p}(\mathbf{w}_{\mathcal{U}} | \mathbf{w}_{\mathcal{S}}) \tilde{p}(\mathbf{w}_{\mathcal{S}}) \prod_{\ell_i \in \mathcal{U}} d(\mathbf{w}(\ell_i)) \prod_{\ell_i \in \mathcal{S}} d(\mathbf{w}(\ell_i)) \\ &= \int \tilde{p}(\mathbf{w}_{\mathcal{S}}) \left(\int \tilde{p}(\mathbf{w}_{\mathcal{U}} | \mathbf{w}_{\mathcal{S}}) \prod_{\ell_i \in \mathcal{U}} d(\mathbf{w}(\ell_i)) \right) \prod_{\ell_i \in \mathcal{S}} d(\mathbf{w}(\ell_i)) = 1. \end{aligned}$$

We now verify the Kolmogorov consistency conditions. Take $\mathcal{L}_{\pi} = \{\ell_{\pi(1)}, \dots, \ell_{\pi(n)}\}$ as any permutation of \mathcal{L} . Then call $\mathcal{U}_{\pi} = \mathcal{L}_{\pi} \setminus \mathcal{S}$ and we get

$$\tilde{p}(\mathbf{w}_{\mathcal{L}_{\pi}}) = \int \tilde{p}(\mathbf{w}_{\mathcal{U}_{\pi}} | \mathbf{w}_{\mathcal{S}}) \tilde{p}(\mathbf{w}_{\mathcal{S}}) \prod_{\mathbf{s}_i \in \mathcal{S} \setminus \mathcal{L}_{\pi}} d(\mathbf{w}(\mathbf{s}_i)).$$

Set membership is invariant to permutations, so $\mathcal{U}_{\pi} = \mathcal{L}_{\pi} \setminus \mathcal{S} = \mathcal{L} \setminus \mathcal{S} = \mathcal{U}$. and therefore in no way the order of the \mathcal{L} locations changes $\tilde{p}(\mathbf{w}_{\mathcal{L}})$. Therefore $\tilde{p}(\mathbf{w}_{\mathcal{L}_{\pi}}) = \tilde{p}(\mathbf{w}_{\mathcal{L}})$, i.e.

$$\tilde{p}(\mathbf{w}(\ell_1), \dots, \mathbf{w}(\ell_n)) = \tilde{p}(\mathbf{w}(\ell_{\pi(1)}), \dots, \mathbf{w}(\ell_{\pi(n)})).$$

Next, take another location $\ell_0 \in \mathcal{D}$. Call $\mathcal{L}_1 = \mathcal{L} \cup \{\ell_0\}$. We want to show that $\int \tilde{p}(\mathbf{w}_{\mathcal{L}_1}) d(\mathbf{w}(\ell_0)) = \tilde{p}(\mathbf{w}_{\mathcal{L}})$. If $\ell_0 \in \mathcal{S}$ then $\mathcal{L}_1 \setminus \mathcal{S} = \mathcal{L} \setminus \mathcal{S} = \mathcal{U}$ hence

$$\begin{aligned} \int \tilde{p}(\mathbf{w}_{\mathcal{L}_1}) d(\mathbf{w}(\ell_0)) &= \int \tilde{p}(\mathbf{w}_{\mathcal{L}_1 \setminus \mathcal{S}} | \mathbf{w}_{\mathcal{S}}) \tilde{p}(\mathbf{w}_{\mathcal{S}}) \prod_{\mathbf{s}_i \in \mathcal{S} \setminus \mathcal{L}_1} d(\mathbf{w}(\mathbf{s}_i)) d(\mathbf{w}(\ell_0)) \\ &= \int \tilde{p}(\mathbf{w}_{\mathcal{U}} | \mathbf{w}_{\mathcal{S}}) \tilde{p}(\mathbf{w}_{\mathcal{S}}) \prod_{\mathbf{s}_i \in \mathcal{S} \setminus \mathcal{L}} d(\mathbf{w}(\mathbf{s}_i)) = \tilde{p}(\mathbf{w}_{\mathcal{L}}) \end{aligned}$$

If $\ell_0 \notin \mathcal{S}$

$$\begin{aligned} \int \tilde{p}(\mathbf{w}_{\mathcal{L}_1}) d(\mathbf{w}(\ell_0)) &= \int \left(\int \tilde{p}(\mathbf{w}_{\mathcal{L}_1 \setminus \mathcal{S}} | \mathbf{w}_{\mathcal{S}}) \tilde{p}(\mathbf{w}_{\mathcal{S}}) \prod_{\mathbf{s}_i \in \mathcal{S} \setminus \mathcal{L}_1} d(\mathbf{w}(\mathbf{s}_i)) \right) d(\mathbf{w}(\ell_0)) \\ &= \int \left(\int \tilde{p}(\mathbf{w}_{\mathcal{L} \setminus \mathcal{S} \cup \{\ell_0\}} | \mathbf{w}_{\mathcal{S}}) \tilde{p}(\mathbf{w}_{\mathcal{S}}) \prod_{\mathbf{s}_i \in \mathcal{S} \setminus \mathcal{L}} d(\mathbf{w}(\mathbf{s}_i)) \right) d(\mathbf{w}(\ell_0)) \\ &= \int \left(\int \tilde{p}(\mathbf{w}_{\{\ell_0\}} | \mathbf{w}_{\mathcal{L} \setminus \mathcal{S}}, \mathbf{w}_{\mathcal{S}}) \tilde{p}(\mathbf{w}_{\mathcal{L} \setminus \mathcal{S}} | \mathbf{w}_{\mathcal{S}}) \tilde{p}(\mathbf{w}_{\mathcal{S}}) \prod_{\mathbf{s}_i \in \mathcal{S} \setminus \mathcal{L}} d(\mathbf{w}(\mathbf{s}_i)) \right) d(\mathbf{w}(\ell_0)) \\ &= \int \tilde{p}(\mathbf{w}_{\mathcal{L} \setminus \mathcal{S}} | \mathbf{w}_{\mathcal{S}}) \tilde{p}(\mathbf{w}_{\mathcal{S}}) \prod_{\mathbf{s}_i \in \mathcal{S} \setminus \mathcal{L}} d(\mathbf{w}(\mathbf{s}_i)) \int \tilde{p}(\mathbf{w}_{\{\ell_0\}} | \mathbf{w}_{\mathcal{S}}) d(\mathbf{w}(\ell_0)) \\ &= \int \tilde{p}(\mathbf{w}_{\mathcal{L} \setminus \mathcal{S}} | \mathbf{w}_{\mathcal{S}}) \tilde{p}(\mathbf{w}_{\mathcal{S}}) \prod_{\mathbf{s}_i \in \mathcal{S} \setminus \mathcal{L}} d(\mathbf{w}(\mathbf{s}_i)) \\ &= \tilde{p}(\mathbf{w}_{\mathcal{L}}). \end{aligned}$$

B Meshed Gaussian Process

In graph \mathcal{G} , the number of parents of node \mathbf{v}_j is b_j , i.e. $b_j = |\text{Pa}[\mathbf{v}_j]|$. Therefore $\mathbf{H}_j = \mathbf{C}_{\mathcal{S}_j, \mathcal{S}_{\text{Pa}[j]}} \mathbf{C}_{\mathcal{S}_{\text{Pa}[j]}}^{-1}$ is a $qn_j \times q \sum_{r=1}^{b_j} n_r$ matrix which can be partitioned by column in b_j blocks:

$$\mathbf{H}_j = \left[\mathbf{H}_{j,1} \mid \cdots \mid \mathbf{H}_{j,b_j} \right],$$

whose r th block is of size $q \times q$ and corresponds to the r th block of $\mathbf{w}_{\text{Pa}[j]}$. $\tilde{p}(\mathbf{w}_{\mathcal{S}}) = \prod_{j=1}^M N(\mathbf{w}_j | \mathbf{H}_j \mathbf{w}_{[j]}, \mathbf{R}_j)$ is then proportional to

$$\frac{1}{\prod_{j=1}^M |\mathbf{R}_j|^{\frac{1}{2}}} \exp \left\{ -\frac{1}{2} \sum_{j=1}^M (\mathbf{w}_j - \mathbf{H}_j \mathbf{w}_{\text{Pa}[j]})^\top \mathbf{R}_j^{-1} (\mathbf{w}_j - \mathbf{H}_j \mathbf{w}_{\text{Pa}[j]}) \right\},$$

and we can define \mathbf{H}_j^* as the $qn_j \times qn_S$ matrix such that $\mathbf{w}_j - \mathbf{H}_j \mathbf{w}_{\text{Pa}[j]} = \mathbf{H}_j^* \mathbf{w}_S$. Analogously to \mathbf{H}_j , we can partition \mathbf{H}_j^* by column in M blocks,

$$\mathbf{H}_j^* = \left[\mathbf{H}_{j,1}^* \mid \cdots \mid \mathbf{H}_{j,h}^* \mid \cdots \mid \mathbf{H}_{j,M}^* \right],$$

where each block $h = 1, \dots, M$

$$\mathbf{H}_{j,h}^* = \begin{cases} \mathbf{O}_{qn_h \times qn_h} & \text{if } \mathbf{v}_h \notin \text{Pa}[\mathbf{v}_j] \\ \mathbf{I}_{qn_j \times qn_j} & \text{if } \mathbf{v}_h = \mathbf{v}_j \\ -\mathbf{H}_{j,r} & \text{if } \mathbf{v}_h \text{ is } \mathbf{v}_j\text{'s } r\text{th parent,} \end{cases}$$

and notice that $h > j$ implies $\mathbf{v}_h \notin \text{Pa}[\mathbf{v}_j]$. Then,

$$\sum_{j=1}^M (\mathbf{w}_j - \mathbf{H}_j \mathbf{w}_{\text{Pa}[j]})^\top \mathbf{R}_j^{-1} (\mathbf{w}_j - \mathbf{H}_j \mathbf{w}_{\text{Pa}[j]}) = \sum_{j=1}^M \mathbf{w}_S^\top \mathbf{H}_j^{*\top} \mathbf{R}_j^{-1} \mathbf{H}_j^* \mathbf{w}_S = \mathbf{w}_S^\top \mathbf{H}_S^{*\top} \mathbf{R}_S^{-1} \mathbf{H}_S^* \mathbf{w}_S,$$

where \mathbf{H}_S^* is a $qn_S \times qn_S$ block-matrix whose j th block-row is \mathbf{H}_j^* for $j = 1, \dots, M$, and $\mathbf{R}_S = \text{blockdiag}(\mathbf{R}_1, \dots, \mathbf{R}_M)$. The resulting precision matrix $\tilde{\mathbf{C}}_S^{-1} = \mathbf{H}_S^{*\top} \mathbf{R}_S^{-1} \mathbf{H}_S^*$ is thus a $qn_S \times qn_S$ matrix made of $M \times M$ blocks of sizes $qn_j \times qn_j$ for $j = 1, \dots, M$ such that $\sum n_j = n_S$. We denote its (i, j) block as $\tilde{\mathbf{C}}_S^{-1}(i, j)$ and note that by symmetry its transpose is $(\tilde{\mathbf{C}}_S^{-1}(i, j))^\top = \tilde{\mathbf{C}}_S^{-1}(j, i)$. We find the nonzero blocks of $\tilde{\mathbf{C}}_S^{-1}$ in its block-diagonal, and for $i \neq j$ in the (i, j) block such that \mathbf{v}_i and \mathbf{v}_j are connected in the moralized subgraph $\mathcal{G}^m(\mathbf{A})$ of $\mathcal{G}(\mathbf{A})$ obtained by subsetting \mathcal{G} to the \mathbf{A} nodes and linking those that share a child; in other words for all $\mathbf{v}_i, \mathbf{v}_j \in \mathbf{A}$, either (1) $\mathbf{v}_i \in \text{Pa}[\mathbf{v}_j]$ or vice-versa, or (2) there exists \mathbf{v}_h such that $\{\mathbf{v}_i, \mathbf{v}_j\} \subset \text{Pa}[\mathbf{v}_h]$. The sparsity of $\tilde{\mathbf{C}}_S^{-1}$ thus depends on the specific choice for $\mathcal{G}(\mathbf{A})$ and the corresponding moralized graph $\mathcal{G}^m(\mathbf{A})$ (Cowell et al., 1999). Suppose that $\mathcal{G}(\mathbf{A})$ has M vertices and induces $\mathcal{G}^m(\mathbf{A})$ whose number of undirected edges is $k \ll M(M-1)/2$. This means that the number of “missing” connections is $\ell = M(M-1)/2 - k$. Assume that each subset \mathcal{S}_j of the reference set is of size $n_j = n/M$, so that each block of $\tilde{\mathbf{C}}_S^{-1}$ is of size $qn/M \times qn/M$. Then the precision matrix will have $\ell \left(\frac{qn}{M}\right)^2$ zeros. Given a partition for \mathcal{D} which induces a partition for \mathcal{S} into M disjoint subsets, one can thus increase sparsity in the precision matrix at \mathcal{S} by choosing a graph which induces a lower ℓ . But note that all nodes linked to locations outside the reference set, i.e. any $\mathbf{v} \in \mathbf{B}$, do not influence the precision matrix $\tilde{\mathbf{C}}_S^{-1}$.

The covariance function of the new process will then be defined using $\tilde{\mathbf{C}}$: consider two locations ℓ_1 and ℓ_2 in \mathcal{D} . If both are in \mathcal{S} then $\ell_1 = \mathbf{s}_i$, $\ell_2 = \mathbf{s}_j$ for some i, j , and $Cov_{\tilde{\mathbf{p}}}(\mathbf{w}(\mathbf{s}_i), \mathbf{w}(\mathbf{s}_j)) = \tilde{\mathbf{C}}_{\mathbf{s}_i, \mathbf{s}_j}$. If $\ell_1 = \mathbf{u}_i \in \mathcal{U}$ with parents $\text{Pa}[\eta(\mathbf{u}_i)]$ and $\ell_2 = \mathbf{s}_j \in \mathcal{S}$ then

$$\begin{aligned} Cov_{\tilde{\mathbf{p}}}(\mathbf{w}(\mathbf{u}_i), \mathbf{w}(\mathbf{s}_j)) &= E_{\tilde{\mathbf{p}}}(Cov_{\tilde{\mathbf{p}}}(\mathbf{w}(\mathbf{u}_i), \mathbf{w}(\mathbf{s}_j) \mid \mathbf{w}_{\mathcal{S}})) + Cov_{\tilde{\mathbf{p}}}(E_{\tilde{\mathbf{p}}}(\mathbf{w}(\mathbf{u}_i) \mid \mathbf{w}_{\mathcal{S}}), E_{\tilde{\mathbf{p}}}(\mathbf{w}(\mathbf{s}_j) \mid \mathbf{w}_{\mathcal{S}})) \\ &= 0 + Cov_{\tilde{\mathbf{p}}}(\mathbf{H}_{\mathbf{u}_i} \mathbf{w}_{[\mathbf{u}_i]}, \mathbf{w}(\mathbf{s}_j)) \\ &= \mathbf{H}_{\mathbf{u}_i} \tilde{\mathbf{C}}_{\mathcal{S}_{[\mathbf{u}_i]}, \mathbf{s}_j} \end{aligned}$$

Finally if $\ell_1 = \mathbf{u}_i \in \mathcal{U}$ and $\ell_2 = \mathbf{u}_j \in \mathcal{U}$ then

$$Cov_{\tilde{\mathbf{p}}}(\mathbf{w}(\mathbf{u}_i), \mathbf{w}(\mathbf{u}_j)) = \delta_{(\mathbf{u}_i = \mathbf{u}_j)} \mathbf{R}_{\mathbf{u}_i} + \mathbf{H}_{\mathbf{u}_i} \tilde{\mathbf{C}}_{\mathcal{S}_{[\mathbf{u}_i]}, \mathcal{S}_{[\mathbf{u}_j]}} \mathbf{H}_{\mathbf{u}_j}^\top.$$

Domain tessellations in meshed GPs will induce discontinuities in the covariance function when evaluated at locations that span different regions; however, $\tilde{\mathbf{C}}_{\ell_1, \ell_2}$ is continuous if ℓ_1 and ℓ_2 are in the same domain region. To show this, we adapt a result from [Datta et al. \(2016\)](#). Let $\|A, B\|$ denote the Euclidean distance matrix at pairs of locations from sets $A, B \subset \mathcal{D} \subset \mathbb{R}^d$. Take $\mathcal{P} = \{\mathcal{D}_1, \dots, \mathcal{D}_M\}$ as the chosen domain partition, and let $\delta\mathcal{P} = \cup_{i=1}^M \delta\mathcal{D}_i$ be the finite union of their boundaries, each of which has measure zero. Then consider $(\ell_1, \ell_2) \in \mathcal{D} \times \mathcal{D} \setminus (\delta\mathcal{P} \cup \{(\ell, \ell) : \ell \in \mathcal{D}\})$ and let $\text{Pa}[\ell_i] = \{\ell \in \mathcal{S} : \eta(\ell) \in \text{Pa}[\eta(\ell_i)]\}$ be the parent set for ℓ_i , i.e. the set of locations that are mapped to parents of the node to which ℓ_i is mapped. Now take $\ell_1 \in \mathcal{D}_j \setminus \mathcal{S}$ and $\ell_2 = \mathbf{s} \in \mathcal{S}_j$ for some $j = 1, \dots, M$ (other cases follow). We assume that $\mathcal{S}_j \subset \mathcal{D}_j \setminus \delta\mathcal{D}_j$ and $\mathbf{C}(\cdot, \cdot)$ is isotropic and continuous. Then $\lim_{\mathbf{h}_i \rightarrow 0} \|\ell_i + \mathbf{h}_i, \text{Pa}[\ell_i + \mathbf{h}_i]\| \rightarrow \|\ell_i, \text{Pa}[\ell_i]\|$ for $i = 1, 2$ and $\lim_{\mathbf{h}_1 \rightarrow 0, \mathbf{h}_2 \rightarrow 0} \|\text{Pa}[\ell_1 + \mathbf{h}_1], \text{Pa}[\ell_2 + \mathbf{h}_2]\| \rightarrow \|\text{Pa}[\ell_1], \text{Pa}[\ell_2]\|$, implying that $\mathbf{H}_{\ell_1 + \mathbf{h}_1} = \mathbf{C}_{\ell_1 + \mathbf{h}_1, \text{Pa}[\ell_1 + \mathbf{h}_1]} \mathbf{C}_{\text{Pa}[\ell_1 + \mathbf{h}_1]}^{-1} \rightarrow \mathbf{C}_{\mathbf{u}, \text{Pa}[\ell_1]} \mathbf{C}_{\text{Pa}[\ell_1]}^{-1} = \mathbf{H}_{\ell_1}$. Taking \mathbf{h}_2 small enough implies $\ell_2 + \mathbf{h}_2 \in \mathcal{D}_j \setminus \mathcal{S}$ thus $\ell_2 \in \text{Pa}[\ell_2 + \mathbf{h}_2]$. Then suppose $\mathbf{s} = \text{Pa}[\ell_2 + \mathbf{h}_2]_1$ i.e. \mathbf{s} is the first element of the parent set of $\ell_2 + \mathbf{h}_2$. Consequently $\mathbf{C}_{\ell_2 + \mathbf{h}_2, \text{Pa}[\ell_2 + \mathbf{h}_2]} \mathbf{C}_{\text{Pa}[\ell_2 + \mathbf{h}_2]}^{-1} \rightarrow \mathbf{e}_1 = (1, 0, \dots, 0)$ and therefore

$$\begin{aligned} \lim_{\mathbf{h}_1 \rightarrow 0, \mathbf{h}_2 \rightarrow 0} \tilde{\mathbf{C}}(\ell_1 + \mathbf{h}_1, \ell_2 + \mathbf{h}_2) &= \mathbf{H}_{\ell_1} \lim_{\mathbf{h}_1 \rightarrow 0, \mathbf{h}_2 \rightarrow 0} \widetilde{Cov}(\mathbf{w}_{\text{Pa}[\ell_1 + \mathbf{h}_1]}, \mathbf{w}_{\text{Pa}[\ell_2 + \mathbf{h}_2]}) \mathbf{e}_1 \\ &= \mathbf{H}_{\ell_1} \lim_{\mathbf{h}_1 \rightarrow 0} \widetilde{Cov}(\mathbf{w}_{\text{Pa}[\ell_1 + \mathbf{h}_1]}, \mathbf{w}(\mathbf{s})) \\ &= \mathbf{H}_{\ell_1} \tilde{\mathbf{C}}_{\text{Pa}[\ell_1], \mathbf{s}} = \tilde{\mathbf{C}}(\ell_1, \ell_2). \end{aligned}$$

C Choosing the tessellation, M , S and \mathcal{U}




Figure 1: Reference sets

Tessellation

Meshed GPs based on tessellations correspond to fixed DAGs constructed from the immediate neighbors bordering each region. For this reason they are unlike “nearest-neighbor” models in which the number of neighbors is a parameter to be fixed. Increasing the parent set in tessellated meshed GPs amounts to changing the tessellation or increasing the size of each region. As a reviewer points out, estimating the unknown covariance parameter may be facilitated by choosing parent sets that include more distant points (see e.g. [Stein, 2014](#)). Doing so in tessellated meshed GPs amounts to choosing a tessellation with elongated regions. With axis-parallel partitioning, one can partition one of the axis in a much smaller number of intervals than the others, resulting in parent sets that are elongated in that directions. Other shapes, such as a “herringbone” tessellation pattern, might be better at capturing variation in more spatial directions.

Multivariate data pose additional challenges, as one can potentially elect to use different tessellations on different margins of the multivariate process. However, considering different tessellations may induce complicated connections in the underlying DAG; this goes against our modeling approach that is to keep a simple, manageable graph structure in order to speed up computations. In cases in which one seeks to maintain good spatial coverage and is willing to make conditional independence assumptions on the variables, we suggest using a single tessellations, but perhaps dropping inter-variable edges across nodes. To clarify, we

may assume that $w_j(\boldsymbol{\ell})$, the j th margin of the multivariate spatial process at location $\boldsymbol{\ell}$, is conditionally independent of $w_{j'}(\boldsymbol{\ell}')$ if $j' \neq j$ and $\boldsymbol{\ell}$ and $\boldsymbol{\ell}'$ are not in the same region. This design choice maintains long range spatial dependence along the same margin of \boldsymbol{w} , but reduces it across different margins.

M – the number of regions

In a meshed GP based on axis-parallel domain partitioning, the number of partitioning regions M is defined as $\prod_{j=1}^{d+1} M_j$ where M_j is the number of intervals into which axis j is partitioned. As seen above, M determines the number of Cholesky factors to compute, as well as the complexity of each factor computation (given by the size of the individual matrices). A small M corresponds to larger regions, which in turn will enlarge the parent sets of each block. For this reason, M should be chosen as small as possible, given the allotted computational budget.

\mathcal{S} and \mathcal{U}

The two sets of locations \mathcal{S} and \mathcal{U} differ in that the former are mapped to nodes \mathbf{A} of graph \mathcal{G} , whereas the latter to \mathbf{B} with no children. Since \mathbf{B} nodes are independent of each other given realizations of \mathbf{A} they provide little information on spatial dependence; in fact, the set \mathcal{S} is akin to the set of knots in a predictive process model. Unlike those models, \mathcal{S} can be chosen to be as large as the set of observations, and \mathcal{U} left empty. Generally, it is possible to choose $\mathcal{S} \supset \mathcal{T}$ to smoothly model spatial dependence across larger regions of unobserved locations, but convergence of the Gibbs sampler may be slower at locations that are very far from the nearest observed ones. Therefore, a convenient choice which is possibly safer for MCMC convergence is to let them coincide, i.e. $\mathcal{S} = \mathcal{T}$.

Matters are less straightforward when redundancies occur. Consider locations $\boldsymbol{t} = (t_1, \dots, t_d) \in \mathcal{D}$ such that $t_j = t_j^* \delta_j$ where $t_j^* = 1, \dots, N_j$ and $\delta_j \in [0, 1/N_j]$ for $j = 1, \dots, d$. Denote the set of such locations as \mathcal{T}^* ; this is a $N_1 \times \dots \times N_d$ grid of equally-spaced locations. If observed locations are on a quantized grid of coordinates, i.e. there are $\{t_j^*, \delta_j\}_{j=1}^d$ such that $\mathcal{T} \subset \mathcal{T}^*$, choosing $\mathcal{S} = \mathcal{T}$ may be inefficient as it may reduce the number of redundant \mathbf{R}_j 's and thus increase computation time. Instead, choosing \mathcal{S} such that $\mathcal{T} \subset \mathcal{S} \subset \mathcal{T}^*$ (possibly $\mathcal{S} = \mathcal{T}^*$) may be much more efficient.

In Figure 1 we show how a small subsample of a time slice of the Serengeti data. In that case, we extend \mathcal{S} to the whole domain (Fig. 1, center) or to extend it just enough to cover all observed locations (Fig. 1, right). For predictions, in the former case we use samples from the full conditional distribution $p(\mathbf{w}_j \mid \mathbf{w}_{-j}, \mathcal{D})$. In the latter case, we map empty blocks to nodes \mathbf{B} : predictions thus easily proceed in parallel since nodes \mathbf{B} are independent of each other given nodes in \mathbf{A} , with $\text{Pa}[\mathbf{b}]$ being composed of the nearest \mathbf{A} nodes along each axis-parallel direction.

D Application: Q-MGP on multivariate response

We consider the multivariate regression model we previously defined:

$$\mathbf{y}(\boldsymbol{\ell}) = \mathbf{X}(\boldsymbol{\ell})^\top \boldsymbol{\beta} + \mathbf{Z}(\boldsymbol{\ell})^\top \mathbf{w}(\boldsymbol{\ell}) + \boldsymbol{\varepsilon}(\boldsymbol{\ell}), \quad (1)$$

with $\boldsymbol{\ell} \in \mathbb{R}^d$ and $\boldsymbol{\varepsilon}(\boldsymbol{\ell}) \stackrel{iid}{\sim} N(0, \mathbf{D})$ and $\mathbf{D} = \text{diag}(\tau_1^2, \dots, \tau_l^2)$, and where $\mathbf{y}(\boldsymbol{\ell}) \in \mathbb{R}^l$ is the multivariate point-referenced outcome. In this section, we specifically consider the case $l = q$, $\mathbf{Z}(\boldsymbol{\ell}) = I_q$, and drop \mathbf{X} and $\boldsymbol{\beta}$ from the model, resulting in

$$\mathbf{y}(\boldsymbol{\ell}) = \mathbf{w}(\boldsymbol{\ell}) + \boldsymbol{\varepsilon}(\boldsymbol{\ell}), \quad (2)$$

where some of the q components of $\mathbf{y}(\boldsymbol{\ell})$ may be unobserved. We construct a cross-covariance function based on latent distances among the q variables. To do so, we set

$$C(\mathbf{h}, \mathbf{v}) = \frac{\exp\{-\phi\|\mathbf{h}\| / \exp\{\frac{1}{2}\beta \log(1 + \alpha\|\mathbf{v}\|)\}\}}{\exp\{\beta \log(1 + \alpha\|\mathbf{v}\|)\}},$$

and for $j = 1, \dots, q$, $C_j(\mathbf{h}) = \exp\{-\phi_j\|\mathbf{h}\|\}$. We then model $\mathbf{w}(\boldsymbol{\ell})$ as a Q-MGP with base cross-covariance

$$C_{ij}(\mathbf{h}, \mathbf{v}) = \begin{cases} \sigma_{i1}^2 C(\mathbf{h}, \mathbf{v}) + \sigma_{i2}^2 C_i(\mathbf{h}) & \text{if } i = j \\ \sigma_{i1} \sigma_{j1} C(\mathbf{h}, \mathbf{v}) & \text{if } i \neq j, \end{cases} \quad (3)$$

which is a valid cross-covariance derived from eq. (7) in [Apanasovich and Genton \(2010\)](#).

D.1 Synthetic data

We generate multivariate outcomes with $q = 3$ at $n = 6,400$ spatial locations on a regular grid on $\mathcal{D} = [0, 1]^2$ for a total sample size of 19,200, sampling the latent effects from a

Gaussian Process with cross-covariance (3) and setting the parameters as

$$\begin{aligned}
 \sigma_{11} &= 1 & \sigma_{21} &= 1.5 & \sigma_{31} &= 0.5 \\
 \sigma_{12} &= 0.1 & \sigma_{22} &= 0.05 & \sigma_{32} &= 0.1 \\
 \phi_1 &= 0.5 & \phi_2 &= 2 & \phi_3 &= 1 \\
 \alpha &= 1 & \beta &= 0.9 & \phi &= 5 \\
 \|\mathbf{v}_{12}\| &= 1 & \|\mathbf{v}_{13}\| &= 1 & \|\mathbf{v}_{23}\| &= 0.05,
 \end{aligned}$$

with measurement errors fixed at $\tau_1^2 = 0.01, \tau_2^2 = 0.05, \tau_3^2 = 0.1$. For predictions, we set $y_1(\ell) = \text{NA}$ if $\ell \in [0.2, 0.7] \times [0.4, 0.7]$, $y_2(\ell) = \text{NA}$ if $\ell \in [0.4, 0.9] \times [0.2, 0.6]$, and $y_3(\ell) = \text{NA}$ if $\ell \in [0.2, 0.7] \times [0.1, 0.5]$ be observed at all 10,000 spatial locations. This setup will highlight the advantages of a multivariate approach given the small latent distance between the second and third margin (given by $\|\mathbf{v}_{23}\|$), and the non-separability of the covariance function on the variables, as given by β . We fit a Q-MGP model on $M = 25^2$ regions. We compare the predictions of this multivariate specification with those obtained from univariate Q-MGP models applied on each margin independently. Domain partitioning for the univariate models along each margin was chosen in order to match the total computing time of the multivariate model. We run MCMC for 7000 iterations, of which 5000 discarded as burn-in and thinning the chain with a 1:1 ratio resulting in an approximated posterior sample of size 1000. The results are reported in Table 1, whereas Figure 2 shows the data along with the results from the two approaches.

	MAE	RMSE	Coverage	Run Time
Q-MGP (multivariate)	0.4439	0.3200	96.39	1.04 min
Q-MGP (univariate)	0.5056	0.4297	95.14	1.11 min

Table 1: Multivariate and univariate Q-MGPs compared on multivariate synthetic data.

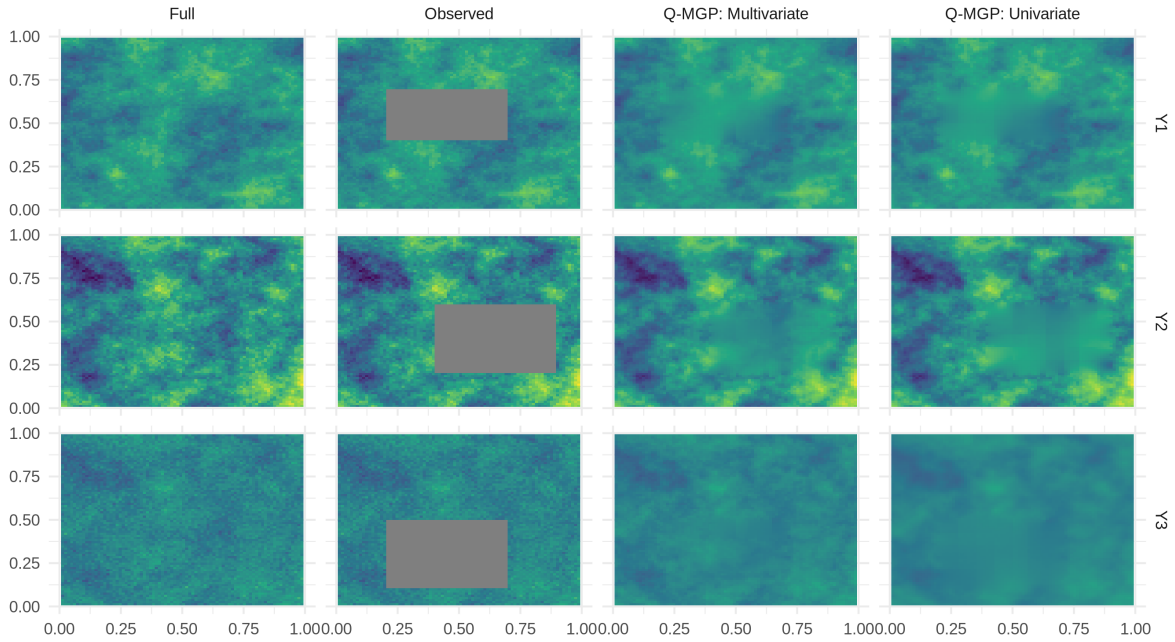


Figure 2: Simulated data and predictions from multivariate and univariate Q-MGPs.

D.2 Spatial multivariate analysis of NDVI in the Serengeti region

We consider the Serengeti dataset previously described; in this section, we build a spatial dataset by considering the NDVI measurement on December 17, 2016, along with elevation and soil moisture at the same spatial locations. The resulting dataset corresponds to $q = 3$ variables and $n = 1,000,000$ spatial locations, for a total size of 3,000,000; cloud cover obstructed the sensor view, resulting in NDVI being missing at 322,555 spatial locations. The goal of is thus to fill those gaps. Unlike in the previous analysis, we are not interested in mapping the varying effect of covariates on the outcome. Instead, we are modeling all variables jointly, aiming to recover the conditional distribution of NDVI given the others. We ran a total of 20000 MCMC iterations, thinning the last 5000 at a 4:1 ratio to eventually obtain a posterior sample of size 1000. Run-time averaged about 1 second per iteration using 50 threads.

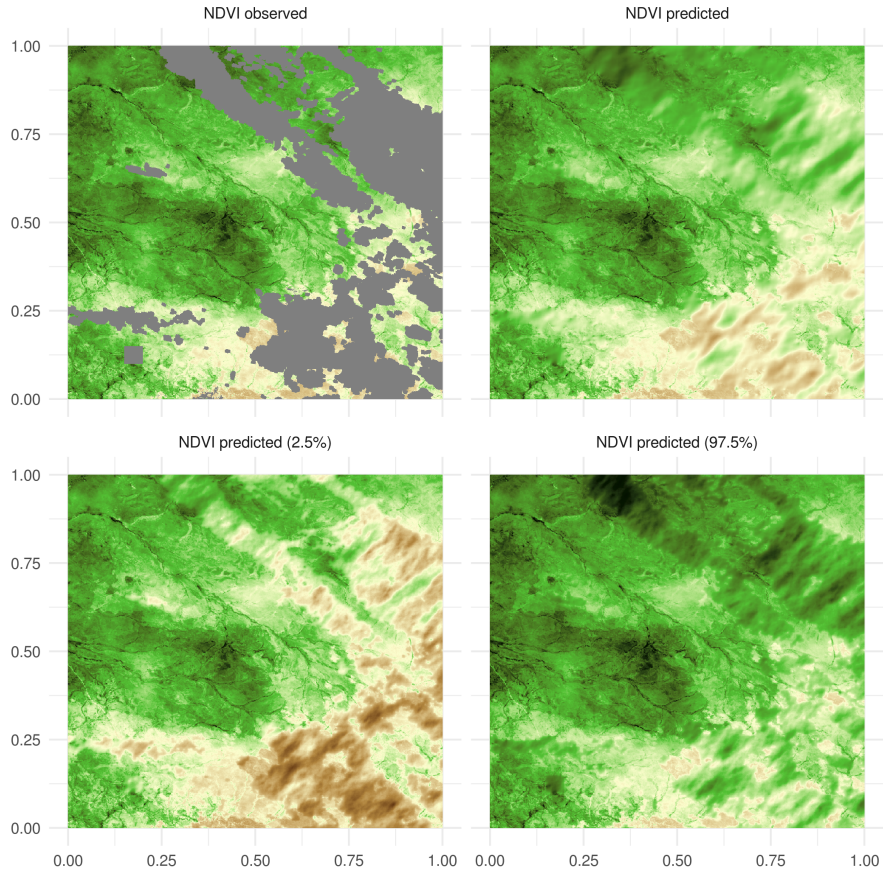


Figure 3: NDVI data and multivariate Q-MGP predictions for the Serengeti region on 2016-12-17. 2.5% and 97.5% quantiles are in the bottom row.

θ					
σ_{11}	3.1889	(3.1965, 3.2049)	σ_{12} 0.0081	(0.0081, 0.0082)	
σ_{13}	0.3015	(0.3019, 0.3023)	σ_{21}	7.1617	(7.1846, 7.2075)
σ_{22}	0.2697	(0.2701, 0.2705)	σ_{23}	0.2731	(0.2738, 0.2747)
ϕ_1	26.5273	(26.6811, 26.8183)	ϕ_2	3.4674	(3.4781, 3.4867)
ϕ_3	0.7512	(0.7520, 0.7527)	α	3.1094	(3.1153, 3.1215)
β	0.0024	(0.0024, 0.0024)	ϕ	1.0615	(1.0632, 1.0651)
v_{12}	0.6281	(0.6292, 0.6306)			
v_{13}	0.2997	(0.3008, 0.3019)	v_{23}	1.0079	(1.0092, 1.0107)

Table 2: Posterior means with 95% credible intervals for the covariance parameters.

E Comparisons with methods in Heaton et al. (2019)

In recent work, Heaton et al. (2019) have reviewed and compared 13 state-of-the-art models for large spatial datasets in a predictive challenge involving (1) simulated and (2) real-world spatial data (daytime land surface temperatures as measured by the Terra instrument onboard the MODIS satellite on August 4, 2016, Level-3 data). The two datasets are available at github.com/finnlindgren/heatoncomparison. The total number of locations is the same in both cases ($n_{\text{all}} = 150,000$), with the goal of predicting outcomes when missing. Both datasets include $n_{\text{train}} = 105,569$ available locations; the test set is of size $n_{\text{test}} = 44,431$ for the simulated data, $n_{\text{test}} = 42,740$ for the MODIS data due to cloud cover.

We estimate two MGP models for each dataset by partitioning the spatial domain into $M = 1500$ (resp. 375) rectangular regions for an average block size of 100 (resp. 400) spatial locations, and fix \mathcal{G} as a cubic mesh. We assign blocks to \mathcal{S} to cover the observed locations; the remaining ones will be used for prediction as in the right plot of Figure 1. Finally, we use an exponential base covariance function. We ran our Gibbs sampler using $\sigma^2 \sim \text{Inv.Gamma}(2.01, 1)$, $\tau^2 \sim \text{Inv.Gamma}(2.01, 1)$, $\phi \sim U[1/10, 30]$ as priors, and respectively 10, 1, 10 as starting values. A log-Normal proposal function was used for the Metropolis update of ϕ . We ran a total of 6,000 Monte Carlo iterations, with 4,000 dropped as burn-in, and thinning 2:1 the remaining ones to obtain a sample of size 1,000 that approximates the joint posterior distribution $p(\mathbf{w}, \phi, \sigma^2, \tau^2 \mid \mathbf{y})$ which we then use to obtain predictions. In both cases, the algorithm ran on 40 threads. We report the results of our analysis at the top of Table 3 along with results from select models tested in Heaton et al. (2019) to facilitate comparisons, whereas prediction plots for both datasets are in Figure 4.

Refer to the cited work for a more in-depth overview. We selected: nearest-neighbor Gaussian process (NNGP), conjugate (Finley et al., 2019) and response (Datta et al., 2016) algorithms; the multiresolution approximation (MRA) model of Katzfuss (2017); a stochastic partial differential equations (SPDE) approach estimated via integrated nested Laplace approximations (Rue et al., 2009; Lindgren et al., 2011); metakriging (Guhaniyogi and Banerjee, 2018).

In terms of predictive performance, coverage, and computation time, both implemented models are competitive with the top-ranking ones in Heaton et al. (2019). In particular, our




Figure 4: Spatial data and predicted temperatures for Q-MGP models. Top: grey areas correspond to missing observations. Middle rows: spatial predictions. Bottom: the true outcomes. Left: MODIS data; Right: simulated data.

Simulated data	MAE	RMSE	Coverage	Run Time	N. Cores	Code Lang.
Q-MGP, $M = 1500$	0.6270	0.8721	95.46	16.4 min	40	C++ (Armadillo)
Q-MGP, $M = 375$	0.6136	0.8407	95.62	133.2 min	40	C++ (Armadillo)
NNGP Conjugate	0.65	0.88	96	1.99 min	10	C++
NNGP Response	0.65	0.88	96	45.56 min	10	C++
MRA	0.61	0.83	93	13.57 min	1	Matlab
SPDE	0.62	0.86	100	138.34 min	2	R (in1a)
Metakriging	0.74	97	99	2888.89 min	30	R (spBayes)

MODIS data	MAE	RMSE	Coverage	Run Time	N. Cores	Code Lang.
Q-MGP, $M = 1500$	1.1151	1.5598	95.14	16.5 min	40	C++ (Armadillo)
Q-MGP, $M = 375$	1.0729	1.5034	95.20	133.7 min	40	C++ (Armadillo)
NNGP Conjugate	1.21	1.64	95	2.06 min	10	C++
NNGP Response	1.24	1.68	94	42.85 min	10	C++
MRA	1.33	1.85	92	15.61 min	1	Matlab
SPDE	1.10	1.53	97	120.33 min	2	R (in1a)
Metakriging	2.08	2.50	89	2888.52 min	30	R (spBayes)

Table 3: Comparative performance of Q-MGP and select methods in [Heaton et al. \(2019\)](#). All values related to methods other than Q-MGPs have been taken from Tables 2 and 3 of [Heaton et al. \(2019\)](#). Refer to the cited article for additional comparisons on the same datasets.

approach (with $M = 1500$) is much faster than Bayesian models estimated via MCMC – note that *NNGP Conjugate* uses cross-validation to select suitable covariance parameters. The slower alternative with $M = 375$ – corresponding to a much coarser partitioning, i.e. larger regions – was possible by caching the redundant matrix operations given that the data are on a regular lattice. We chose $M = 375$ by targeting a total computing time similar to the SPDE model, which achieved the lowest predictive error in the MODIS dataset. Computing times should not differ dramatically on a dual-CPU Intel Xeon server as in [Heaton et al. \(2019\)](#).

F C++ implementation of Q-MGP in meshgp

All applications were run on a 512GB-memory, dual-CPU server with two Intel Xeon E5-2699v3 processors, each with 18 cores at 2.3Ghz each, running Debian linux and R version 3.6 linked to the OpenBLAS libraries. Source code for implementing Q-MGP models is available as R package `meshgp` (github.com/mkln/meshgp). The `meshgp` package is written in Armadillo ([Sanderson and Curtin 2016](#), a C++ library for linear algebra) which can easily be linked to high performance libraries. Parallelization of all high-dimensional operations is implemented via OpenMP ([Dagum and Menon, 1998](#)). R packages `Rcpp` ([Eddelbuettel and François, 2011](#)) and `RcppArmadillo` ([Eddelbuettel and Sanderson, 2014](#)) are used for interfacing the C++ source with R.

G Caching algorithm

Input: $\Gamma = \{\Gamma_i\}_{i=1}^g$: a collection of matrices of size $n_{\Gamma_i} \times d$ for $i = 1, \dots, g$. Each row is denoted $\Gamma_i[\mathbf{r}, :]$ for $r = 1, \dots, n_{\Gamma_i}$.

for $i \in \{1, \dots, g\}$ **do**

- Sort Γ_i using column 1; resolve ties using columns 2 to d .
- Calculate $\Gamma_{0,i}$ as the matrix of size $n_{\Gamma_i} \times d$ such that

$$\Gamma_{0,i}[\mathbf{r}, :] = \Gamma_i[1, :] \text{ for } r = 1, \dots, n_{\Gamma_i}.$$
- Calculate $\tilde{\Gamma}_i = \Gamma_i - \Gamma_{0,i}$.

end

Initialize: $\Delta_\Gamma = \{\delta_1, \dots, \delta_g\}$ where $\delta_i = i$ for all $i = 1, \dots, g$.

for $i \in \{1, \dots, g\}$ **do**

- for** $j \in \{1, \dots, i\}$ **do**
 - if** $\tilde{\Gamma}_i = \tilde{\Gamma}_j$ (*element by element*) **then**
 - | Set $\delta_i = j$.

end

end

Output: Δ_Γ : a dictionary with g keys and g^* unique values. If f is a function such that $f(\Gamma + t) = f(\Gamma)$ then $f(\Gamma_{\delta_i}) = f(\Gamma_j)$ for all i such that $\delta_i = j$.

Algorithm 1: Caching algorithm.

Algorithm 1 can be used by taking $\Gamma = \{(\mathcal{S}_j, \mathcal{S}_{[j]})\}_{j=1:M}$ for $\mathbf{C}_{\mathcal{S}_j, \mathcal{S}_{[j]}}$, $\Gamma = \{\mathcal{S}_j\}_{j=1:M}$ for $\mathbf{C}_{\mathcal{S}_j}$, and $\Gamma = \{\mathcal{S}_{[j]}\}_{j=1:M}$ for $\mathbf{C}_{\mathcal{S}_{[j]}}^{-1}$. The resulting dictionaries can be used to build a dictionary $\Delta_{\mathbf{R}}$ for $\{\mathbf{R}_j\}_{j=1:M}$ with M^* unique values. Each iteration of the Gibbs sampler becomes cheaper if the number of unique values in $\Delta_{\mathbf{R}}$ is $M^* \ll M$.

H Compute time comparisons with NNGPs

Figure 5 shows the time-per-iteration of Q-MGP, on a fully-observed slice of the Serengeti data, $n = 1,000,000$, choosing $M \in \{150^2, 250^2\}$, and with or without caching. $M = 250^2$ corresponds to blocks with 16 locations (all blocks are of dimension 4×4), whereas $M = 150^2$ to blocks of dimension 36, 42, 49 corresponding to squares of size 6×6 , 6×7 or 7×6 , 7×7 , respectively. Caching is optional; in this case we showcase the speed differential between a favorable case (regular lattice) and the worst case (completely irregularly-spaced

observations) scenarios by disabling caching.




Figure 5: Computation time per iteration, in logarithmic scale, of Q-MGP and NNGP at different configurations.

We compare our results with the NNGP sequential sampler of [Datta et al. \(2016\)](#) as implemented in the `spNNGP` package for R, choosing $m \in \{10, 15\}$. This fully Bayesian NNGP recovers the spatial random effects and is also estimated via a Gibbs sampler, but unlike a Q-MGP, it does not use caching or sample \mathbf{w} in parallel. Computing times of Q-MGP models without caching are comparable to NNGP models; the Q-MGP model with block size 16 is about 30% faster than the NNGP model with 10 neighbors.

The optional caching feature results in a much smaller computational penalty when considering large regions. The large-region model is 25% slower than the small-region model when caching is activated, but 250% slower if caching is disabled because many more large matrix inverses are computed. The caching feature of Q-MGP models allowed us to design the best-performing model for the real-world application of [Heaton et al. \(2019\)](#), see [Table 3](#).

I Effective sample size of MCMC posterior samples

Markov-chain Monte Carlo algorithms output correlated samples from the posterior distribution, and their efficiency is decreased when successive samples are highly autocorrelated. The




Figure 6: Effective sample size (ESS) of Q-MGP and NNGP for different values of the spatial exponential decay parameter ϕ .

effective sample size (ESS) is the size of an independent sample equivalent to the correlated sample at hand. Highly autocorrelated Markov chains result in lower ESS. This scenario arises in geostatistical settings with large-scale spatial dependency, as the high-dimensional latent spatial random effects are highly correlated at nearby locations. For this reason, Gibbs samplers such as the one proposed in this article and the sequential NNGP sampler of [Datta et al. \(2016\)](#) are expected to be negatively affected by slowly-decaying spatial correlation.

With this in mind, we now compare the parallel sampler of Q-MGPs to a sequential NNGP sampler (from R package `spNNGP`, [Finley et al. 2020](#)) in terms of ESS of the latent random effects. We generate $\mathbf{y} = \mathbf{Z}\mathbf{w} + \boldsymbol{\varepsilon}$ on a regular 80×80 grid where $\mathbf{Z} = \mathbf{I}_n$, $n = 6400$, the spatial process is $\mathbf{w} \sim GP(0, \mathbf{C})$, $\mathbf{C}(\|\mathbf{h}\|) = \sigma^2 \exp\{-\phi\|\mathbf{h}\|\}$ for $\mathbf{h} = \boldsymbol{\ell}' - \boldsymbol{\ell}$, $\boldsymbol{\varepsilon}(\boldsymbol{\ell}) \sim N(0, \tau^2)$ for all $\boldsymbol{\ell}$ and using parameters $\sigma^2 = 1$ and $\phi \in \{0.01, 0.025, 0.05, .1, .2, .4, .8, 1.6, 3.2, 6.4, 12.8\}$, and nugget $\tau^2 = 0.01$. The goal is to sample \mathbf{w} a posteriori. We assign priors $\phi \sim U[.01, 30]$, $\sigma^2 \sim Inv.G.(2.01, 1)$, $\tau^2 \sim Inv.G.(2.01, 1)$ and use the true values of $\{\sigma^2, \phi, \tau^2\}$ as starting values for the Gibbs sampler of both NNGP and Q-MGP. The starting value of \mathbf{w} was set to a vector of zeros of size n . We let the Markov chain run for a total of 2000 iterations, discarding the first 1000. The NNGP model used $m = 10$ neighbors, whereas the Q-MGP model used $M = 270$ regions obtained by axis-parallel partitioning the spatial domain along the two dimensions into 15 and 18 intervals, respectively. We chose $M = 270$ to achieve about the same computation time when caching was disabled (actual Q-MGP run times were about 2.6 to 3 times faster than NNGP-sequential once caching was turned on). Figure 6 shows the effective sample size of the resulting approximated posterior sample of \mathbf{w} , averaged across all

spatial locations. While both models evidence a degradation of performance with long-range spatial dependence (small ϕ), the NNGP model shows significantly worse performance. With low spatial decay, the Q-MGP model exhibits effective sample sizes up to 2.5 times larger than the equivalent NNGP model. The smaller effective sample size of NNGPs implies a longer effective runtime to achieve similar performance. Coupled with speed advantages due to caching, we estimate Q-MGP models to result in an effective improvement in compute times up to one order of magnitude over sequential NNGPs.

J Using MGPs with $d = 1$



Figure 7: Graph \mathcal{G} for interval-MGP in unidimensional domains.

MGPs can be implemented on unidimensional domain by splitting the real line into disjoint intervals, then choosing a finite reference set of locations $\mathcal{S} \subset \mathfrak{R}$, linking each of them to a reference node based on the intervals they falls into, and using the DAG (e.g. see Figure 7) to establish a process at all other locations. Grouping the reference locations is not an appealing strategy if the chosen covariance exhibits the screening effect (see e.g. Stein (2002) for a more general discussion), in which case there may be numerical errors in calculating \mathbf{R}_j . In fact, suppose $\ell < \ell'$ for all $\ell \in \mathcal{S}_j$ and $\ell \in \mathcal{S}_{j'}$ and $\mathbf{a} \rightarrow \mathbf{a}'$ in \mathcal{G} . If \mathbf{C} exhibits the screening effect, $p(\mathbf{w}_{\mathcal{S}'} | \mathbf{w}_{\mathcal{S}}) = p(\mathbf{w}_{\mathcal{S}'} | \mathbf{w}(\ell^*))$ where $\ell^* = \max_{\ell}(\mathcal{S})$. Therefore, one must either choose a valid covariance function, or reduce the size of the reference sets to 1.

K Recovering missing pixels of an animated GIF image

Lion GIF	MAE	RMSE	Coverage (90%)
Q-MGP, $M = 1050$	0.0341	0.0715	93.28
Gapfill	0.0480	0.0923	62.76

Table 4: Comparative performance of Q-MGP and Gapfill on recovering missing pixels of an animated GIF image.

We compare Q-MGPs with Gapfill (Gerber et al., 2018) on non-Gaussian data in the form of an animated GIF image. The original GIF image collects 30 frames, each of size 200×250 , for a total data size of 1.5 million locations. We subsample the data size to $n_{\text{all}} = 94,500$ by downsampling each of the 30 frames at a 16:1 ratio. We simulate cloud cover by placing random clouds in 5 of the 30 frames, covering all but 10 locations in 5 of 30 frames, and covering 50% random locations in 5 of the 30 frames; 15 frames were thus untouched. See Figure 8. The resulting dataset includes $n_{\text{obs}} = 67,363$ non-missing observations and $n_{\text{test}} = 27,137$ missing ones to be predicted. We implement a Q-MGP model with $M = 1050$ found by partitioning the spatial coordinates into 5 and 21 intervals, and time in 10 intervals. We run the Gibbs sampler for 15000 iterations; we discard the first 10000 and keep one every five of the remaining 5000 to obtain a sample of size 1000 from the approximated posterior distribution, which we use for predictions. We used the same base covariance, prior distributions, and starting values as in previous sections, for a total run time of 35 minutes (0.14s/iteration) on 11 threads. As seen in Table 4 and from a visual inspection of Figure 8, Q-MGP outperforms Gapfill.




Figure 8: From the top: full original image frames to be recovered, predictions by Q-MGP and Gapfill, and observed image frames, at three select times (of 30 total) of the GIF image.

References

- Apanasovich, T. V. and Genton, M. G. (2010). Cross-covariance functions for multivariate random fields based on latent dimensions. *Biometrika*, 97:15–30. doi:10.1093/biomet/asp078. 7
- Cowell, R. G., Dawid, A. P., Lauritzen, S. L., and Spiegelhalter, D. J. (1999). *Probabilistic Networks and Expert Systems*. Springer-Verlag, New York. 3
- Dagum, L. and Menon, R. (1998). OpenMP: an industry standard api for shared-memory programming. *Computational Science & Engineering, IEEE*, 5(1):46–55. 14
- Datta, A., Banerjee, S., Finley, A. O., and Gelfand, A. E. (2016). Hierarchical nearest-neighbor gaussian process models for large geostatistical datasets. *Journal of the American Statistical Association*, 111:800–812. doi:10.1080/01621459.2015.1044091. 4, 11, 16, 17
- Eddelbuettel, D. and François, R. (2011). Rcpp: Seamless R and C++ integration. *Journal of Statistical Software*, 40(8):1–18. doi:10.18637/jss.v040.i08. 14
- Eddelbuettel, D. and Sanderson, C. (2014). RcppArmadillo: Accelerating R with high-performance C++ linear algebra. *Computational Statistics and Data Analysis*, 71:1054–1063. doi:10.1016/j.csda.2013.02.005. 14
- Finley, A. O., Datta, A., and Banerjee, S. (2020). R package for Nearest Neighbor Gaussian Process models. arXiv:2001.09111. 17
- Finley, A. O., Datta, A., Cook, B. D., Morton, D. C., Andersen, H. E., and Banerjee, S. (2019). Efficient algorithms for Bayesian nearest neighbor Gaussian processes. *Journal of Computational and Graphical Statistics*, 28:401–414. doi:10.1080/10618600.2018.1537924. 11
- Gerber, F., Furrer, R., Schaepman-Strub, G., de Jong, R., and Schaepman, M. E. (2018). Predicting missing values in spatio-temporal remote sensing data. *IEEE Transactions on Geoscience and Remote Sensing*, 56(5):2841–2853. doi:10.1109/TGRS.2017.2785240. 19
- Guhaniyogi, R. and Banerjee, S. (2018). Meta-kriging: Scalable bayesian modeling and inference for massive spatial datasets. *Technometrics*, 60(4):430–444. doi:10.1080/00401706.2018.1437474. 11

- Heaton, M. J., Datta, A., Finley, A. O., Furrer, R., Guinness, J., Guhaniyogi, R., Gerber, F., Gramacy, R. B., Hammerling, D., Katzfuss, M., Lindgren, F., Nychka, D. W., Sun, F., and Zammit-Mangion, A. (2019). A case study competition among methods for analyzing large spatial data. *Journal of Agricultural, Biological and Environmental Statistics*, 24(3):398–425. [doi:10.1007/s13253-018-00348-w](https://doi.org/10.1007/s13253-018-00348-w). 11, 13, 14, 16
- Katzfuss, M. (2017). A multi-resolution approximation for massive spatial datasets. *Journal of the American Statistical Association*, 112:201–214. [doi:10.1080/01621459.2015.1123632](https://doi.org/10.1080/01621459.2015.1123632). 11
- Lindgren, F., Rue, H., and Lindström, J. (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B*, 73:423–498. [doi:10.1111/j.1467-9868.2011.00777.x](https://doi.org/10.1111/j.1467-9868.2011.00777.x). 11
- Rue, H., Martino, S., and Chopin, N. (2009). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of the Royal Statistical Society: Series B*, 71:319–392. [doi:10.1111/j.1467-9868.2008.00700.x](https://doi.org/10.1111/j.1467-9868.2008.00700.x). 11
- Sanderson, C. and Curtin, R. (2016). Armadillo: a template-based C++ library for linear algebra. *Journal of Open Source Software*, 1:26. 14
- Stein, M. L. (2002). The screening effect in kriging. *The Annals of Statistics*, 30(1):298–323. [doi:10.1214/aos/1015362194](https://doi.org/10.1214/aos/1015362194). 18
- Stein, M. L. (2014). Limitations on low rank approximations for covariance matrices of spatial data. *Spatial Statistics*, 8:1–19. [doi:doi:10.1016/j.spasta.2013.06.003](https://doi.org/10.1016/j.spasta.2013.06.003). 5