

UCLA

Research Reports

Title

High-dimensional Multivariate Geostatistics: A Conjugate Bayesian Matrix-Normal Approach

Permalink

<https://escholarship.org/uc/item/58c1r34h>

Authors

Zhang, Lu

Banerjee, Sudipto

Finley, Andrew O

Publication Date

2021-02-05

High-dimensional Multivariate Geostatistics: A Conjugate Bayesian Matrix-Normal Approach

LU ZHANG

UCLA DEPARTMENT OF BIOSTATISTICS

Lu.Zhang@ucla.edu

SUDIPTO BANERJEE

UCLA DEPARTMENT OF BIOSTATISTICS

sudipto@ucla.edu

ANDREW O. FINLEY

MICHIGAN STATE UNIVERSITY DEPARTMENTS OF FORESTRY AND GEOGRAPHY

finleya@msu.edu

May 20, 2020

Abstract

Joint modeling of spatially-oriented dependent variables are commonplace in the environmental sciences, where scientists seek to estimate the relationships among a set of environmental outcomes accounting for dependence among these outcomes and the spatial dependence for each outcome. Such modeling is now sought for very large data sets where the variables have been measured at a very large number of locations. Bayesian inference, while attractive for accommodating uncertainties through their hierarchical structures, can become computationally onerous for modeling massive spatial data sets because of their reliance on iterative estimation algorithms. This manuscript develops a conjugate Bayesian framework for analyzing multivariate spatial data using analytically tractable posterior distributions that do not require iterative algorithms. We discuss differences between modeling the multivariate response itself as a spatial process and that of modeling a latent process. We illustrate the computational and inferential benefits of these models using simulation studies and real data analyses for a Vege Indices dataset with observed locations numbering in the millions.

Key words: Conjugate Bayesian multivariate regression; Multivariate spatial processes; Matrix-variate normal and inverse-Wishart distributions; Nearest-Neighbor Gaussian processes.

I. INTRODUCTION

Analysis for environmental data sets often require joint modeling of multiple spatially dependent variables accounting for dependence among the variables and the spatial association for each variable. For point-referenced variables, multivariate Gaussian processes (GPs) serve as versatile tools for joint modeling of spatial variables (see, e.g., Schabenberger and Gotway, 2004; Cressie and Wikle, 2015; Banerjee et al., 2014, and references therein). However, for a dataset with n observed locations, fitting a GP based spatial model typically requires floating point operations (flops) and memory requirements of the order $\sim n^3$ and $\sim n^2$, respectively. This is challenging when n is large. This “Big Data” problem has received much attention in the literature and a comprehensive review is beyond the scope of this article; see, e.g., Banerjee (2017); Heaton et al. (2019); Sun et al. (2011) for a review and comparison of scalable modeling methods. Much of the aforementioned literature for scalable models focused on univariate spatial processes, i.e., assuming only one response for each location.

Multivariate processes (see, e.g., Genton and Kleiber, 2015; Salvaña and Genton, 2020; Le and Zidek, 2006, and references therein), has received relatively limited developments in the context of massive data. Bayesian models are attractive for inference on multivariate spatial processes because they can accommodate uncertainties in the process parameters more flexibly through their hierarchical structure. Multivariate spatial interpolation using conjugate Bayesian modeling can be found in Brown et al. (1994); Le et al. (1997); Sun et al. (1998); Le et al. (2001); Gamerman and Moreira (2004), but these methods do not address the challenges encountered in massive data sets. More flexible methods for joint modeling, including spatial factor models, have been investigated in Bayesian contexts (see, e.g. Schmidt and Gelfand, 2003; Ren and Banerjee, 2013; Taylor-Rodriguez et al., 2019), but these methods have focused upon delivering full Bayesian inference through iterative algorithms such as Markov chain Monte Carlo (MCMC).

Our current contribution is an augmented Bayesian multivariate linear model framework that accommodates conjugate distribution theory, similar to Gamerman and Moreira (2004), but that can scale up to massive data sets with locations numbering in the millions. More specifically, we embed the Nearest-Neighbor Gaussian process (NNGP) (Datta et al., 2016a) within our conjugate Bayesian framework. We will consider two classes of models. The first is obtained by modeling the spatially dependent variables jointly as a multivariate spatial process, while the second models a latent multivariate spatial process in a hierarchical setup. We refer to the former as the “response” model and the latter as the “latent” model and we explore some properties of these models.

The remainder of our paper is arranged as follows. Section II develops a conjugate Bayesian multivariate spatial regression framework using Matrix-Normal and Inverse-Wishart prior distributions. We first develop two classes of models, response models and latent models using Gaussian spatial processes, in Section I. Subsequently, in Section II we develop scalable versions of these models using the Nearest Neighbor Gaussian process (NNGP). We develop NNGP response models and NNGP latent models in this conjugate Bayesian framework. A cross-validation algorithm to fix certain hyperparameters in these models is presented in Section III and some theoretical attributes of these models are presented in Section IV. Section III present some simulation experiments, while Section IV analyzes a massive Normalized Difference Vegetation Index data with a few million locations. Finally, Section V concludes the manuscript with some discussion.

II. BAYESIAN MULTIVARIATE GEOSTATISTICAL MODELING

I. Conjugate Multivariate Spatial Models

Conjugate Multivariate Response Model Let $\mathbf{y}(\mathbf{s}) = (y_1(\mathbf{s}), \dots, y_q(\mathbf{s}))^\top \in \mathbb{R}^q$ be a $q \times 1$ vector of outcomes at location $\mathbf{s} \in \mathcal{D} \subset \mathbb{R}^d$ and $\mathbf{x}(\mathbf{s}) = (x_1(\mathbf{s}), \dots, x_p(\mathbf{s}))^\top \in \mathbb{R}^p$ be a corresponding $p \times 1$ vector of explanatory variables. Conditional on these explanatory variables, the response is assumed to follow a multivariate Gaussian process,

$$\mathbf{y}(\mathbf{s}) \sim \text{GP}(\boldsymbol{\beta}^\top \mathbf{x}(\mathbf{s}), \mathbf{C}(\cdot, \cdot)); \quad \mathbf{C}(\mathbf{s}, \mathbf{s}') = [\rho_\psi(\mathbf{s}, \mathbf{s}') + (\alpha^{-1} - 1)\delta_{\mathbf{s}=\mathbf{s}'}]\boldsymbol{\Sigma}, \quad (1)$$

where the mean of $\mathbf{y}(\mathbf{s})$ is $\boldsymbol{\beta}^\top \mathbf{x}(\mathbf{s})$, $\boldsymbol{\beta}$ is a $p \times q$ matrix of regression coefficients and $\mathbf{C}(\mathbf{s}, \mathbf{s}') = \{\text{cov}\{y_i(\mathbf{s}), y_j(\mathbf{s}')\}\}$ is a $q \times q$ cross-covariance matrix (Genton and Kleiber, 2015) whose (i, j) -th element is the covariance between $y_i(\mathbf{s})$ and $y_j(\mathbf{s}')$. The cross-covariance matrix is defined for each pair of locations and is further specified as a multiple of a nonspatial positive definite matrix $\boldsymbol{\Sigma}$. The multiplication factor is a function of the two locations and is composed of two components: a spatial correlation function, $\rho_\psi(\mathbf{s}, \mathbf{s}')$, which introduces spatial dependence between the outcomes through hyperparameters ψ , and a micro-scale adjustment $(1/\alpha - 1)\delta_{\mathbf{s}=\mathbf{s}'}$, where $\delta_{\mathbf{s}=\mathbf{s}'} = 1$ if $\mathbf{s} = \mathbf{s}'$ and is 0 if $\mathbf{s} \neq \mathbf{s}'$, and $\alpha \in (0, 1]$ is a scalar parameter representing the overall strength of the spatial variability as a proportion of the total variation.

The covariance among the elements of $\mathbf{y}(\mathbf{s})$ within a location \mathbf{s} is given by the elements of $\mathbf{C}(\mathbf{s}, \mathbf{s}) = (1/\alpha)\boldsymbol{\Sigma}$ suggesting the interpretation of $\boldsymbol{\Sigma}$ as the within-location (nonspatial) dependence among the outcomes adjusted by a scale of $1/\alpha$ to accommodate additional variation at local scales. The interpretation of α is equivalent to the ratio of the ‘‘partial sill’’ to the ‘‘sill’’ in classical geostatistics. For example, in the special case that $\boldsymbol{\Sigma} = \sigma^2 \mathbf{I}_q$, $\text{cov}\{y_i(\mathbf{s}), y_j(\mathbf{s}')\} = \sigma^2 \rho(\mathbf{s}, \mathbf{s}') + \sigma^2(1/\alpha - 1)\delta_{\mathbf{s}=\mathbf{s}'}$, which shows that $\sigma^2(1/\alpha - 1) = \tau^2$ is the variance of micro-scale processes (or the ‘‘nugget’’), so that $\alpha = \sigma^2 / (\sigma^2 + \tau^2)$ is the ratio of the spatial variance (partial sill) to the total variance (sill). A similar interpretation for α results in the univariate setting with $q = 1$.

Let $\mathcal{S} = \{\mathbf{s}_1, \dots, \mathbf{s}_n\} \subset \mathcal{D}$ be a set of n locations yielding observations on $\mathbf{y}(\mathbf{s})$. Then $\mathbf{Y} = \mathbf{y}(\mathcal{S}) = [\mathbf{y}(\mathbf{s}_1) : \dots : \mathbf{y}(\mathbf{s}_n)]^\top$ is $n \times q$ and $\mathbf{X} = \mathbf{x}(\mathcal{S}) = [\mathbf{x}(\mathbf{s}_1) : \dots : \mathbf{x}(\mathbf{s}_n)]^\top$ is the corresponding $n \times p$ matrix of explanatory variables observed over \mathcal{S} . We will assume that \mathbf{X} has full column rank ($= p < n$). The likelihood emerging from (1) is $\mathbf{Y} | \boldsymbol{\beta}, \boldsymbol{\Sigma} \sim \text{MN}_{n,q}(\mathbf{X}\boldsymbol{\beta}, \boldsymbol{\mathcal{K}}, \boldsymbol{\Sigma})$, where MN denotes the Matrix-Normal distribution defined in Ding and Cook (2014), i.e.,

$$\text{MN}_{n,q}(\mathbf{Y} | \mathbf{X}\boldsymbol{\beta}, \boldsymbol{\mathcal{K}}, \boldsymbol{\Sigma}) = \frac{\exp\left(-\frac{1}{2} \text{tr}\left[\boldsymbol{\Sigma}^{-1}(\mathbf{Y} - \mathbf{X}\boldsymbol{\beta})^\top \boldsymbol{\mathcal{K}}^{-1}(\mathbf{Y} - \mathbf{X}\boldsymbol{\beta})\right]\right)}{(2\pi)^{np/2} |\boldsymbol{\Sigma}|^{n/2} |\boldsymbol{\mathcal{K}}|^{p/2}}, \quad (2)$$

where tr denotes trace, $\boldsymbol{\mathcal{K}} = \boldsymbol{\rho}_\psi + (\alpha^{-1} - 1)\mathbf{I}_n$ and $\boldsymbol{\rho}_\psi = \{\rho_\psi(\mathbf{s}_i, \mathbf{s}_j)\}$ is the $n \times n$ spatial correlation matrix. A conjugate Bayesian model is obtained by a Matrix-Normal-Inverse-Wishart (MNIW) prior on $\{\boldsymbol{\beta}, \boldsymbol{\Sigma}\}$, which we denote as

$$\text{MNIW}(\boldsymbol{\beta}, \boldsymbol{\Sigma} | \boldsymbol{\mu}_\beta, \mathbf{V}_r, \boldsymbol{\Psi}, \nu) = \text{IW}(\boldsymbol{\Sigma} | \boldsymbol{\Psi}, \nu) \times \text{MN}_{p,q}(\boldsymbol{\beta} | \boldsymbol{\mu}_\beta, \mathbf{V}_r, \boldsymbol{\Sigma}), \quad (3)$$

where $\text{IW}(\boldsymbol{\Sigma} | \cdot, \cdot)$ is the inverted-Wishart distribution with parameters ν and $\boldsymbol{\Psi}$ describing the degrees of freedom and the scale matrix (see section 3.6 in Gelman et al., 2013). The MNIW family is a conjugate prior with respect to the likelihood (2) and, for any fixed values of α , ψ and the hyperparameters in the prior density, we obtain the posterior density

$$p(\boldsymbol{\beta}, \boldsymbol{\Sigma} | \mathbf{Y}) \propto \text{MNIW}(\boldsymbol{\beta}, \boldsymbol{\Sigma} | \boldsymbol{\mu}_\beta, \mathbf{V}_r, \boldsymbol{\Psi}, \nu) \times \text{MN}_{n,q}(\mathbf{Y} | \mathbf{X}\boldsymbol{\beta}, \boldsymbol{\mathcal{K}}, \boldsymbol{\Sigma}) \propto \text{MNIW}(\boldsymbol{\mu}^*, \mathbf{V}^*, \boldsymbol{\Psi}^*, \nu^*), \quad (4)$$

where

$$\begin{aligned}\mathbf{V}^* &= (\mathbf{X}^\top \mathcal{K}^{-1} \mathbf{X} + \mathbf{V}_r^{-1})^{-1}, \quad \boldsymbol{\mu}^* = \mathbf{V}^* (\mathbf{X}^\top \mathcal{K}^{-1} \mathbf{Y} + \mathbf{V}_r^{-1} \boldsymbol{\mu}_\beta), \\ \boldsymbol{\Psi}^* &= \boldsymbol{\Psi} + \mathbf{Y}^\top \mathcal{K}^{-1} \mathbf{Y} + \boldsymbol{\mu}_\beta^\top \mathbf{V}_r^{-1} \boldsymbol{\mu}_\beta - \boldsymbol{\mu}^{*\top} \mathbf{V}^{*-1} \boldsymbol{\mu}^*, \quad \text{and } v^* = v + n.\end{aligned}\quad (5)$$

Direct sampling from the MNIW posterior distribution in (4) is achieved by first sampling $\boldsymbol{\Sigma} \sim \text{IW}(\boldsymbol{\Psi}^*, v^*)$ and then sampling one draw of $\boldsymbol{\beta} \sim \text{MN}_{p,q}(\boldsymbol{\mu}^*, \mathbf{V}^*, \boldsymbol{\Sigma})$ for each draw of $\boldsymbol{\Sigma}$. The resulting pairs $\{\boldsymbol{\beta}, \boldsymbol{\Sigma}\}$ will be samples from (4). Since this scheme draws directly from the posterior distribution, the sample is exact and does not require burn-in or convergence.

Turning to predictions, let $\mathcal{U} = \{\mathbf{u}_1, \dots, \mathbf{u}_{n'}\}$ be a finite set of locations where we intend to predict or impute the value of $\mathbf{y}(\mathbf{s})$ based upon an observed $n' \times p$ design matrix $\mathbf{X}_{\mathcal{U}} = [\mathbf{x}(\mathbf{u}_1) : \dots : \mathbf{x}(\mathbf{u}_{n'})]^\top$ for \mathcal{U} . If $\mathbf{Y}_{\mathcal{U}} = [\mathbf{y}(\mathbf{u}_1) : \dots : \mathbf{y}(\mathbf{u}_{n'})]^\top$ is the $n' \times q$ matrix of predictive random variables, then the conditional predictive distribution is

$$\begin{aligned}p(\mathbf{Y}_{\mathcal{U}} | \mathbf{Y}, \boldsymbol{\beta}, \boldsymbol{\Sigma}) &= \text{MN}_{n',q}(\mathbf{Y}_{\mathcal{U}} | \mathbf{X}_{\mathcal{U}} \boldsymbol{\beta} + \boldsymbol{\rho}_\psi(\mathcal{U}, \mathcal{S}) \mathcal{K}^{-1} [\mathbf{Y} - \mathbf{X} \boldsymbol{\beta}], \\ &\quad \boldsymbol{\rho}_\psi(\mathcal{U}, \mathcal{U}) + (\alpha^{-1} - 1) \mathbf{I}_{n'} - \boldsymbol{\rho}_\psi(\mathcal{U}, \mathcal{S}) \mathcal{K}^{-1} \boldsymbol{\rho}_\psi(\mathcal{S}, \mathcal{U}), \boldsymbol{\Sigma}),\end{aligned}\quad (6)$$

where $\boldsymbol{\rho}_\psi(\mathcal{U}, \mathcal{S}) = \{\rho_\psi(\mathbf{u}_i, \mathbf{s}_j)\}$ is $n' \times n$ and $\boldsymbol{\rho}_\psi(\mathcal{S}, \mathcal{U}) = \{\rho_\psi(\mathbf{s}_i, \mathbf{u}_j)\} = \boldsymbol{\rho}_\psi(\mathcal{U}, \mathcal{S})^\top$. Predictions can also be directly carried out in posterior predictive fashion, where we sample from

$$\begin{aligned}p(\mathbf{Y}_{\mathcal{U}} | \mathbf{Y}) &= \int \text{MN}_{n',q}(\mathbf{X}_{\mathcal{U}} \boldsymbol{\beta} + \boldsymbol{\rho}_\psi(\mathcal{U}, \mathcal{S}) \mathcal{K}^{-1} [\mathbf{Y} - \mathbf{X} \boldsymbol{\beta}], \\ &\quad \boldsymbol{\rho}_\psi(\mathcal{U}, \mathcal{U}) + (\alpha^{-1} - 1) \mathbf{I}_{n'} - \boldsymbol{\rho}_\psi(\mathcal{U}, \mathcal{S}) \mathcal{K}^{-1} \boldsymbol{\rho}_\psi(\mathcal{S}, \mathcal{U}), \boldsymbol{\Sigma}) \\ &\quad \times \text{MNIW}(\boldsymbol{\mu}^*, \mathbf{V}^*, \boldsymbol{\Psi}^*, v^*) d\boldsymbol{\beta} d\boldsymbol{\Sigma}.\end{aligned}\quad (7)$$

Sampling from (7) is achieved by drawing one $\mathbf{Y}_{\mathcal{U}}$ from (6) for each posterior draw of $\{\boldsymbol{\beta}, \boldsymbol{\Sigma}\}$.

Conjugate Multivariate Latent Model We now discuss a conjugate Bayesian model for a latent process. Consider the spatial regression model

$$\mathbf{y}(\mathbf{s}) = \boldsymbol{\beta}^\top \mathbf{x}(\mathbf{s}) + \boldsymbol{\omega}(\mathbf{s}) + \boldsymbol{\epsilon}(\mathbf{s}), \quad \mathbf{s} \in \mathcal{D}, \quad (8)$$

where $\boldsymbol{\omega}(\mathbf{s}) \sim \text{GP}(\mathbf{0}_{q \times 1}, \rho_\psi(\cdot, \cdot) \boldsymbol{\Sigma})$ is a $q \times 1$ multivariate latent process with cross-covariance matrix $\rho_\psi(\mathbf{s}, \mathbf{s}') \boldsymbol{\Sigma}$ and $\boldsymbol{\epsilon}(\mathbf{s}) \stackrel{iid}{\sim} \text{N}(\mathbf{0}_{q \times 1}, (\alpha^{-1} - 1) \boldsymbol{\Sigma})$ captures micro-scale variation. The ‘‘proportionality’’ assumption for the variance of $\boldsymbol{\epsilon}(\mathbf{s})$ will allow us to derive analytic posterior distributions using conjugate priors.

The latent process $\boldsymbol{\omega}(\mathbf{s})$ captures the underlying spatial pattern and holds specific interest in many applications. Let $\boldsymbol{\omega} = \boldsymbol{\omega}(\mathcal{S}) = [\boldsymbol{\omega}(\mathbf{s}_1) : \dots : \boldsymbol{\omega}(\mathbf{s}_n)]^\top$ be $n \times q$. The parameter space with the latent process is $\{\boldsymbol{\beta}, \boldsymbol{\omega}, \boldsymbol{\Sigma}\}$. Letting $\boldsymbol{\gamma}^\top = [\boldsymbol{\beta}^\top, \boldsymbol{\omega}^\top]$ be $q \times (p + n)$, we assume that $\{\boldsymbol{\gamma}, \boldsymbol{\Sigma}\} \sim \text{MNIW}(\boldsymbol{\mu}_\gamma, \mathbf{V}_\gamma, \boldsymbol{\Psi}, v)$, where $\boldsymbol{\mu}_\gamma^\top = [\boldsymbol{\mu}_\beta^\top, \mathbf{0}_{q \times n}]$ and $\mathbf{V}_\gamma = \text{blockdiag}\{\mathbf{V}_r, \boldsymbol{\rho}_\psi(\mathcal{S}, \mathcal{S})\}$. The posterior density is

$$\begin{aligned}p(\boldsymbol{\gamma}, \boldsymbol{\Sigma} | \mathbf{Y}) &\propto \text{MNIW}(\boldsymbol{\gamma}, \boldsymbol{\Sigma} | \boldsymbol{\mu}_\gamma, \mathbf{V}_\gamma, \boldsymbol{\Psi}, v) \times \text{MN}_{n,q}(\mathbf{Y}_{n \times q} | [\mathbf{X} : \mathbf{I}_n] \boldsymbol{\gamma}, (\alpha^{-1} - 1) \mathbf{I}_n, \boldsymbol{\Sigma}) \\ &\propto \text{MNIW}(\boldsymbol{\gamma}, \boldsymbol{\Sigma} | \boldsymbol{\mu}_\gamma^*, \mathbf{V}^*, \boldsymbol{\Psi}^*, v^*),\end{aligned}\quad (9)$$

where

$$\begin{aligned}\mathbf{V}^* &= \begin{bmatrix} \frac{\alpha}{1-\alpha} \mathbf{X}^\top \mathbf{X} + \mathbf{V}_r^{-1} & \frac{\alpha}{1-\alpha} \mathbf{X}^\top \\ \frac{\alpha}{1-\alpha} \mathbf{X} & \boldsymbol{\rho}_\psi^{-1}(\mathcal{S}, \mathcal{S}) + \frac{\alpha}{1-\alpha} \mathbf{I}_n \end{bmatrix}^{-1}, \quad \boldsymbol{\mu}_\gamma^* = \mathbf{V}^* \begin{bmatrix} \frac{\alpha}{1-\alpha} \mathbf{X}^\top \mathbf{Y} + \mathbf{V}_r^{-1} \boldsymbol{\mu}_\beta \\ \frac{\alpha}{1-\alpha} \mathbf{Y} \end{bmatrix}, \\ \boldsymbol{\Psi}^* &= \boldsymbol{\Psi} + \frac{\alpha}{1-\alpha} \mathbf{Y}^\top \mathbf{Y} + \boldsymbol{\mu}_\beta^\top \mathbf{V}_r^{-1} \boldsymbol{\mu}_\beta - \boldsymbol{\mu}_\gamma^{*\top} \mathbf{V}^{*-1} \boldsymbol{\mu}_\gamma^* \quad \text{and } v^* = v + n.\end{aligned}\quad (10)$$

For prediction on a set of location \mathcal{U} , we can estimate the unobserved latent process $\omega_{\mathcal{U}} = \omega(\mathcal{U}) = [\omega(\mathbf{u}_1) : \dots : \omega(\mathbf{u}_{n'})]^\top$ and the response $\mathbf{Y}_{\mathcal{U}}$ through

$$p(\mathbf{Y}_{\mathcal{U}}, \omega_{\mathcal{U}} | \mathbf{Y}) = \int \text{MN}_{n',q}(\mathbf{Y}_{\mathcal{U}} | \mathbf{X}_{\mathcal{U}}\boldsymbol{\beta} + \omega_{\mathcal{U}}, (\alpha^{-1} - 1)\mathbf{I}_{n'}, \boldsymbol{\Sigma}) \times \text{MN}_{n',q}(\omega_{\mathcal{U}} | \mathbf{M}_{\mathcal{U}}\boldsymbol{\omega}, \mathbf{V}_{\omega_{\mathcal{U}}}, \boldsymbol{\Sigma}) \\ \times \text{MNIW}(\boldsymbol{\gamma}, \boldsymbol{\Sigma} | \boldsymbol{\mu}_{\boldsymbol{\gamma}}^*, \mathbf{V}^*, \boldsymbol{\Psi}^*, \nu^*) d\boldsymbol{\gamma}d\boldsymbol{\Sigma}, \quad (11)$$

where $\mathbf{M}_{\mathcal{U}} = \boldsymbol{\rho}_{\psi}(\mathcal{U}, \mathcal{S})\boldsymbol{\rho}_{\psi}^{-1}(\mathcal{S}, \mathcal{S})$ and $\mathbf{V}_{\omega_{\mathcal{U}}} = \boldsymbol{\rho}_{\psi}(\mathcal{U}, \mathcal{U}) - \boldsymbol{\rho}_{\psi}(\mathcal{U}, \mathcal{S})\boldsymbol{\rho}_{\psi}^{-1}(\mathcal{S}, \mathcal{S})\boldsymbol{\rho}_{\psi}(\mathcal{S}, \mathcal{U})$. Posterior predictive inference proceeds by sampling one draw of $\omega_{\mathcal{U}} \sim \text{MN}_{n',q}(\omega_{\mathcal{U}} | \mathbf{M}_{\mathcal{U}}\boldsymbol{\omega}, \mathbf{V}_{\omega_{\mathcal{U}}}, \boldsymbol{\Sigma})$ for each posterior draw of $\{\boldsymbol{\gamma}, \boldsymbol{\Sigma}\}$ and then one draw of $\mathbf{Y}_{\mathcal{U}} \sim \text{MN}(\mathbf{X}_{\mathcal{U}}\boldsymbol{\beta} + \omega_{\mathcal{U}}, (\alpha^{-1} - 1)\mathbf{I}_{n'}, \boldsymbol{\Sigma})$ for each drawn $\{\omega_{\mathcal{U}}, \boldsymbol{\gamma}, \boldsymbol{\Sigma}\}$.

II. Scalable Conjugate Bayesian Multivariate Models

Conjugate multivariate response NNGP model A conjugate Bayesian modeling framework is appealing for massive spatial data sets because the posterior distribution of the parameters are available in closed form circumventing the need for MCMC algorithms. The key computational bottleneck for Bayesian estimation of spatial process models concerns the computation and storage involving \mathcal{K}^{-1} in (5). The required matrix computations require $\mathcal{O}(n^3)$ flops and $\mathcal{O}(n^2)$ storage when \mathcal{K} is $n \times n$ and dense. While conjugate models reduce computational expenses by enabling direct sampling from closed-form posterior and posterior predictive distributions, the computation and storage of \mathcal{K} is still substantial for massive datasets.

One approach to obviate the overwhelming computations is to develop a sparse alternative for \mathcal{K}^{-1} in (5). One such approximation that has generated substantial recent attention in the spatial literature is an approximation due to Vecchia (Vecchia, 1988). Consider the spatial covariance matrix $\mathcal{K} = \boldsymbol{\rho}_{\psi} + \delta_{\mathbf{s}=\mathbf{s}'}\mathbf{I}_n$ in (2). This is a dense $n \times n$ matrix with no apparent exploitable structure. Instead, we specify a sparse Cholesky representation

$$\mathcal{K}^{-1} = (\mathbf{I} - \mathbf{A}_{\mathcal{K}})^\top \mathbf{D}_{\mathcal{K}}^{-1} (\mathbf{I} - \mathbf{A}_{\mathcal{K}}), \quad (12)$$

where $\mathbf{D}_{\mathcal{K}}$ is a diagonal matrix and $\mathbf{A}_{\mathcal{K}}$ is a sparse lower-triangular matrix with 0 along the diagonal and with no more than a fixed small number m of nonzero entries in each row of $\mathbf{A}_{\mathcal{K}}$. The diagonal entries of $\mathbf{D}_{\mathcal{K}}^{-1}$ and the nonzero entries of $\mathbf{A}_{\mathcal{K}}$ are obtained from the conditional variance and conditional expectations for a Gaussian process with covariance function $\rho_{\psi}(\mathbf{s}, \mathbf{s}')$. To be precise, we consider a fixed order of locations in \mathcal{S} and define $N_m(\mathbf{s}_i)$ to be the set comprising at most m neighbors of \mathbf{s}_i among locations $\mathbf{s}_j \in \mathcal{S}$ such that $j < i$. The (i, j) -th entry of $\mathbf{A}_{\mathcal{K}}$ is 0 whenever $\mathbf{s}_j \notin N_m(\mathbf{s}_i)$. If $j_1 < j_2 < \dots < j_m$ are the m column indices indicating the nonzero entries in the i -th row of $\mathbf{A}_{\mathcal{K}}$, then the (i, j_k) -th element of $\mathbf{A}_{\mathcal{K}}$ is set equal to the k -th element of the $1 \times m$ vector $\mathbf{a}_i^\top = \boldsymbol{\rho}_{\psi}(\mathbf{s}_i, N_m(\mathbf{s}_i))\boldsymbol{\rho}_{\psi}(N_m(\mathbf{s}_i), N_m(\mathbf{s}_i))^{-1}$. The (i, i) -th diagonal element of $\mathbf{D}_{\mathcal{K}}^{-1}$ is given by $\rho_{\psi}(\mathbf{s}_i, \mathbf{s}_i) - \mathbf{a}_i^\top \boldsymbol{\rho}_{\psi}(N_m(\mathbf{s}_i), \mathbf{s}_i)$. Repeating these calculations for each row completes the construction of $\mathbf{A}_{\mathcal{K}}$ and $\mathbf{D}_{\mathcal{K}}^{-1}$ and yields a sparse \mathcal{K}^{-1} in (12). This construction can be performed in parallel and requires storage or computation of at most $m \times m$ matrices, where $m \ll n$, costing $\mathcal{O}(n)$ flops and storage. Further algorithmic details about this construction can be found in Finley et al. (2019).

Based on Section I, the posterior distribution $\boldsymbol{\beta}, \boldsymbol{\Sigma} | \mathbf{Y}$ follows $\text{MNIW}(\boldsymbol{\mu}^*, \mathbf{V}^*, \boldsymbol{\Psi}^*, \nu^*)$ where $\{\boldsymbol{\mu}^*, \mathbf{V}^*, \boldsymbol{\Psi}^*, \nu^*\}$ are given in (5). With the sparse representation of \mathcal{K}^{-1} in (12), the process of obtaining posterior inference for $\{\boldsymbol{\beta}, \boldsymbol{\Sigma}\}$ only involves steps with storage and computational requirement in $\mathcal{O}(n)$.

The predictions on the unobserved locations $\mathcal{U} = \{\mathbf{u}_1, \dots, \mathbf{u}_{n'}\}$ is also simplified as follows. We extend the definition of $N_m(\mathbf{s}_i)$'s to arbitrary locations by defining $N_m(\mathbf{u}_i)$ to be the set of m nearest neighbors of \mathbf{u}_i from \mathcal{S} . Furthermore, we assume that $\mathbf{y}(\mathbf{u})$ and $\mathbf{y}(\mathbf{u}')$ are conditionally independent of each other given $\mathbf{Y} = \mathbf{y}(\mathcal{S})$ and the other model parameters. Thus, for any $\mathbf{u}_i \in \mathcal{U}$, we have

$$\mathbf{y}(\mathbf{u}_i) | \mathbf{Y}, \boldsymbol{\beta}, \boldsymbol{\Sigma} \sim \mathbf{N}(\boldsymbol{\beta}^\top \mathbf{x}(\mathbf{u}_i) + [\mathbf{Y} - \mathbf{X}\boldsymbol{\beta}]^\top \tilde{\mathbf{a}}_i, \tilde{d}_i \boldsymbol{\Sigma}), \quad i = 1, \dots, n', \quad (13)$$

where $\tilde{\mathbf{a}}_i$ is an $n \times 1$ vector with m non-zero elements. If $N_m(\mathbf{u}_i) = \{\mathbf{s}_{j_k}\}_{k=1}^m$, then

$$\begin{aligned} (\{\tilde{\mathbf{a}}_i\}_{j_1}, \dots, \{\tilde{\mathbf{a}}_i\}_{j_m}) &= \boldsymbol{\rho}_\psi(\mathbf{u}_i, N_m(\mathbf{u}_i)) \{\boldsymbol{\rho}_\psi(N_m(\mathbf{u}_i), N_m(\mathbf{u}_i)) + (\alpha^{-1} - 1)\mathbf{I}_m\}^{-1}, \\ \tilde{d}_i &= \alpha^{-1} - \boldsymbol{\rho}_\psi(\mathbf{u}_i, N_m(\mathbf{u}_i)) [\boldsymbol{\rho}_\psi(N_m(\mathbf{u}_i), N_m(\mathbf{u}_i)) + (\alpha^{-1} - 1)\mathbf{I}_m]^{-1} \boldsymbol{\rho}_\psi(N_m(\mathbf{u}_i), \mathbf{u}_i). \end{aligned} \quad (14)$$

If $\tilde{\mathbf{A}} = [\tilde{\mathbf{a}}_1 : \dots : \tilde{\mathbf{a}}_{n'}]^\top$ and $\tilde{\mathbf{D}} = \text{diag}(\{\tilde{d}_i\}_{i=1}^{n'})$, then the conditional predictive density for $\mathbf{Y}_{\mathcal{U}}$ is

$$\mathbf{Y}_{\mathcal{U}} | \mathbf{Y}, \boldsymbol{\beta}, \boldsymbol{\Sigma} \sim \text{MN}(\mathbf{X}_{\mathcal{U}}\boldsymbol{\beta} + \tilde{\mathbf{A}}[\mathbf{Y} - \mathbf{X}\boldsymbol{\beta}], \tilde{\mathbf{D}}, \boldsymbol{\Sigma}). \quad (15)$$

Since the posterior distribution of $\{\boldsymbol{\beta}, \boldsymbol{\Sigma}\}$ and the conditional predictive distribution of $\mathbf{Y}_{\mathcal{U}}$ are both available in closed form, direct sampling from the posterior predictive distribution is straightforward. A detailed algorithm for obtaining the posterior inference on parameter set $\{\boldsymbol{\beta}, \boldsymbol{\Sigma}\}$ and the posterior prediction over a new set of location \mathcal{U} is given as below.

Algorithm 1: Obtaining posterior inference of $\{\boldsymbol{\beta}, \boldsymbol{\Sigma}\}$ and predictions on \mathcal{U} for conjugate multivariate response NNGP

1. Construct \mathbf{V}^* , $\boldsymbol{\mu}^*$, $\boldsymbol{\Psi}^*$ and ν^* :
 - (a) Compute \mathbf{L}_r the Cholesky decomposition of \mathbf{V}_r
 - (b) Compute $\mathbf{D}\mathbf{I}\mathbf{A}\mathbf{X} = \mathbf{D}_k^{-\frac{1}{2}}(\mathbf{I} - \mathbf{A}_k)\mathbf{X}$ and $\mathbf{D}\mathbf{I}\mathbf{A}\mathbf{Y} = \mathbf{D}_k^{-\frac{1}{2}}(\mathbf{I} - \mathbf{A}_k)\mathbf{Y}$
 - Construct \mathbf{A}_K and \mathbf{D}_K as described, for example, in Finley et al. (2019) $\mathcal{O}(nm^3)$
 - Compute $\mathbf{D}\mathbf{I}\mathbf{A}\mathbf{X} = \mathbf{D}_k^{-\frac{1}{2}}(\mathbf{I} - \mathbf{A}_k)\mathbf{X}$ and $\mathbf{D}\mathbf{I}\mathbf{A}\mathbf{Y} = \mathbf{D}_k^{-\frac{1}{2}}(\mathbf{I} - \mathbf{A}_k)\mathbf{Y}$ $\mathcal{O}(n(m+1)(p+q+2))$
 - (c) Obtain \mathbf{V}^* , $\boldsymbol{\mu}^*$ and $\boldsymbol{\Psi}^*$
 - Compute $\mathbf{V}^* = (\mathbf{D}\mathbf{I}\mathbf{A}\mathbf{X}^\top \mathbf{D}\mathbf{I}\mathbf{A}\mathbf{X} + \mathbf{V}_r^{-1})^{-1}$ and its Cholesky decomposition \mathbf{L}_{v^*} $\mathcal{O}(np^2)$
 - Compute $\boldsymbol{\mu}^* = \mathbf{V}^*(\mathbf{D}\mathbf{I}\mathbf{A}\mathbf{X}^\top \mathbf{D}\mathbf{I}\mathbf{A}\mathbf{Y} + \mathbf{V}_r^{-1}\boldsymbol{\mu}_\beta)$ $\mathcal{O}(npq)$
 - Compute $\boldsymbol{\Psi}^* = \boldsymbol{\Psi} + \mathbf{D}\mathbf{I}\mathbf{A}\mathbf{Y}^\top \mathbf{D}\mathbf{I}\mathbf{A}\mathbf{Y} + (\mathbf{L}_r^{-1}\boldsymbol{\mu}_\beta)^\top (\mathbf{L}_r^{-1}\boldsymbol{\mu}_\beta) - (\mathbf{L}_{v^*}^{-1}\boldsymbol{\mu}^*)^\top (\mathbf{L}_{v^*}^{-1}\boldsymbol{\mu}^*)$ $\mathcal{O}(nq^2)$
 - Compute $\nu^* = \nu + n$
 2. Generate posterior samples $\{\mathbf{Y}_{\mathcal{U}}^{(l)}\}_{l=1}^L$ on a new set \mathcal{U} given $\mathbf{X}_{\mathcal{U}}$
 - (a) Construct $\tilde{\mathbf{A}}$ and $\tilde{\mathbf{D}}$ as described in (14) $\mathcal{O}(n'm^3)$
 - (b) For l in $1 : L$
 - i. Sample $\boldsymbol{\Sigma}^{(l)} \sim \text{IW}(\boldsymbol{\Psi}^*, \nu^*)$
 - ii. Sample $\boldsymbol{\beta}^{(l)} \sim \text{MN}(\boldsymbol{\mu}^*, \mathbf{V}^*, \boldsymbol{\Sigma}^{(l)})$
 - Calculate Cholesky decomposition of $\boldsymbol{\Sigma}^{(l)}$, $\boldsymbol{\Sigma}^{(l)} = \mathbf{L}_{\boldsymbol{\Sigma}^{(l)}} \mathbf{L}_{\boldsymbol{\Sigma}^{(l)}}^\top$
 - Sample $\mathbf{u} \sim \text{MN}(\mathbf{0}, \mathbf{I}_p, \mathbf{I}_q)$ (i.e. $\text{vec}(\mathbf{u}) \sim \text{MVN}(\mathbf{0}, \mathbf{I}_{pq})$)
 - Generate $\boldsymbol{\beta}^{(l)} = \boldsymbol{\mu}^* + \mathbf{L}_{v^*} \mathbf{u} \mathbf{L}_{\boldsymbol{\Sigma}^{(l)}}^\top$
 - iii. Sample $\mathbf{Y}_{\mathcal{U}}^{(l)} \sim \text{MN}(\mathbf{X}_{\mathcal{U}}\boldsymbol{\beta}^{(l)} + \tilde{\mathbf{A}}[\mathbf{Y} - \mathbf{X}\boldsymbol{\beta}^{(l)}], \tilde{\mathbf{D}}, \boldsymbol{\Sigma}^{(l)})$
 - Sample $\mathbf{u} \sim \text{MN}(\mathbf{0}, \mathbf{I}_{n'}, \mathbf{I}_q)$.
 - Generate $\mathbf{Y}_{\mathcal{U}}^{(l)} = \mathbf{X}_{\mathcal{U}}\boldsymbol{\beta}^{(l)} + \tilde{\mathbf{A}}[\mathbf{Y} - \mathbf{X}\boldsymbol{\beta}^{(l)}] + \tilde{\mathbf{D}}^{\frac{1}{2}} \mathbf{u} \mathbf{L}_{\boldsymbol{\Sigma}^{(l)}}^\top$ $\mathcal{O}((n' + n)pq + n'(q^2 + mq))$
-

Conjugate multivariate latent NNGP model Bayesian estimation for the conjugate multivariate latent model is more challenging because inference is usually sought on the (high-dimensional) latent process itself. In particular, the calculations involved in \mathbf{V}^* in (9) are often too expensive for large data sets even when the precision matrix $\rho_{\psi}^{-1}(\mathcal{S}, \mathcal{S})$ is sparse. Here, the latent process $\omega(\mathbf{s})$ in (8) follows a multivariate Gaussian process so that its realizations over \mathcal{S} follows $\omega \sim \text{MN}(\mathbf{O}_{n \times q}, \tilde{\rho}, \Sigma)$, where $\tilde{\rho}$ is the Vecchia approximation of $\rho_{\psi}(\mathcal{S}, \mathcal{S})$. Hence, $\tilde{\rho}^{-1} = (\mathbf{I} - \mathbf{A}_{\rho})^{\top} \mathbf{D}_{\rho}^{-1} (\mathbf{I} - \mathbf{A}_{\rho})$, where \mathbf{A}_{ρ} and \mathbf{D}_{ρ} are constructed analogous to $\mathbf{A}_{\mathcal{K}}$ and $\mathbf{D}_{\mathcal{K}}$ in (12) with \mathcal{K} replaced by $\rho_{\psi}(\mathcal{S}, \mathcal{S})$. This corresponds to modeling $\omega(\mathbf{s})$ with a Nearest-Neighbor Gaussian Process (NNGP) (see, e.g., Datta et al., 2016a,b; Banerjee, 2017, for details).

The posterior distribution of $\{\gamma, \Sigma\}$ follows a Matrix-Normal distribution similar to (9), but with $\rho_{\psi}(\mathcal{S}, \mathcal{S})^{-1}$ in (10) replaced by its Vecchia approximation $\tilde{\rho}_{\psi}(\mathcal{S}, \mathcal{S})$. We will solve the linear system $\mathbf{X}^{*\top} \mathbf{X}^* \boldsymbol{\mu}^* = \mathbf{X}^{*\top} \mathbf{Y}^*$ for $\boldsymbol{\mu}^*$, compute $\{\Psi^*, \nu^*\}$ and generate posterior samples of Σ from $\text{IW}(\Psi^*, \nu^*)$. Posterior samples of γ are obtained by generating $\boldsymbol{\eta} \sim \text{MN}(\mathbf{O}, \mathbf{I}_{2n+p}, \Sigma)$, solving $\mathbf{X}^{*\top} \mathbf{X}^* \mathbf{v} = \mathbf{X}^{*\top} \boldsymbol{\eta}$ for \mathbf{v} and then obtaining posterior samples of γ from $\gamma = \boldsymbol{\mu}^* + \mathbf{v}$.

However, sampling $\{\gamma, \Sigma\}$ is still challenging for massive data sets, where we seek to minimize storage and operations with large matrices. Here we introduce a useful representation. Let \mathbf{V}_{ρ} be a non-singular square matrix such that $\rho_{\psi}^{-1}(\mathcal{S}, \mathcal{S}) = \mathbf{V}_{\rho}^{\top} \mathbf{V}_{\rho}$ where we write $\mathbf{V}_{\rho} = \mathbf{D}_{\rho}^{-1/2} (\mathbf{I} - \mathbf{A}_{\rho})$. We treat the prior of γ as additional ‘‘observations’’ and recast $p(\mathbf{Y}, \gamma | \Sigma) = p(\mathbf{Y} | \gamma, \Sigma) \times p(\gamma | \Sigma)$ into an augmented linear model

$$\underbrace{\begin{bmatrix} \sqrt{\frac{\alpha}{1-\alpha}} \mathbf{Y} \\ \mathbf{L}_r^{-1} \boldsymbol{\mu}_{\beta} \\ \mathbf{0} \end{bmatrix}}_{\mathbf{Y}^*} = \underbrace{\begin{bmatrix} \sqrt{\frac{\alpha}{1-\alpha}} \mathbf{X} & \sqrt{\frac{\alpha}{1-\alpha}} \mathbf{I}_n \\ \mathbf{L}_r^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{V}_{\rho} \end{bmatrix}}_{\mathbf{X}^*} \underbrace{\begin{bmatrix} \boldsymbol{\beta} \\ \boldsymbol{\omega} \end{bmatrix}}_{\boldsymbol{\gamma}} + \underbrace{\begin{bmatrix} \eta_1 \\ \eta_2 \\ \eta_3 \end{bmatrix}}_{\boldsymbol{\eta}}, \quad (16)$$

where \mathbf{L}_r is the Cholesky decomposition of \mathbf{V}_r , and $\boldsymbol{\eta} \sim \text{MN}(\mathbf{0}, \mathbf{I}_{2n+p}, \Sigma)$. With a flat prior for $\boldsymbol{\beta}$, \mathbf{L}_r^{-1} degenerates to \mathbf{O} and does not contribute to the linear system. The expression in (10) can now be simplified as follows

$$\begin{aligned} \mathbf{V}^* &= (\mathbf{X}^{*\top} \mathbf{X}^*)^{-1}, \quad \boldsymbol{\mu}^* = (\mathbf{X}^{*\top} \mathbf{X}^*)^{-1} \mathbf{X}^{*\top} \mathbf{Y}^*, \\ \Psi^* &= \Psi + (\mathbf{Y}^* - \mathbf{X}^* \boldsymbol{\mu}^*)^{\top} (\mathbf{Y}^* - \mathbf{X}^* \boldsymbol{\mu}^*), \quad \nu^* = \nu + n. \end{aligned} \quad (17)$$

Following developments in Zhang et al. (2019) for the univariate case, one can efficiently generate posterior samples through a conjugate gradient algorithm exploiting the sparsity of \mathbf{V}_{ρ} . The sampling process for γ will be scalable when there is a sparse precision matrix $\rho_{\psi}^{-1}(\mathcal{S}, \mathcal{S})$. It is also possible to construct \mathbf{V}^* and $\boldsymbol{\mu}^*$ in (17) using $\rho_{\psi}^{-1}(\mathcal{S}, \mathcal{S})$ instead of \mathbf{V}_{ρ} . We refer to Zhang et al. (2019) for further details of this construction. We provide a detailed algorithm of the conjugate multivariate latent NNGP model below, where we implement a ‘‘Sparse Equations and Least Squares’’ (LSMR) algorithm (Fong and Saunders, 2011) to solve the linear system $\mathbf{X}^{*\top} \mathbf{X}^* \boldsymbol{\mu}^* = \mathbf{X}^{*\top} \mathbf{Y}^*$ and $\mathbf{X}^{*\top} \mathbf{X}^* \mathbf{v} = \mathbf{X}^{*\top} \boldsymbol{\eta}$ needed to generate γ . LSMR is a conjugate-gradient type algorithm for solving sparse linear equations $\mathbf{A} \mathbf{x} = \mathbf{b}$ where the matrix \mathbf{A} may be square or rectangular. The matrix $\mathbf{A} := \mathbf{X}^*$ is a sparse tall matrix. LSMR only requires storing \mathbf{X}^* , \mathbf{Y}^* and $\boldsymbol{\eta}^*$ and, unlike the conjugate gradient algorithm, avoids $\mathbf{X}^{*\top} \mathbf{X}^*$, $\mathbf{X}^{*\top} \mathbf{Y}^*$ and $\mathbf{X}^{*\top} \boldsymbol{\eta}^*$. LSMR also tends to produce more stable estimates than conjugate gradient. We have also tested a variety of conjugate gradient methods and preconditioning methods, where we have observed that their performances varied across different data sets. The LSMR without conditioning showed a relatively good performance for the latent models. Therefore, we choose LSMR without preconditioning for our current illustrations.

Posterior predictive inference will adapt from (11) for scalable models. After sampling $\{\gamma, \Sigma\}$, we sample one draw of $\omega_{\mathcal{U}} \sim \omega_{\mathcal{U}} | \gamma, \Sigma \sim \text{MN}([\mathbf{0}_{n' \times p}, \tilde{\mathbf{A}}] \gamma, \tilde{\mathbf{D}}, \Sigma)$ for each sampled $\{\gamma, \Sigma\}$, where $\tilde{\mathbf{A}} = [\tilde{\mathbf{a}}_1 : \dots : \tilde{\mathbf{a}}_n]^\top$, $\tilde{\mathbf{D}} = \text{diag}(\{\tilde{d}_i\}_{i=1}^n)$ with

$$\begin{aligned} (\{\tilde{\mathbf{a}}_i\}_{j_1}, \dots, \{\tilde{\mathbf{a}}_i\}_{j_m}) &= \rho_\psi(\mathbf{u}_i, N_m(\mathbf{u}_i)) \rho_\psi^{-1}(N_m(\mathbf{u}_i), N_m(\mathbf{u}_i)), \\ \tilde{d}_i &= 1 - \rho_\psi(\mathbf{u}_i, N_m(\mathbf{u}_i)) \rho_\psi^{-1}(N_m(\mathbf{u}_i), N_m(\mathbf{u}_i)) \rho_\psi(N_m(\mathbf{u}_i), \mathbf{u}_i). \end{aligned} \quad (18)$$

Finally, for each sampled $\{\beta, \omega_{\mathcal{U}}, \Sigma\}$ we make one draw of $\mathbf{Y}_{\mathcal{U}} \sim \text{MN}(\mathbf{X}_{\mathcal{U}} \beta + \omega_{\mathcal{U}}, (\alpha^{-1} - 1) \mathbf{I}_{n'}, \Sigma)$. The following provides details of the algorithm for predictive inference.

Algorithm 2: Obtaining posterior inference of $\{\gamma, \Sigma\}$ and predictions on set \mathcal{U} for conjugate multivariate latent NNGP

1. Construct \mathbf{X}^* and \mathbf{Y}^* in (16)
 - (a) \mathbf{L}_r^{-1} and $\mathbf{L}_r^{-1} \mu_\beta$
 - Compute the Cholesky decomposition of $\mathbf{V}_r, \mathbf{L}_r$
 - Compute \mathbf{L}_r^{-1} and $\mathbf{L}_r^{-1} \mu_\beta$
 - (b) \mathbf{V}_ρ
 - Construct \mathbf{A}_ρ and \mathbf{D}_ρ as described, for example, in Finley et al. (2019) $\mathcal{O}(nm^3)$
 - Compute $\mathbf{V}_\rho = \mathbf{D}_\rho^{-\frac{1}{2}} (\mathbf{I} - \mathbf{A}_\rho)$ $\mathcal{O}(n(m+1))$
 - (c) Construct \mathbf{X}^* and \mathbf{Y}^*
 2. Obtain μ^*, Ψ^* and ν^* .
 - (a) Obtain $\mu^* = [\mu_1^* : \dots : \mu_q^*]$
 - Solve μ_i^* from $\mathbf{X}^* \mu_i^* = \mathbf{Y}_i^*$ by LSMR for $i = 1, \dots, q$.
 - (b) Obtain Ψ^* and ν^*
 - Generate $\mathbf{u} = \mathbf{Y}^* - \mathbf{X}^* \mu^*$ $\mathcal{O}(n(1 + (p+m+1)q))$
 - Compute $\Psi^* = \Psi + \mathbf{u}^\top \mathbf{u}$ $\mathcal{O}(nq^2)$
 - Compute $\nu^* = \nu + n$
 3. Generate posterior samples of $\{\gamma^{(l)}, \Sigma^{(l)}\}_{l=1}^L$. For l in $1 : L$
 - (a) Sample $\Sigma^{(l)} \sim \text{IW}(\Psi^*, \nu^*)$
 - (b) Sample $\gamma^{(l)} \sim \text{MN}(\mu^*, \mathbf{V}^*, \Sigma^{(l)})$
 - Sample $\mathbf{u} \sim \text{MN}(\mathbf{0}, \mathbf{I}_{2n+p}, \mathbf{I}_q)$ $\mathcal{O}(2nq)$
 - Calculate Cholesky decomposition of $\Sigma^{(l)}, \Sigma^{(l)} = \mathbf{L}_{\Sigma^{(l)}} \mathbf{L}_{\Sigma^{(l)}}^\top$
 - Generate $\boldsymbol{\eta} = \mathbf{u} \mathbf{L}^{(l)\top} = [\boldsymbol{\eta}_1 : \dots : \boldsymbol{\eta}_q]$ $\mathcal{O}(2nq^2)$
 - Solve \mathbf{v}_i from $\mathbf{X}^* \mathbf{v}_i = \boldsymbol{\eta}_i$ by LSMR for $i = 1, \dots, q$.
 - Generate $\gamma^{(l)} = \mu^* + \mathbf{v}$ with $\mathbf{v} = [\mathbf{v}_1 : \dots : \mathbf{v}_q]$ $\mathcal{O}(nq)$
 4. Generate posterior samples of $\{\mathbf{Y}_{\mathcal{U}}^{(l)}\}$ on a new set \mathcal{U} given $\mathbf{X}_{\mathcal{U}}$.
 - (a) Construct $\tilde{\mathbf{A}}$ and $\tilde{\mathbf{D}}$ using (18) $\mathcal{O}(n'm^3)$
 - (b) For l in $1 : L$
 - i. Sample $\omega_{\mathcal{U}}^{(l)} \sim \text{MN}([\mathbf{0}_{n' \times p}, \tilde{\mathbf{A}}] \gamma^{(l)}, \tilde{\mathbf{D}}, \Sigma^{(l)})$
 - Sample $\mathbf{u} \sim \text{MN}(\mathbf{0}, \mathbf{I}_{n'}, \mathbf{I}_q)$ $\mathcal{O}(n'q)$
 - Generate $\omega_{\mathcal{U}}^{(l)} = [\mathbf{0}_{n' \times p}, \tilde{\mathbf{A}}] \gamma^{(l)} + \tilde{\mathbf{D}}^{\frac{1}{2}} \mathbf{u} \mathbf{L}_{\Sigma^{(l)}}^\top$ $\mathcal{O}(n'mq + n'q^2)$
 - ii. Sample $\mathbf{Y}_{\mathcal{U}}^{(l)} | \omega_{\mathcal{U}}^{(l)}, \gamma^{(l)}, \Sigma^{(l)} \sim \text{MN}(\mathbf{X}_{\mathcal{U}} \beta + \omega_{\mathcal{U}}^{(l)}, (\alpha^{-1} - 1) \mathbf{I}_{n'}, \Sigma)$
 - Sample $\mathbf{u} \sim \text{MN}(\mathbf{0}, \mathbf{I}_{n'}, \mathbf{I}_q)$ $\mathcal{O}(n'q)$
 - Generate $\mathbf{Y}_{\mathcal{U}}^{(l)} = \mathbf{X}_{\mathcal{U}} \beta + \omega_{\mathcal{U}}^{(l)} + (\alpha^{-1} - 1) \mathbf{u} \mathbf{L}_{\Sigma^{(l)}}^\top$ $\mathcal{O}(n'pq + n'q^2)$
-

III. Cross-validation for Conjugate Multivariate NNGP Models

Conjugate Bayesian multivariate regression models will depend upon fixing hyperparameters in the model. Here, we apply a K -fold cross-validation algorithm for choosing $\{\psi, \alpha\}$. This algorithm is a straightforward generalization of the univariate algorithm in (Finley et al., 2019). We run the conjugate models for each point $\{\psi, \alpha\}$ on a grid and choose the value that produces the least magnitude of root mean square prediction error. The inference on that point is then presented. This is appealing for scalable Gaussian process models that, for any fixed $\{\psi, \alpha\}$, can deliver posterior inference at new locations requiring storage and flops in $\mathcal{O}(n)$.

Algorithm 3: Cross-validation of tuning ϕ, α for conjugate multivariate response or latent NNGP model

1. Split \mathcal{S} into K folds, and build neighbor index.
 - Split \mathcal{S} into K folds $\{\mathcal{S}_k\}_{k=1}^K$. We use \mathcal{S}_{-k} to denote \mathcal{S} without the locations in \mathcal{S}_k .
 - Build nearest neighbors for $\{\mathcal{S}_{-k}\}_{k=1}^K$
 - Find the collection of nearest neighbor set for \mathcal{S}_k among \mathcal{S}_{-k} for $k = 1, \dots, K$.
 2. (For response NNGP) Fix ϕ and α , obtain posterior mean of β after removing the k^{th} fold of the data:
 - Use step 1 in Algorithm 1 to obtain $\hat{\beta}_k$ by taking \mathcal{S} to be \mathcal{S}_{-k} and μ^* to be $\hat{\beta}_k$.
 (For latent NNGP) Fix ϕ and α , obtain posterior mean of $\gamma_k = \{\beta, \omega(\mathcal{S}_{-k})\}$ after removing the k^{th} fold of the data:
 - Use step 1-2 in Algorithm 3 to obtain $\hat{\gamma}_k$ by taking \mathcal{S} to be \mathcal{S}_{-k} and μ^* to be $\hat{\gamma}_k$.
 3. (For response NNGP) Predict posterior means of $\mathbf{y}(\mathcal{S}_k)$
 - Construct matrix $\tilde{\mathbf{A}}$ through (14) by taking \mathcal{S} to be \mathcal{S}_{-k} and \mathcal{U} to be \mathcal{S}_k .
 - According to (15), the predicted posterior mean of $\mathbf{y}(\mathcal{S}_k)$ follows

$$\hat{\mathbf{y}}(\mathcal{S}_k) = \mathbb{E}[\mathbf{y}(\mathcal{S}_k) | \mathbf{y}(\mathcal{S}_{-k})] = \mathbf{x}(\mathcal{S}_k)\hat{\beta}_k + \tilde{\mathbf{A}}[\mathbf{y}(\mathcal{S}_{-k}) - \mathbf{x}(\mathcal{S}_{-k})\hat{\beta}_k]$$
 (For latent NNGP) Predict posterior means of $\mathbf{y}(\mathcal{S}_k)$
 - Construct matrix $\tilde{\mathbf{A}}$ by taking \mathcal{S} to be \mathcal{S}_{-k} and \mathcal{U} to be \mathcal{S}_k .
 - The predicted posterior mean of $\mathbf{y}(\mathcal{S}_k)$ follows

$$\hat{\mathbf{y}}(\mathcal{S}_k) = \mathbb{E}[\mathbf{y}(\mathcal{S}_k) | \mathbf{y}(\mathcal{S}_{-k})] = \mathbb{E}_\omega[\mathbb{E}_{\mathbf{y}}[\mathbf{y}(\mathcal{S}_k) | \omega(\mathcal{S}_{-k}), \mathbf{y}(\mathcal{S}_{-k})]] = [\mathbf{x}(\mathcal{S}_k), \tilde{\mathbf{A}}]\hat{\gamma}_k$$
 4. Root Mean Square Predictive Error (RMSPE) over K folds
 - Initialize $e = 0$
 - for (k in $1 : K$)
 - for (\mathbf{s}_i in \mathcal{S}_k)
 - $e = e + \|\mathbf{y}(\mathbf{s}_i) - \hat{\mathbf{y}}(\mathbf{s}_i)\|^2$
 5. Cross validation for choosing ψ and α
 - Repeat steps (2) - (4) for all candidate values of ψ and α
 - Choose ψ_0 and α_0 as the value that minimizes the average RMSPE
-

IV. Comparison of Response and Latent Models

Modeling the response as an NNGP produces a different model from modeling the latent process as an NNGP. In the former, Vecchia's approximation to the joint density of the response yields a sparse precision matrix for the response. In the latter, it is the precision matrix of the realizations of the latent process that is sparse. This has been discussed in Datta et al. (2016a) and also explored in greater generality by Katzfuss and Guinness (2017). Comparisons based on the Kullback-Leibler divergence (KL-D) between the NNGP based models and their parent full GP models reveal that the latent NNGP model tends to be closer to the full GP than the response NNGP. A proof of such a result is provided by Katzfuss and Guinness (2017), but this result holds only in the context of an augmented directed acyclical graphical model with nodes comprising the response and the latent variables. However, if we compute the KL-D between the NNGP models and their full GP counterparts in terms of the collapsed or marginal distribution for \mathbf{Y} , then it is theoretically possible for the response model to be closer to the full GP.

Here we provide a simple example where a response NNGP model outperforms a latent NNGP model on a collapsed space. Assume the observed location set is $\mathcal{S} = \{\mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_3\}$, $\omega(\mathcal{S})$ has covariance matrix $\sigma^2 \mathbf{R}$ with correlation matrix

$$\mathbf{R} = \begin{bmatrix} 1 & \rho_{12} & \rho_{13} \\ \rho_{12} & 1 & \rho_{23} \\ \rho_{13} & \rho_{23} & 1 \end{bmatrix}. \quad (19)$$

Let us suppress the connection between knots \mathbf{s}_1 and \mathbf{s}_3 in the directed acyclic graph corresponding to the finite realization of the NNGP on \mathcal{S} . Then the covariance matrix of the response NNGP model Σ_R and that of the latent NNGP model Σ_l have the following forms:

$$\Sigma_R = \sigma^2 \begin{bmatrix} 1 + \delta^2 & \rho_{12} & \frac{\rho_{12}\rho_{23}}{1 + \delta^2} \\ \rho_{12} & 1 + \delta^2 & \rho_{23} \\ \frac{\rho_{12}\rho_{23}}{1 + \delta^2} & \rho_{23} & 1 + \delta^2 \end{bmatrix}, \quad \Sigma_l = \sigma^2 \begin{bmatrix} 1 + \delta^2 & \rho_{12} & \rho_{12}\rho_{23} \\ \rho_{12} & 1 + \delta^2 & \rho_{23} \\ \rho_{12}\rho_{23} & \rho_{23} & 1 + \delta^2 \end{bmatrix}, \quad (20)$$

where $\delta^2 = \frac{\tau^2}{\sigma^2}$ is the noise-to-signal ratio with τ^2 as the variance of the noise process $\epsilon(s)$. Since R is positive-definite, we must have

$$1 - (\rho_{12}^2 + \rho_{13}^2 + \rho_{23}^2) + 2\rho_{12}\rho_{13}\rho_{23} > 0, \quad 1 - \rho_{12}^2 > 0. \quad (21)$$

It is easy to show that Σ_R and Σ_l are also positive-definite. If $\rho_{13} = \frac{\rho_{12}\rho_{23}}{1 + \delta^2}$, then the KL-D from the response NNGP model to the true model always equals zero, which is no more than the KL-D from the latent NNGP model to the true model. If $\rho_{13} = \rho_{12}\rho_{23}$, then the KL-D of the latent NNGP model to the true model always equals zero, which reverses the relationship. Numerical examples can be found in https://luzhangstat.github.io/notes/KL-D_com.html

Still, our simulations indicate that the latent NNGP model tends to outperform the response NNGP model in approximating their parent GP based models. This is consistent with the theoretical result of (Katzfuss and Guinness, 2017) and also with our intuition: the presence of the latent process should certainly improve the goodness of fit of the model. Without loss of generality, our discussion here considers the univariate case, but the argument applies to the multivariate setting as well. Let $\{y(\mathbf{s}) : \mathbf{s} \in \mathcal{D}\}$ be the process of interest over $\mathcal{D} \subset \mathbb{R}^d, d \in \mathbb{N}^+$, and let $y(\mathbf{s}) = \omega(\mathbf{s}) + \epsilon(\mathbf{s})$ for some latent spatial GP $\omega(\mathbf{s})$ and white noise process $\epsilon(\mathbf{s})$. A response NNGP model specifies the NNGP on $y(\mathbf{s})$, while a latent NNGP model assumes that $\omega(\mathbf{s})$ follows the NNGP.

Let the covariance matrix of $\mathbf{y} = y(\mathcal{S})$ of the parent GP based models be $\mathbf{C} + \tau^2\mathbf{I}$, where \mathbf{C} is the covariance matrix of the latent process $\omega(\mathcal{S})$. Consider the Vecchia approximation of the precision matrices \mathbf{C}^{-1} and $\mathbf{K}^{-1} = \{\mathbf{C} + \tau^2\mathbf{I}\}^{-1}$:

$$\text{Vecchia}(\mathbf{C}^{-1}) = \tilde{\mathbf{C}}^{-1}, \text{Vecchia}(\mathbf{K}^{-1}) = \tilde{\mathbf{K}}^{-1}. \quad (22)$$

The covariance matrix of $y(\mathcal{S})$ from the latent NNGP model is $\tilde{\mathbf{C}} + \tau^2\mathbf{I}$, while the precision matrix of $y(\mathcal{S})$ from the response NNGP model is $\tilde{\mathbf{K}}^{-1}$. We denote the error matrix of the Vecchia approximation of \mathbf{C}^{-1} by \mathbf{E} . We assume that \mathbf{E} is small so that $\tilde{\mathbf{C}}^{-1}$ approximates \mathbf{C}^{-1} well. With the same observed location \mathcal{S} and the fixed number of nearest neighbors, the error matrix of the Vecchia approximation of \mathbf{K}^{-1} is believed to be close to \mathbf{E} , i.e.,

$$\mathbf{C}^{-1} = \tilde{\mathbf{C}}^{-1} + \mathbf{E}; \mathbf{K}^{-1} = \tilde{\mathbf{K}}^{-1} + \mathcal{O}(\mathbf{E}). \quad (23)$$

Representing the precision matrices of $y(\mathcal{S})$ of the parent GP based model and the latent NNGP model by

$$\begin{aligned} (\mathbf{C} + \tau^2\mathbf{I})^{-1} &= \mathbf{C}^{-1} - \mathbf{C}^{-1}\mathbf{M}^{-1}\mathbf{C}^{-1}, \mathbf{M} = \mathbf{C}^{-1} + \tau^{-2}\mathbf{I}, \\ (\tilde{\mathbf{C}} + \tau^2\mathbf{I})^{-1} &= \tilde{\mathbf{C}}^{-1} - \tilde{\mathbf{C}}^{-1}\mathbf{M}^{*-1}\tilde{\mathbf{C}}^{-1}, \mathbf{M}^* = \tilde{\mathbf{C}}^{-1} + \tau^{-2}\mathbf{I}, \end{aligned} \quad (24)$$

we find that the difference between the precision metrics over the collapsed space for the parent NNGP and for the latent NNGP model is

$$\begin{aligned} (\mathbf{C} + \tau^2\mathbf{I})^{-1} - (\tilde{\mathbf{C}} + \tau^2\mathbf{I})^{-1} &= \mathbf{C}^{-1} - \mathbf{C}^{-1}\mathbf{M}^{-1}\mathbf{C}^{-1} - \tilde{\mathbf{C}}^{-1} + \tilde{\mathbf{C}}^{-1}\mathbf{M}^{*-1}\tilde{\mathbf{C}}^{-1} \\ &= \underbrace{\mathbf{E} - \mathbf{E}\mathbf{M}^{-1}\tilde{\mathbf{C}}^{-1} - \tilde{\mathbf{C}}^{-1}\mathbf{M}^{-1}\mathbf{E} - \tilde{\mathbf{C}}^{-1}(\mathbf{M}^{-1} - \mathbf{M}^{*-1})\tilde{\mathbf{C}}^{-1}}_{\mathbf{B}} - \underbrace{\mathbf{E}\mathbf{M}^{-1}\mathbf{E}}_{\mathcal{O}(\mathbf{E}^2)} \end{aligned}$$

Representing \mathbf{B} in terms of $\tilde{\mathbf{C}}^{-1}$, \mathbf{M}^* and \mathbf{E} , where \mathbf{E} is assumed to be nonsingular, we find

$$\begin{aligned} \mathbf{B} &= \mathbf{E} - \mathbf{E}\mathbf{M}^{*-1}\tilde{\mathbf{C}}^{-1} + \mathbf{E}\mathbf{M}^{*-1}(\mathbf{E}^{-1} + \mathbf{M}^{*-1})^{-1}\mathbf{M}^{*-1}\tilde{\mathbf{C}}^{-1} - \tilde{\mathbf{C}}^{-1}\mathbf{M}^{*-1}\mathbf{E} \\ &\quad + \tilde{\mathbf{C}}^{-1}\mathbf{M}^{*-1}(\mathbf{E}^{-1} + \mathbf{M}^{*-1})^{-1}\mathbf{M}^{*-1}\mathbf{E} + \tilde{\mathbf{C}}^{-1}\mathbf{M}^{*-1}(\mathbf{E}^{-1} + \mathbf{M}^{*-1})^{-1}\mathbf{M}^{*-1}\tilde{\mathbf{C}}^{-1}. \end{aligned} \quad (25)$$

Using the familiar Woodbury matrix identity and the expansion $(\mathbf{I} + \mathbf{X})^{-1} = \sum_{n=0}^{\infty} \{-\mathbf{X}\}^n$, we find

$$\begin{aligned} (\mathbf{E}^{-1} + \mathbf{M}^{*-1})^{-1}\mathbf{M}^{*-1} &= \{\mathbf{M}^*(\mathbf{E}^{-1} + \mathbf{M}^{*-1})\}^{-1} = \{\mathbf{M}^*\mathbf{E}^{-1} + \mathbf{I}\}^{-1} \\ &= \mathbf{I} - \{\mathbf{I} + \mathbf{E}\mathbf{M}^{*-1}\}^{-1} = \mathbf{I} - \{\mathbf{I} - \mathbf{E}\mathbf{M}^{*-1} + \mathcal{O}(\mathbf{E}^2)\} \\ &= \mathbf{E}\mathbf{M}^{*-1} + \mathcal{O}(\mathbf{E}^2). \end{aligned}$$

Using the above equations and excluding the terms of order $\mathcal{O}(\mathbf{E}^2)$ in the expression of \mathbf{B} , the leading term in the difference is

$$\mathbf{B} = (\mathbf{I} - \tilde{\mathbf{C}}^{-1}\mathbf{M}^{*-1})\mathbf{E}(\mathbf{I} - \mathbf{M}^{*-1}\tilde{\mathbf{C}}^{-1}) = (\mathbf{I} + \tau^2\tilde{\mathbf{C}}^{-1})^{-1}\mathbf{E}(\mathbf{I} + \tau^2\tilde{\mathbf{C}}^{-1})^{-1}. \quad (26)$$

Using the spectral decomposition $(\mathbf{I} + \tau^2\tilde{\mathbf{C}}^{-1}) = \mathbf{P}^\top(\mathbf{I} + \tau^2\mathbf{D})\mathbf{P}$, where \mathbf{P} is orthogonal and \mathbf{D} is diagonal with positive elements on the diagonal, we obtain

$$\begin{aligned} \|\mathbf{B}\|_F &= \|\mathbf{P}^\top(\mathbf{I} + \tau^2\mathbf{D})^{-1}\mathbf{P}\mathbf{E}\mathbf{P}^\top(\mathbf{I} + \tau^2\mathbf{D})^{-1}\mathbf{P}\|_F = \|(\mathbf{I} + \tau^2\mathbf{D})^{-1}\mathbf{P}\mathbf{E}\mathbf{P}^\top(\mathbf{I} + \tau^2\mathbf{D})^{-1}\|_F \\ &\leq \|\mathbf{P}\mathbf{E}\mathbf{P}^\top\|_F = \|\mathbf{E}\|_F, \end{aligned} \quad (27)$$

where $\|\cdot\|_F$ denotes the Frobenius matrix norm. The inequality also holds for the absolute value of the determinant and p norms. And the equality holds if and only if $\tau^2 = 0$ when the difference is the same as the error matrix for response NNGP model. Thus, the latent model tends to shrink the error from the Vecchia approximation, which explains the expected superior performance of the latent NNGP model over the response NNGP model based on KL-Ds.

III. SIMULATION

Our proposed models were implemented in Julia 1.2.0 (Bezanson et al., 2017). All models were run on a Linux environment (Ubuntu 18.04.2 LTS), with 32 Gbytes of random-access memory and 1 Intel Core i7-7700K CPU @ 4.20GHz processor with 4 cores each and 2 threads per core - totaling 8 possible threads for use in parallel. Model diagnostics and other posterior summaries were implemented within the Julia statistical environment and R 3.6.1. Each model was compared in terms of the posterior inference of parameters (posterior mean and 95% confidence interval), root mean squared predict error (RMSPE = $n^{-1} \sum_{i=1}^n ((y_j(\mathbf{s}_i) - \hat{y}_j(\mathbf{s}_i))^2)^{\frac{1}{2}}, j = 1, \dots, q$), mean squared error of intercept-centered latent processes (MSEL = $n^{-1} \sum_{i=1}^n ((\omega_j(\mathbf{s}_i) + \beta_{1j} - \hat{\omega}_j(\mathbf{s}_i) - \hat{\beta}_{1j})^2)^{\frac{1}{2}}, j = 1, \dots, q$), prediction interval coverage (CVG; the percent of intervals containing the true value), interval coverage for intercept-centered latent process of observed response (CVGL), mean continuous rank probability score (MCRPS = $n'^{-1} \sum_{i=1}^{n'} \text{CRPS}_j(u_i), j = 1, \dots, q$, where $\text{CRPS}_j(\mathbf{u}_i)$ is the CRPS of j -th response on held location \mathbf{u}_i see Gneiting and Raftery (2007)), and run time. To calculate $\text{CRPS}_j(\mathbf{u}_i)$, we assumed the associated predictive distribution was well approximated by a Gaussian distribution with mean centered at the predicted value $\hat{y}_j(\mathbf{u}_i)$ and standard deviation equal to the predictive standard error $\hat{\sigma}_j(\mathbf{u}_i)$, $\text{CRPS}_j(\mathbf{u}_i) = \hat{\sigma}_j(\mathbf{u}_i) [1/\sqrt{\pi} - 2\varphi(z_{ij}) - z_{ij}(2\Phi(z_{ij}) - 1)]$ with $z_{ij} = (y_j(\mathbf{u}_i) - \hat{y}_j(\mathbf{u}_i))/\hat{\sigma}_j(\mathbf{u}_i)$, φ and Φ denoting the probability density function and the cumulative distribution function of a standard Gaussian variable. All NNGP models in this section specified at most $m = 10$ nearest neighbors.

We simulated $\mathbf{y}(\mathbf{s})$ using model (8) with $q = 2, p = 2$ over 1200 randomly generated locations inside a unit square. The design matrix \mathbf{X} consisted of a column of 1's and a single predictor generated from a standard normal distribution. An exponential covariance function with decay ϕ was used to model $\rho_\psi(\cdot, \cdot)$ in (8), i.e., $\rho_\psi(\mathbf{s}', \mathbf{s}'') = \exp(-\phi \|\mathbf{s}' - \mathbf{s}''\|)$, for $\mathbf{s}', \mathbf{s}'' \in \mathcal{D}$, with $\|\mathbf{s}' - \mathbf{s}''\|$ be the L^2 norm of $\mathbf{s}' - \mathbf{s}''$ and $\psi = \phi$. The value of parameters are listed in table 1. We withheld 200 locations to assess predictive performance for conjugate models and benchmark models. NNGP based BSLMC model was also tested here for a comparison.

We assigned a flat prior for β , the prior of Σ was set to follow $\text{IW}(\Psi, \nu)$ with $\Psi = \text{diag}([1.0, 1.0])$ and $\nu = 3$. The candidate values for $\{\phi, \alpha\}$ used in cross-validation algorithm were a 25 by 25 grid over $[2.12, 26.52] \times [0.8, 0.99]$. The posterior inference of conjugate response and latent NNGP models were based on 500 samples. The run times for conjugate models include the time for choosing hyper-parameters through cross-validation and the time for the sampling process. We summarize posterior inference for regression coefficients $\beta = \{\beta_{ij}\}_{i=1, j=1}^{p, q}$, covariance of measurement error (labeled as $\text{cov}(\epsilon)$ in summary table), covariance across different latent processes (labeled as $\text{cov}(\omega)$ in summary table) and hyperparameters $\{\phi, \alpha\}$ in Table 1.

Table 1 lists the parameter estimates and performance metrics of the candidate models. The posterior inference of regression slopes $\{\beta_{21}, \beta_{22}\}$ are close among two models. The 95% confidence intervals of the intercepts $\{\beta_{11}, \beta_{12}\}$ all include the actual value.

The interpolated map of the recovered posterior mean of latent processes (figure 1) capture the patterns of the underlying latent processes. The conjugate NNGP models all yielded close RMSPEs and MCRPSs. The CVG and CVGL are close to 0.95, supporting reliable inference from conjugate NNGP models. The two models finished within a minute. The simulation example shows that fitting a conjugate model is a pragmatic method for quick inference in multivariate spatial data analysis.

Table 1: Simulation 1 study summary table: posterior mean (2.5%, 97.5%) percentiles

	True	Conj resp	Conj latent
β_{11}	1.0	1.391 (0.814, 1.902)	1.459 (0.865, 2.057)
β_{12}	-1.0	0.813 (0.344, 1.286)	0.734(0.201, 1.276)
β_{21}	-2.0	-1.978 (-2.114, -1.841)	-1.979 (-2.121, -1.842)
β_{22}	2.0	2.076 (1.952, 2.21)	2.082 (1.961, 2.208)
$\text{cov}(\epsilon)_{11}$	0.222	0.226 (0.205, 0.248)	0.231 (0.212, 0.252)
$\text{cov}(\epsilon)_{12}$	-0.111	-0.113(-0.129, -0.099)	-0.115(-0.128, -0.103)
$\text{cov}(\epsilon)_{22}$	0.167	0.172 (0.158, 0.188)	0.175 (0.16, 0.189)
$\text{cov}(\omega)_{11}$	1.234	-	1.208 (1.148, 1.268)
$\text{cov}(\omega)_{12}$	-0.701	-	-0.705(-0.75, -0.658)
$\text{cov}(\omega)_{22}$	1.077	-	1.077 (1.023, 1.131)
ϕ	6.0	8.220	7.204
α	0.9	0.863	0.871
RMSPE ^a	-	[0.727, 0.602, 0.668]	[0.723, 0.6, 0.664]
MSEL	-	-	[0.112, 0.112, 0.103]
CVG ^a	-	[0.935, 0.955, 0.945]	[0.925, 0.95, 0.9375]
CVGL ^a	-	-	[0.957, 0.945, 0.951]
MCRPS ^a	-	[-0.408, -0.336, -0.372]	[-0.405, -0.334, -0.37]
time(s)	-	[12, 1] ^b	[17, 1] ^b

^a[response 1, response 2, all responses]

^b[time for cross-validation, time for sampling]

^c[time for MCMC sampling, time for recovering β and predictions]

IV. VEGETATION INDICES DATA ANALYSIS

We implement all proposed models on a real dataset to test their performances in a realistic analysis scenario. Our dataset comprises Vege Indices data and landcover data (see Ramon Solano et al., 2010; Sulla-Menashe and Friedl, 2018, for further details). The Vege Indices data records the standard Normalized Difference Vegetation Index (NDVI) and Enhanced Vegetation Index (EVI). These two indices are robust, empirical measures of vegetation activity at the land surface, that are studied for an understanding of the global distribution of vegetation types as well as their biophysical and structural properties and spatial/temporal variations (Ramon Solano et al., 2010). Provided along with the two vegetation indexes are red reflectance, near-infrared (NIR) reflectance, blue reflectance mid-infrared (MIR) reflectance, view zenith angle, sun zenith angle and relative azimuth angle. All data were mapped to Euclidean coordinates using the Sinusoidal (SIN) grid projection. We chose zone *h08v05* which covers 11,119,505 to 10,007,555 meters south of the prime meridian and 3,335,852 to 4,447,802 meters north of the equator. The land region in zone *h08v05* is the western United States. We generated a dummy variable for no vegetation or urban area through the 2016 landcover data, and took it along with the intercept as the explanatory variables in the analysis. All other data were measured through MODIS satellite over a 16-days period from 2016.04.06 to 2016.04.21. Some variables were rescaled and transformed in exploratory data analysis for the sake of better model fitting. The datasets were downloaded using the R package *MODIS*, and the code for exploratory data analysis is provided on https://github.com/LuZhangstat/Conj_Multi_NNGP.

There are 3,115,934 observed locations. We chose transformed NDVI ($\log(\text{NDVI} + 1)$) labeled




Figure 1: Interpolated maps of (a) & (c) the true generated latent processes and the posterior means of the spatial latent process ω from the (b) & (d) conjugate latent NNGP model. The NNGP based models were all fit using $m = 10$ nearest neighbors.

as NDVI) and red reflectance (red refl) as responses. Bayesian linear models were fitted for comparison. All NNGP based models specified at most $m = 10$ nearest neighbors. We randomly held out 1% of observed locations and then held all responses over region 10,400,000 to 10,300,000 meters south of the prime meridian and 3,800,000 to 3,900,000 meters north of the equator to examine the predictive performance of models on randomly missing locations and a missing region. There were in total 67,132 locations held for prediction. Figure 2a illustrates the map of the transformed NDVI data. The white square region within the Continent is the region held out for prediction.

Posterior inference from our conjugate models were based on 500 independent samples from the posterior distribution. Recall that the samples are directly drawn from the conjugate posterior distribution and, hence, there is no need to monitor convergence of these samples. The priors for all parameters, except the decay, follow those in the simulation section. We recursively shrink the domain and grid of candidate values $\{\phi, \alpha\}$ through repeatedly using cross-validation algorithms for fixing parameters. The recorded run time for running the cross-validation algorithms, therefore, varied a lot across different models. The number of threads used in the cross-validation algorithm for conjugate models and response NNGP models with misalignment were equal to the number of folds. The remaining part of all the code were run with single thread.

The results for the conjugate models are listed in Table 2. Consistent with the related background, the regression coefficients of the index of no vegetation or urban area show relatively low biomass (low NDVI) and high red reflectance over no vegetation or urban area. The inference of the covariance of the noise and non-spatial covariance of the latent process shows a negative association between the residuals and latent processes of transformed NDVI and red reflectance, which satisfies the underlying relationship between two responses. The maps of the latent processes recovered by conjugate latent NNGP shown in Figure 2 also support this relationship.

Model performances were compared in terms of RMSPE, CVG, MCRPS and run time. The spatial models, unsurprisingly, greatly improved predictive accuracy. Conjugate Bayesian spatial models effected 35% shrinkage over the (non-spatial) Bayesian linear model in the magnitude of RMSPE. The performance in terms of the CVG is similar among all the models, but all the spatial models provided more accurate predictions than the Bayesian linear models based on MCRPS. Visual inspections of the predictive surfaces based on conjugate response NNGP model are shown in Figure 2. Notably, the proposed methods smooth out the predictions in the held-out region.

Posterior sampling for the conjugate response and latent models cost 1.8 and 18.88 minutes, respectively, which is impressive given our sample sizes of around 3 million locations. The run time for both the cross-validation algorithm and sampling for conjugate models is appealing for such massive datasets.

Table 2: Real data analysis summary table 1: posterior mean (2.5%, 97.5%) percentiles

	Bayesian linear model	conj response	conj latent
intercept ₁	0.25144 (0.25131, 0.25158)	0.1023(0.0822, 0.1223)	0.240729 (0.240723, 0.240736)
intercept ₂	0.13951 (0.13944, 0.13958)	0.2218(0.2094, 0.2338)	0.144277 (0.144273, 0.144281)
no vege or urban area ₁	-0.13375(-0.13425, -0.13329)	-8.010e-3(-8.233e-3, -7.796e-3)	-8.025e-3 (-8.050e-3, -8.001e-3)
no vege or urban area ₂	6.026e-2(6.002e-2, 6.052e-2)	4.381e-3 (4.261e-3, 4.514e-3)	4.390e-3 (4.376e-3, 4.402e-3)
cov(ϵ) ₁₁	1.599e-2 (1.596e-2, 1.602e-2)	3.493e-5 (3.487e-5, 3.499e-5)	3.125e-5 (3.120e-5, 3.130e-5)
cov(ϵ) ₁₂	-6.494e-3 (-6.505e-3, -6.483e-3)	-1.214e-5 (-1.217e-5, -1.212e-5)	-1.086e-5 (-1.089e-5, -1.085e-5)
cov(ϵ) ₂₂	3.656e-3 (3.651e-3, 3.662e-3)	1.090e-5(1.089e-5, 1.092e-5)	9.760e-6 (9.745e-6, 9.776e-6)
cov(ω) ₁₁	-	7.776e-2 (7.764e-2, 7.789e-2)	1.7192e-2 (1.7190e-2, 1.7193e-2)
cov(ω) ₁₂	-	-2.703e-2 (-2.709e-2, -2.697e-2)	-7.0307e-3(-7.0314e-3, -7.03e-3)
cov(ω) ₂₂	-	2.428e-2(2.424e-2, 2.432e-2)	3.8897e-3 (3.8893e-3, 3.8901e-3)
ϕ	-	17.919 ($\alpha = 0.999551$)	20.1755 ($\alpha = 0.999551$)
RMSPE ^a	[0.09899 0.04932 0.07821]	[0.05707 0.03187 0.04622]	[0.0503 0.02572 0.03995]
MCRPS ^a	[-0.05588 -0.02818 -0.04203]	[-0.03301 -0.0188 -0.02591]	[-0.0314 -0.01748 -0.02444]
CVG ^a	[0.9664 0.9847 0.9755]	[0.9756 0.9707 0.9732]	[0.9764 0.9715 0.974]
time(mins) ^b	-	[1012.18, 1.8]	[270.28, 18.88]

^a[1st response transformed NDVI, 2nd response red reflectance, all responses]

^b[time for cross-validation, time for generating 500 samples]

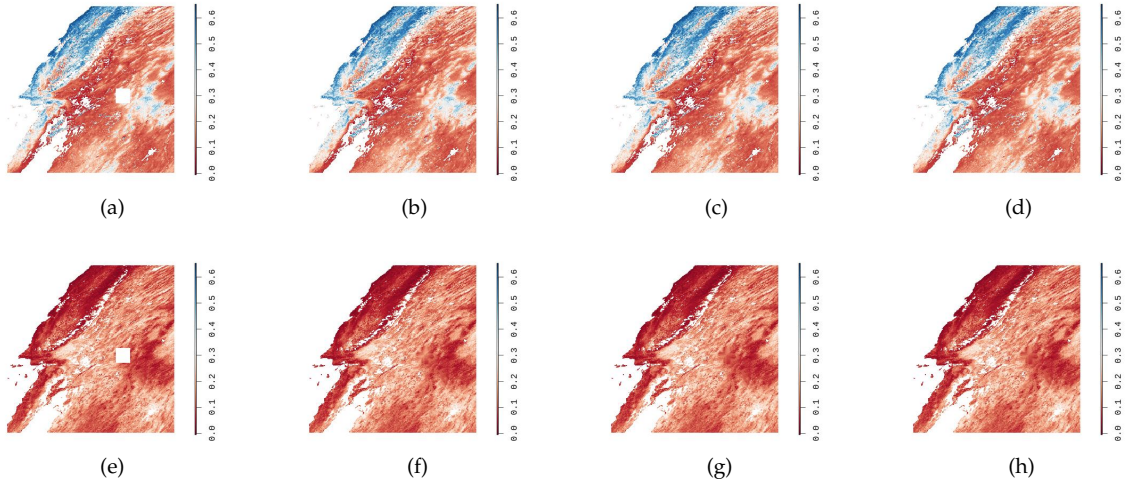


Figure 2: Colored NDVI and red reflectance images (first and second row respectively) of western United States (zone h08v05). Maps of raw data (a) & (e), raw data with predictions fitted by NNGP based conjugate response model (b) & (f), raw data with predictions fitted by NNGP based conjugate latent model (c) & (g) and the posterior mean of the intercept-centered latent process recovered from NNGP based conjugate latent model (d) & (h).

V. SUMMARY AND DISCUSSION

We have presented a conjugate Bayesian multivariate spatial regression model using Matrix-Normal and Inverse-Wishart distributions. A specific contribution is to embed the latent spatial process within an augmented Bayesian multivariate regression to obtain posterior inference for the high-dimensional latent process with stochastic uncertainty quantification. For scalability to massive spatial datasets—our examples here comprise locations in the millions—we adopt the increasingly popular Vecchia approximation and, more specifically, the NNGP models that render savings in terms of storage and floating point operations. We present elaborate simulation experiments to test the performance of different models using datasets exhibiting different behaviors. Our conjugate modeling framework fixes hyperparameters using a K -fold cross-validation approach. While our analysis is based upon fixing these hyperparameters, the subsequent inference obtained is seen to be effective in capturing the features of the generating latent process (in our simulation experiments) and is orders of magnitude faster than iterative alternatives at such massive scales as ours. We also applied our models, and compared them, in our analysis of an NDVI dataset. The scalability of our approach is guaranteed when univariate scalable model can exploit a tractable precision or covariance matrix. Our approach can, therefore, incorporate other methods such as multiresolution approximation (MRA) and more general Vecchia-type of approximations (see, e.g. Katzfuss and Guinness, 2017).

Future work can extend and adapt this framework to univariate and multivariate spatiotemporal modeling. A modification is to use a dynamic nearest-neighbor Gaussian process (DNNGP) (Datta et al., 2016b) instead of the NNGP in our models, which dynamically learns about space-time neighbors rather than fixing them. We can also develop conjugate Bayesian modeling frameworks for spatially-varying coefficient models, where the regression coefficients β are themselves random fields capturing the spatially-varying impact of predictors on the vector of outcomes. While

conceptually straightforward, their actual implementation at massive scales will require substantial development.

Developments in scalable statistical models must be accompanied by explorations in high performance computing. While the algorithms presented here are efficient in terms of storage and flops, they have been implemented on modest hardware. Implementations exploiting Graphical Processing Units (GPUs) and parallel CPUs can be further explored. For the latent NNGP models, the algorithms relied upon sparse solvers such as conjugate gradients and LSMR matrix algorithms. Adapting such libraries to GPUs and other high performance computing hardware will need to be explored and tested further in the context of our spatial Gaussian process models.

SUPPORTING INFORMATION

The work of the first and second authors were supported, in part, by federal grants NSF/DMS 1513654, NSF/IIS 1562303 and NIH/NIEHS 1R01ES027027

REFERENCES

- Banerjee, S. (2017). "High-Dimensional Bayesian Geostatistics." *Bayesian Analysis*, 12: 583–614.
- Banerjee, S., Carlin, B. P., and Gelfand, A. E. (2014). *Hierarchical modeling and analysis for spatial data*. CRC Press, Boca Raton, FL.
- Bezanson, J., Edelman, A., Karpinski, S., and Shah, V. B. (2017). "Julia: A fresh approach to numerical computing." *SIAM review*, 59(1): 65–98.
URL <https://doi.org/10.1137/141000671>
- Brown, P. J., Le, N. D., and Zidek, J. V. (1994). "Multivariate spatial interpolation and exposure to air pollutants." *Canadian Journal of Statistics*, 22(4): 489–509.
- Cressie, N. and Wikle, C. K. (2015). *Statistics for spatio-temporal data*. John Wiley & Sons, Hoboken, NJ.
- Datta, A., Banerjee, S., Finley, A. O., and Gelfand, A. E. (2016a). "Hierarchical Nearest-Neighbor Gaussian Process Models for Large Geostatistical Datasets." *Journal of the American Statistical Association*, 111: 800–812.
URL <http://dx.doi.org/10.1080/01621459.2015.1044091>
- Datta, A., Banerjee, S., Finley, A. O., Hamm, N. A. S., and Schaap, M. (2016b). "Non-separable Dynamic Nearest-Neighbor Gaussian Process Models for Large spatio-temporal Data With an Application to Particulate Matter Analysis." *Annals of Applied Statistics*, 10: 1286–1316.
URL <http://dx.doi.org/10.1214/16-A0AS931>
- Ding, S. and Cook, R. D. (2014). "Dimension folding PCA and PFC for matrix-valued predictors." *Statistica Sinica*, 24(1): 463–492.
- Finley, A. O., Datta, A., Cook, B. C., Morton, D. C., Andersen, H. E., and Banerjee, S. (2019). "Efficient algorithms for Bayesian Nearest Neighbor Gaussian Processes." *Journal of Computational and Graphical Statistics*, 28(2): 401–414.
- Fong, D. C.-L. and Saunders, M. (2011). "LSMR: An iterative algorithm for sparse least-squares problems." *SIAM Journal on Scientific Computing*, 33(5): 2950–2971.

-
- Gamerman, D. and Moreira, A. R. (2004). "Multivariate spatial regression models." *Journal of multivariate analysis*, 91(2): 262–281.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B. (2013). *Bayesian Data Analysis, 3rd Edition*. Chapman & Hall/CRC Texts in Statistical Science. Chapman & Hall/CRC.
- Genton, M. G. and Kleiber, W. (2015). "Cross-covariance functions for multivariate geostatistics." *Statistical Science*, 147–163.
- Gneiting, T. and Raftery, A. E. (2007). "Strictly proper scoring rules, prediction, and estimation." *Journal of the American Statistical Association*, 102(477): 359–378.
- Heaton, M. J., Datta, A., Finley, A. O., Furrer, R., Guinness, J., Guhaniyogi, R., Gerber, F., Gramacy, R. B., Hammerling, D., Katzfuss, M., et al. (2019). "A case study competition among methods for analyzing large spatial data." *Journal of Agricultural, Biological and Environmental Statistics*, 24(3): 398–425.
- Katzfuss, M. and Guinness, J. (2017). "A General Framework for Vecchia Approximations of Gaussian Processes." *arXiv preprint arXiv:1708.06302*.
- Le, N., Sun, L., and Zidek, J. V. (2001). "Spatial prediction and temporal backcasting for environmental fields having monotone data patterns." *Canadian Journal of Statistics*, 29(4): 529–554.
- Le, N. D., Sun, W., and Zidek, J. V. (1997). "Bayesian multivariate spatial interpolation with data missing by design." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 59(2): 501–510.
- Le, N. D. and Zidek, J. V. (2006). *Statistical analysis of environmental space-time processes*. Springer Science & Business Media.
- Ramon Solano, R., Didan, K., Jacobson, A., and Huete, A. (2010). "Modis Vegetation Index User's Guide." *The University of Arizona: Tucson, AZ, USA*.
- Ren, Q. and Banerjee, S. (2013). "Hierarchical factor models for large spatially misaligned datasets: A low-rank predictive process approach." *Biometrics*, 69: 19–30.
- Salvaña, M. L. O. and Genton, M. G. (2020). "Nonstationary cross-covariance functions for multivariate spatio-temporal random fields." *Spatial Statistics*, 100411.
- Schabenberger, O. and Gotway, C. A. (2004). *Statistical Methods for Spatial Data Analysis*. Chapman and Hall/CRC Press, Boca Raton, FL, first edition.
- Schmidt, A. M. and Gelfand, A. E. (2003). "A Bayesian coregionalization approach for multivariate pollutant data." *Journal of Geophysical Research: Atmospheres*, 108(D24).
- Sulla-Menashe, D. and Friedl, M. A. (2018). "User guide to collection 6 MODIS land cover (MCD12Q1 and MCD12C1) product." *USGS: Reston, VA, USA*, 1–18.
- Sun, W., Le, N. D., Zidek, J. V., and Burnett, R. (1998). "Assessment of a Bayesian multivariate interpolation approach for health impact studies." *Environmetrics: The official journal of the International Environmetrics Society*, 9(5): 565–586.

-
- Sun, Y., Li, B., and Genton, M. (2011). "Geostatistics for large datasets." In Montero, J., Porcu, E., and Schlather, M. (eds.), *Advances And Challenges In Space-time Modelling Of Natural Events*, 55–77. Berlin Heidelberg: Springer-Verlag.
- 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(3): 1155–1180.
- Vecchia, A. V. (1988). "Estimation and Model Identification for Continuous Spatial Processes." *Journal of the Royal Statistical society, Series B*, 50: 297–312.
- Zhang, L., Datta, A., and Banerjee, S. (2019). "Practical Bayesian Modeling and Inference for Massive Spatial Datasets On Modest Computing Environments." *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 12(3): 197–209.
URL <https://doi.org/10.1002/sam.11413>