

# **UCLA**

## **Research Reports**

### **Title**

Practical Bayesian Modeling and Inference for Massive Spatial Datasets On Modest Computing Environments

### **Permalink**

<https://escholarship.org/uc/item/6mr2986t>

### **Authors**

Zhang, Lu  
Datta, Abhirup  
Banerjee, Sudipto

### **Publication Date**

2018-11-07

### **Data Availability**

The data associated with this publication are available upon request.

Peer reviewed

**ARTICLE TYPE**

# Practical Bayesian Modeling and Inference for Massive Spatial Datasets On Modest Computing Environments<sup>†</sup>

Lu Zhang<sup>1</sup> | Abhirup Datta<sup>2</sup> | Sudipto Banerjee<sup>\*3</sup><sup>1</sup>Department of Biostatistics, University of California Los Angeles, California, USA<sup>2</sup>Department of Biostatistics, Johns Hopkins University, Maryland, USA<sup>3</sup>Department of Biostatistics, University of California Los Angeles, California, USA**Correspondence**

\*Sudipto Banerjee, UCLA Department of Biostatistics. Email: sudipto@ucla.edu

**Present Address**

650 Charles E. Young Drive South, Los Angeles, CA 90095-1772.

**Summary**

With continued advances in Geographic Information Systems and related computational technologies, statisticians are often required to analyze very large spatial datasets. This has generated substantial interest over the last decade, already too vast to be summarized here, in scalable methodologies for analyzing large spatial datasets. Scalable spatial process models have been found especially attractive due to their richness and flexibility and, particularly so in the Bayesian paradigm, due to their presence in hierarchical model settings. However, the vast majority of research articles present in this domain have been geared toward innovative theory or more complex model development. Very limited attention has been accorded to approaches for easily implementable scalable hierarchical models for the practicing scientist or spatial analyst. This article devises massively scalable Bayesian approaches that can rapidly deliver inference on spatial process that are practically indistinguishable from inference obtained using more expensive alternatives. A key emphasis is on implementation within very standard (modest) computing environments (e.g., a standard desktop or laptop) using easily available statistical software packages without requiring message-parsing interfaces or parallel programming paradigms. Key insights are offered regarding assumptions and approximations concerning practical efficiency.

**KEYWORDS:**

Bayesian inference, Gaussian processes, Latent spatial processes, Nearest-neighbor Gaussian processes

## 1 | INTRODUCTION

Rapidly increasing usage and growing capabilities of Geographic Information Systems (GIS) have spawned considerable research in modeling and analyzing spatial datasets in diverse disciplines including, but not limited to, environmental sciences, economics, biometry and so on [see, e.g., 13, 7, 2]. Much of spatial modeling is carried out within the familiar hierarchical modeling paradigm,

$$[\text{data} \mid \text{process}] \times [\text{process} \mid \text{parameters}] \times [\text{parameters}]. \quad (1)$$

For point-referenced data sets, where spatial locations are indexed by coordinates on a map, the “process” is modeled as a spatial random field over the domain of interest and the observations are treated as a finite realization of this random field. The Gaussian process (GP) is, perhaps, the most conspicuous of process specifications and offers flexibility and richness in modeling. The GP’s popularity as a modeling tool is enhanced due to their extensibility to multivariate and spatial-temporal geostatistical settings,

although we do not pursue such generalizations in this article. They also provide comparatively greater theoretical tractability among spatial processes [32].

Fitting GPs incur onerous computational costs that severely hinders their implementation for large datasets. The key bottleneck stems from the massive spatial covariance matrix present in the multivariate normal density for the finite realizations of the GP. For irregularly situated spatial locations, as is common in geostatistics, these matrices are typically dense and carry no exploitable structure to facilitate computations. Even for a modestly large number of points ( $\approx 50,000$  or greater), the computational demands become prohibitive for a modern computer and preclude inference from GP models.

A substantial literature exists on methodologies for massive spatial datasets and it is already too vast to be summarized here [see, e.g., 1, 18, and references therein]. Some are more amenable than others to the hierarchical setup in (1). Even within the hierarchical paradigm, there is already a burgeoning literature on massively scalable spatial process models. There are two pressing issues facing the practicing spatial analyst. The first is to analyze massive amounts of spatial data on “modest” computing environments such as standard desktop or laptop architectures. By “modest computing environments”, we mean environments such as R (<http://cran.r-project.org>) or STAN (<http://mc-stan.org>) that do not require knowledge of lower-level languages or distributed programming paradigms. The second pressing issue is that of *full inference* that subsumes parameter estimation, spatial prediction of the outcome, and estimation of the underlying latent process. Yet the size of the datasets easily exceed the CPU memory available for computing, which means that we need to rely upon statistical models that will enable analysis with the available memory.

Some scalable processes such as the multi-resolution predictive process models proposed by [20] or the nearest-neighbor Gaussian process (NNGP) models by [8] can be programmed in modest computing environments to estimate parameters and predict outcomes, but not necessarily infer on the latent process efficiently. [20] does not address this, while [8] and [9] implement high-dimensional Gibbs sampling algorithms that had to be run for several iterations on a high-performance computing environment to yield adequate convergence due to high autocorrelations. Other approaches such as Gaussian Markov random field (GMRF) approximations to spatial processes [29, 23] use Integrated Nested Laplace Approximations (INLA) for computing the marginal distribution of the process at given locations. These approximations can be implemented on standard environments for a variety of spatial models using the R-INLA software ([www.r-inla.org](http://www.r-inla.org)), but may be difficult to program for new unconventional process models within (1).

This article outlines strategies for achieving fully model-based Bayesian inference including parameter estimation, response surface predictions and interpolation of the latent spatial process for massive spatial datasets on modest computing environments. To achieve this goal, we need a massively scalable spatial process that will be able to estimate (1) by obviating the memory obstacles. Here, there are a few choices that are well-suited for (1) all of whom seem to be competitive based upon the recent “contest” paper by [18], but we opt for the sparsity-inducing Nearest-neighbor Gaussian process (NNGP) primarily because of its ease of use and also because of its easier accessibility through the `spnngp` package available from [cran.r-project.org/web/packages/spnngp](http://cran.r-project.org/web/packages/spnngp) (see Section 2).

In fact, [11] outline several strategies for estimating NNGP models, including a conjugate Bayesian NNGP response model for exact inference without requiring MCMC that was fitted to a dataset with approximately 5 million locations in a matter of seconds. However, the approach did not accommodate the latent process and, hence, did not generalize to (1). For the latent process models of [8], [11] used the collapsed model after integrating out the latent effects and exploited permutation-based sparse Cholesky methods embedded within MCMC algorithms. This still proved too expensive for massive datasets in the order of  $10^6$  locations. Our current contribution lies in casting the latent process models of [8] within a conjugate Bayesian framework for exact inference avoiding MCMC. Like [11] we will exploit conjugacy in conjunction with cross-validatory estimation of process parameters, but we need to be careful that the model formulation and computations will not require loading large data objects into memory at any point. In addition to the latent process, we also show how the response model, also analyzed by [11] on much more sophisticated computing environments, can be executed on very modest architectures. The details of the conjugate Bayesian formulation and algorithms for their effective implementation constitute the novelty of this paper.

The remainder of the paper evolves as follows. Section 2 gives a brief review of nearest-neighbor Gaussian process and NNGP based models. Section 3 discussed the conjugate NNGP based models. We propose the conjugate latent NNGP model, and provide the algorithms for practical implementation. The simulation study is in Section 4 for discussing the performance of the proposed models. A data analysis on sea surface temperature with over 2.5 million locations is conducted in Section 5. Finally, we conclude with some discussion in Section 6.

## 2 | THE NEAREST-NEIGHBOR GAUSSIAN PROCESS

The computational burden in GP models arises from the  $n \times n$  covariance matrix  $C_\theta(S, S)$ , where  $S = \{s_1, s_2, \dots, s_n\}$  is the set of observed locations and  $n$  is large. The  $(i, j)$ -th element of this matrix is the value of a spatial covariance function evaluated at locations  $s_i$  and  $s_j$ . Spatial covariance functions in general do not produce exploitable structures in the resulting matrices. One effective approach to achieve efficient computations is to replace  $C_\theta(S, S)$  with an approximate  $\tilde{C}_\theta(S, S)$  such that the inverse of  $\tilde{C}_\theta(S, S)$  is sparse. There are multiple options, but notable among them are approximations based upon Gaussian Markov random fields or GMRFs [see, e.g., 28, 29] that yield computationally efficient sparse representations. An alternative approach exploits an idea familiar in graphical models or Bayesian networks [see, e.g., 22, 5, 24] that has also been exploited by [35], [33] and [34] to construct composite likelihoods for inference. Datta et al. [8 9] extended this idea to construct a Nearest Neighbor Gaussian Process (NNGP) for modeling large spatial data. NNGP is a well defined Gaussian Process that yields finite dimensional Gaussian densities with sparse precision matrices. It delivers massive scalability both in terms of parameter estimation and spatial prediction or “kriging”.

### 2.1 | Response NNGP model

Consider modeling a point-referenced outcome as a partial realization of a Gaussian process,  $\{y(s) : s \in D\} \sim GP(m_\theta(s), C_\theta(\cdot, \cdot))$  on a spatial domain  $D \in \mathfrak{R}^d$ . The mean and covariance functions are assumed to be determined by one or more parameters in a set  $\theta$ . The finite-dimensional distribution for the  $n \times 1$  vector  $y(S)$  with elements  $y(s_i)$  is multivariate normal with mean  $m_\theta(S)$  and covariance matrix  $C_\theta(S, S)$ . As a directed acyclic graph (DAG) [5], the joint density is  $p(y(S)) = \prod_{i=1}^n p(y(s_i) | y(\text{Pa}[s_i]))$ , where  $\text{Pa}[s_1]$  is the empty set and  $\text{Pa}[s_i] = \{s_1, s_2, \dots, s_{i-1}\}$  for  $i = 2, 3, \dots, n-1$  is the set of parent nodes with directed edges to  $s_i$ . [35] suggested approximating the multivariate normal likelihood by shrinking  $\text{Pa}[s_i]$  from the set of all nodes preceding  $s_i$  to a much smaller subset of locations preceding  $s_i$  that are among the  $m$  (a fixed small number) nearest neighbors of  $s_i$  based upon their Euclidean distance. [8] extended that notion to arbitrary points in the domain by defining

$$\text{Pa}[s] = \begin{cases} \text{empty set} & \text{if } s = s_1, \\ \{s_1, s_2, \dots, s_{i-1}\} & \text{if } s \in S \text{ and } i = 1, 2, \dots, m, \\ m \text{ closest points to } s \text{ among } \{s_1, s_2, \dots, s_{i-1}\} & \text{if } s \in S \text{ and } i > m, \\ m \text{ closest points to } s \text{ among } S & \text{if } s \notin S. \end{cases}$$

for any arbitrary point  $s$  in the domain, where  $m$  is the fixed number of nearest neighbors. This results in another multivariate Gaussian density

$$p(y(S)) = N(y(S) | m_\theta(S), C_\theta(S, S)) \approx N(y(S) | m_\theta(S), \tilde{C}_\theta(S, S)), \quad (2)$$

where  $\tilde{C}_\theta(S, S)^{-1} = (I - A_S)^\top D_S^{-1} (I - A_S)$  is sparse,  $A_S$  is sparse and strictly lower triangular with  $A_S(i, i) = 0$  for  $i = 1, 2, \dots, n$  and at most  $m$  non-zero entries in each row, and  $D_S$  is diagonal whose elements are the conditional variances  $\text{var}\{y(s_i) | y(\text{Pa}[s_i])\}$  based upon the full GP model, i.e.,  $D_S(1, 1) = C_\theta(s_1, s_1)$  and  $D_S(i, i) = C_\theta(s_i, s_i) - C_\theta(s_i, \text{Pa}[s_i])C_\theta(\text{Pa}[s_i], \text{Pa}[s_i])^{-1}C_\theta(\text{Pa}[s_i], s_i)$  for  $i = 2, \dots, n$ . Turning to the structure of  $A_S$ , all its elements are completely determined from  $C_\theta(S, S)$ . Its first row, i.e.,  $A_S(1, \cdot)$  has all zeroes. For the  $i+1$ -th row, the nonzero entries appear in the positions indexed by  $\text{Pa}[s_{i+1}]$  and are obtained as row vectors,

$$A_S(i+1, \text{Pa}[s_{i+1}]) = C_\theta(s_{i+1}, \text{Pa}[s_{i+1}])C_\theta(\text{Pa}[s_{i+1}], \text{Pa}[s_{i+1}])^{-1}.$$

The nonzero entries in each row of  $A_S$  are precisely the “kriging” weights of  $y(s_i)$  based upon the values of  $y(s)$  at neighboring locations, i.e.,  $\text{Pa}[s_i]$  [6]. The  $\tilde{C}_\theta(S, S)$ , constructed as above, is called an NNGP approximation to  $C_\theta(S, S)$ .

With the above definition of  $\text{Pa}[s]$ , we can express the partial realizations of an NNGP as a linear model. Let  $S$  be the set of the  $n$  observed locations as defined earlier (and  $n$  is assumed to be large) and let  $U = \{u_1, u_2, \dots, u_{n'}\}$  be a set of  $n'$  arbitrary locations where we wish to predict  $y(s)$ . Then,

$$\underbrace{\begin{bmatrix} y(S) \\ y(U) \end{bmatrix}}_y = \underbrace{\begin{bmatrix} m_\theta(S) \\ m_\theta(U) \end{bmatrix}}_{m_\theta} + \underbrace{\begin{bmatrix} A(S) \\ A(U) \end{bmatrix}}_A (y(S) - m_\theta(S)) + \underbrace{\begin{bmatrix} \eta(S) \\ \eta(U) \end{bmatrix}}_\eta, \quad (3)$$

where  $\eta \sim N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} D(S) & O \\ O & D(U) \end{bmatrix}\right)$ ,  $D(U)$  is  $n' \times n'$  diagonal and  $A(U)$  is sparse  $n' \times n$  formed by extending the definitions of  $D(S)$  and  $A(S)$  as

$$\begin{aligned} D(U)(i, i) &= C_\theta(u_i, u_i) - C_\theta(u_i, \text{Pa}[u_i])C_\theta(\text{Pa}[u_i], \text{Pa}[u_i])^{-1}C_\theta(\text{Pa}[u_i], u_i), \\ A(U)(i, \text{Pa}[u_i]) &= C_\theta(u_i, \text{Pa}[u_i])C_\theta(\text{Pa}[u_i], \text{Pa}[u_i])^{-1}. \end{aligned} \quad (4)$$

Each row of  $A(U)$  has exactly  $m$  nonzero entries corresponding to the column indices in  $\text{Pa}[u_i]$ . The above structure implies that  $y(s)$  and  $y(s')$  are conditionally independent for any two points  $s$  and  $s'$  outside of  $S$ , given  $y(S)$ . The parameters  $\theta$  will be estimated from the data  $y(S)$  and predictions will be carried out using the conditional distribution of  $y(U)$  given  $y(S)$ . In a Bayesian setting,  $\theta$  will be sampled from its posterior distribution  $p(\theta | y(S))$ ,

$$p(\theta) \times \left( \prod_{i=1}^n \frac{1}{\sqrt{D(S)(i, i)}} \right) \times \exp \left\{ -\frac{1}{2} z_\theta(S)^\top (I - A(S)^\top) D(S)^{-1} (I - A(S)) z_\theta(S) \right\}, \quad (5)$$

where  $z_\theta(S) = y(S) - m_\theta(S)$  and  $p(\theta)$  is the prior distribution for  $\theta$ .

Consider a specific example with the covariance function  $C_\theta(s, s') = \sigma^2 \exp(-\phi \|s - s'\|) + \tau^2 \delta_{s=s'}$ , where  $\delta_{s=s'}$  is equal to one if  $s = s'$  and 0 otherwise, and  $m_\theta = x(s)^\top \beta$  is a linear regression with spatial predictors  $x(s)$  and corresponding slope vector  $\beta$ . Then  $\theta = \{\beta, \sigma^2, \phi, \tau^2\}$  and one choice of priors could be

$$p(\theta) \propto U(\phi | a_\phi, b_\phi) \times IG(\sigma^2 | a_\sigma, b_\sigma) \times IG(\tau^2 | a_\tau, b_\tau) \times N(\beta | \mu_\beta, V_\beta),$$

where we are using standard notations for the above distributions as, e.g., in Gelman et al. [14]. The parameter space for this model is not high-dimensional and MCMC algorithms such as Gibbs sampling in conjunction with random-walk Metropolis (RWM) or HMC can be easily implemented. Other approximate algorithms such as Variational Bayes or INLA can also be used.

Once the parameter estimates (i.e., posterior samples) are obtained from (5) we can carry out predictive inference for  $y(U)$  from the posterior predictive distribution

$$p(y(U) | y(S)) = \int p(y(U) | y(S), \theta) p(\theta | y(S)) d\theta = E_{\theta | y(S)} [N(y(U) | \mu_\theta(U | \cdot), D(U))], \quad (6)$$

where  $p(y(U) | y(S), \theta)$  is an  $n'$ -dimensional multivariate normal distribution with mean  $\mu_\theta(U | \cdot) = m_\theta(U) + A(U)(y(S) - m_\theta(S))$  and conditional covariance matrix  $D(U)$ . Since  $D(U)$  is diagonal, it is easy to sample from  $p(y(U) | y(S), \theta)$ . For each  $\theta$  sampled from (5), we sample an  $n'$ -dimensional vector  $y(U)$  from  $p(y(U) | y(S), \theta)$ . The resulting  $y(U)$ 's are samples from (6). The NNGP exploits the conditional independence between the elements of  $y(U)$ , given  $y(S)$  and  $\theta$ , to achieve efficient posterior predictive sampling for  $y(U)$ . This assumption of conditional independence is not restrictive as the samples from (6) are not independent. In fact, the marginal covariance matrix of  $y(U)$ , given  $\theta$  only, is  $A(U)\tilde{C}_\theta(S, S)A(U)^\top + D(U)$ , which is clearly not diagonal.

## 2.2 | Latent NNGP model

Rather than model the outcome as an NNGP, as was done for the response model in the preceding subsection, one could use the NNGP as a prior for the latent process [8]. In fact, as discussed in Section 4 of [8], the response model does not strictly follow the paradigm in (1) and it is not necessarily possible to carry out inference on a latent or residual spatial process after accounting for the mean.

A more general setting envisions a spatial regression model at any location  $s$

$$y(s) = m_\theta(s) + w(s) + \epsilon(s), \quad \epsilon(s) \stackrel{iid}{\sim} N(0, \tau^2), \quad (7)$$

where, usually,  $m_\theta(s) = x(s)^\top \beta$  and  $w(s)$  is a latent spatial process capturing spatial dependence. Using definitions analogous to Section 2.1, we assume  $\{w(s) : s \in D\} \sim NNGP(0, \tilde{C}_\theta(\cdot, \cdot))$ , which means that for any  $S$  and  $U$ , as constructed in (3),  $w \equiv w(S \cup U)$  will have a zero-centered multivariate normal law with covariance matrix  $(I - A)^{-1} D (I - A)^{-\top}$ . The posterior distribution to be sampled from is now given by

$$p(\theta) \times N(w | 0, \tilde{C}_\theta(S, S)) \times \exp \left\{ -\frac{1}{2\tau^2} \sum_{i=1}^n (y(s_i) - m_\theta(s_i) - w(s_i))^2 \right\}. \quad (8)$$

It is easier to sample from (5) than from (8) since the parameter space in the latter includes the high-dimensional random vector  $w$  in addition to  $\theta$ . One option is to integrate out  $w$  from (8) which yields the posterior

$$p(\theta) \times \exp \left\{ -\frac{1}{2} \sum_{i=1}^n z_{\theta}(S)^{\top} (\tilde{C}_{\theta}(S, S) + \tau^2 I_n)^{-1} z_{\theta}(S) \right\}, \quad (9)$$

where  $p(\theta)$  and  $z_{\theta}(S)$  are as defined for (5). The parameter space has collapsed from  $\{\theta, w\}$  to  $\theta$ , so (9) is called the collapsed version of (8). However, the matrix  $(\tilde{C}_{\theta}(S, S) + \tau^2 I)^{-1}$  is still high-dimensional, not sparse or computationally tractable as  $\tilde{C}_{\theta}(S, S)^{-1}$ . Efficient computations will require sophisticated sparse-Cholesky algorithms in conjunction with parallel programming paradigms such as message-parsing-interfaces (MPI) or some variant thereof. Since our definition of ‘‘modest’’ does not extend to such environments, we will turn to conjugate models in the next section.

### 3 | CONJUGATE BAYESIAN MODEL

The response NNGP and latent NNGP models outlined in Sections 2.1 and 2.2, respectively, will still require iterative simulation methods such as MCMC for full Bayesian inference. Conjugate models, i.e., using conjugate priors, can provide exact Bayesian inference by exploiting analytic forms for the posterior distributions. While some specific assumptions are needed, these models are much faster to implement even for massive datasets. Here we develop conjugate NNGP models using the tractable Normal Inverse-Gamma (NIG) family of conjugate priors. We formulate a *conjugate response model* (also formulated in [11]) and is available in the spNNGP package from [cran.r-project.org/web/packages/spNNGP](http://cran.r-project.org/web/packages/spNNGP) and a new conjugate latent NNGP model. These are conjugate versions of the models described in Sections 2.1 and 2.2. We especially focus on the conjugate latent NNGP model and show how it can exploit sparsity by sampling from latent spatial processes over massive numbers of locations efficiently using a conjugate gradient algorithm for solving large sparse systems.

#### 3.1 | The NIG conjugate prior family

Let the spatial linear regression model be specified as

$$y(S) = X\beta + w(S) + \epsilon(S) \quad (10)$$

where  $y(S)$ ,  $w(S)$  and  $\epsilon(S)$  are the realization of the corresponding processes defined in (7) over the  $n$  observed locations  $S = \{s_1, \dots, s_n\}$ ,  $X$  is the  $n \times p$  matrix of regressors with  $i$ -th row being a  $1 \times p$  vector of regressors,  $x(s_i)^{\top}$  at location  $s_i \in S$ . Henceforth, we suppress the dependence of  $y$ ,  $w$ ,  $\epsilon$  and their covariance matrix on  $S$  when this will not lead to confusion. Assume that  $w \sim N(0, \sigma^2 C)$ ,  $\epsilon \sim N(0, \delta^2 \sigma^2 I_n)$ , where  $C$  and  $\delta^2 = \frac{\tau^2}{\sigma^2}$  are known. Let  $\gamma^{\top} = [\beta^{\top}, w^{\top}]$ ,  $\mu_{\gamma}^{\top} = [\mu_{\beta}^{\top}, O^{\top}]$  and  $V_{\gamma} = \begin{bmatrix} V_{\beta} & O \\ O & C \end{bmatrix}$ . The Normal-Inverse-Gamma (NIG) density yields a convenient conjugate prior,

$$p(\gamma, \sigma^2) = NIG(\gamma, \sigma^2 | \mu_{\gamma}, V_{\gamma}, a, b) = N(\gamma | \mu_{\gamma}, \sigma^2 V_{\gamma}) \times IG(\sigma^2 | a, b). \quad (11)$$

The posterior distribution of the parameters, up to proportionality, is

$$p(\gamma, \sigma^2 | y) \propto NIG(\gamma, \sigma^2 | \mu_{\gamma}, V_{\gamma}, a_{\sigma}, b_{\sigma}) \times N(y | [X : I_n]\gamma, \delta^2 \sigma^2 I_n). \quad (12)$$

The joint posterior distribution is of the form  $NIG(\mu^*, V^*, a^*, b^*)$ , where

$$\begin{aligned} y^* &= \frac{1}{\delta} y, & X^* &= \begin{bmatrix} \frac{1}{\delta} X & \frac{1}{\delta} I_n \end{bmatrix}, \\ \mu^* &= [V_{\gamma}^{-1} + X^{*\top} X^*]^{-1} (V_{\gamma}^{-1} \mu_{\gamma} + X^{*\top} y^*), \\ V^* &= [V_{\gamma}^{-1} + X^{*\top} X^*]^{-1}, \\ a^* &= a_{\sigma} + \frac{n}{2}, \\ b^* &= b_{\sigma} + \frac{1}{2} [\mu_{\gamma}^{\top} V_{\gamma} \mu_{\gamma} + y^{*\top} y^* - \mu^{*\top} V^{*-1} \mu^*]. \end{aligned} \quad (13)$$

The marginal posterior distribution of  $\sigma^2$  follows an  $IG(a^*, b^*)$  and the marginal posterior distribution of  $\gamma$  can be identified as a multivariate  $t$ -distribution with mean  $\mu^*$ , variance  $\frac{b^*}{a^*} V^*$  and degree of freedom  $2a^*$  (i.e.  $MVS-t_{2a^*}(\mu^*, \frac{b^*}{a^*} V^*)$ ). Exact Bayesian inference is carried out by sampling directly from the joint posterior density: we sample  $\sigma^2$  from  $IG(a^*, b^*)$  and then, for each sampled  $\sigma^2$ , we draw  $\gamma$  from its conditional posterior density  $N(\mu^*, \sigma^2 V^*)$ . This yields posterior samples from (12). Furthermore,

note that once the posterior samples of  $\sigma^2$  are obtained, we can obtain samples from  $p(\tau^2 | y)$  by simply multiplying the sampled  $\sigma^2$ s with  $\delta^2$ . Thus, posterior samples are obtained without recourse to MCMC or other iterative algorithms.

### 3.2 | Conjugate response NNGP model

Finley et al. [11] formulated a conjugate NNGP model for the response model described in Section 2.1. This is formed by integrating out  $w(S)$  from (10) and applying an NNGP approximation to the marginal covariance matrix of  $y(S)$ . The model can be cast as a conjugate Bayesian linear regression model

$$p(\beta, \sigma^2 | y) \propto NIG(\beta, \sigma^2 | \mu_\beta, V_\beta, a_\sigma, b_\sigma) \times N(y | X\beta, \sigma^2 \tilde{K}), \quad (14)$$

where  $\tilde{K}$  is the NNGP approximation of  $K = C + \delta^2 I$ ,  $C$  and  $\delta^2$  are as described in Section 3.13.1. Also,  $\tilde{K}^{-1} = \sigma^2(I - A(S)^\top)D(S)^{-1}(I - A(S))$  with  $A(S)$  and  $D(S)$  as described in Section 2.12.1. We will refer to (14) as the conjugate response NNGP model. Note that this model can estimate  $\{\beta, \sigma^2\}$  and also impute the outcome at unknown locations, but does not permit inference on the latent process  $w(\cdot)$  [also see Section 4 of 8]. We address this shortcoming with a new conjugate latent NNGP model in the next section.

### 3.3 | Conjugate latent NNGP model

The conjugate models in Section 3.1 works for any covariance matrix  $C$ . Here, we derive a conjugate latent NNGP model that will subsume inference on  $w(\cdot)$ . We rewrite the covariance matrix  $\tilde{C}_\theta(S, S)$  in section 2.22.2 for  $w(S)$  as  $\sigma^2 \tilde{M}_\phi$  with fixed parameter  $\phi$ . Note that  $\tilde{M}_\phi$  is the NNGP approximation of the dense matrix  $M$ , where  $C = \sigma^2 M$ . Specifically,  $\tilde{M}_\phi^{-1} = (I - A_M)^\top D_M^{-1}(I - A_M)$ , where  $A_M$  and  $D_M$  depend only on  $\phi$ . We recast the model as

$$\underbrace{\begin{bmatrix} \frac{1}{\delta} y \\ L_\beta^{-1} \mu_\beta \\ 0 \end{bmatrix}}_{y_*} = \underbrace{\begin{bmatrix} \frac{1}{\delta} X & \frac{1}{\delta} I_n \\ L_\beta^{-1} & O \\ O & D_M^{-\frac{1}{2}}(I - A_M) \end{bmatrix}}_{X_*} \underbrace{\begin{bmatrix} \beta \\ w \end{bmatrix}}_{\gamma} + \underbrace{\begin{bmatrix} \eta_1 \\ \eta_2 \\ \eta_3 \end{bmatrix}}_{\eta} \quad (15)$$

where  $L_\beta$  is the Cholesky decomposition of the  $p \times p$  matrix  $V_\beta$ , and  $\eta \sim N(0, \sigma^2 I_{2n+p})$ . The joint posterior distribution of  $\gamma$  and  $\sigma^2$  follows an NIG distribution

$$p(\gamma, \sigma^2 | y) = NIG(\gamma, \sigma^2 | \hat{\gamma}, (X_*^\top X_*)^{-1}, a_*, b_*) \quad (16)$$

where  $\hat{\gamma} = (X_*^\top X_*)^{-1} X_*^\top y_*$ ,  $a_* = a_\sigma + \frac{n}{2}$  and  $b_* = b_\sigma + \frac{1}{2}(y_* - X_* \hat{\gamma})^\top (y_* - X_* \hat{\gamma})$ . Evaluating the posterior mean of  $\gamma$  involves solving  $X_*^\top X_* \hat{\gamma} = X_*^\top y_*$ , which requires  $\mathcal{O}(\frac{1}{3}(n+p)^3)$  flops. However, when  $p \ll n$ , the structure of  $X_*$  ensures low storage complexity. Also,  $X_*^\top X_* =$

$$\begin{bmatrix} \frac{1}{\delta^2} X^\top X + L_\beta^{-\top} L_\beta^{-1} & \frac{1}{\delta^2} X^\top \\ \frac{1}{\delta^2} X & \frac{1}{\delta^2} I_n + (I_n - A_M)^\top D_M^{-1}(I_n - A_M) \end{bmatrix} \quad (17)$$

Since  $(I_n - A_M)$  has less than  $n(m+1)$  nonzero elements and each of its row has at most  $m+1$  nonzero elements, the storage of the  $n \times n$  matrix  $(I_n - A_M)^\top D_M^{-1}(I_n - A_M)$  is less than  $n(m+1)^2$ , and the computational complexity is less than  $nm + n(m+1)^2$ .

This sparsity in  $X_*^\top X_*$  can be exploited by a conjugate gradient solver [see, e.g., 16] which transforms the problem of solving a linear system into minimizing a quadratic function. Instead of direct methods such as Cholesky decompositions, conjugate gradient methods produce successive iterations that converge to the exact solution in at most  $n$  steps. The computational cost per iteration here is  $\mathcal{O}(n)$  due to the sparsity of  $X_*^\top X_*$ . Since the algorithm converges in super-linear time, a sufficiently good approximation is often obtained in a few iterations [3], hence the performance of the conjugate gradient algorithm will be competitive when  $n$  is large. This enables posterior sampling of the latent process  $w(S)$  in high-dimensional settings. The algorithm for sampling  $\{\gamma, \sigma^2\}$  from (16) using the conjugate gradient method is given below.

---

**Algorithm 1:** Sample  $\{\gamma, \sigma^2\}$  from conjugate latent NNGP model

---

1. Fixing  $\phi$  and  $\delta^2$ , obtain  $L_\beta^{-1} \mu_\beta$  and  $L_\beta^{-1}$ :

- Compute a Cholesky decomposition of  $V_\beta$  to get  $L_\beta$

$\mathcal{O}(p^3)$

- Compute  $L_\beta^{-1}$  and  $L_\beta^{-1}\mu_\beta$   $\mathcal{O}(p^2)$
- 2. Obtain the posterior mean for  $\gamma$ :
  - Construct  $A_M$  and  $D_M$  as described, for example, in Finley et al. [11]  $\mathcal{O}(nm^3)$
  - Construct  $X_*$  and  $Y_*$  from (15)  $\mathcal{O}(nm)$
  - Calculate  $X_*^\top X_*$  and  $X_*^\top y_*$   $\mathcal{O}(n(m+1)^2)$
  - Use conjugate gradient to solve  $X_*^\top X_* \hat{\gamma} = X_*^\top y_*$
- 3. Obtain posterior samples of  $\sigma^2$ 
  - Calculate  $a_*$  and  $b_*$  as given below (16)  $\mathcal{O}(n(m+4+p))$
  - Sample  $\sigma^2$  from  $IG(a_*, b_*)$
- 4. Obtain posterior samples of  $\gamma$ 
  - Generate  $u \sim N(0, \sigma^2 I_{2n+p})$
  - Calculate  $v$  by solving  $X_*^\top X_* v = X_*^\top u$  using conjugate gradient
  - Obtain  $\gamma = \hat{\gamma} + v$   $\mathcal{O}(n)$

It is readily seen that the  $v$  in step 4 follows a Gaussian distribution with variance  $\sigma^2(X_*^\top X_*)^{-1}$ . Note that Algorithm 1 implements the conjugate gradient method for an  $n+p$ -dimensional linear system in steps 2 and 4. Since  $X_*$  and  $y_*$  depend only on  $\{\phi, \delta^2\}$ , the linear equation in step 2 only need to be solved once for each choice of  $\{\phi, \delta^2\}$ .

After formulating the posterior distribution of the latent process and the remain process parameters, the posterior distributions of the replicated values  $y_{rep}(S)$  over the observed locations along with  $\sigma^2$  follow an NIG:

$$\begin{aligned} p(y_{rep}(S), \sigma^2 | y) &= \int N(y_{rep}(S) | [X : I_n]\gamma, \sigma^2 \delta^2 I_n) \times NIG(\gamma, \sigma^2 | \hat{\gamma}, (X_*^\top X_*)^{-1}, a_*, b_*) d\gamma \\ &= NIG(y_{rep}(S), \sigma^2 | \mu_{rep}, V_{rep}, a_{rep}, b_{rep}), \end{aligned} \quad (18)$$

where we have used the common assumption that the replicated data are independent of the observed data, given the model parameters, i.e.,  $p(y_{rep}(S) | \gamma, \sigma^2, y) = p(y_{rep}(S) | \gamma, \sigma^2)$ , and

$$\begin{aligned} \mu_{rep} &= [X \ I_n] \hat{\gamma}, \\ V_{rep} &= [X \ I_n] (X_*^\top, X_*)^{-1} \begin{bmatrix} X_*^\top \\ I_n \end{bmatrix} + \delta^2 I_n, \\ a_{rep} &= a_*, \quad b_{rep} = b_*. \end{aligned} \quad (19)$$

According to Section 3.1, the posterior distribution of  $y_{rep}(S)$  follows a multivariate  $t$ -distribution  $MVS-t_{2a_*}(\mu_{rep}, \frac{b_*}{a_*} V_{rep})$ . Sampling from the posterior distribution of  $y_{rep}(S)$  directly requires the expensive Cholesky decomposition of an  $n \times n$  matrix  $V_{rep}$ . However, the Cholesky decomposition can be avoided once we obtain the posterior samples of  $\gamma$  and  $\sigma^2$ . By the first equation in (18), the sampling of  $y_{rep}(S)$  follows a Gaussian distribution with mean  $[X : I_n]\gamma$  and a diagonal variance matrix. In the next section, we illustrate how the derived posterior distributions  $p(\gamma, \sigma^2 | y)$  (16) helps to formulate the posterior distribution of the response and latent process over the unobserved locations in close form, and discuss their sampling algorithms in the practical implementation.

### 3.4 | Posterior predictive inference for conjugate latent NNGP

We extend the predictive inference for the response NNGP model in Section 2.1 to the conjugate latent NNGP model. Assume  $w(U)$  and  $y(U)$  are the realization of the latent process and the response process over the  $n'$  locations  $U = \{u_1, \dots, u_{n'}\}$  where we wish to predict. let  $C_\theta(\cdot, \cdot)$  be the covariance function for the latent process  $w(s)$  in (7),  $Pa[u_i]$  be the nearest neighbors of  $i$ th location in  $U$  as defined in section 2.1. Define  $A_u = A(U)$  and  $D_u = \frac{1}{\sigma^2} D(U)$  where  $A(U)$  and  $D(U)$  are constructed by (4). Here,  $\sigma^2$  refers to the variance of the latent process  $w(s)$ , and  $A_u$  and  $D_u$  are defined in the way that they only depend on fixed parameter  $\phi$ . According to the definition of NNGP process over the whole domain given in section 2, the joint distribution of  $w(U)$  and  $\gamma, \sigma^2$  given  $y(S)$  follows:

$$\begin{aligned} p(w(U), \gamma, \sigma^2 | y(S)) &= N(w(U) | [O : A_u]\gamma, \sigma^2 D_u) \\ &\quad \times NIG(\gamma, \sigma^2 | \hat{\gamma}, (X_*^\top X_*)^{-1}, a_*, b_*) \end{aligned} \quad (20)$$



Marginalizing the joint distribution (20) over  $\gamma$  and  $\sigma^2$ , the posterior distribution of  $w(U)$  can be identified as a multivariate  $t$ -distribution:

$$w(U) | y(S) \sim \text{MVS-}t_{2a_*} \left( \mu_{wu}, \frac{b_*}{a_*} V_{wu} \right) \quad (21)$$

where

$$\mu_{wu} = [O : A_u] \hat{\gamma}, \quad V_{wu} = [O : A_u] (X_*^\top X_*)^{-1} \begin{bmatrix} O \\ A_u^\top \end{bmatrix} + D_u.$$

It is straightforward to see that the joint posterior distribution of  $\{y(U), w(U), \gamma, \sigma^2\}$  is

$$p(y(U), w(U), \gamma, \sigma^2 | y(S)) = N(y(U) | X(U)\beta + w(U), \sigma^2 \delta^2 I_{n'}) \times p(w(U), \gamma, \sigma^2 | y(S)), \quad (22)$$

which is the product of the conditional distribution of  $y(U)$  from the spatial linear regression model (7) and the posterior distribution (20). It can be shown that the posterior distribution of the predictive process  $y(U)$  and  $\sigma^2$  follows an NIG after marginalizing out  $\gamma$  and  $w(U)$ , and the posterior distribution of  $y(U)$  follows a multivariate  $t$ -distribution:

$$y(U) | y(S) \sim \text{MVS-}t_{2a_*} \left( \mu_{yu}, \frac{b_*}{a_*} V_{yu} \right) \quad (23)$$

where

$$\mu_{yu} = [X(U) : A_u] \hat{\gamma} \quad \text{and} \quad V_{yu} = [X(U) : A_u] (X_*^\top X_*)^{-1} \begin{bmatrix} X(U)^\top \\ A_u^\top \end{bmatrix} + \delta^2 I_{n'} + D_u.$$

Sampling  $w(U)$   $y(U)$  from their posterior distribution requires taking Cholesky decomposition of matrix  $V_{wu}$  and  $V_{yu}$ . Since the matrix  $(X_*^\top X_*)^{-1}$  is involved in the calculation, the required computation power is expensive and the calculation quickly become forbidden when the number of locations to predict is large. Rather than direct sampling, we recommend using a two stage sampling method based on the joint distribution (20) and (22) in this subsection. First, obtain the posterior samples  $\{\gamma^{(l)}, \sigma^{2(l)}\}_{l=1}^L$ . Then generate the posterior samples of  $w(U)$  through  $w(U)^{(l)} \sim N([O : A_u] \gamma^{(l)}, \sigma^{2(l)} D_u)$  for  $l = 1, \dots, L$ . Finally use  $y(U)^{(l)} \sim N(X(U)\beta^{(l)} + w(U)^{(l)}, \delta^2 \sigma^{2(l)})$  to generate the posterior samples of  $y(U)$ .

### 3.5 | Inference of $\phi$ and $\delta^2$

Algorithm 1 provides the exact posterior sampling of the process parameters after specifying  $\phi$  and  $\delta^2$ . This motivates us to estimate all the process parameters by first obtaining the inference of a small set of parameters  $\phi$  and  $\delta^2$ , then implement Algorithm 1 to sample  $\{\gamma, \sigma^2\}$ . When we fix  $\phi$  and  $\delta^2$  at a point estimator (i.e.  $\arg \max \{p(\phi, \delta^2 | y)\}$ ), the conjugate latent NNGP model becomes a special case of fitting latent NNGP model with Empirical Bayes method.

Here we propose a  $K$ -folder cross-validation algorithm for picking a point estimate of  $\{\phi, \delta^2\}$  of the conjugate Latent NNGP. We first split the data randomly into  $K$  folds and denote the  $k$ -th folder of the observed locations  $S[k]$ , whereas  $S[-k]$  denotes the observed locations without  $S[k]$ . Then we fit the predictive mean  $E[y(S[k]) | y(S[-k])]$  by the posterior distribution given in (23). We use the Root Mean Square Predictive Error (RMSPE) (Yeniay and Goktas [36]) to select  $\phi$  and  $\delta^2$  from a grid of candidate values. Algorithm 2 describes the  $K$ -fold cross-validation of choosing  $\phi, \delta^2$  for conjugate latent NNGP model:

**Algorithm 2:** Cross-validation of tuning  $\phi, \delta^2$  for conjugate latent NNGP model

1. Split the data into  $K$  folds, and build neighbor index.
  - Build nearest neighbors for  $S[-k]$
  - Find the collection of nearest neighbor set for  $S[k]$  among  $S[-k]$ .
2. Fix  $\phi$  and  $\delta^2$ , Obtain the posterior mean for  $\gamma_k = \{\beta, w(S[-k])\}$  after removing the  $k^{\text{th}}$  fold of the data:
  - Use step 1-2 in Algorithm 1 to obtain  $\hat{\gamma}_k$
3. Predicting posterior means of  $y(S[k])$ 
  - Construct matrix  $A_u$  for  $S[k]$
  - According to (23), the predicted posterior mean follows  $\hat{y}(S[k]) = E[y(S[k]) | y(S[-k])] = [X(U) : A_u] \hat{\gamma}$
4. Root Mean Square Predictive Error (RMSPE) over  $K$  folds
  - Initialize  $e = 0$
  - for ( $k$  in  $1 : K$ )
  - for ( $s_i$  in  $S[k]$ )
  - $e = e + (y(s_i) - \hat{y}(s_i))^2$

### 5. Cross validation for choosing $\phi$ and $\delta^2$

- Repeat steps (2) - (4) for all candidate values of  $\phi$  and  $\delta^2$
- Choose  $\phi_0$  and  $\delta_0$  as the value that minimizes the average RMSPE

The main computational burden lies in step 1 in Algorithm 2. However, step 1 serves as a pre-calculation for the whole cross-validation since it only need to be calculated for once. We recommend using KD-tree algorithm provided in R package `spNNGP` [10] to build the nearest neighbor matrices. Step 2 dominates the computational requirement in Algorithm 2 after the pre-calculation, which calls Algorithm 1 for  $k$  times for each choice of  $\{\phi, \delta^2\}$ .

An alternative choice for picking point estimation is to carry out the cross-validation with the conjugate response NNGP model. The practical advantage here is that the function `spConjNNGP` within the `spNNGP` package in R can be used to carry out the cross-validation with the conjugate response NNGP model in (14). The algorithm behind `spConjNNGP` is exactly linear in  $n$  and highly efficient for implementation. Empirical studies show that response NNGP performance similarly with the latent process model. While through a Directed Acyclic Graph (DAG) built on an augmented latent space(parameter space including latent process), [21] show that the Kullback-Leibler divergence(KL-D) Gneiting and Raftery [15] from the latent NNGP model to the full GP model is no more than that of the response NNGP model. Which suggests that the latent NNGP model is better than the response NNGP model. This raises the concern of using response NNGP model instead of latent NNGP model for model fitting. However, the claim in [21] based on KL-Ds on an augmented space (parameter space with latent process  $w$ ) is not guaranteed to hold on a collapsed space (parameter space without latent process  $w$ ). We provide an example to discuss this claim in the next subsection.

### 3.6 | Comparison of response NNGP and latent NNGP models

Consider (7) over locations  $S = \{s_1, s_2, s_3\}$ . Also assume that  $m_\theta(S) = 0$  and that  $w(S)$  has covariance matrix  $\sigma^2 C$ , where

$$C = \begin{bmatrix} 1 & \rho_{12} & \rho_{13} \\ \rho_{12} & 1 & \rho_{23} \\ \rho_{13} & \rho_{23} & 1 \end{bmatrix}. \quad (24)$$

Let us suppress the dependence between  $s_2$  and  $s_3$  in the DAG corresponding to the finite realization of the NNGP on  $S$ . Then the covariance matrix of  $y(S)$  of the response NNGP model,  $\Sigma_R$ , and that of the latent NNGP model,  $\Sigma_I$ , have the following forms:

$$\Sigma_R = \sigma^2 \begin{bmatrix} 1 + \delta^2 & \rho_{12} & \frac{\rho_{12}\rho_{23}}{1+\delta^2} \\ \rho_{12} & 1 + \delta^2 & \rho_{23} \\ \frac{\rho_{12}\rho_{23}}{1+\delta^2} & \rho_{23} & 1 + \delta^2 \end{bmatrix}, \quad \Sigma_I = \sigma^2 \begin{bmatrix} 1 + \delta^2 & \rho_{12} & \rho_{12}\rho_{23} \\ \rho_{12} & 1 + \delta^2 & \rho_{23} \\ \rho_{12}\rho_{23} & \rho_{23} & 1 + \delta^2 \end{bmatrix}, \quad (25)$$

By the sufficient and necessary condition of  $C$  being positive-definite:

$$1 - (\rho_{12}^2 + \rho_{13}^2 + \rho_{23}^2) + 2\rho_{12}\rho_{13}\rho_{23} > 0, \quad 1 - \rho_{12}^2 > 0, \quad (26)$$

it is easy to show that  $\Sigma_R$  and  $\Sigma_I$  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. This example shows that whether a latent NNGP model is a better approximation to its parent GP model than a response NNGP model on a collapsed space depends on the covariance function of the full GP.

## 4 | SIMULATION STUDY

We use a simulation study in this section to discuss the performance of the aforementioned models in Sections 2 and 3. Algorithm 1 were programmed in R which calls the Rstan environment [31] for building matrix  $A_M$  and  $D_M$ . The conjugate gradient solver for sparse linear systems was implemented through `RcppEigen` [4].The nearest-neighbor sets were built using the `spConjNNGP` function in the `spNNGP` package. All simulations were conducted on a OS High sierra system (version 10.13.4) with 16GB RAM and one 3.1 GHz Intel-Core i7 processors.

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

	True	Full GP	NNGP	Conj LNNGP
$\beta_0$	1	1.07(0.72, 1.42)	1.10 (0.74, 1.43)	1.06 (0.76, 1.46)
$\beta_1$	-5	-4.97 (-5.02, -4.91)	-4.97 (-5.02, -4.91)	-4.97 (-5.02, -4.91)
$\sigma^2$	2	1.94 (1.63, 2.42)	1.95 (1.63, 2.41)	1.94 (1.77, 2.12)
$\tau^2$	0.2	0.14 (0.07, 0.23)	0.15 (0.06, 0.24)	0.17 (0.16, 0.19)
$\phi$	16	19.00 (13.92, 23.66)	18.53 (14.12, 24.17)	17.65
KL-D	-	4.45(1.16, 9.95)	5.13(1.66, 11.39)	3.58(1.27, 8.56)
MSE(w)	-	297.45(231.62, 444.79)	303.38(228.18, 429.54)	313.28 (258.96, 483.75)
RMSPE	-	0.94	0.94	0.94
time(s)	-	2499 + 23147	109.5	12 + 0.6

#### 4.1 | Univariate simulation study

We generated data using the spatial regression model in (7) over a set of  $n = 1200$  spatial locations within a unit square. The true values of the parameters generating the data are supplied in Table 1. The size of the data set was kept moderate to permit comparisons with the expensive full GP models. The model had an intercept and a single predictor  $x(s)$  generated from a standard normal distribution. An exponential covariance function was used to generate the data.

Candidate models for fitting the data included full Gaussian process based model (labeled as full GP in Table 1), a latent NNGP model with  $m = 10$  neighbors and a conjugate latent NNGP model with  $m = 10$  neighbors. These models were trained using  $n = 1000$  of the 1200 observed locations. And the remaining 200 observations were withheld to assess predictive performance. The full Gaussian process base model was implemented with function *spLM* in R package *spBayes*. The latent NNGP model was conducted with function *spNNGP* in R package *spNNGP*. The fixed parameters  $\{\phi, \delta^2\}$  for the conjugate latent NNGP model were picked through the  $k$ -th folder cross-validation algorithm (Algorithm 2). And the choice from *spConjNNGP* coincide with the cross-validation for the conjugate latent NNGP model.

The intercept and slope parameters  $\beta$  were assigned improper flat priors. The spatial decay  $\phi$  was modeled using a fairly wide uniform prior  $U(2.2, 220)$ . We use Inverse-Gamma priors  $IG(2, b)$  (mean  $b$ ) for the nugget ( $\tau^2$ ) and the partial sill ( $\sigma^2$ ) in order to compare the conjugate Bayesian models with other models. The shape parameter was fixed at 2 and the scale parameter was set from the empirical estimate provided by the variogram using the *geoR* package [27]. The parameter estimates and performance metrics are provided in Table 1. The summaries for the full Gaussian process based model and the latent NNGP model were based on 1 MCMC chain with 20,000 iterations. The number of iterations was taken to be large enough to guarantee the convergence of the MCMC chains. We took the first half of the MCMC chains as burn-in. The inference from the conjugate latent NNGP model were based on 300 samples. 300 samples is sufficient for the conjugate latent NNGP model since the conjugate model provides independent samples from the exact posterior distribution. We don't need extra memory for burn-in, and the samples from the conjugate model are more efficient than that from MCMC algorithms.

All models were assessed by the Kullback-Leibler divergence (labeled KL-D; Gneiting and Raftery [15]) and the out-of-sample root mean squared prediction error (RMSPE) (Yeniay and Goktas [36]). The KL-D between true distribution  $Q$  and fitted distribution  $P_\theta$  is measured by:

$$d(P_\theta, Q) = \frac{1}{2} \{ \text{tr}(\Sigma_P^{-1} \Sigma_Q) - \log \det(\Sigma_P^{-1} \Sigma_Q) + (\mu_P - \mu_Q)' \Sigma_P^{-1} (\mu_P - \mu_Q) - n \} \quad (27)$$

where  $P_\theta$  and  $Q$  define Gaussian distributions on  $\mathfrak{R}^n$  with mean vectors  $\mu_P$  and  $\mu_Q$ , respectively, and covariance matrices  $\Sigma_P$  and  $\Sigma_Q$ , respectively. The KL-D in Table 1 are on the collapsed space  $\theta = \{\beta, \sigma^2, \tau^2, \phi\}$ . We estimated the KL-D by the empirical estimator:

$$E_{\theta|y(S)}(d(P_\theta, Q)) \approx \frac{1}{L} \sum_{i=1}^L d(P_{\theta_{(i)}}, Q), \quad (28)$$

where  $\theta_{(i)}, i = 1, \dots, L$  are  $L$  samples from the posterior distribution of  $\theta$ . We also present the 95% credible intervals for  $d(P_\theta, Q)$  in Table 1. The predicted outcome at any withheld location  $s_0$  was estimated as

$$\hat{y}(s_0) = E[\tilde{y}(s_0) | y(S)] \approx \frac{1}{L} \sum_{i=1}^L \tilde{y}_{\theta_{(i)}}(s_0), \quad (29)$$

where  $\tilde{y}_{\theta_{(i)}}(s_0) \sim p(y(s_0) | y(S), \theta_{(i)})$  and  $p(\cdot | y(S), \theta_{(i)})$  is the likelihood for the respective model. These were used to calculate the RMSPE using the 200 hold-out values. We randomly picked 300 out of the 10000 samples from the post burn-in MCMC chains for calculating the KL-D and RMSPE. The  $y(s_0)$  for full Gaussian process based and the latent NNGP model are sampled by function *spPredict*. For the purpose of assessing the performance of recovering spatial latent process, we also report the Mean Squared Error (MSE) with respect to the true values of the spatial latent process ( $MSE(w)$ ) over the observed locations in the simulation. The KL-D,  $MSE(w)$  and RMSPE metrics reveal that the NNGP provides a highly competitive alternative to the full Gaussian process based model.

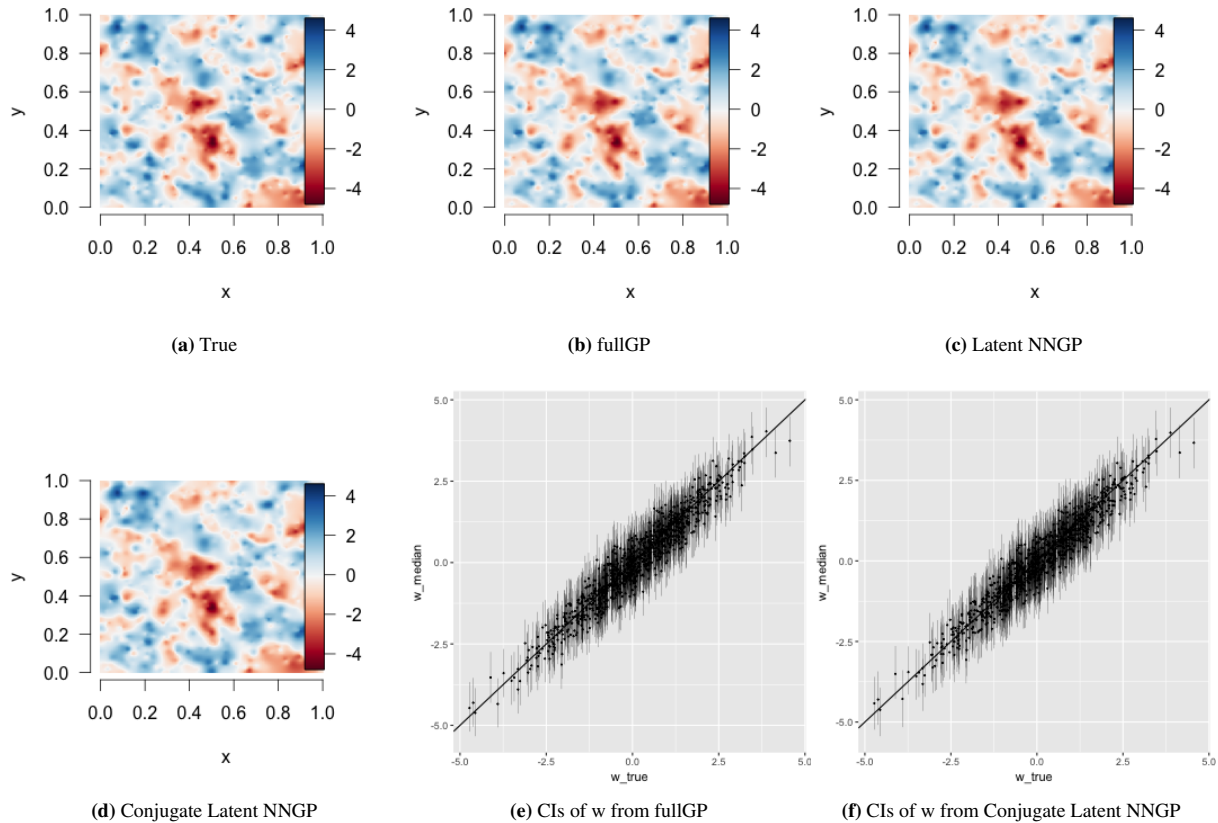
Table 1 lists the parameter estimates and performance metrics for the candidate models. The posterior inference of the regression coefficients  $\beta$  are close for all three models. While the posterior estimates of  $\{\sigma^2, \tau^2, \phi\}$  are similar for full Gaussian process based model and latent NNGP model but, somewhat expectedly, different from the conjugate latent NNGP model. The 95% confidence interval for  $\sigma^2$  and  $\tau^2$  are narrower since we fix the parameter  $\phi, \delta^2$ . The KL-Ds on the parameter space  $\{w, \beta, \tau^2\}$  show that the conjugate latent NNGP provides reliable inference for the latent process and the regression coefficients. The same RMSPE across all three models also support that conjugate latent NNGP is comparable with full Gaussian process based model in prediction. The latent NNGP model is 200 times faster than the full Gaussian process based model, while the conjugate latent NNGP model use one tenth of the time required for the latent NNGP model to obtain similar inference on the regression coefficients and latent process. Notice that the time for the sampling of the 300 samples after fixing the parameter  $\phi$  and  $\delta^2$  in the conjugate latent NNGP model is less than one second. And the conjugate latent NNGP spare the effect of testing the tuning parameters in MCMC algorithm. Based on KL-D and RMSPE, the conjugate latent NNGP models emerge as highly competitive alternatives to latent NNGP models for prediction and inference on the latent process.

Figure 1 shows interpolated surfaces from the simulation example: 1 (a) shows an interpolated map of the “true” spatial latent process  $w$ , 1 (b)–(d) are maps of the posterior means of the latent process using a full GP model, a latent NNGP model and a conjugate latent NNGP model, respectively. Figure 1 (e)–(f) present the 95% confidence intervals for  $w$  from a full GP model and a conjugate latent NNGP model. The recovered spatial residual surfaces are almost indistinguishable, and are comparable to the true interpolated surface of  $w(s)$ . Notice that the posterior mean of  $w$  of the conjugate latent NNGP model can be theoretically calculated by the  $\hat{\gamma}$  in (16). Thus the posterior samples of the latent process  $w$  is only required for measuring uncertainty. Figure 1 f provides the 95% confidence interval for all latent process  $w$  from the conjugate latent NNGP model. There are 955 out of 1000 95% confidence intervals successfully include the true value. This is comparable to the full Gaussian process based model (fig 1 e) which has 946 out of 1000 95% confidence intervals covering the true value.

## 5 | SEA SURFACE TEMPERATURE ANALYSIS

Global warming continues to be an ongoing concern among scientists. In order to develop conceptual and predictive global models, NASA monitors temperature and other atmospheric properties of the Earth regularly by two Moderate Resolution Imaging Spectroradiometer (MODIS) instruments in Aqua and Terra platforms. There is an extensive global satellite-based database processed and maintained by NASA. Details of the data can be found in <http://modis-atmos.gsfc.nasa.gov/index.html>. In particular, inferring on processes generating sea surface temperatures (SST) are of interest to atmospheric scientists studying exchange of heat, momentum, and water vapor between the atmosphere and ocean. Our aforementioned development will enable scientists to analyze large spatially-indexed datasets using a Bayesian geostatistical model easily implementable on modest computing platforms.

Model-based inference is obtained rapidly using the conjugate latent NNGP model and, based on simulation studies, will be practically indistinguishable from MCMC-based output from more general NNGP specification. The dataset we analyze here consists of 2,827,252 spatially indexed observations of sea surface temperature (SST) collected between June 18-26, 2017, the data covers the ocean from longitude  $-140^\circ$  to  $0^\circ$  and from latitude  $0^\circ$  to  $60^\circ$ . Among the 2,827,252 observations,  $n = 2,544,527$  (90%) were used for model fitting and the rest were withheld to assess predictive performance of the candidate models. Figure 2 depicts an interpolated map of the observed SST records over training locations. The temperatures are color-coded from shades of blue indicating lower temperatures, primarily seen in the higher latitudes, to shades of red indicating high temperatures. The missing data are colored by yellow and the gray part refers to land. To understand trends across the coordinates, we used sinusoidally projected coordinates (scaled to 1000km units) as explanatory variables. An exponential spatial covariance function was used for all models. Further model specifications included non-informative flat priors for the

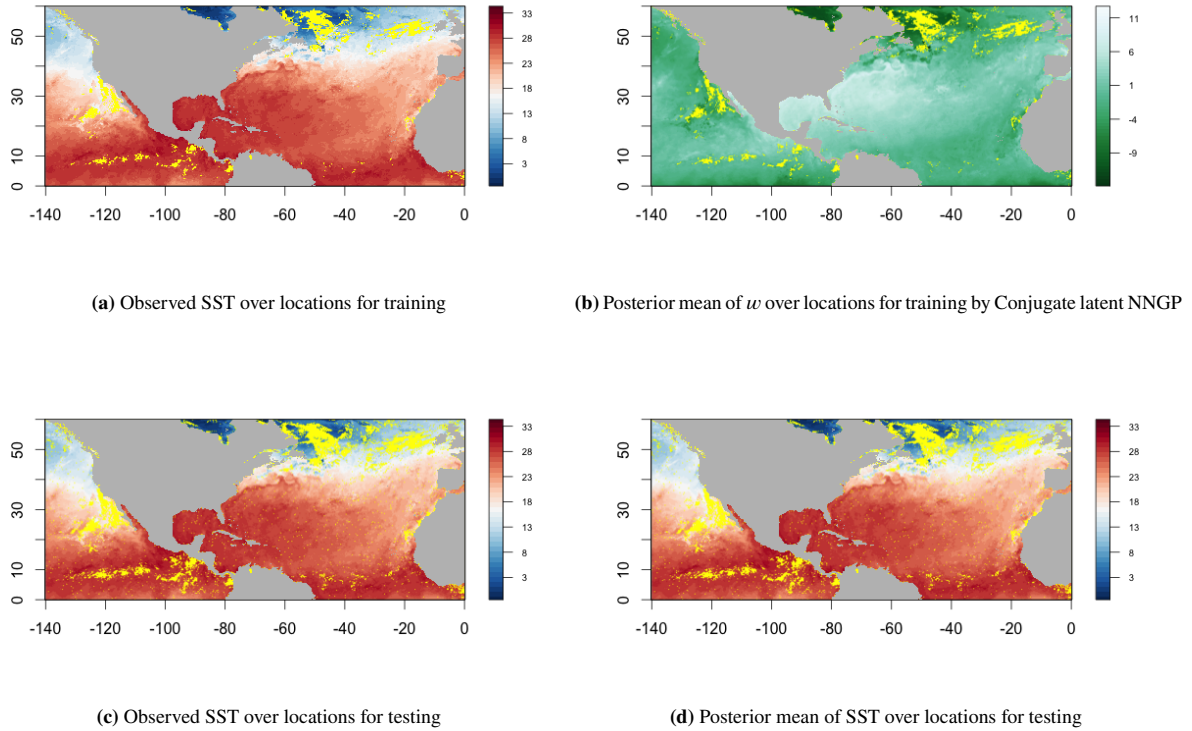


**FIGURE 1** Interpolated maps of (a) the true generated surface, the posterior means of the spatial latent process  $w(s)$  for (b) the full Gaussian Process (Full GP), (c) the latent NNGP and (d) the conjugate latent NNGP. The 95% confidence intervals for  $w$  from (e) the full GP and (f) the conjugate latent NNGP. The models in (c), and (d) were all fit using  $m = 10$  nearest neighbors.

intercept and regression coefficients, inverse-gamma priors for  $\tau^2$  and  $\sigma^2$  with shape parameter 2 and scale parameter equaling the respective estimates from an empirical variogram.

We fit the conjugate Bayesian model with fixed  $\phi$  and  $\delta^2$  using the algorithm 1 in Section 3.3 with  $m = 10$  nearest neighbors. We implement Algorithm 2 to choose the values of  $\{\phi, \delta^2\}$  at  $\phi = 7$ ,  $\delta^2 = 0.001$ . Figure 2 b shows the posterior means for the latent process of the conjugate latent NNGP model. The temperatures are color-coded from light green indicating high temperatures to dark of green indicating low temperatures. The map of the latent process  $w$  indicates lower temperature on the east coast and higher temperature on the west coast. At the same time, we observed high temperature at center of the map. These features coincide with the ocean current, suggesting that the ocean current plays an important role in the sea surface temperature.

Parameter estimates along with their estimated 95% credible intervals and performance metrics for candidate models are shown in Table 2. The RMSPE for a non-spatial linear regression model, conjugate latent NNGP model were 1.13, 0.31, respectively. Compared to the spatial models, the non-spatial models have substantially higher values of RMSPE, which suggest that coordinates alone do not adequately capture the spatial structure of SST. The fitted SST map over the withheld locations (Fig 2 d) using conjugate latent NNGP model is almost indistinguishable from the real SST map (Fig 2 c). All the inference from the conjugate latent NNGP model are based on 300 samples. The sampling process took 2367 seconds. In average, the posterior mean of the latent process  $w$  can be obtained within 20 seconds.



**FIGURE 2** Notes: (a) Observed SST over locations for training (b) Posterior mean of  $w$  over locations for training by Conjugate latent NNGP (c) Posterior mean of SST over locations for testing (d) Posterior mean of SST over locations for testing The land is colored by gray, locations in the ocean without observations are colored by yellow.

**TABLE 2** Real data analysis summary table. Parameter Posterior summary mean (2.5, 97.5) percentiles

	Non-spatial	Conjugate latent NNGP <sup>1</sup>
$\beta_0$	31.92(31.91, 31.92)	31.43 (31.28, 31.59)
$\beta_1$	0.12 (0.12, 0.12)	0.07 (0.05, 0.09)
$\beta_2$	-3.07 (-3.07, -3.07)	-3.03 (-3.08, -2.99)
$\sigma^2$	–	3.95 (3.94, 3.95)
$\phi$	–	7.00
$\tau^2$	11.44 (11.43, 11.46)	$3.95e^{-3}$ ( $3.94e^{-3}$ , $3.95e^{-3}$ )
RMSPE	1.13	0.31

## 6 | CONCLUSIONS AND FUTURE WORK

This article has attempted to address some practical issues encountered by scientists and statisticians in the hierarchical modeling and analysis for very large geospatial datasets. Building upon some recent work on nearest-neighbor Gaussian processes for massive spatial data, we build conjugate Bayesian spatial regression models and propose strategies for rapidly deliverable inference on modest computing environments equipped with user-friendly and readily available software packages. In particular, we have demonstrated how judicious use of a conjugate latent NNGP model can be effective for estimation and uncertainty quantification of latent (underlying) spatial processes. This provides an easily implementable practical alternative to computationally onerous Bayesian computing approaches. All the computations done in the paper were implemented on a standard desktop using R and Stan. The article intends to contribute toward innovations in statistical practice rather than novel methodologies.

It is important to recognize that the conjugate Bayesian models outlined here are not restricted to the NNGP. Any spatial covariance structure that leads to efficient computations can, in principle, be used. There are a number of recently proposed approaches that can be adopted here. These include, but are not limited to, multi-resolution approaches [e.g., 26, 25, 20], covariance tapering and its use in full-scale approximations [e.g., 12, 30, 19], and stochastic partial differential equation approximations [23], among several others [see, e.g., 1, and references therein].

With regard to the NNGP specifically, our choice was partially dictated by its easy implementation in R using the `spNNGP` package and in Stan as described in <http://mc-stan.org/users/documentation/case-studies/nngp.html>. The NNGP is built upon a very effective likelihood approximation [35, 33], which has also been explored recently by several authors in a variety of contexts [34, 17]. [17] provides empirical evidence about Vecchia's approximation outperforming other alternate methods, but also points out some optimal methods for permuting the order of the spatial locations before constructing the model. His methods for choosing the order of the locations can certainly be executed prior to implementing the models proposed in this article. Finally, an even more recent article by [21] proposes further extensions of the Vecchia approximation, but its practicability for massive datasets on modest computing environments with easily available software packages is yet to be ascertained.

## ACKNOWLEDGEMENTS

The authors wish to thank Dr. Michael Betancourt, Dr. Bob Carpenter and Dr. Aki Vehtari of the STAN Development Team for useful guidance regarding the implementation of non-conjugate NNGP models in Stan for full Bayesian inference. The work of the first and third authors was supported, in part, by federal grants NSF/DMS 1513654, NSF/IIS 1562303 and NIH/NIEHS 1R01ES027027.

## SUPPLEMENTARY MATERIAL

All computer programs implementing the examples in this article can be found in the public domain and downloaded from <https://github.com/LuZhangstat/ConjugateNNGP>.



## APPENDIX

### References

- [1] Banerjee, S., 2017: High-dimensional bayesian geostatistics. *Bayesian Analysis*, **12**, 583–614.
- [2] Banerjee, S., B. P. Carlin, and A. E. Gelfand, 2014: *Hierarchical modeling and analysis for spatial data*. Crc Press.
- [3] Banerjee, S. and A. Roy, 2014: *Linear algebra and matrix analysis for statistics*. CRC Press.
- [4] Bates, D. and D. Eddelbuettel, 2013: Fast and elegant numerical linear algebra using the RcppEigen package. *Journal of Statistical Software*, **52**, no. 5, 1–24.  
URL <http://www.jstatsoft.org/v52/i05/>
- [5] Bishop, C., 2006: *Pattern Recognition and Machine Learning*. Springer-Verlag, New York, NY.
- [6] Chilés, J. and P. Delfiner, 1999: *Geostatistics: Modeling Spatial Uncertainty*. John Wiley: New York.
- [7] Cressie, N. and C. K. Wikle, 2015: *Statistics for spatio-temporal data*. John Wiley & Sons.
- [8] Datta, A., S. Banerjee, A. O. Finley, and A. E. Gelfand, 2016: 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>

- [9] Datta, A., S. Banerjee, A. O. Finley, N. A. S. Hamm, and M. Schaap, 2016: 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-AOAS931>
- [10] Finley, A., A. Datta, and S. Banerjee, 2017: *spNNGP: Spatial Regression Models for Large Datasets using Nearest Neighbor Gaussian Processes*. R package version 0.1.1.  
URL <https://CRAN.R-project.org/package=spNNGP>
- [11] Finley, A. O., A. Datta, B. C. Cook, D. C. Morton, H. E. Andersen, and S. Banerjee, 2018: Applying nearest neighbor gaussian processes to massive spatial data sets: Forest canopy height prediction across tanana valley alaska. *Journal of Computational and Graphical Statistics (in press)*; *arXiv preprint arXiv:1702.00434*.
- [12] Furrer, R., M. G. Genton, and D. Nychka, 2006: Covariance tapering for interpolation of large spatial datasets. *Journal of Computational and Graphical Statistics*, **15**, 503–523.
- [13] Gelfand, A. E., P. Diggle, P. Guttorp, and M. Fuentes, 2010: *Handbook of spatial statistics*. CRC press.
- [14] Gelman, A., J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin, 2013: *Bayesian Data Analysis, 3rd Edition*. Chapman & Hall/CRC Texts in Statistical Science, Chapman & Hall/CRC.
- [15] Gneiting, T. and A. E. Raftery, 2007: Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, **102**, no. 477, 359–378.
- [16] Golub, G. H. and C. F. Van Loan, 2012: *Matrix Computations, 4th Edition*. Johns Hopkins University Press.
- [17] Guinness, J., 2016: Permutation methods for sharpening gaussian process approximations. *arXiv preprint arXiv:1609.05372*.
- [18] Heaton, M., A. Datta, A. Finley, R. Furrer, R. Guhaniyogi, F. Gerber, D. Hammerling, M. Katzfuss, F. Lindgren, D. Nychka, and A. Zammit-Mangion, 2017: Methods for analyzing large spatial data: A review and comparison. *arXiv:1710.05013*.  
URL <https://arxiv.org/abs/1710.05013>
- [19] Katzfuss, M., 2013: Bayesian nonstationary modeling for very large spatial datasets. *Environmetrics*, **24**, 189—200.
- [20] — 2017: A multi-resolution approximation for massive spatial datasets. *Journal of the American Statistical Association*, **112**, 201–214, doi:10.1080/01621459.2015.1123632.  
URL <http://dx.doi.org/10.1080/01621459.2015.1123632>
- [21] Katzfuss, M. and J. Guinness, 2017: A general framework for vecchia approximations of gaussian processes. *arXiv preprint arXiv:1708.06302*.
- [22] Lauritzen, S. L., 1996: *Graphical Models*, Clarendon Press, Oxford, United Kingdom.
- [23] Lindgren, F., H. Rue, and J. Lindstrom, 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 (Statistical Methodology)*, **73**, no. 4, 423–498, doi:10.1111/j.1467-9868.2011.00777.x.  
URL <http://dx.doi.org/10.1111/j.1467-9868.2011.00777.x>
- [24] Murphy, K., 2012: *Machine Learning: A probabilistic perspective*. The MIT Press, Cambridge, MA.
- [25] Nychka, D., S. Bandyopadhyay, D. Hammerling, F. Lindgren, and S. Sain, 2015: A multiresolution gaussian process model for the analysis of large spatial datasets. *Journal of Computational and Graphical Statistics*, **24**, no. 2, 579–599, doi:10.1080/10618600.2014.914946.  
URL <http://dx.doi.org/10.1080/10618600.2014.914946>
- [26] Nychka, D., C. Wikle, and J. A. Royle, 2002: Multiresolution models for nonstationary spatial covariance functions. *Statistical Modelling*, **2**, no. 4, 315–331.



- [27] Ribeiro Jr, P. J. and P. J. Diggle, 2012: geoR: a package for geostatistical analysis. R package version 1.7-4.  
URL <https://cran.r-project.org/web/packages/geoR>
- [28] Rue, H. and L. Held, 2005: *Gaussian Markov Random Fields : Theory and Applications*, Chapman & Hall/CRC, Boca Raton, FL. Monographs on statistics and applied probability.  
URL <http://opac.inria.fr/record=b1119989>
- [29] Rue, H., S. Martino, and N. Chopin, 2009: Approximate bayesian inference for latent gaussian models by using integrated nested laplace approximations. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **71**, no. 2, 319–392, doi:10.1111/j.1467-9868.2008.00700.x.  
URL <http://dx.doi.org/10.1111/j.1467-9868.2008.00700.x>
- [30] Sang, H. and J. Z. Huang, 2012: A full scale approximation of covariance functions for large spatial data sets. *Journal of the Royal Statistical society, Series B*, **74**, 111–132.
- [31] Stan Development Team, 2016: *RStan: the R interface to Stan*. R package version 2.14.1.  
URL <http://mc-stan.org/>
- [32] Stein, M. L., 1999: *Interpolation of Spatial Data: Some Theory for Kriging*, Springer. Firstnd ed.
- [33] Stein, M. L., Z. Chi, and L. J. Welty, 2004: Approximating likelihoods for large spatial data sets. *Journal of the Royal Statistical society, Series B*, **66**, 275–296.
- [34] Stroud, J. R., M. L. Stein, and S. Lysen, 2017: Bayesian and maximum likelihood estimation for gaussian processes on an incomplete lattice. *Journal of Computational and Graphical Statistics*, **26**, 108–120.  
URL <http://dx.doi.org/10.1080/10618600.2016.1152970>
- [35] Vecchia, A. V., 1988: Estimation and model identification for continuous spatial processes. *Journal of the Royal Statistical society, Series B*, **50**, 297–312.
- [36] Yeniay, O. and A. Goktas, 2002: A comparison of partial least squares regression with other prediction methods. *Hacetetepe Journal of Mathematics and Statistics*, **31**, no. 99, 99–101.