

Incorporating Subsampling into Bayesian Models for High-Dimensional Spatial Data

Sudipto Saha* and Jonathan R. Bradley†

Abstract. Additive spatial statistical models with weakly stationary process assumptions have become standard in spatial statistics. However, one disadvantage of such models is the computation time, which rapidly increases with the number of data points. The goal of this article is to apply an existing subsampling strategy to standard spatial additive models and to derive the spatial statistical properties. We call this strategy the “spatial data subset model” (SDSM) approach, which can be applied to big datasets in a computationally feasible way. Our approach has the advantage that one does not require any additional restrictive model assumptions. That is, computational gains increase as model assumptions are removed when using our model framework. This provides one solution to the computational bottlenecks that occur when applying methods such as Kriging to “big data”. We provide several properties of this new spatial data subset model approach in terms of moments, sill, nugget, and range under several sampling designs. An advantage of our approach is that it subsamples without throwing away data, and can be implemented using datasets of any size that can be stored. We present the results of the spatial data subset model approach on simulated datasets, and on a large dataset consists of 150,000 observations of daytime land surface temperatures measured by the MODIS instrument onboard the Terra satellite.

Keywords: Bayesian hierarchical model, big data, curse of dimensionality, subsampling method, Gibbs-within-composite sampler, Markov chain Monte Carlo, spatial data.

1 Introduction

A usual inferential task for spatially referenced data is to predict a spatial process at observed and unobserved locations (i.e. spatial prediction). The classical approach to perform spatial prediction is known as “Kriging”. Taking a Bayesian approach for parameter estimation and spatial prediction is a well-established and widely used approach in spatial statistics (e.g., see More et al., 1989; Omre and Halvorsen, 1989; Cressie, 1993; Handcock and Stein, 1993; Banerjee et al., 2014 for standard references). Classical spatial statistical models assume that a latent spatial process is weakly stationary with known covariogram, e.g. Matérn, exponential etc. (Banerjee et al., 2014). Although Bayesian prediction in this setting is possible and provides accurate spatial predictions for small datasets, the processing time rapidly increases with the increase in the size of

arXiv: 2305.13221

*Department of Statistics, Florida State University, 117 N. Woodward Ave., Tallahassee, FL 32306, ssaha2@fsu.edu

†Corresponding author. Department of Statistics, Florida State University, 117 N. Woodward Ave., Tallahassee, FL 32306, bradley@stat.fsu.edu

dataset. This is a well-known problem that has motivated several methodologies (Vecchia, 1988; Nychka, 2000; Cressie and Johannesson, 2008; Rue et al., 2009; Xu et al., 2013; Nychka et al., 2015; Rue et al., 2017; Vignés et al., 2017). In this article, we provide a computationally feasible solution to analyze big data using a classical Bayesian spatial model with standard weakly stationary or non-stationary process assumptions. In particular, we develop a new “spatial data subset model” approach by applying the use of a recently introduced “data subset model” approach (Bradley, 2021) to spatial data. We also provide inferential properties of the data in terms of its moments, and spatial properties in terms of sill, nugget and range, under two sampling designs that are natural for spatial data.

We extend the use of data subset model (Bradley, 2021) and develop the spatial data subset model that allows one to implement classical Bayesian spatial models for analyzing big data quickly. To achieve this, we redefine the data model of a classical Bayesian spatial model using subsamples, without imposing any additional model assumptions. This is different than what has been done in the existing literature. For example, Byers and Gill (2022) presented a similar work where they used a Bootstrap Random Spatial Sampling (BRSS) method to estimate the full Bayesian-Kriging model over large samples. They independently generated B bootstrap subsamples of size n , and estimated model parameters independently for each of B bootstraps. In our proposed model, we redefine the full data model of the classical Bayesian spatial model to be semi-parametric using subsamples and we add a subset model to the hierarchical structure, rather than estimating model parameters independently for each bootstrap subsamples or approximating the full data model using the subsamples. This is the primary difference between our exact approach, and the other exact and approximate approaches (Vecchia, 1988; Stein et al., 2004; Gunawan et al., 2017; Quiroz et al., 2019; Katzfuss and Guinness, 2021; Byers and Gill, 2022). Moreover, adding a subset model to the hierarchical structure itself provides the flexibility to choose a sampling method (from the survey sampling literature) inside the model.

There are many existing approaches for subsampling spatial data. One approach is to divide the spatial domain into overlapping subblocks which are obtained by moving a subsampling window over sequences of increasing index sets (Guan et al., 2004), and draw the subsamples from these subblocks. One disadvantage of this approach is that the overlapping subblocks can introduce unnecessary complexity into the model. To overcome this issue, a common approach is to divide the spatial domain into several partitions or subshapes (Possolo, 1991; Sherman and Carlstein, 1994; Lahiri, 1999; Ekström and Luna, 2004). One can partition a spatial domain using one of many partitioning methods. For example, equal area partition (Sang et al., 2011), hierarchical clustering based partition (Heaton et al., 2017), tree based partition (Konomi et al., 2014), partition based on cluster centroids (Knorr-Held and Raßer, 2000; Kim et al., 2005), partition based on mixture model (Neelon et al., 2014) etc. However, all these spatial partitioning approaches assume independence across partitions (blockwise independence), which can be a strong assumption (Bradley et al., 2015) in certain scenarios. To avoid any additional assumption, we choose the sampling method for our spatial data subset mode from the survey sampling literature (Lohr, 2010; Nassiuma, 2001). In particular, we consider simple designs, like simple random sampling (SRS) and stratified

random sampling, that are natural for spatial data. Cluster sampling is another natural choice for spatial data. However, in the survey literature it has been shown to be less efficient than stratified random sampling (Lohr, 2010), and hence, we do not consider it here.

Possolo (1991) is an example of a method that uses a single subsample of size n and discards $(N - n)$ spatial observations, where N is the size of the entire dataset. Of course this is a standard data pre-processing step for whenever a method is not scalable (e.g., a subset of the data is used). Several in the more general Gaussian process literature have given this strategy a name, i.e., the “Subset of Data” (SoD) method (Lawrence et al., 2002; Seeger, 2003; Keerthi and Chu, 2005; Chalupka et al., 2013; Hayashi et al., 2020; Liu et al., 2020). Our proposed method is different from using a single subsample, as the Data Subset Model approach iteratively re-samples the data, and if one runs the MCMC long enough, all data is used for inference. We define such method that makes use of all N observations without discarding them as “fully scalable”. To our knowledge, there is no such Bayesian spatial model to analyze large spatial data that is fully scalable (for large enough length of the Markov chain) to datasets of any size that can be stored.

Outside of subsampling there are several other strategies to make a Gaussian Process (GP) scalable in Bayesian and non-Bayesian contexts. For example, reduced rank methodologies reduce the parameter space to speed up computations (Cressie and Johannesson, 2008). Sparse GPs similarly reduce the parameter space via sparsity (e.g., see Besag (1974); Besag et al. (1991) in the spatial context, and Hensman et al. (2013) in a more general GP context). Approximate Bayesian strategies approximate the likelihood to aid with computations, for example, the Vecchia approximation (Katzfuss and Guinness, 2021) writes the likelihood as a product and truncate the product to fewer terms, and integrated nested Laplace approximations (INLA, Rue et al., 2009) use Gaussian approximations of the marginal posterior to aid with inference. Some apply spatial statistical models to spatial partitions or for a given prediction location adaptively to create a partition based on nearest neighborhoods (e.g., see Heaton et al., 2019, and the local GP approach of Gramacy and Apley, 2015 for examples). For a broad discussion on the available strategies in spatial statistics see Bradley (2021), Heaton et al. (2019) and in the generic GP context see Liu et al. (2020). In general, all of the fully Bayesian strategies involve additional parametric assumptions, or approximations. The data subset model, does not add assumptions, but rather removes assumptions; that is, the hold-out sample in the data subset model approach is assumed to follow its true nonparametric distribution instead of some parametric density enforced to have sparsity or dimension reduction structure. Furthermore the data subset model approach does not involve approximations used in approximate Bayesian strategies.

Our proposed spatial data subset model has a hyperparameter, subsample size (n), that can be set to a value to achieve a computational goal, rather than determining an optimal subsample size. There are several ways to choose an optimal subsample size. For example, one can choose an optimal subsample size by minimizing the mean square error of variance estimation of the blocked sample variance (Politis and Romano, 1993), or by minimizing the combined expansions for the bias and the variance parts of the

subsample variance estimator (Nordman and Lahiri, 2004). A limitation of both of these approaches is that they both demand an approximation to the full likelihood. However, in our model, we *redefine* the full likelihood to be semi-parametric (or more flexible) based on the subsamples instead of approximating it. So, we keep the subsample size as a hyperparameter for the users to choose based on their computational goal, which makes our model flexible to demand. This is particularly exciting because our model scales with the subsample design, and not with the size of the entire dataset. Hence, one can scale our model to a dataset of any size that can be stored by choosing the subsample size small enough.

While n can be chosen to achieve a computational goal, this is not necessarily a positive. Ideally, we would choose n for inference, and choosing n in this manner, is to make an admission that a given model is not fully scalable. The main goal of this article is to apply the data subset approach to standard spatial models (i.e., kriging). It is well known that kriging is not fully scalable, which leads to a natural restriction on n , since kriging is known to break down computationally when $n \geq 10,000$ (e.g., see Rulli re et al., 2018), and can be extremely slow at $n = 10,000$. In our empirical results we often see that $n \ll N$ produces similar answers compared to kriging based on N . This is consistent with results in Bradley (2021) and results in the SoD literature (Chalupka et al., 2013, §5).

The choice of the sampling method affects the fundamental properties of the model. In particular, we are interested in the changes in the first and second moments, and the spatial properties of the data as we change the sampling design. Hence, we present a detailed discussion on the inferential, as well as the spatial properties of the data under two sampling strategies – SRS and stratified random sampling. That is, we investigate several spatial properties like sill, nugget and (effective) range (Cressie, 1993; Banerjee et al., 2014), including the mean, variance and covariance (Casella and Berger, 2021) of the data for SRS and stratified random sampling. One exciting feature of our model is that our model retains the stationarity of the true data model, i.e. our model becomes weakly/intrinsically stationary or non-stationary when the true data model is respectively weakly/intrinsically stationary or non-stationary.

The remainder of the article is organized as follows. In Section 2, we provide an extensive review of the existing data subset model in a generic GP setting. In Section 3, we present the proposed spatial data subset model, along with the inferential and spatial properties under the mentioned two sampling designs and the MCMC implementation. In Section 4, we demonstrate the performance of our model by providing an illustration through a simulated dataset, followed by a more extensive simulation study and an analysis on the scalability of the model. We demonstrate that the performance of our model gets better as we increase the number of subsamples. We also show that the proposed model is fully scalable (for large enough length of the Markov chain) to the datasets of any size that can be stored. In Section 5, we apply our model to the benchmark dataset from Heaton et al. (2019), that consists of Land Surface Temperature (LST) captured by the Terra satellite. We find that our semi-parametric version of classical Bayesian spatial models provides competitive inferential and computational results as compared to several existing strategies (Heaton et al., 2019). Finally, a concluding dis-

ussion is presented in Section 6. We provide the proof of all the technical results in the supplementary material (Saha and Bradley, 2024) for convenience.

2 Review

In this section, we review general subsampling strategies. In particular, we review the “data subset approach” in the generic GP Bayesian setting from Bradley (2021), and the SoD method (Chalupka et al., 2013).

2.1 Review of Data Subset Model

Let, an N -dimensional data vector be $\mathbf{y} = (Y_1, \dots, Y_N)'$, the latent process vector be $\boldsymbol{\nu}$, and a generic real-valued parameter vector be $\boldsymbol{\theta}$. In general, a Bayesian hierarchical model can be written as the product of the following conditional and marginal distributions:

$$\begin{aligned} \text{Data Model: } & \prod_{i=1}^N f(Y_i | \boldsymbol{\nu}, \boldsymbol{\theta}) \\ \text{Process Model: } & f(\boldsymbol{\nu} | \boldsymbol{\theta}) \\ \text{Parameter Model: } & f(\boldsymbol{\theta}). \end{aligned} \tag{2.1}$$

Bradley (2021) considered N to be so large that estimating $\boldsymbol{\nu}$ and $\boldsymbol{\theta}$ using (2.1) (referred as “full parametric model”) directly is not possible. To solve this problem, Bradley (2021) proposed to replace the data model in (2.1) with the data subset model (DSM) in the following way:

$$\begin{aligned} \text{Data Subset Model: } & \left\{ \prod_{i=1}^N f(Y_i | \boldsymbol{\nu}, \boldsymbol{\theta}, \boldsymbol{\delta})^{\delta_i} \right\} \frac{m(\mathbf{1}_N, \mathbf{y})}{m(\boldsymbol{\delta}, \mathbf{y}_\delta)} \\ \text{Process Model: } & f(\boldsymbol{\nu} | \boldsymbol{\theta}) \\ \text{Parameter Model: } & f(\boldsymbol{\theta}) \\ \text{Subset Model: } & \Pr(\boldsymbol{\delta} | n), \end{aligned} \tag{2.2}$$

where $\boldsymbol{\delta} = (\delta_1, \dots, \delta_N)'$ be the N -dimensional random vector consists of ones and zeros, $\mathbf{1}_N$ (and $\mathbf{0}_N$) be the N -dimensional vector of ones (and zeros), and $\mathbf{y}_\delta = (Y_i : \delta_i = 1)'$ be the n -dimensional random vector since they chose exactly n number of ones in the $\boldsymbol{\delta}$ vector by setting $\sum_{i=1}^N \delta_i = n$. The subset model determines which n observations to include in the likelihood through δ_i ($= 0$ or 1), which is treated as a parameter in the hierarchical model. The marginal distribution, $m(\boldsymbol{\delta}, \mathbf{y}_\delta)$, is defined as

$$m(\boldsymbol{\delta}, \mathbf{y}_\delta) = \int \int \left\{ \prod_{i=1}^N f(Y_i | \boldsymbol{\nu}, \boldsymbol{\theta}, \boldsymbol{\delta})^{\delta_i} \right\} f(\boldsymbol{\nu} | \boldsymbol{\theta}) f(\boldsymbol{\theta}) d\boldsymbol{\nu} d\boldsymbol{\theta}, \tag{2.3}$$

where $m(\boldsymbol{\delta}, \mathbf{y}_\delta)$ is the marginal distribution of \mathbf{y}_δ from the full parametric model. $m(\mathbf{1}_N, \mathbf{y})$ is the marginal distribution of the data associated with the full paramet-

ric model. Hence, the data subset model in (2.2) re-weights the likelihood evaluated at $n(\ll N)$ observations by the ratio of marginal densities, $m(\mathbf{1}_N, \mathbf{y})/m(\boldsymbol{\delta}, \mathbf{y}_\delta)$.

Bradley (2021) established several important theoretical results of the model in (2.2). In particular, the model in (2.2) is proper given that the model in (2.1) is proper, and the marginal distributions of data \mathbf{y} are the same for the Bayesian models in (2.1) and (2.2). Thus, the Bayesian models in (2.1) and (2.2) have the same expression of $f(\boldsymbol{\nu}|\boldsymbol{\theta})$, $f(\boldsymbol{\theta})$ and $m(\mathbf{1}_N, \mathbf{y})$. This implies that, marginally, the distributional assumptions of the data, process, and parameters are the same between (2.2) and (2.1). Sufficiency is a key property in Bayesian analysis, and Bradley (2021) showed that \mathbf{y}_δ is a partial sufficient statistic for $\boldsymbol{\nu}$ and $\boldsymbol{\theta}$. Another important result is the data vector \mathbf{y} and the $\boldsymbol{\delta}$ vector are independent. This result is particularly important from a computational perspective, as it leads to an efficient sampler that we develop further in the spatial setting. In particular, to sample from the posterior $f(\boldsymbol{\nu}, \boldsymbol{\theta}, \boldsymbol{\delta}|\mathbf{y})$, one first samples from $\Pr(\boldsymbol{\delta}|n)$ and then samples from the posterior using subsample \mathbf{y}_δ (i.e., $p(\boldsymbol{\nu}, \boldsymbol{\theta}|\mathbf{y}_\delta)$). Note that, the data are continually resampled in this sampling scheme, which is different from many existing subsampling strategies that use a single subsample. Bradley (2021) explored this sampler in truly high-dimensional settings (i.e., a dataset of size 10 Gigabytes and a simulated dataset of size 100 million).

A semi-parametric interpretation of the model in (2.2) in the generic Bayesian GP setting is also discussed in Bradley (2021), where the following model is considered.

$$\begin{aligned}
 \text{Data Model: } & w_\delta(\mathbf{y}_{-\delta}) \prod_{\{i:\delta_i=1\}} f(Y_i|\boldsymbol{\nu}, \boldsymbol{\theta}, \boldsymbol{\delta}) \\
 \text{Process Model: } & f(\boldsymbol{\nu}|\boldsymbol{\theta}) \\
 \text{Parameter Model: } & f(\boldsymbol{\theta}) \\
 \text{Subset Model: } & \Pr(\boldsymbol{\delta}|n), \tag{2.4}
 \end{aligned}$$

where $w_\delta(\cdot)$ is the true unknown unparameterized pdf for the $(N - n)$ -dimensional random vector $\mathbf{y}_{-\delta}$ for a given $\boldsymbol{\delta}$. Bradley (2021) referred this model as the semi-parametric full model (SFM). When one assumes that \mathbf{y} and $\boldsymbol{\delta}$ are independent in (2.4), then Bradley (2021) showed that (2.4) has many of the same properties as (2.2). Specifically, \mathbf{y}_δ is a partially sufficient statistic for $\boldsymbol{\nu}$ and $\boldsymbol{\theta}$. Furthermore, Bradley (2021) established that $f_{DSM}(\boldsymbol{\nu}, \boldsymbol{\theta}|\mathbf{y}, \boldsymbol{\delta}, n) = f_{SFM}(\boldsymbol{\nu}, \boldsymbol{\theta}|\mathbf{y}, \boldsymbol{\delta}, n)$ and $f_{DSM}(\boldsymbol{\nu}, \boldsymbol{\theta}|\mathbf{y}, n) = f_{SFM}(\boldsymbol{\nu}, \boldsymbol{\theta}|\mathbf{y}, n)$, where f_{DSM} and f_{SFM} be the posterior distributions under the model in (2.2) and (2.4), respectively. These two results are particularly interesting because together they show that the posterior distributions of DSM and SFM are equivalent when one adds the assumption that \mathbf{y} and $\boldsymbol{\delta}$ are independent in (2.4).

2.2 Review of Subset of Data Method

It is important to understand that the data subset model does not use a single subsample defined by $\boldsymbol{\delta}$. This is because $\boldsymbol{\delta}$ is actively resampled in the Gibbs-within-composite sampler. This is different from the setting where one makes use of a single subsample. Some in the general GP literature call this approach the Subset of Data (SoD) method

(Lawrence et al., 2002; Seeger, 2003; Keerthi and Chu, 2005; Chalupka et al., 2013; Hayashi et al., 2020; Liu et al., 2020). The SoD method simply uses a single subsample of size n ($n \ll N$) for inference. For reviews of SoD see Chalupka et al. (2013) and Liu et al. (2020). The difference between SoD and the Data Subset Approach is important because the data subset model scales to higher dimensions in a similar way that SoD scales to higher dimension without discarding $N - n$ observations. We say a method is “fully scalable” if it can be implemented on N observations *without* discarding any observations. Thus, the SoD method is not fully scalable by definition, and the data subset model approach is fully scalable for large enough G , where G is the length of the Markov chain. That is, as one continually samples from the Gibbs-within-composite sampler, as long as each observation has a non-zero probability of being selected, all data will be used when G is large enough.

Resampling in the SoD literature is sometimes based on n randomly selected points, which has computational complexity of $\mathcal{O}(n)$, since one does not need to scan through all N data-points to produce the subsample. Others in the SoD literature have considered clustering algorithms, which are on the order of $\mathcal{O}(nN)$ or $\mathcal{O}(N \log(n))$ (Chalupka et al., 2013). A fully Bayesian data subset model strategy is more restrictive here, since the data cannot be used to select the n points from N (such a strategy would yield an empirical Bayesian approach with unchecked variability). As a result, we make use of random selected points and do not make use of any clustering algorithms.

3 Methodology

Suppose, we observe an N -dimensional spatial data vector $\mathbf{y} = (Y(\mathbf{s}_1), \dots, Y(\mathbf{s}_N))'$, where the locations are $\{\mathbf{s}_1, \dots, \mathbf{s}_N\} \in D \subset \mathfrak{R}^d$. Let, the corresponding N -dimensional latent process of interest be $\boldsymbol{\nu} = (\nu(\mathbf{s}_1), \dots, \nu(\mathbf{s}_N))' \in \mathbb{R}^N$, the $N \times p$ dimensional matrix of known covariates be $\mathbf{X} \in \mathbb{R}^{N \times p}$, and the p -dimensional unknown regression parameter be $\boldsymbol{\beta} \in \mathbb{R}^p$. The additive model of the classical Bayesian spatial model is defined as,

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\nu} + \boldsymbol{\epsilon}, \quad (3.1)$$

where $\epsilon(\mathbf{s}_i) \stackrel{iid}{\sim} \mathcal{N}(0, \tau^2)$, $\forall i = 1, \dots, N$ is the white noise and $\mathbf{w} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\nu}$ is known as the latent process. Hence, the data model of classical Bayesian spatial model refers to the conditional distribution $f(\mathbf{y}|\boldsymbol{\nu}, \boldsymbol{\theta}) = \prod_{i=1}^N f(Y(\mathbf{s}_i)|\boldsymbol{\nu}, \boldsymbol{\theta})$, where $\boldsymbol{\theta}$ be a real-valued parameter vector. However, we consider N to be so large that estimating $\boldsymbol{\nu}$ and $\boldsymbol{\theta}$ using classical Bayesian spatial model directly is not possible in real time. As a solution, we extend the data subset model (Bradley, 2021) by incorporating it into the classical Bayesian spatial model and the resulting model is known as spatial data subset model.

3.1 Bayesian Hierarchical Model

We write the data subset model (Bradley, 2021) incorporating into the classical Bayesian spatial model, where the resulting joint spatial data subset model (SDSM) is propor-

tional to the product of the following conditional and marginal distributions,

$$\begin{aligned}
& \text{Data Subset Model: } w_{\delta}(\mathbf{y}_{-\delta}) \left\{ \prod_{\{i:\delta(\mathbf{s}_i)=1\}} f(Y(\mathbf{s}_i)|\boldsymbol{\nu}, \boldsymbol{\theta}) \right\}, \\
& \text{Process Model: } \boldsymbol{\nu}|\sigma^2, \phi \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{H}(\phi)), \\
& \text{Parameter Model 1: } \boldsymbol{\beta}|\sigma_{\beta}^2 \sim \mathcal{N}(\mathbf{0}, \sigma_{\beta}^2 \mathbf{I}_p), \\
& \text{Parameter Model 2: } \tau^2 \sim \mathcal{IG}(a_{\tau}, b_{\tau}), \\
& \text{Parameter Model 3: } \sigma^2 \sim \mathcal{IG}(a_{\sigma}, b_{\sigma}), \\
& \text{Parameter Model 4: } \sigma_{\beta}^2 \sim \mathcal{IG}(a_{\beta}, b_{\beta}), \\
& \text{Parameter Model 5: } \phi \sim \pi(\phi), \\
& \text{Subset Model: } \boldsymbol{\delta} \sim \Pr(\boldsymbol{\delta}|n),
\end{aligned} \tag{3.2}$$

where $f(Y(\mathbf{s}_i)|\boldsymbol{\nu}, \boldsymbol{\theta})$ is the probability density function (pdf) of a $\mathcal{N}(\mathbf{x}(\mathbf{s}_i)' \boldsymbol{\beta} + \nu(\mathbf{s}_i), \tau^2)$ distribution, $\boldsymbol{\delta} = (\delta(\mathbf{s}_1), \dots, \delta(\mathbf{s}_N))'$ is an N -dimensional random vector consists of ones and zeros, $\mathbf{y}_{-\delta} = (Y(\mathbf{s}_i) : \delta(\mathbf{s}_i) = 0)'$ consists of all such $Y(\mathbf{s}_i)$ where $\delta(\mathbf{s}_i) = 0$ for a given $\boldsymbol{\delta}$, and $w_{\delta}(\mathbf{y}_{-\delta})$ is a probability density that represents the true unknown unparameterized pdf for $\mathbf{y}_{-\delta}$. We define the real-valued parameter vector $\boldsymbol{\theta} = (\boldsymbol{\beta}, \tau^2, \sigma^2, \sigma_{\beta}^2, \phi)'$. A general structure of an element of the matrix $\mathbf{H}(\phi) = \{h_{ij}(\phi) : i, j = 1, \dots, N\} \in \mathbb{R}^{N \times N}$ is $h_{ij}(\phi) = \rho_c(\mathbf{s}_i - \mathbf{s}_j; \phi)$, where ρ_c is a valid correlation function on \mathfrak{R}^2 indexed by the parameter ϕ . $h_{ij}(\phi)$ and the density function $\pi(\phi)$ both depends on the choice of the function ρ_c . For example, if we choose exponential covariogram, i.e. $h_{ij}(\phi) = \exp(-\phi \|\mathbf{s}_i - \mathbf{s}_j\|)$, then a Gamma distribution might seem sensible for the choice of $\pi(\phi)$ (Banerjee et al., 2014). $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ is shorthand for a multivariate Normal distribution with mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. Similarly, $\mathcal{IG}(a, b)$ is used as shorthand for inverse Gamma distribution with shape parameter a and scale parameter b . Notice that, the hyperparameters of the \mathcal{IG} distributions are different across τ^2 , σ^2 and σ_{β}^2 . It is important to mention that one has the flexibility to consider other expressions of the spatial data model. For example, one can marginalize across $\boldsymbol{\nu}$. This would lead to one fewer block update of parameters in the sampler. We choose to include $\boldsymbol{\nu}$, as it leads to conjugate updates, which helps with mixing of Markov chains (Banerjee et al., 2014). Additionally, one has the flexibility to include other covariograms. We often use the exponential covariogram for illustration.

The model in (3.2) shows that for a given $\boldsymbol{\delta}$, $\delta(\mathbf{s}_i) = 1$ includes $Y(\mathbf{s}_i)$ in the expression of the parametric portion of the likelihood (proportionally), and $\delta(\mathbf{s}_i) = 0$ removes $Y(\mathbf{s}_i)$ from the expression of the parametric portion of the likelihood (proportionally). We sample $\boldsymbol{\delta}$ from $\Pr(\boldsymbol{\delta}|n)$ for a given n such that $\sum_{i=1}^N \delta(\mathbf{s}_i) = n \ll N$, where n is the ‘‘subsample size’’. Therefore, for any sample of $\boldsymbol{\delta}$, only n number of $Y(\mathbf{s}_i)$ would be included in the expression of likelihood (proportionally). This further means that, the spatial data subset model only requires an n -dimensional data vector $\mathbf{y}_{\delta} = (Y(\mathbf{s}_i) : \delta(\mathbf{s}_i) = 1)'$ (at a time), instead of the larger N -dimensional data vector \mathbf{y} , to estimate $\boldsymbol{\nu}$ and $\boldsymbol{\theta}$. We choose n in a way that the model in (3.2) achieves computational goals (Bradley, 2021). We specify $\Pr(\boldsymbol{\delta}|n)$ in a manner that it produces simple random samples (SRS) without replacement and stratified random samples as described in Section 1.

We assume that $\boldsymbol{\delta}$ is independent of \mathbf{y} , and Bradley (2021) showed that such a model exists. This assumption leads to a very efficient composite sampling scheme for posterior inference, where one first samples from $\Pr(\boldsymbol{\delta}|\mathbf{y}, n) = \Pr(\boldsymbol{\delta}|n)$, via independence, and then samples from $\Pr(\boldsymbol{\nu}, \boldsymbol{\theta}|\mathbf{y}, \boldsymbol{\delta}, n)$. The conditional distribution $\Pr(\boldsymbol{\nu}, \boldsymbol{\theta}|\mathbf{y}, \boldsymbol{\delta}, n)$ only includes $n \ll N$ observations, and hence, one can choose n small enough so that sampling from $\Pr(\boldsymbol{\nu}, \boldsymbol{\theta}|\mathbf{y}, \boldsymbol{\delta}, n)$ is efficient. This composite sampling scheme iteratively re-samples the data, so that every observation is included provided that the composite sampling procedure runs long enough. In practice, sampling from $\Pr(\boldsymbol{\nu}, \boldsymbol{\theta}|\mathbf{y}, \boldsymbol{\delta}, n)$ can require MCMC, which leads to a Gibbs-within-Composite sampler. We provide more discussion on this in Section 3.4.

3.2 Properties of Spatial Data Subset Model when the Sampling Design is SRS Without Replacement

The properties of a classical Bayesian spatial model are well-established in numerous literature (e.g., see Cressie, 1993; Banerjee et al., 2014, among several others). Enforcing $\boldsymbol{\delta}$ into the model does not change the process and parameter models, and hence it does not change the inherent properties of the latent process. But it does change the properties of the data since we are subsampling a smaller number of data points from the entire dataset. Moreover, the properties vary with the choice of sampling technique for $\boldsymbol{\delta}$. This leads us to the following propositions when $\Pr(\boldsymbol{\delta}|n)$ is defined by simple random samples (SRS) without replacement.

In nonparametric statistics it is often assumed that the true measure that generated $\mathbf{y}_{-\delta}$ is unknown and not parameterized (Hollander et al., 2013, pg. 2). We assume that the true (unknown) distribution is absolutely continuous so that it has density w , which is also unknown. Essentially, we assume that any given n -dimensional “training” dataset \mathbf{y}_{δ} follows a parametric model, and the corresponding holdout $\mathbf{y}_{-\delta}$ is distributed according to its true unknown data generating measure as one would in nonparametric statistics. We call the mean computed using w the “true mean”, the variance computed using w the “true variance”, and the covariance computed using w the “true covariance”.

Proposition 3.1. *When $\Pr(\boldsymbol{\delta}|n)$ is defined by simple random samples without replacement, the expectation, variance and covariance of the data $Y(\mathbf{s}_i)$ for $i = 1, \dots, N$ given the real-valued parameter vector $\boldsymbol{\theta} = (\boldsymbol{\beta}, \tau^2, \sigma^2, \sigma_{\beta}^2, \phi)'$ are given by*

$$\begin{aligned} \mathbb{E}[Y(\mathbf{s}_i)|\boldsymbol{\theta}] &= p\{\mathbf{x}(\mathbf{s}_i)'\boldsymbol{\beta}\} + (1-p)\tilde{\mu}(\mathbf{s}_i) = p\{\mathbf{x}(\mathbf{s}_i)'\boldsymbol{\beta} - \tilde{\mu}(\mathbf{s}_i)\} + \tilde{\mu}(\mathbf{s}_i), \\ \text{Var}(Y(\mathbf{s}_i)|\boldsymbol{\theta}) &= p(\tau^2 + \sigma^2) + (1-p)\tilde{\sigma}^2 + p(1-p)\{\mathbf{x}(\mathbf{s}_i)'\boldsymbol{\beta} - \tilde{\mu}(\mathbf{s}_i)\}^2 \\ &= p(\tau^2 + \sigma^2 - \tilde{\sigma}^2) + p(1-p)\{\mathbf{x}(\mathbf{s}_i)'\boldsymbol{\beta} - \tilde{\mu}(\mathbf{s}_i)\}^2 + \tilde{\sigma}^2, \\ \text{Cov}(Y(\mathbf{s}_i), Y(\mathbf{s}_j)|\boldsymbol{\theta}) &= a\tilde{C}(\mathbf{s}_i, \mathbf{s}_j) + (b + p^2)\sigma^2 h_{ij}(\phi) \\ &\quad + b\{\mathbf{x}(\mathbf{s}_i)'\boldsymbol{\beta} - \tilde{\mu}(\mathbf{s}_i)\}\{\mathbf{x}(\mathbf{s}_j)'\boldsymbol{\beta} - \tilde{\mu}(\mathbf{s}_j)\}, \end{aligned}$$

where $p = \frac{n}{N}$, $a = \left(\frac{n}{N} \frac{n-1}{N-1} - 2\frac{n}{N} + 1\right)$, $b = \frac{n}{N} \left(\frac{n-1}{N-1} - \frac{n}{N}\right)$, $\tilde{\mu}(\mathbf{s}_i)$ is the true mean of $Y(\mathbf{s}_i)$, $\tilde{\sigma}^2$ is the true variance of $Y(\mathbf{s}_i)$, and $\tilde{C}(\mathbf{s}_i, \mathbf{s}_j)$ is the true covariance between $Y(\mathbf{s}_i)$ and $Y(\mathbf{s}_j)$.

Proof: See Appendix A.1 of the supplementary material (Saha and Bradley, 2024).

Proposition 3.1 is useful to understand the characteristics of the spatial data subset model in (3.2) from the perspective of its first and second moments. Notice that, the expectation of $Y(\mathbf{s}_i)|\boldsymbol{\theta}$ is a weighted sum of the parametric mean and the true mean (shown in the first expression of $\mathbb{E}[Y(\mathbf{s}_i)|\boldsymbol{\theta}]$). The second expression of $\mathbb{E}[Y(\mathbf{s}_i)|\boldsymbol{\theta}]$ shows that the true mean of $Y(\mathbf{s}_i)|\boldsymbol{\theta}$ is biased by the weighted difference between the parametric mean and the true mean. Furthermore, from the expression of the expectation of $Y(\mathbf{s}_i)|\boldsymbol{\theta}$ we obtain

$$\mathbb{E}[Y(\mathbf{s}_i)|\boldsymbol{\theta}] = \begin{cases} \tilde{\mu}(\mathbf{s}_i) & n = 0 \\ \mathbf{x}(\mathbf{s}_i)'\boldsymbol{\beta} & n = N. \end{cases} \quad (3.3)$$

Equation (3.3) shows that $Y(\mathbf{s}_i)|\boldsymbol{\theta}$ achieves the true mean when we do not subsample, and achieves the parametric mean when we subsample all observations. Proposition 3.1 provides two expressions for the variance. In the first expression, the variance of $Y(\mathbf{s}_i)|\boldsymbol{\theta}$ is a weighted sum of the parametric variance and the true variance with an additional component. The second expression of $\text{Var}(Y(\mathbf{s}_i)|\boldsymbol{\theta})$ shows that the true variance of $Y(\mathbf{s}_i)|\boldsymbol{\theta}$ is biased by two components – (a) the weighted difference between the parametric variance and the true variance, and (b) the weighted square of the difference between the parametric mean and the true mean. However, we still obtain

$$\text{Var}(Y(\mathbf{s}_i)|\boldsymbol{\theta}) = \begin{cases} \tilde{\sigma}^2 & n = 0 \\ \tau^2 + \sigma^2 & n = N. \end{cases} \quad (3.4)$$

Again we see from Equation (3.4) that the variance of $Y(\mathbf{s}_i)|\boldsymbol{\theta}$ attains the true variance when we do not subsample, and attains the parametric variance when we subsample all observations. Lastly, Proposition 3.1 also provides the covariance between $Y(\mathbf{s}_i)$ and $Y(\mathbf{s}_j)$ given $\boldsymbol{\theta}$, which is a weighted sum of the true covariance, the parametric covariance, and the cross-product between the difference between the parametric and true means. Furthermore, we have

$$\text{Cov}(Y(\mathbf{s}_i), Y(\mathbf{s}_j)|\boldsymbol{\theta}) = \begin{cases} \tilde{C}(\mathbf{s}_i, \mathbf{s}_j) & n = 0 \\ \sigma^2 h_{ij}(\phi) & n = N. \end{cases} \quad (3.5)$$

Equation 3.5 shows that the covariance between $Y(\mathbf{s}_i)$ and $Y(\mathbf{s}_j)$ given $\boldsymbol{\theta}$ also attains the true covariance when we do not subsample, and attains the parametric covariance when we subsample all observations. Notice that, the joint spatial data subset model from (3.2) becomes the true unparameterized model, $w_\delta(\mathbf{y})$, when $n = 0$, and becomes the full parametric classical Bayesian spatial model, $\prod_{i=1}^N f(Y(\mathbf{s}_i), \boldsymbol{\nu}, \boldsymbol{\theta})$, when $n = N$. Hence, Equations (3.3), (3.4) and (3.5) exactly show the characteristics that we expect from the joint spatial data subset model derived from (3.2) under the conditions of no subsampling and subsampling all observations.

Typically, the spatial properties of a data are considered for the “de-trended” data process that has a zero mean. Hence, we define the de-trended version of $Y(\mathbf{s}_i)$, $\forall i = 1, \dots, N$ given $\boldsymbol{\theta}$ and $\boldsymbol{\delta}$ as

$$Z(\mathbf{s}_i)|\boldsymbol{\theta}, \boldsymbol{\delta} = \begin{cases} Y(\mathbf{s}_i) - \mathbf{x}(\mathbf{s}_i)'\boldsymbol{\beta} & \delta(\mathbf{s}_i) = 1 \\ Y(\mathbf{s}_i) - \tilde{\mu}(\mathbf{s}_i) & \delta(\mathbf{s}_i) = 0. \end{cases} \quad (3.6)$$

In a similar manner, we have that \tilde{C} is the covariance of the true de-trended data process $Y(\mathbf{s}_i) - \tilde{\mu}(\mathbf{s}_i), \forall i = 1, \dots, N$. This leads to Proposition 3.2.

Proposition 3.2 (Spatial Property). *Under SRS without replacement, if the true de-trended data process model is weakly and/or intrinsically stationary, then the de-trended data model defined in (3.6) is respectively weakly and/or intrinsically stationary. If the true de-trended data process model is non-stationary, then the de-trended data model defined in (3.6) is also non-stationary.*

Proof: See Appendix A.1 of the supplementary material (Saha and Bradley, 2024).

Proposition 3.2 implies that the de-trended data model defined in (3.6) retains the spatial characteristics of the true de-trended data process model, since the covariance between $Z(\mathbf{s}_i)$ and $Z(\mathbf{s}_j)$ given $\boldsymbol{\theta}$ depends on the true covariance between $Y(\mathbf{s}_i) - \tilde{\mu}(\mathbf{s}_i)$ and $Y(\mathbf{s}_j) - \tilde{\mu}(\mathbf{s}_j)$ in the following way,

$$\text{Cov}(Z(\mathbf{s}_i), Z(\mathbf{s}_j) | \boldsymbol{\theta}) = a\tilde{C}(\mathbf{s}_i, \mathbf{s}_j) + (b + p^2)\sigma^2 h_{ij}(\phi), \quad (3.7)$$

where $\tilde{C}(\mathbf{s}_i, \mathbf{s}_j)$ is the covariance of the true de-trended data process model, and p, a, b are defined in Proposition 3.1 (see Appendix A.1 of the supplementary material (Saha and Bradley, 2024) for the proof of Equation (3.7)). Notice that, the covariance between $Z(\mathbf{s}_i)$ and $Z(\mathbf{s}_j)$ given $\boldsymbol{\theta}$ is also a weighted sum of the true and parametric covariances of the de-trended data. In the case where the true de-trended data process model is intrinsically stationary, we get the variogram as

$$2\gamma(\mathbf{s}_i - \mathbf{s}_j) = 2\{p(\tau^2 + \sigma^2) + (1 - p)\tilde{\sigma}^2\} - 2\{a\tilde{C}(\mathbf{s}_i, \mathbf{s}_j) + (b + p^2)\sigma^2 h_{ij}(\phi)\}, \quad (3.8)$$

where $p, a, b, \tilde{\sigma}^2$ are defined in Proposition 3.1, $\mathbf{s}_i - \mathbf{s}_j$ is the spatial lag between \mathbf{s}_i and \mathbf{s}_j , and $2\gamma(\mathbf{s}_i - \mathbf{s}_j)$ is called the variogram (see Appendix A.1 of the supplementary material (Saha and Bradley, 2024) for the proof of variogram). This helps us to visualize the spatial properties of the model in (3.2). If we assume that the true de-trended data process model is weakly stationary, i.e. if the true covariance $\tilde{C}(\mathbf{s}_i, \mathbf{s}_j)$ depends on the lag between \mathbf{s}_i and \mathbf{s}_j , then we have the usual definitions of the sill, nugget and range,

$$\begin{aligned} \text{Sill} &= \lim_{\|\mathbf{d}\| \rightarrow \infty} \gamma(\mathbf{d}) = p(\tau^2 + \sigma^2) + (1 - p)\tilde{\sigma}^2, \\ \text{Nugget} &= \lim_{\|\mathbf{d}\| \rightarrow 0^+} \gamma(\mathbf{d}) = [p\tau^2 + (p - b - p^2)\sigma^2] + [(1 - p - a)\tilde{\sigma}^2 + a\tilde{\tau}^2], \\ \text{Range} &= \inf \{\mathbf{d} : a\tilde{C}(\mathbf{d}) + (b + p^2)\sigma^2 \rho_c(\mathbf{d}; \phi) = 0\}, \end{aligned} \quad (3.9)$$

where $\mathbf{d} = \mathbf{s}_i - \mathbf{s}_j$ is the spatial lag and $\tilde{\tau}^2$ is the nugget of the true de-trended data process. Notice that, the nugget is dependent on both parametric nugget (τ^2) and the true nugget ($\tilde{\tau}^2$). Refer to Appendix A.1 of the supplementary material (Saha and Bradley, 2024) for the proof of sill, nugget and range. Equation (3.9) reproduces the sill, nugget, and range of the (true) parametric model when $(n = 0) n = N$. Thus, we again achieve a type of balance between the true model and parametric model.

3.3 Properties of Spatial Data Subset Model when the Sampling Design is Stratified Random Sampling

In the case of stratified random sampling, the properties of the spatial data subset model discussed in Section 3.2 change, which is important to explore. We start with the moment properties when $\Pr(\delta|n)$ is defined by stratified random samples.

Proposition 3.3. *Let, D_1, \dots, D_R represent R disjoint strata such that the spatial domain $D = \cup_{r=1}^R D_r$. Let, $\Pr(\delta|n)$ be defined by stratified random samples with R strata. Then the expectation, variance and covariance of the data $Y(\mathbf{s}_i)$ for $i = 1, \dots, N$ given the real-valued parameter vector $\boldsymbol{\theta} = (\boldsymbol{\beta}, \tau^2, \sigma^2, \sigma_{\beta}^2, \phi)'$ are given by*

When $\mathbf{s}_i \in D_r$ for $r = 1, \dots, R$,

$$\begin{aligned}\mathbb{E}[Y(\mathbf{s}_i)|\boldsymbol{\theta}] &= p_r\{\mathbf{x}(\mathbf{s}_i)'\boldsymbol{\beta}\} + (1 - p_r)\tilde{\mu}(\mathbf{s}_i) = p_r\{\mathbf{x}(\mathbf{s}_i)'\boldsymbol{\beta} - \tilde{\mu}(\mathbf{s}_i)\} + \tilde{\mu}(\mathbf{s}_i), \\ \text{Var}(Y(\mathbf{s}_i)|\boldsymbol{\theta}) &= p_r(\tau^2 + \sigma^2) + (1 - p_r)\tilde{\sigma}^2 + p_r(1 - p_r)\{\mathbf{x}(\mathbf{s}_i)'\boldsymbol{\beta} - \tilde{\mu}(\mathbf{s}_i)\}^2 \\ &= p_r(\tau^2 + \sigma^2 - \tilde{\sigma}^2) + \tilde{\sigma}^2 + p_r(1 - p_r)\{\mathbf{x}(\mathbf{s}_i)'\boldsymbol{\beta} - \tilde{\mu}(\mathbf{s}_i)\}^2;\end{aligned}$$

When $\mathbf{s}_i, \mathbf{s}_j \in D_r$ for $r = 1, \dots, R$,

$$\begin{aligned}\text{Cov}(Y(\mathbf{s}_i), Y(\mathbf{s}_j)|\boldsymbol{\theta}) &= a_r\tilde{C}(\mathbf{s}_i, \mathbf{s}_j) + (b_r + p_r^2)\sigma^2 h_{ij}(\phi) \\ &\quad + b_r\{\mathbf{x}(\mathbf{s}_i)'\boldsymbol{\beta} - \tilde{\mu}(\mathbf{s}_i)\}\{\mathbf{x}(\mathbf{s}_j)'\boldsymbol{\beta} - \tilde{\mu}(\mathbf{s}_j)\};\end{aligned}$$

When $\mathbf{s}_i \in D_r, \mathbf{s}_j \in D_t$ and $r \neq t$ for $r, t \in \{1, \dots, R\}$,

$$\text{Cov}(Y(\mathbf{s}_i), Y(\mathbf{s}_j)|\boldsymbol{\theta}) = (p_r p_t - p_r - p_t + 1)\tilde{C}(\mathbf{s}_i, \mathbf{s}_j) + p_r p_t \sigma^2 h_{ij}(\phi),$$

where $p_r = \frac{n_r}{N_r}, p_t = \frac{n_t}{N_t}, a_r = \left(\frac{n_r}{N_r} \frac{n_r - 1}{N_r - 1} - 2 \frac{n_r}{N_r} + 1\right), b_r = \frac{n_r}{N_r} \left(\frac{n_r - 1}{N_r - 1} - \frac{n_r}{N_r}\right)$, n_r is the number of subsamples drawn from stratum r with size N_r , n_t is the number of subsamples drawn from stratum t with size N_t , and $\tilde{\mu}(\mathbf{s}_i), \tilde{\sigma}^2, \tilde{C}(\mathbf{s}_i, \mathbf{s}_j)$ are defined in Proposition 3.1.

Proof: See Appendix A.2 of the supplementary material (Saha and Bradley, 2024).

Notice that, the expectation of $Y(\mathbf{s}_i)|\boldsymbol{\theta}$ is a weighted sum of the parametric mean and the true mean (shown in the first expression of $\mathbb{E}[Y(\mathbf{s}_i)|\boldsymbol{\theta}]$) like it was for SRS without replacement. Also, the variance of $Y(\mathbf{s}_i)|\boldsymbol{\theta}$ is a weighted sum of the parametric variance and the true variance with an additional component as shown in the first expression of $\text{Var}(Y(\mathbf{s}_i)|\boldsymbol{\theta})$ (similar to SRS without replacement). On the other hand, the covariance between $Y(\mathbf{s}_i)$ and $Y(\mathbf{s}_j)$ depends on whether \mathbf{s}_i and \mathbf{s}_j are in the same stratum or not. In both the cases, the covariance is a weighted sum of the parametric covariance and the true covariance (plus a weighted cross-product between the difference between the parametric and true means, when \mathbf{s}_i and \mathbf{s}_j are in same stratum). Furthermore, we arrive at the same equations as (3.3), (3.4) and (3.5) from Proposition 3.3 under the conditions $n_r = n_t = 0, n_r = N_r$ and $n_t = N_t$. This implies that even if the sampling design is different, the properties of the model in (3.2) are fully non-parametric when we do not subsample and fully parametric when we subsample all observations.

It is also noticeable that p_r (and p_t) is dependent on the number of subsamples drawn from stratum r (and t) and size of the corresponding stratum. If the sizes of all strata are same, and equal number of subsamples are drawn from each stratum, then the probability of sampling one data point using SRS without replacement is same as the probability of sampling one data point using stratified random sampling, since $p = n/N = \frac{n/R}{N/R} = n_r/N_r = p_r$ for $r = 1, \dots, R$ (same relation can be found for p_t). So, p_r (or p_t) is different than p when the sizes of strata are different or unequal number of subsamples are drawn from each stratum. Thus, the expectation and variance of $Y(\mathbf{s}_i)|\boldsymbol{\theta}$ in Propositions 3.1 and 3.3 are the same when equal number of subsamples are drawn from equal sized strata. However, the covariances between $Y(\mathbf{s}_i)$ and $Y(\mathbf{s}_j)$ given $\boldsymbol{\theta}$ under two sampling techniques are unequal.

Proposition 3.4 (Spatial Property). *Under stratified random sampling, if the true de-trended data process model is weakly and/or intrinsically stationary, then the de-trended data model defined in (3.6) is respectively weakly and/or intrinsically stationary. Similarly, if the true de-trended data process model is non-stationary, then the de-trended data model defined in (3.6) is also non-stationary.*

Proof: See Appendix A.2 of the supplementary material (Saha and Bradley, 2024).

Proposition 3.4 also implies that the de-trended data model defined in (3.6) always retains the stationary properties of the true data process model. Under stratified random sampling, the covariance between $Z(\mathbf{s}_i)$ and $Z(\mathbf{s}_j)$ given $\boldsymbol{\theta}$ becomes

$$\text{Cov}(Z(\mathbf{s}_i), Z(\mathbf{s}_j)|\boldsymbol{\theta}) = \begin{cases} a_r \tilde{C}(\mathbf{s}_i, \mathbf{s}_j) + (b_r + p_r^2) \sigma^2 h_{ij}(\phi) & \text{when } \mathbf{s}_i, \mathbf{s}_j \in D_r \\ a_{rt} \tilde{C}(\mathbf{s}_i, \mathbf{s}_j) + p_r p_t \sigma^2 h_{ij}(\phi) & \text{when } \mathbf{s}_i \in D_r, \mathbf{s}_j \in D_t, r \neq t, \end{cases} \quad (3.10)$$

where $a_{rt} = p_r p_t - p_r - p_t + 1$, $\tilde{C}(\mathbf{s}_i, \mathbf{s}_j)$ is the covariance of the true de-trended data process model, and p_r, a_r, b_r, p_t are defined in Proposition 3.3. Equation (3.10) (see Appendix A.2 of the supplementary material (Saha and Bradley, 2024) for the proof) shows that under stratified random sampling, the covariance between $Z(\mathbf{s}_i)$ and $Z(\mathbf{s}_j)$ given $\boldsymbol{\theta}$ is again a weighted sum of the true and parametric covariances of the de-trended data, but the weights change based on \mathbf{s}_i and \mathbf{s}_j are in same stratum or not. If the true data process model is assumed to be intrinsically stationary, then we get two different variograms depending on whether \mathbf{s}_i and \mathbf{s}_j are in the same stratum or not. When $\mathbf{s}_i, \mathbf{s}_j \in D_r$ for $r = 1, \dots, R$, we obtain the variogram as

$$2\gamma(\mathbf{s}_i - \mathbf{s}_j) = 2\{p_r(\tau^2 + \sigma^2) + (1 - p_r)\tilde{\sigma}^2\} - 2\{a_r \tilde{C}(\mathbf{s}_i, \mathbf{s}_j) + (b_r + p_r^2)\sigma^2 h_{ij}(\phi)\}, \quad (3.11)$$

where p_r, a_r and b_r are defined in Proposition 3.3. When $\mathbf{s}_i \in D_r, \mathbf{s}_j \in D_t$ and $r \neq t$ for $r, t \in \{1, \dots, R\}$, the variogram becomes

$$2\gamma(\mathbf{s}_i - \mathbf{s}_j) = (p_r + p_t)(\tau^2 + \sigma^2) + (2 - p_r - p_t)\tilde{\sigma}^2 - 2\{a_{rt} \tilde{C}(\mathbf{s}_i, \mathbf{s}_j) + p_r p_t \sigma^2 h_{ij}(\phi)\}, \quad (3.12)$$

where $a_{rt} = p_r p_t - p_r - p_t + 1$. The proof of Equations (3.11) and (3.12) are shown in Appendix A.2 of the supplementary material (Saha and Bradley, 2024). If the true de-trended data process model is assumed to be weakly stationary, i.e. if the true covariance

$\tilde{C}(\mathbf{s}_i, \mathbf{s}_j)$ depends on the lag between \mathbf{s}_i and \mathbf{s}_j , then we have the usual definitions of the sill, nugget and range,

$$\begin{aligned}
\text{Sill} &= \begin{cases} p_r(\tau^2 + \sigma^2) + (1 - p_r)\tilde{\sigma}^2 & \mathbf{s}_i, \mathbf{s}_j \in D_r \\ \frac{p_r+p_t}{2}(\tau^2 + \sigma^2) + \left(1 - \frac{p_r+p_t}{2}\right)\tilde{\sigma}^2 & \mathbf{s}_i \in D_r, \mathbf{s}_j \in D_t, r \neq t, \end{cases} \\
\text{Nugget} &= \begin{cases} [p_r\tau^2 + (p_r - b_r - p_r^2)\sigma^2] + [(1 - p_r - a_r)\tilde{\sigma}^2 + a_r\tilde{\tau}^2] & \mathbf{s}_i, \mathbf{s}_j \in D_r \\ \left[\frac{p_r+p_t}{2}\tau^2 + \frac{p_r+p_t-2p_r p_t}{2}\sigma^2\right] + \left[\frac{p_r+p_t-2p_r p_t}{2}\tilde{\sigma}^2 + a_{rt}\tilde{\tau}^2\right] & \mathbf{s}_i \in D_r, \mathbf{s}_j \in D_t, \\ & r \neq t, \end{cases} \\
\text{Range} &= \begin{cases} \inf \{ \mathbf{d} : a_r\tilde{C}(\mathbf{d}) + (b_r + p_r^2)\sigma^2\rho_c(\mathbf{d}; \phi) = 0 \} & \mathbf{s}_i, \mathbf{s}_j \in D_r \\ \inf \{ \mathbf{d} : a_{rt}\tilde{C}(\mathbf{d}) + p_r p_t \sigma^2 \rho_c(\mathbf{d}; \phi) = 0 \} & \mathbf{s}_i \in D_r, \mathbf{s}_j \in D_t, r \neq t, \end{cases}
\end{aligned} \tag{3.13}$$

where $\mathbf{d} = \mathbf{s}_i - \mathbf{s}_j$ is the spatial lag and $\tilde{\tau}^2$ is the nugget of the true de-trended data process. This is to emphasize that the spatial characteristics change based on whether \mathbf{s}_i and \mathbf{s}_j belong to same stratum or not (see Appendix A.2 of the supplementary material (Saha and Bradley, 2024) for the proof of sill, nugget and range). Equation (3.13) reproduces the sill, nugget, and range of the (true) parametric model when ($n_r = 0$ and $n_t = 0$) $n_r = N_r$ and $n_t = N_t$. Thus, we again achieve a type of balance between the true model and parametric model.

Propositions 3.2 and 3.4 show that the SDSM assumes the data is stationary when w implies a stationary process, and the SDSM assumes the data is nonstationary when w implies a nonstationary process. This consequence is particularly important for motivating the use of the SDSM. There are times when one would just prefer to implement a standard model. This is because the properties of standard models (e.g., weak stationarity, sill, nugget, and range) are well-known and easier to communicate to the general public and collaborators in different fields. However, it is difficult to apply these models in high-dimensional settings because of computational problems and the assumptions on weak stationarity may be suspect in high-dimensions. Our proposed method allows one to fit a traditional Kriging model (to each subsample) in a way that respects the stationary properties of the entire spatial dataset (via our propositions), without throwing away data, and in a way that is fully scalable.

One important point to mention is that we can not use these propositions in practice. This is because w is the true nonparametric density associated with the measure that generated the realization $\{Y(\mathbf{s}_i)\}$ and is generally unknown (a common assumption in nonparametric statistics, see, Hollander et al., 2013). The main purpose of these propositions is that it aids with interpretation (i.e., interpretation of model assumptions), which subsequently, motivates the use of our SDSM model. Specifically, we are assuming the subsample (i.e., all $Y(\mathbf{s}_i)$ such that $\delta(\mathbf{s}_i) = 1$) follows a parametric model, and the holdout data (i.e., all $Y(\mathbf{s}_i)$ such that $\delta(\mathbf{s}_i) = 0$) is assumed to be distributed according to the actual data generating mechanism that produced it (assuming it has a density), which is unknown, and this is true for every split of the data into training and holdout. Moreover, we can strike a balance between these models through our choice of n (i.e., when $n = N$ we obtain a parametric model, and with $n = 0$ we obtain a nonparametric model).

One can plot the sill, nugget, and range associated with the covariance function given by $cov(Y(\mathbf{s}_i), Y(\mathbf{s}_j)|\boldsymbol{\theta}, \delta(\mathbf{s}_i) = 1, \delta(\mathbf{s}_j) = 1) = \sigma^2 h_{ij}(\phi)$, where $h_{ij}(\phi)$ is the covariogram evaluated at $\mathbf{s}_i - \mathbf{s}_j$. Note that $cov(Y(\mathbf{s}_i), Y(\mathbf{s}_j)|\boldsymbol{\theta}, \delta(\mathbf{s}_i) = 1, \delta(\mathbf{s}_j) = 1)$ is different than $cov(Y(\mathbf{s}_i), Y(\mathbf{s}_j)|\boldsymbol{\theta})$. One cannot plot the sill, nugget, and range associated with the covariance function $cov(Y(\mathbf{s}_i), Y(\mathbf{s}_j)|\boldsymbol{\theta})$, which contains the unknown w in its expression. The fact that $cov(Y(\mathbf{s}_i), Y(\mathbf{s}_j)|\boldsymbol{\theta}, \delta(\mathbf{s}_i) = 1, \delta(\mathbf{s}_j) = 1)$ is different from $cov(Y(\mathbf{s}_i), Y(\mathbf{s}_j)|\boldsymbol{\theta})$, restricts our interpretation of the sill, nugget, and range. Traditionally, these parameters describe a covariance function at any two locations in the spatial domain. For example, the traditional interpretation of the sill and nugget describes the covariance function of the response at any two locations in the spatial domain as the spatial lag goes to infinity and zero, respectively. In our model, these interpretations hold only for the locations in the training set (i.e., $\delta(\mathbf{s}_i) = \delta(\mathbf{s}_j) = 1$). That is, the sill and nugget in our model describe the covariance function of the response at any two locations with $\delta(\mathbf{s}_i) = \delta(\mathbf{s}_j) = 1$ as the spatial lag goes to infinity and zero, respectively.

3.4 MCMC Implementation

The full-conditional distributions of $\boldsymbol{\nu}_\delta, \boldsymbol{\nu}_A, \boldsymbol{\beta}, \tau^2, \sigma^2, \sigma_\beta^2$, and ϕ are derived (see Appendix B of the supplementary material; Saha and Bradley, 2024) using a standard proportionality argument, and we reiterate that for any given $\boldsymbol{\delta}$ the full-conditional distribution is only implemented using a single dataset of size $n \ll N$. The prior distribution for all the variance parameters ($\tau^2, \sigma^2, \sigma_\beta^2$) are chosen to be a non-informative inverse Gamma distribution, $\mathcal{IG}(a, b)$, among different choices (Gelman et al., 1995; Daniels, 1999; Gelman, 2006; Banerjee et al., 2014). A standard choice for the prior of $\boldsymbol{\beta}$ is independent Normal distribution with mean 0 and variance σ_β^2 (Banerjee et al., 2014). Learning the range parameter, ϕ , has been a challenge in spatial statistics over decades (Berger et al., 2001; Paulo, 2005; Pilz and Spöck, 2008; Kazianka and Pilz, 2012). Berger et al. (2001) investigated that common choices of prior distribution for ϕ , like a constant prior $\pi(\phi) = 1$ or independent Jeffreys prior typically result in an improper posterior distribution for $\boldsymbol{\theta} = (\boldsymbol{\beta}, \tau^2, \sigma^2, \sigma_\beta^2, \phi)'$. As a solution, they suggested that a Jeffrey rule prior or a reference prior for ϕ yields a proper posterior distribution. It is also often suggested to choose an informative prior for ϕ (Banerjee et al., 2014) such as, a uniform distribution over a specified small interval, or a discrete uniform distribution over a specific finite set of points. We specify $\pi(\phi)$ as a discrete uniform distribution with a support of a finite set of points.

To sample from the joint distribution of random variables say X and Y , one can first sample Y from its density $f(Y)$ and then sample X from the conditional density $f(X|Y)$. This type of sampler is well-known (e.g., see Dunn and Shultis, 2022; Van Dyk and Park, 2008), and we call this a “composite sampler” borrowing terminology from the composite likelihood literature (e.g., see Varin et al., 2011). In our case, we first sample from $p(\boldsymbol{\delta}|\mathbf{y}, n) = p(\boldsymbol{\delta}|n)$ and then sample from $p(\boldsymbol{\nu}, \boldsymbol{\theta}|\mathbf{y}, \boldsymbol{\delta}, n)$ using a Gibbs sampler. We call this a “Gibbs-within-composite sampler”, since a Gibbs sampling step is used within our composite sampler. This sampling approach is needed to avoid specifying the value of the true density w . The density w does not contain $\boldsymbol{\theta}$ and $\boldsymbol{\nu}$, and hence is a proportionality constant when deriving the full-conditionals for $\boldsymbol{\theta}$ and $\boldsymbol{\nu}$. As a

result, those full-conditionals do not contain w in its expression. We also assume $\boldsymbol{\delta}$ is independent of the data \mathbf{y} , where note, Bradley (2021) showed that there exists Bayesian hierarchical models that allow $\boldsymbol{\delta}$ to be independent of \mathbf{y} . Consequently $p(\boldsymbol{\delta}|\mathbf{y}, n) = p(\boldsymbol{\delta}|n)$, which does not contain w in its expression. This allows for one to adopt a Gibbs-within-composite sampler strategy to avoid specifying the true density for the holdout $\mathbf{y}_{-\delta}$.

Bradley (2021) suggested to use a Gibbs sampler, which is reasonable to do when the full-conditional distribution for $\boldsymbol{\delta}$ has the following property, $p(\boldsymbol{\delta}|\boldsymbol{\nu}, \boldsymbol{\theta}, \mathbf{y}, n) = p(\boldsymbol{\delta}|n)$; that is, if $\boldsymbol{\delta}$ is independent of $\boldsymbol{\nu}$, $\boldsymbol{\theta}$ and \mathbf{y} . When $p(\boldsymbol{\delta}|\boldsymbol{\nu}, \boldsymbol{\theta}, \mathbf{y}, n) = p(\boldsymbol{\delta}|n)$, then the Gibbs-within-composite sampler becomes a Gibbs sampler. This assumption is reasonable when n is sufficiently large, so that the distribution $p(\boldsymbol{\nu}, \boldsymbol{\theta}|\mathbf{y}, \boldsymbol{\delta}, n)$ is roughly constant over different values of $\boldsymbol{\delta}$, which occurs when the spatial statistical model is posterior consistent. This can be seen through the following heuristic,

$$\begin{aligned}
 p(\boldsymbol{\delta}|\boldsymbol{\nu}, \boldsymbol{\theta}, \mathbf{y}, n) &= \frac{p(\boldsymbol{\delta}, \boldsymbol{\nu}, \boldsymbol{\theta}|\mathbf{y}, n)}{p(\boldsymbol{\nu}, \boldsymbol{\theta}|\mathbf{y}, n)} = \frac{p(\boldsymbol{\delta}|\mathbf{y}, n)p(\boldsymbol{\nu}, \boldsymbol{\theta}|\boldsymbol{\delta}, \mathbf{y}, n)}{p(\boldsymbol{\nu}, \boldsymbol{\theta}|\mathbf{y}, n)} \\
 &= \frac{p(\boldsymbol{\delta}|\mathbf{y}, n)p(\boldsymbol{\nu}, \boldsymbol{\theta}|\boldsymbol{\delta}, \mathbf{y}, n)}{\sum_{\boldsymbol{\delta}^*} p(\boldsymbol{\delta}^*|\mathbf{y}, n)p(\boldsymbol{\nu}, \boldsymbol{\theta}|\boldsymbol{\delta}^*, \mathbf{y}, n)} \\
 &= \frac{p(\boldsymbol{\delta}|\mathbf{y}, n)}{\sum_{\boldsymbol{\delta}^*} p(\boldsymbol{\delta}^*|\mathbf{y}, n) \{p(\boldsymbol{\nu}, \boldsymbol{\theta}|\boldsymbol{\delta}^*, \mathbf{y}, n)/p(\boldsymbol{\nu}, \boldsymbol{\theta}|\boldsymbol{\delta}, \mathbf{y}, n)\}} \\
 &\approx \frac{p(\boldsymbol{\delta}|\mathbf{y}, n)}{\sum_{\boldsymbol{\delta}^*} p(\boldsymbol{\delta}^*|\mathbf{y}, n)} = p(\boldsymbol{\delta}|\mathbf{y}, n) = p(\boldsymbol{\delta}|n),
 \end{aligned} \tag{3.14}$$

when $\lim_{n \rightarrow \infty} \frac{p(\boldsymbol{\nu}, \boldsymbol{\theta}|\boldsymbol{\delta}^*, \mathbf{y}, n)}{p(\boldsymbol{\nu}, \boldsymbol{\theta}|\boldsymbol{\delta}, \mathbf{y}, n)} = 1$ if the spatial statistical model is posterior consistent. In practice, one needs to check that $p(\boldsymbol{\nu}, \boldsymbol{\theta}|\boldsymbol{\delta}^*, \mathbf{y}, n)$ approaches $p(\boldsymbol{\nu}, \boldsymbol{\theta}|\boldsymbol{\delta}, \mathbf{y}, n)$ if n is sufficiently large, which can be checked empirically by comparing the replicates from the Gibbs sampler shown in Algorithm 1 to the replicates from the Gibbs-within-composite sampler. We provide examples of this empirical investigation for both our simulations and application in Appendix C of the supplementary material (Saha and Bradley, 2024). In both studies we find it reasonable to use a Gibbs sampler, but emphasize that this assumption should be checked empirically (or theoretically) in practice.

4 Simulations

In this section, we use simulation to demonstrate the predictive and computational performance of the proposed spatial data subset model in (3.2) under different model specifications. The simulation setup is provided in Section 4.1. We specify $\Pr(\boldsymbol{\delta}|n)$ in a way that it generates (a) SRS without replacement, and (b) stratified random samples. We illustrate the results generated by the spatial data subset model under both sampling techniques in Section 4.2. We also compare the spatial data subset model to the full model by demonstrating a simulation study on a smaller dataset in Section 4.3. Finally, we give a detailed analysis on the scalability of our proposed model in Section 4.4.

Algorithm 1: Implementation of the model in (3.2) for a specific n .

1 Define $n \ll N$, G , g_0 (burn-in), and a set of prediction locations $A \subset D$.
2 Initialize $\theta^{[0]} = (\beta^{[0]}, \tau^{2[0]}, \sigma^2[0], \sigma_\beta^2[0], \phi^{[0]})'$.
3 **for** $g = 1 : G$ **do**
4 Sample $\delta^{[g]} \sim \text{Pr}(\delta|n)$.
5 Set the n -dimensional vector $\mathbf{y}_\delta^{[g]} = (Y(\mathbf{s}_i) : \delta(\mathbf{s}_i)^{[g]} = 1)'$.
6 Set the $n \times p$ matrix $\mathbf{X}_\delta^{[g]} = (\mathbf{x}(\mathbf{s}_i) : \delta(\mathbf{s}_i)^{[g]} = 1)'$.
7 Calculate the $n \times n$ matrix $\mathbf{H}_\delta(\phi^{[g-1]}) = (h_{jl}(\phi^{[g-1]}) : j, l \in \{i : \delta(\mathbf{s}_i)^{[g]} = 1\})$.
8 Sample $\boldsymbol{\nu}_\delta^{[g]} = (\nu(\mathbf{s}_i)^{[g]} : \delta(\mathbf{s}_i)^{[g]} = 1)'$ from

$$\mathcal{N}\left\{\left(\mathbf{I}_n + \frac{\tau^{2[g-1]}}{\sigma^2[g-1]} \mathbf{H}_\delta^{-1}(\phi^{[g-1]})\right)^{-1} \left(\mathbf{y}_\delta^{[g]} - \mathbf{X}_\delta^{[g]} \boldsymbol{\beta}^{[g-1]}\right), \left(\frac{1}{\tau^{2[g-1]}} \mathbf{I}_n + \frac{1}{\sigma^2[g-1]} \mathbf{H}_\delta^{-1}(\phi^{[g-1]})\right)^{-1}\right\}.$$

9 Sample $\boldsymbol{\beta}^{[g]}$ from

$$\mathcal{N}\left\{\left(\mathbf{X}_\delta^{[g]'} \mathbf{X}_\delta^{[g]} + \frac{\tau^{2[g-1]}}{\sigma_\beta^2[g-1]} \mathbf{I}_p\right)^{-1} \mathbf{X}_\delta^{[g]'} \left(\mathbf{y}_\delta^{[g]} - \boldsymbol{\nu}_\delta^{[g]}\right), \left(\frac{1}{\tau^{2[g-1]}} \mathbf{X}_\delta^{[g]'} \mathbf{X}_\delta^{[g]} + \frac{1}{\sigma_\beta^2[g-1]} \mathbf{I}_p\right)^{-1}\right\}.$$

10 Sample $\tau^{2[g]}$ from

$$\mathcal{IG}\left\{a_\tau + \frac{n}{2}, b_\tau + \frac{1}{2} \left(\mathbf{y}_\delta^{[g]} - \mathbf{X}_\delta^{[g]} \boldsymbol{\beta}^{[g]} - \boldsymbol{\nu}_\delta^{[g]}\right)' \left(\mathbf{y}_\delta^{[g]} - \mathbf{X}_\delta^{[g]} \boldsymbol{\beta}^{[g]} - \boldsymbol{\nu}_\delta^{[g]}\right)\right\}.$$

11 Sample $\sigma^2[g]$ from $\mathcal{IG}\left\{a_\sigma + \frac{n}{2}, b_\sigma + \frac{1}{2} \boldsymbol{\nu}_\delta^{[g]'} \mathbf{H}_\delta^{-1}(\phi^{[g-1]}) \boldsymbol{\nu}_\delta^{[g]}\right\}$.
12 Sample $\sigma_\beta^2[g]$ from $\mathcal{IG}\left\{a_\beta + \frac{p}{2}, b_\beta + \frac{1}{2} \boldsymbol{\beta}^{[g]'} \boldsymbol{\beta}^{[g]}\right\}$.
13 Sample $\phi^{[g]}$ from its full-conditional distribution, where $\pi(\phi)$ is discrete uniform.
14 **end**
15 **Prediction Steps:**
16 **for** $g = g_0 : G$ **do**
17 Calculate $\mathbf{H}_m(\phi^{[g-1]}) = (h_{jl}(\phi^{[g-1]}) : j \in A, l \in \{i : \delta(\mathbf{s}_i)^{[g]} = 1\})$.
18 Calculate $\mathbf{H}_A(\phi^{[g-1]}) = (h_{jl}(\phi^{[g-1]}) : j, l \in A)$.
19 **Parallel computation at each** $i \in A$:
20 Set \mathbf{h}'_{mi} = i -th row of $\mathbf{H}_m(\phi^{[g-1]})$.
21 Set h_{Aii} = i -th diagonal element of $\mathbf{H}_A(\phi^{[g-1]})$.
22 Sample $\nu(\mathbf{s}_i)^{[g]}$ from

$$\mathcal{N}\left\{\mathbf{h}'_{mi} \mathbf{H}_\delta^{-1}(\phi^{[g-1]}) \boldsymbol{\nu}_\delta^{[g]}, \sigma^2[g-1] \left[h_{Aii} - \mathbf{h}'_{mi} \mathbf{H}_\delta^{-1}(\phi^{[g-1]}) (\mathbf{h}'_{mi})'\right]\right\}.$$

23 Compute $W(\mathbf{s}_i)^{[g]} = \mathbf{x}(\mathbf{s}_i)' \boldsymbol{\beta}^{[g]} + \nu(\mathbf{s}_i)^{[g]}$.
24 **end**
25 **end**
26 **end**

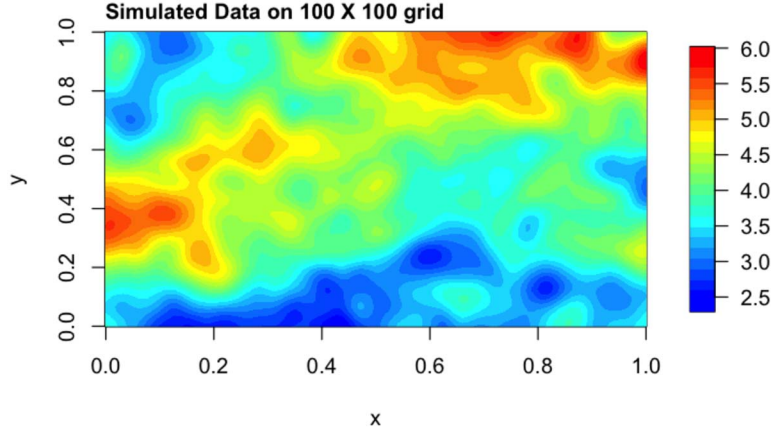


Figure 4.1: Image plot of the simulated spatial data on a 100×100 grid over $[0, 1] \times [0, 1]$. This data is simulated based on the simulation setup mentioned in Section 4.1.

4.1 Simulation Setup

We fix the spatial domain $D \subset \mathbb{R}^2$, where D is a 100×100 grid over $[0, 1] \times [0, 1]$. Thus, there are 10,000 spatial locations in total, all equally spaced. First, we simulate \mathbf{X} , which is a $10,000 \times 2$ dimensional matrix, from an independent $\mathcal{U}(0, 1)$ distribution. Then, we simulate the process model $\boldsymbol{\nu}$, a vector of size 10,000, from a non-stationary model using `NSconvo_sim` function from `convoSPAT` library in R, with true $\boldsymbol{\beta} = (2, 3)'$ and exponential covariance function, and calculate $\mathbf{w} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\nu}$. Finally, we simulate the data vector \mathbf{y} , a vector of size 10,000, from the latent process and choose τ^2 so that we have a signal-to-noise-ratio of 3 ($\text{SNR} = \text{Var}(\mathbf{w})/\text{Var}(\boldsymbol{\epsilon})$). Figure 4.1 shows the simulated spatial data \mathbf{y} based on this simulation setup. We also enforce 20% missing values at random, and 10% are enforced to be missing on an arbitrary 25×40 grid (total number of missing locations = 3,000). In practice, missing data may occur over large regions similar to our 25×40 grid specification. We consider two specifications of $\Pr(\boldsymbol{\delta}|n)$: (a) SRS without replacement, and (b) stratified random sampling. For the latter part, we define four strata: $\{\mathbf{s}_i : x \leq 0.5, y \leq 0.5\}$, $\{\mathbf{s}_i : x \leq 0.5, y > 0.5\}$, $\{\mathbf{s}_i : x > 0.5, y \leq 0.5\}$ and $\{\mathbf{s}_i : x > 0.5, y > 0.5\}$. The same missing data structure is used for both sampling designs.

4.2 An Illustration

We implement the spatial data subset model in (3.2) on the simulated data with missing values defined in Section 4.1. We denote the locations of these missing values as $A \subset D$, and we aim to make prediction at these locations. The covariogram for the model in (3.2) is chosen to be exponential, i.e. $h_{ij}(\phi) = \exp(-\phi\|\mathbf{s}_i - \mathbf{s}_j\|)$. We use standard conjugate prior specifications (see Appendix B of the supplementary material (Saha and Bradley, 2024) for more details). The model in (3.2) is implemented in the Gibbs sampler using Algorithm 1 (since stationarity reaches irrespective of using a Gibbs-within-composite

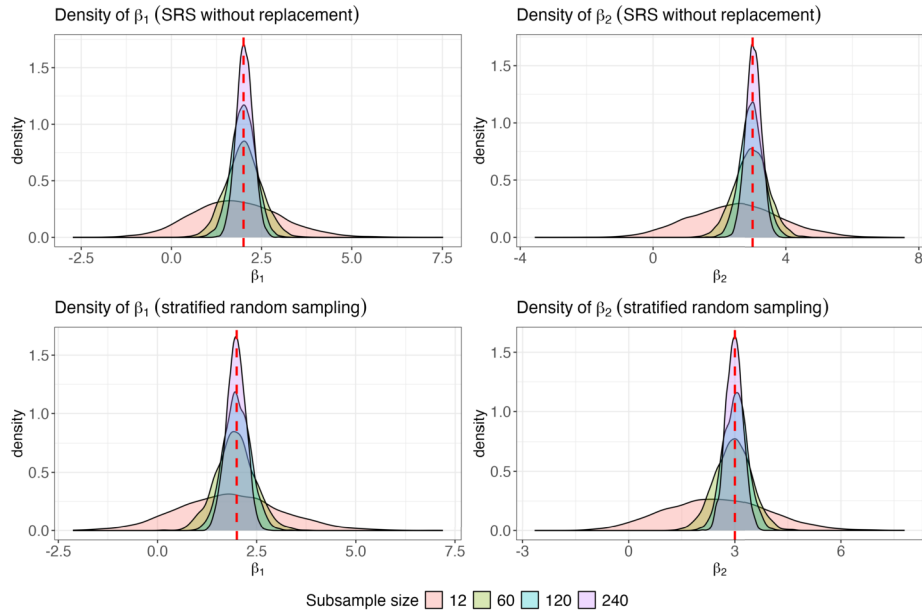


Figure 4.2: We plot the density of MCMC replicates for β_1 (left column) and β_2 (right column) for different subsample sizes when $\Pr(\delta|n)$ is defined by SRS without replacement (top row) and by stratified random sampling (bottom row). The red dashed lines represent the true values $\beta_1 = 2$ (left column) and $\beta_2 = 3$ (right column).

sampler; see Appendix C of the supplementary material (Saha and Bradley, 2024) for a detailed description) with $G = 10,000$ iterations and a burn-in of $g_0 = 2,000$. Finally, we predict the response at each location in A using Algorithm 1. We run the same algorithm (Algorithm 1) when $\Pr(\delta|n)$ is defined by (a) SRS without replacement, and (b) stratified random samples. We use the author’s laptop computer with the following specification to run the algorithm: 11 Gen Intel(R) Core(TM) i7-1165G7 CPU with 2.80GHz and 16 GB RAM.

To investigate the effect of subsample size on the estimates of β , we plot the density of MCMC replicates for β_1 and β_2 for different values of n under both SRS without replacement and stratified random sampling in Figure 4.2. We see that the densities of MCMC replicates for both β_1 and β_2 have high variance, non-zero skewness and platykurtic peak when n is small. As we increase n , the densities become sharper with less variability, and become more symmetric about the corresponding true values of β_1 and β_2 . The densities achieve least variability when $n = 240$. These properties show that the posterior samples of β tend to have less variability with its mean approaching to the true value as n grows, which is consistent with the results in Bradley (2021). We see similar properties within the posterior samples of β across the sampling strategies.

To evaluate the inferential performance of our model, we compute the root mean square error (RMSE) between the process $W(\mathbf{s}_i)$ and the prediction $\widehat{W}(\mathbf{s}_i)$ for each

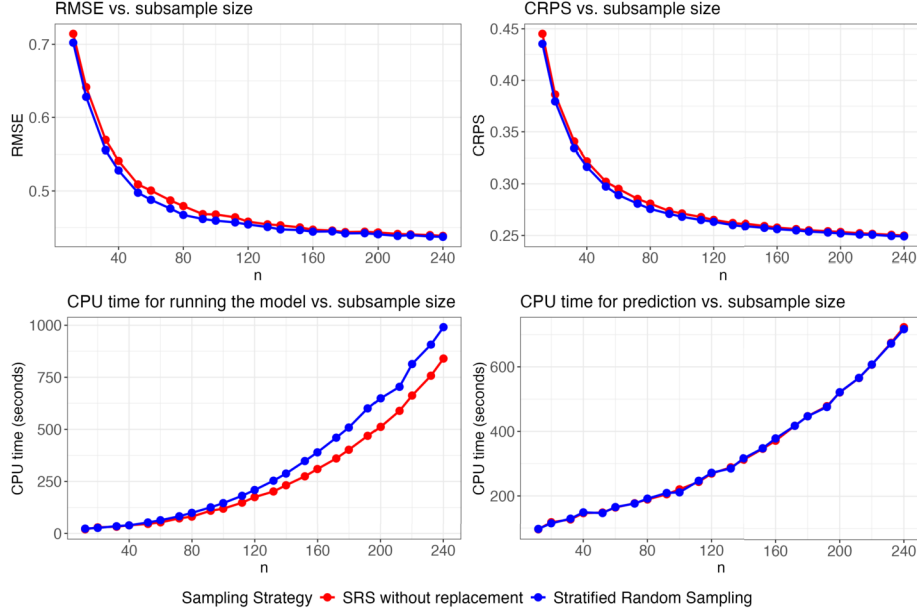


Figure 4.3: We plot the RMSE (top left panel), CRPS (top right panel), CPU time for running the model (bottom left panel), and CPU time for prediction (bottom right panel) at 3,000 missing locations over different subsample size, for both SRS without replacement (red) and stratified random sampling (blue).

$i \in A$, using the following equation,

$$\text{RMSE} = \left\{ \sum_{i \in A} \left[W(\mathbf{s}_i) - \widehat{W}(\mathbf{s}_i) \right]^2 / 3000 \right\}^{1/2}, \quad (4.1)$$

where $W(\mathbf{s}_i)$ represents the true process and $\widehat{W}(\mathbf{s}_i)$ represents the prediction at location \mathbf{s}_i . Since, we know the true process in case of simulated data, it is possible to use (4.1) to calculate the RMSE. To evaluate the computational performance of our model, we calculate the CPU time to run the model in (3.2) and perform predictions. We calculate the CPU time for predictions separately because the CPU time increases linearly with the number of missing locations for same subsample size. We compare RMSE and CPU times across the sampling strategies for δ .

In Figure 4.3, we plot the RMSE, Continuous Rank Probability Score (CRPS; see Gneiting et al., 2005, to incorporate uncertainty quantification in the metric), CPU time to run the model, and CPU time for predictions over different subsample size, when $\text{Pr}(\delta|n)$ is defined by (a) SRS without replacement, and (b) stratified random samples. We choose n to be $\{12, 20, 32, \dots, 220, 232, 240\}$ which are multiples of 4, since there are four strata in this simulation setting, and we draw equal number of samples from each stratum (in case of stratified random sampling). We keep n same for SRS without



Figure 4.4: We plot the RMSE (left panel) and pairwise difference in predictions (right panel) over different subsample size under SRS without replacement (red) and stratified random sampling (blue). The elbow of pairwise difference in predictions is much sharper in both sampling designs.

replacement to make a fair performance comparison between two sampling designs. For both the sampling strategies, we see that the RMSE and CRPS decrease, and the CPU times (for both running the model and prediction) increase as the subsample size increases, and the RMSE (and CRPS) is at its smallest value when $n = 240$. Hence, it can be inferred that if we increase n even further, the RMSE and CRPS are expected to keep on decreasing, but at a very small rate. But at the same time, the CPU times increase with n . It takes roughly 13 minutes for running the model and roughly 11.5 minutes for prediction when $n = 240$ for SRS without replacement. In case of stratified random sampling, it takes roughly 14 minutes for running the model and roughly 12 minutes for prediction when $n = 240$. So, we see a trade-off between the model accuracy and the CPU time as n changes. However, both RMSE and CRPS show an “elbow” pattern (in the top row of Figure 4.3), i.e. the decreasing rate is very less as subsample size increases after the elbow. Therefore, to get the best prediction, one should choose n to be a value after the elbow seen in the RMSE (or in the CRPS). However, it is suggested by Bradley (2021) to simultaneously look at the pairwise difference between $\hat{\mathbf{w}}_n$ and $\hat{\mathbf{w}}_{n+1}$ ($\hat{\mathbf{w}}_n$ is the vector of predictions for subsample size n) for $n = n_1, \dots, n_B$. In Figure 4.4, we plot the pairwise difference between $\hat{\mathbf{w}}_n$ and $\hat{\mathbf{w}}_{n+1}$, or $(\hat{\mathbf{w}}_n - \hat{\mathbf{w}}_{n+1})'(\hat{\mathbf{w}}_n - \hat{\mathbf{w}}_{n+1})$ (in the right panel) along with the RMSE (in the left panel), under both sampling strategies. Notice that, the elbow of pairwise difference (for both sampling) is much sharper than it is for RMSE. Considering the plots, when $\Pr(\delta|n)$ is defined by SRS without replacement (or stratified random samples) one choice of n can be around 100, for which it takes roughly 2 minutes (or 2.5 minutes) to run the model, and roughly 3.6 minutes (or 3.5 minutes) for prediction.




Figure 4.5: We plot the predictions $\{\widehat{W}(\mathbf{s}_i) : i \in A\}$ vs. the true process $\{W(\mathbf{s}_i) : i \in A\}$ for different choices of n , when $\Pr(\delta|n)$ is defined by SRS without replacement. The straight line represents the 45° reference line. The corresponding R^2 value is shown on the title of each plot.

In Figures 4.5 and 4.6, we plot the predictions ($\{\widehat{W}(\mathbf{s}_i) : i \in A\}$) vs. the true process ($\{W(\mathbf{s}_i) : i \in A\}$) for different choices of n , when $\Pr(\delta|n)$ is defined by SRS without replacement and by stratified random sampling respectively. In both figures, we see that the scatter plots are mostly aligned with the 45° reference line even for the smaller values of n . This suggests that the model is able to predict accurately even with small subsample size. We also see an increase in R^2 value as n increases, which does indicate that the predictions become more precise as we increase n . Thus in terms of prediction, we see similar properties for both sampling designs in this simulation study.

4.3 Simulation Study

We now assess the performance of our proposed model over multiple simulated replicates. Since the spatial data subset model is a semi-parametric model, it is important to compare it with the full model using simulation study. By “full model” we mean $n = N$, which is the weakly stationary spatial model. This comparison is particularly interesting, since the full model is misspecified and our spatial data subset model approach has marginal distributions at any given subset (for the holdout) that are correctly specified. Thus, we simulate a smaller dimensional dataset of size 900 on a 30×30 grid over $[0, 1] \times [0, 1]$ the same way as we described in Section 4.1 so that we can implement the full model. Here, we enforce 20% (or, 180) missing values at random locations over the grid. Hence, we generate random subsamples of sizes 10, 20, \dots , 100, 200, \dots , 700 from the remaining 720 data points and predict the response at the 180 missing locations

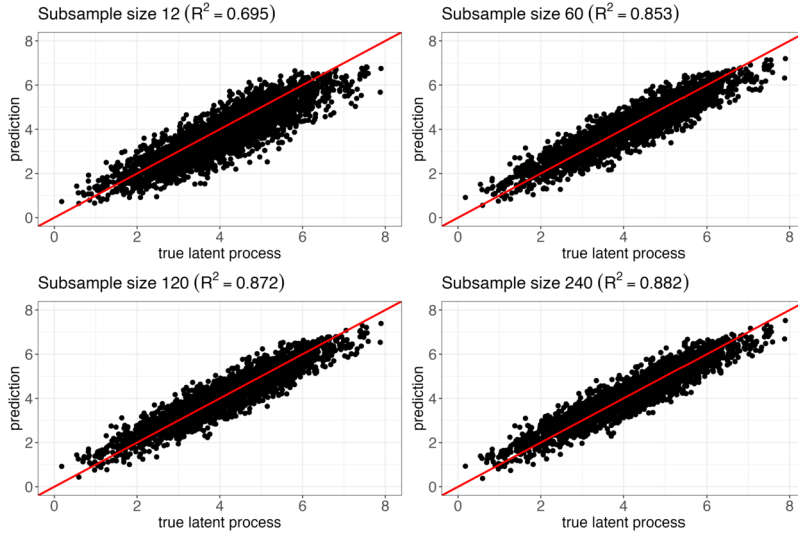


Figure 4.6: We plot the predictions $\{\widehat{W}(s_i) : i \in A\}$ vs. the true process $\{W(s_i) : i \in A\}$ for different choices of n , when $\Pr(\delta|n)$ is defined by stratified random samples. The straight line represents the 45° reference line. The corresponding R^2 value is shown on the title of each plot.

using Algorithm 1 with $G = 5,000$ and $g_0 = 1,000$ under (a) SRS without replacement and (b) stratified random sampling. We also run the full model using all 720 data points and we predict the response at the missing locations using the full model. We compare both models with respect to the RMSE between the true process and the predictions at the missing locations. We generate 10 replicates of each of the mentioned results.

In Figure 4.7, we plot the Boxplots of the 10 replicates of RMSE that we get from our proposed spatial data subset model using Algorithm 1. The results under SRS without replacement and stratified random sampling are shown on the left and right panel respectively. Also, the red horizontal line on both panels represents the average RMSE of the 10 replicates that we get from the full model. As expected, the range of the Boxplots decreases with increase in n . An interesting observation from Figure 4.7 is that the spatial data subset model for $n = 500, 600$ and 700 performs even better than the full model which was trained with all 720 data points. This may be because the data is simulated from a non-stationary model whereas the full model is a stationary model, but our proposed model becomes non-stationary because of its semi-parametric nature.

4.4 Scalability of the Spatial Data Subset Model

In principle, our proposed approach can be applied to a spatial dataset of *any size* that can be stored. A common target for complexity time for spatial models is $\mathcal{O}(N)$, where

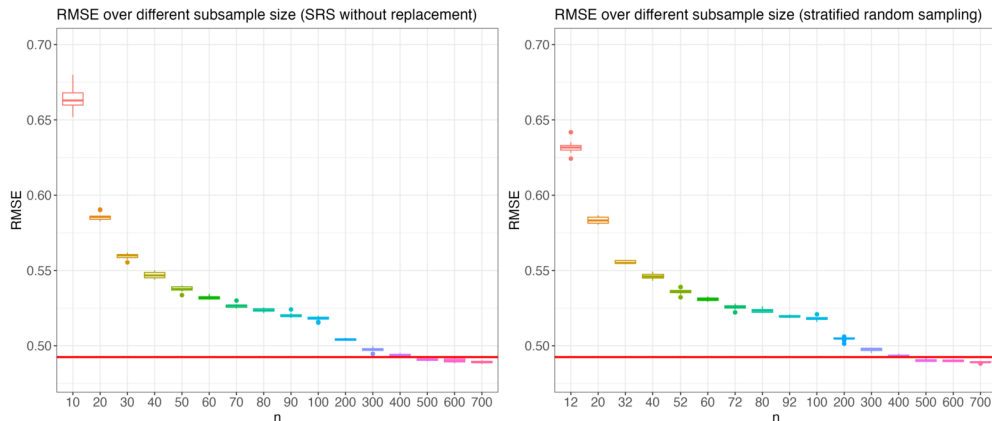


Figure 4.7: We plot the Boxplots of 10 replicates of RMSE that we get from the spatial data subset model for each subsample size, and the average RMSE (the red horizontal line) of 10 replicates that we get from the full model. The left and right panels show the result under SRS without replacement and under stratified random sampling respectively.

N is the size of the dataset (see Cooley and Tukey, 1965; Mallat, 1989; Cressie and Johannesson, 2008 for standard references). In general, this goal of $\mathcal{O}(N)$ still can lead to computational failures, since in the era of big data the size of N is continuously increasing. But, our proposed spatial data subset model (which is also a fully Bayesian model) uses $n \ll N$ data points at a time (and all N for large enough G), which makes it a $\mathcal{O}(n^3G)$ spatial model. Thus, no matter what the value of N is, we can always choose an n scalable in the presence of any data size N . To investigate this, we simulate large datasets with different sizes and we fit the spatial methods for large dataset from a benchmark study in Heaton et al. (2019). Our goal is to find the N such that the methods from Heaton et al. (2019) crash the first author’s laptop or stop with an error message regarding size allocation. We refer to this N as the “crash point”. For this study, we use the author’s laptop computer with the same specifications mentioned in Section 4.2 (11 Gen Intel(R) Core(TM) i7-1165G7 CPU with 2.80GHz and 16 GB RAM).

We simulate large datasets of size N using an efficient spectral simulation method (see Mejía and Rodríguez-Iturbe, 1974; Cressie, 1993, pg. 204; Yang and Bradley, 2021 for more details) with the exponential covariogram with range equals to 3 and spatial domain as the unit square. An independent error component is included with an SNR equals to 3. We also enforce 1,000 missing values to perform predictions. We keep this number small to focus on the computation time on training the model. We fit the spatial methods mentioned in Heaton et al. (2019) one by one to datasets with different sizes until we reach the crash point. In Table 4.1, we list the crash points of each method. We find that the methods crash at different data sizes, and LAGP scales to the largest N at 20M. However, the proposed spatial data subset model is still able to process

Spatial Method	Crash Point
FRK	550 K
Gapfill	8 M
Lattice Kriging	110 K
LAGP	20 M
Metakriging	100 K
MRA	1 M
NNGP	5 M
Partition	600 K
Predictive Process	3 M
SPDE	150 K
Tapering	100 K

Table 4.1: Crash points (value of N at which a method crash) for different spatial methods.




Figure 4.8: We plot the RMSE of predictions at 1,000 missing locations on the left panel, and the total CPU time (in minutes) on the right panel for different subsample sizes on a simulated dataset of size 20M.

that much data for different values of n . We provide the R code that generates the 20M dimensional spatial dataset and fits the SDSM method to it on GitHub at <https://github.com/sudiptosaha/Spatial-Data-Subset-Model>. The code to implement the competing methods to the 20M observations can be found at <https://github.com/finnlindgren/heatoncomparison> (Heaton et al., 2019). In Figure 4.8, we plot the performance of the spatial data subset model on the same 20M data. On the left panel, we plot the RMSE of the predictions at 1,000 missing locations for different subsample sizes. On the right panel, we plot the total CPU time (CPU time to run the model and to make predictions) in minutes for different subsample sizes. We see that, not only our proposed model is able to fit a data of size 20M, it also takes less CPU time (at lower values of n) to run and make predictions, and shows compelling performance, whereas all the other methods crash for the same dataset.

We emphasize that we do not throw away any of 20M observations, and simply only consider $n \ll N$ at each step of our sampler. The crash point can be interpreted as the

value of N such that the corresponding model can be fully scalable, where recall, we say a method is fully scalable if it makes use of all N observations without discarding them. For $N = 20\text{M}$, we have checked that 99.998% observations are used after $G = 106\text{K}$ when implementing in Algorithm 1. What we have observed is that our sampler for posterior replicates of $(\boldsymbol{\theta}, \boldsymbol{\nu})$ reaches its stationary distribution well before $G = 106\text{K}$. That is, it appears that one does not need to be fully scalable in order to appropriately search through the parameter space. In practice, we suggest choosing G based on standard MCMC diagnostic criteria (e.g., trace plots and Gelman-Rubin diagnostics) instead of choosing G to obtain full scalability.

We reiterate that the crash points depend on the system that is used. Thus, if we run the methods on a different system, the crash points will change. But, our proposed method does not crash for any size of dataset for an $n \ll N$ provided N observations can be stored. We have also considered $N = 30\text{M}$ and found that our method can still produce posterior inferences in this case (i.e., the algorithm does not crash). To our knowledge, there is no spatial model other than our proposed spatial data subset model that is fully scalable (for large enough G) to a spatial dataset of any size that can be stored.

5 Application

In this section, we present an analysis of a real dataset that consists of land surface temperatures collected by MODIS, a remote sensing instrument onboard the Terra satellite. We discuss the dataset in detail in Section 5.1, and we provide a detailed analysis of the dataset in Section 5.2.

5.1 The Dataset

On December 18, 1999 the National Aeronautics and Space Administration (NASA) launched the Terra satellite, which is the Earth Observing System (EOS) flagship satellite orbiting at an altitude of 705 kilometers (<https://terra.nasa.gov/>; see Ranson, 2003). The Moderate Resolution Imaging Spectroradiometer (MODIS) is an important remote sensing instrument onboard the Terra satellite. It primarily captures Earth’s atmospheric conditions which helps scientists in monitoring the changes in the biosphere. MODIS measures several different variables including land and sea surface temperatures. In particular, MODIS retrieves the Land Surface Temperature (LST) data daily at 1 km pixels by using the generalized split-window algorithm (Wan and Dozier, 1996), and at 6 km grids by using the day/night algorithm (Wan and Li, 1997). The pair of daytime and nighttime LSTs are recorded by the day/night algorithm in seven Thermal Infrared (TIR) bands. We observe the Level-3 (Wan et al., 2006; Hulley et al., 2016) daytime LSTs as measured by MODIS on August 4, 2016. Heaton et al. (2019) considered this dataset as large spatial data and provided a comparative study among several spatial methods applied on this dataset. The data were originally downloaded from the MODIS reprojection tool web interface (MRTweb), and later made available at a public repository at <https://github.com/finnlindgren/heatoncomparison> (Heaton et al., 2019).

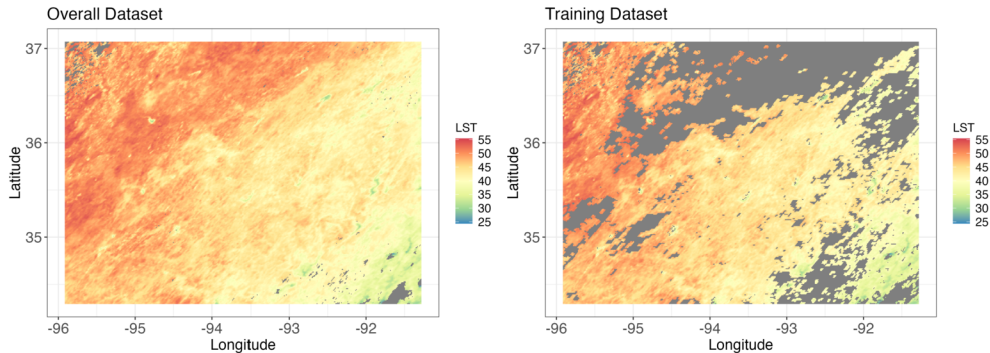


Figure 5.1: We plot the overall LST data on the left panel, and the training LST data on the right panel. The grey region on the left panel represents the missing data. The grey region on the right panel represents the missing data and the test data together.

This benchmark dataset is particularly useful because it allows us to compare our modification (i.e., using a data subset model) of a standard spatial model (i.e., based on covariograms) to a large list of current spatial methods including FRK (Cressie and Johannesson, 2008), Gapfill (Gerber et al., 2018), Lattice Kriging (Nychka et al., 2015), Local Approximate Gaussian Process (LAGP; Gramacy and Apley, 2015), Metakriging (Guhaniyogi and Banerjee, 2018), Multiresolution Approximations (MRA; Jurek and Katzfuss, 2021), Nearest Neighbor Process (NNGP; Datta et al., 2016), Spatial Partitioning (Sang et al., 2011), Predictive Process (Banerjee et al., 2008), SPDE (Lindgren et al., 2011), Covariance Tapering (Furrer et al., 2006), Periodic Embedding (Guinness, 2019), and Multi-scale Vecchia (MSV) approximation (Zhang and Katzfuss, 2022). We also compare our proposed model’s performance with SoD (Chalupka et al., 2013) for subsample sizes 2,000 and 5,000. The author’s laptop computer crashed when we ran SoD with $n = 10,000$, which is not surprising because of kriging’s natural restriction on n (Rulli ere et al., 2018). Our goal is to show that we can take a familiar standard spatial model for low-dimensional stationary spatial data, and modify it with a data subset model to make it produce similar predictions as the state-of-the-art methods. This is particularly exciting, as the spatial data subset model can be scaled to effectively any dataset that can be stored, and hence, similar predictive performance on a well-known benchmark dataset would suggest that our approach would be reasonable to consider it in extremely high-dimensional settings.

The dataset consists of LST observations on a 500×300 grid (150,000 observations in total) ranging longitude values from -95.91153 to -91.28381 , and latitude values from 34.29519 to 37.06811 . The ranges of longitude and latitude, and also the date on which the LST is measured are chosen because of the sparse cloud cover over the region on that day. We have 1.1% missing data due to cloud cover, leaving 148,309 out of 150,000 observed values to be used. We use the same train and test dataset split as it is given in Heaton et al. (2019). We keep 105,569 observations in the training set, and 42,740 observations in the out-of-sample test set. In Figure 5.1, we plot the entire dataset on the left panel where the 1.1% missing observations are represented by the grey region.

n	MAE	RMSE	CRPS	INT	CVG	CPU Time to run the model (min.)	CPU Time for prediction (min.)
96	2.35	2.71	1.63	11.37	0.80	0.42	31.18
208	2.25	2.59	1.62	14.73	0.72	1.85	115.75
512	2.13	2.47	1.60	19.48	0.59	16.16	457.56
1024	2.00	2.38	1.38	9.75	0.88	98.14	2304.19
2000	1.93	2.32	1.35	10.41	0.83	468.26	4183.77

Table 5.1: Model performance and CPU time on LST dataset over different subsample size.

We also plot the training dataset on the right panel where the grey region represents the missing data and the test data.

5.2 Analysis of Daytime LST Data Captured By MODIS

We implement Algorithm 1 (since stationarity reaches irrespective of using a Gibbs-within-composite sampler; see Appendix C of the supplementary material (Saha and Bradley, 2024) for a detailed description) on the training dataset, where we set $G = 2,000$, $g_0 = 800$, discrete uniform distribution with support $\{0.001, 0.0002, \dots, 0.1\}$ for the prior of ϕ and $\Pr(\delta|n)$ is defined by stratified random samples with 16 equal sized strata. Notice that we have 42,740 locations in the test set (and an additional 1,691 locations for the missing data) and the CPU time increases with the large number of prediction locations. The algorithm takes comparatively longer time for prediction, which also increases as the subsample size n increases. We run Algorithm 1 for different values of n . For each n , we generate equal number of random samples ($n/16$) from each of the 16 strata, and we predict the LST values on the grey region shown on the right panel of Figure 5.1. We also calculate mean absolute error (MAE), root mean squared error (RMSE), continuous rank probability score (CRPS; see Gneiting et al., 2005), interval score (INT; see Gneiting and Raftery, 2007) and prediction interval coverage (CVG) at $\alpha = 0.05$ on the out-of-sample test dataset. We keep track of the CPU times to run/fit the model and to perform prediction separately for each n . In Table 5.1, we show these metrics for different values of n . The values of n are chosen in a way that they are multiples of 16 (number of strata) so that we can draw equal number of samples from each stratum. We see gradual decrease in MAE, RMSE and CRPS as n increases, but simultaneously the CPU times for both running the model and prediction increase. It takes very less time to run/fit the model than to predict because of the large test set. We also check trace plots and credible intervals informally. We do not find any sign of lack of convergence, and the credible intervals decrease as n increases.

In Figure 5.2, we plot the actual LST dataset on the left panel. Note that, the 1.1% missing data are represented by the grey region in the plot. On the right panel, we filled in the grey region of the “training dataset” plot of Figure 5.1 (represents the test data and missing data together) with the predicted LSTs, from the spatial data subset

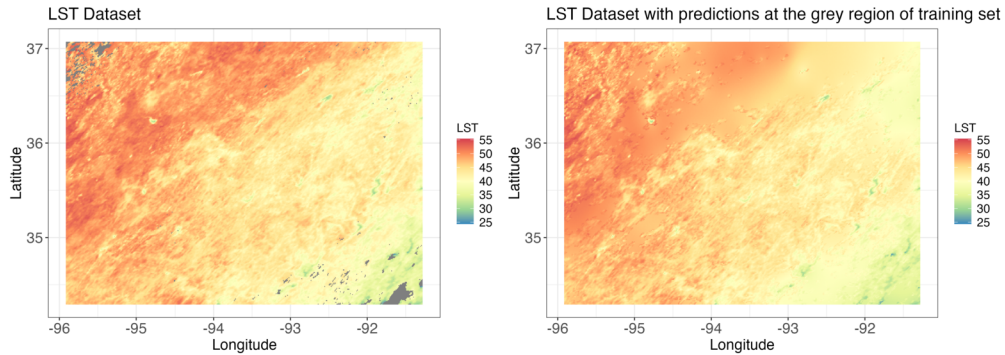


Figure 5.2: We plot the actual LST dataset on the left panel. We filled in the grey region of training set with the predictions, from the spatial data subset model, on those locations with $n = 2,000$ and plot that on the right panel.

model, on those locations using $n = 2,000$. For a given computational goal, one can choose a different value of n and that will produce a different set of predictions. We see that the model correctly identifies the spatial property of the data, i.e. the nearby locations have similar value than the locations at far.

In Table 5.2, we show a comparison of the performance of our proposed spatial data subset model (SDSM) when $n = 2,000$ with all the methods listed in Section 5.1 (methods described in Heaton et al., 2019; Zhang and Katzfuss, 2022 and Chalupka et al., 2013). While our method does not consistently dominate all other methods in all criteria, it is competitive. We find that in terms of RMSE, our proposed model does perform better than FRK, Metakriging, Predictive Process, Tapering and SoD with just $n = 2,000$. In terms of both MAE and CRPS, our proposed model outperforms FRK, Metakriging, Predictive Process and SoD. It is also noticeable that in terms of INT, our proposed model outperforms FRK, Gapfill, LAGP, Metakriging, Partition, Predictive Process and SoD. However, in terms of CVG, it outperforms only FRK and Gapfill, and it matches the performance of LAGP and Predictive Process. These results are particularly exciting for spatial statisticians. Our model allows one to implement a traditional spatial model (on subsamples) so that it can be implemented on high-dimensional data, respecting the spatial properties of the entire dataset, and is competitive with several state-of-the-art methods used in high-dimensional setting.

The total run time for our proposed model with $n = 2,000$ is 4,652 minutes. But, the run time depends on number of cores, RAM size, CPU speed and number of jobs running in the system. The author's laptop computer has different specification (mentioned in Section 4.2) than what it was used for the methods in Heaton et al. (2019) and Zhang and Katzfuss (2022), which is why the run time is not comparable. Moreover, in Section 4.4, we show that our proposed method is fully scalable (for large enough G) to spatial datasets of any size that can be stored, and the CPU time stays fixed (for an n) as N increases. There is no other method that is scalable at this extent. Depending on the system used, other methods will crash for large enough N , but the spatial data subset

Method	Reference	MAE	RMSE	CRPS	INT	CVG	Run time (min)	Cores used
SDSM	–	1.93	2.32	1.35	10.41	0.83	4652.03	1
FRK	Cressie and Johannesson (2008)	1.96	2.44	1.44	14.08	0.79	2.32	1
Gapfill	Gerber et al. (2018)	1.33	1.86	1.17	34.78	0.36	1.39	40
Lat-ticeKrig	Nychka et al. (2015)	1.22	1.68	0.87	7.55	0.96	27.92	1
LAGP	Gramacy and Apley (2015)	1.65	2.08	1.17	10.81	0.83	2.27	40
Metakriging	Guhaniyogi and Banerjee (2018)	2.08	2.50	1.44	10.77	0.89	2888.52	30
MRA	Jurek and Katzfuss (2021)	1.33	1.85	0.94	8.00	0.92	15.61	1
NNGP	Datta et al. (2016)	1.21	1.64	0.85	7.57	0.95	2.06	10
Partition	Sang et al. (2011)	1.41	1.80	1.02	10.49	0.86	79.98	55
Pred. Proc.	Banerjee et al. (2008)	2.15	2.64	1.55	15.51	0.83	160.24	10
SPDE	Lindgren et al. (2011)	1.10	1.53	0.83	8.85	0.97	120.33	2
Tapering	Furrer et al. (2006)	1.87	2.45	1.32	10.31	0.93	133.26	1
Period. Emb.	Guinness (2019)	1.29	1.79	0.91	7.44	0.93	9.81	1
MSV	Zhang and Katzfuss (2022)	1.11	1.42	–	7.32	0.88	–	–
SoD with $n = 2000$	Chalupka et al. (2013)	50.48	75.37	32.64	145.8	0.97	796.95	1
SoD with $n = 5000$	Chalupka et al. (2013)	35.36	57.52	22.90	105.86	0.88	4822.23	1

Table 5.2: Comparison of spatial data subset model (when $n = 2,000$) with the other methods.

model will not. Although our proposed method does not outperform all the methods in terms of performance, it is competitive and it has a big advantage in terms of scalability.

Each method in Table 5.2 used different computational resources. Thus, one should be careful not to use the reported CPU times to rank value the methods computationally, since each method could be made faster by capitalizing on improved computational resources. For example, in Table 5.2 we provide the number of cores used, and not all methods were implemented on more than one core. Moreover, SDSM uses one core and was implemented on a different computer than the other methods. In Table 5.1, metrics are provided for implementations of SDSM in our application as we vary the choice of n . In Table 5.1, we see that SDSM is considerably faster when $n = 512$ (CPU time of roughly 32 minutes) than when $n = 2,000$ (CPU time of 4652 minutes). Previous applications of the data subset model approach in Bradley (2021) suggest that the rate of increase of SDSM’s CPU time, as a function of n , would be smaller if more computational resources (e.g., cores) were available. Thus, the use of additional computational resources (e.g., cores) would allow one to choose larger values of n while taking on a smaller computational cost.

The results from Section 4.4 are particularly important for evaluating SDSM relative to the state-of-the-art methods. In particular, when N is so large that it is beyond the largest crash-point of the state-of-the-art methods, SDSM is the *only* approach that is possible to compute that does not throw away data. This is because SDSM is fully scalable in the limit. Moreover, the method that throws away data, i.e. SoD, performed considerably worse than SDSM in terms of inferential metrics in Table 5.2.

There are times when one would just prefer to implement a standard spatial statistical model. This is because the properties of standard models (e.g., weak stationarity, sill, nugget, and range) are well-known and easy to communicate to the general public and collaborators in different fields. Bayesian Kriging is the traditional spatial model that we focus on in this article. However, in the big N case, it is unclear if the standard model assumptions are appropriate for such high-dimensional spatial data with spatial domains with complicated terrains, and would ideally like to use all N observations. Propositions 3.1–3.4, show that when the data is generated from a model with non-stationary covariances the SDSM is non-stationary, and when the data is generated from a model with stationary covariances the SDSM assumes a stationary covariance. Thus, the assumptions on the spatial properties of this SDSM model are well aligned with the data, we obtain inferences on parameters that arise in traditional spatial models, and use all N observations in a fully Bayesian context. However, it is unclear how well this SDSM performs in practice. In this section, we show that the SDSM performs similar (i.e., does not produce the worst value of a metric, but not the best) to several of the state-of-the-art methods. Thus, one is able to obtain reasonable inferential results (i.e., the metrics considered are not always the best, but are not always the worst) when assume a traditional model for subsamples of the data (scaling it to large N), without necessarily throwing out data, and while still respecting the true spatial properties of the data.

6 Discussion

In this article, we propose a new spatial data subset model (SDSM) that allows one to implement the classical Bayesian spatial model on big data in a computationally feasible way. We do this by redefining the full likelihood of the data model using subsamples without imposing any additional restrictive assumptions. Also, we insert more levels to the Bayesian hierarchical model to introduce subsampling inside the model. We see that the subsampling changes the spatial properties of the dataset that is being used in the model. These properties are also dependent on the subsampling strategy. Specifically, we provide theoretical results which show that if the true spatial model is stationary (or non-stationary) then the spatial data subset model in (3.2) is also stationary (or non-stationary) under both SRS without replacement and stratified random sampling. We also provide moment results under SRS without replacement and stratified random sampling, and compare and contrast the models. Code to implement the SDSM is made available at <https://github.com/sudiptosaha/Spatial-Data-Subset-Model>.

In our simulation study, we have found that the prediction errors decrease as n (subsample size) increases, but the CPU time increases simultaneously. This suggests

that one should choose a value of n in a way that it achieves the computational goal (for example, producing inference in 2 minutes) while getting reasonable predictions. This performance is particularly exciting since the CPU time does not scale with N (total number of observations), but rather n , which we can choose. The decrease in RMSE have shown an elbow pattern which is why we have plotted the pairwise difference in prediction as the RMSE decreases gradually. This helps us to choose n at which the prediction error starts decreasing at a slower rate. Furthermore, we have analyzed the scatter plots between the true process and the predictions, and saw that the predictions get closer to the 45° reference line as n increases. We have also done a simulation study where we compared the performance of our model with the full model using a smaller dataset. It turns out that our proposed model performs even better than the full model for large enough n , which may be due to the semi-parametric nature of our model. The biggest advantage of our proposed model is scalability. Using simulation, we were able to show that our model is scalable to such large value of N where other methods fail. We also emphasize that our proposed model is fully scalable (for large enough G) to a dataset of any size that can be stored and there is no model out there which is fully scalable to any size of the data.

In practice, a model is sometimes chosen simply because it is a traditional choice, and in spatial statistics a Gaussian process with a stationary covariogram is considered a textbook model (Banerjee et al., 2014; Cressie and Wikle, 2015). However, this traditional spatial model does not scale naturally to high-dimensions and does not allow for non-stationary covariance functions. Our proposed approach allows spatial statisticians to apply traditional Bayesian spatial statistical models in a way that (i) is fully scalable in the limit and (ii) removes assumptions from the traditional spatial model to allow for the possibility of non-stationary covariance functions. This is achieved through the data subset model framework, which allows one to assume a traditional spatial model for all data with $\delta(\mathbf{s}_i) = 1$ such that $\sum_i \delta(\mathbf{s}_i) = n \ll N$. This is exciting solution for those who wish to adopt the more traditional model assumptions because of its ability to easily explain the spatial properties. However, it is important to clarify that while Items (i) and (ii) are attractive features of our model, it does not mean the proposed model will dominate estimation, prediction, and computational metrics. In our application, we demonstrate that the proposed model performs similar to competing methods, but not the best (i.e., our proposed method did not perform the best nor did it perform worst for any of metrics). In particular, we illustrate the use of our model to a daytime Land Surface Temperature (LST) data captured by the MODIS instrument onboard the Terra satellite. We have found that our model outperforms several methods in terms of MAE, RMSE, CRPS, INT, and CVG with $n = 2,000$. We emphasize that although our proposed model is not the best, we are able to take a standard spatial model for low-dimensional stationary spatial data and modify it to be competitive for large data and to be able to model even non-stationary data.

Although our proposed model does produce competitive results, it depends on the choice of n . As n changes, the model performance also changes. Thus, the choice of n is important and an open problem in our proposed framework. In our empirical results, we see that posterior samples tend to have less variability as the subsample size n grows. This is consistent with the results from Bradley (2021). Consequently, the best choice

of n is the largest possible value that can be computed. In this article, we focus on applying the data subset model to kriging, which is well-known to be computationally feasible when n is on the order of thousands (e.g., see Rulli re et al., 2018, where $n = 2,000$ appears reasonable). When implementing our model in practice, one should simply choose an $n = 2,000$, or if more computational resources are available than our laptop, $n = 10,000$. In this article we make use of elbow plots simply to demonstrate the sensitivity of the choice of n , and we see that the “elbow” is often quite small (in the simulations elbow appears at $n = 300$). Of course, in general, when one does not know what n is practical a priori through historical usage, we suggest what is commonly used in the SoD literature (Chalupka et al., 2013). That is, an “elbow plot” can be used, where one selects an n from a sequence that “passes the elbow”.

Our proposed spatial data subset model is built on the classical Bayesian spatial model. However, one can take the similar approach of applying the subsampling strategy on some other spatial model. One important thing to note is that the data subset model approach works on the Bayesian framework. So, one has to choose a Bayesian framework of a spatial model to build a model using this strategy. Therefore, building different spatial models on the data subset model framework can be a potential topic for future research.

Supplementary Material

Supplementary Materials to “Incorporating Subsampling into Bayesian Models for High-Dimensional Spatial Data” (DOI: [10.1214/24-BA1426SUPP](https://doi.org/10.1214/24-BA1426SUPP); .pdf). Proof of the propositions, related results, and the full-conditional distributions are presented. Also, an empirical investigation of the posterior consistency is demonstrated both for the simulated data and LST data.

References

- Banerjee, S., Carlin, B., and Gelfand, A. (2014). *Hierarchical Modeling and Analysis for Spatial Data, Second Edition*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Taylor & Francis. URL <https://books.google.com/books?id=zNLhAAQBAJ> MR3362184. 1, 4, 8, 9, 15, 32
- Banerjee, S., Gelfand, A. E., Finley, A. O., and Sang, H. (2008). “Gaussian predictive process models for large spatial data sets.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(4): 825–848. MR2523906. doi: <https://doi.org/10.1111/j.1467-9868.2008.00663.x>. 27, 30
- Berger, J. O., De Oliveira, V., and Sans , B. (2001). “Objective Bayesian analysis of spatially correlated data.” *Journal of the American Statistical Association*, 96(456): 1361–1374. MR1946582. doi: <https://doi.org/10.1198/016214501753382282>. 15
- Besag, J. (1974). “Spatial interaction and the statistical analysis of lattice systems.” *Journal of the Royal Statistical Society: Series B (Methodological)*, 36(2): 192–225. MR0373208. 3

- Besag, J., York, J., and Mollié, A. (1991). “Bayesian image restoration, with two applications in spatial statistics.” *Annals of the institute of statistical mathematics*, 43: 1–20. MR1105822. doi: <https://doi.org/10.1007/BF00116466>. 3
- Bradley, J. R. (2021). “An Approach to Incorporate Subsampling Into a Generic Bayesian Hierarchical Model.” *Journal of Computational and Graphical Statistics*, 30(4): 889–905. MR4356593. doi: <https://doi.org/10.1080/10618600.2021.1923518>. 2, 3, 4, 5, 6, 7, 8, 9, 16, 19, 21, 30, 32
- Bradley, J. R., Cressie, N., and Shi, T. (2015). “Comparing and selecting spatial predictors using local criteria.” *Test*, 24(1): 1–28. MR3314567. doi: <https://doi.org/10.1007/s11749-014-0415-1>. 2
- Byers, J. S. and Gill, J. (2022). “Applied Geospatial Bayesian Modeling in the Big Data Era: Challenges and Solutions.” *Mathematics*, 10(21). URL <https://www.mdpi.com/2227-7390/10/21/4116> 2
- Casella, G. and Berger, R. L. (2021). *Statistical inference*. Cengage Learning. MR1051420. 4
- Chalupka, K., Williams, C. K., and Murray, I. (2013). “A framework for evaluating approximation methods for Gaussian process regression.” *Journal of Machine Learning Research*, 14: 333–350. MR3033334. 3, 4, 5, 7, 27, 29, 30, 33
- Cooley, J. W. and Tukey, J. W. (1965). “An algorithm for the machine calculation of complex Fourier series.” *Mathematics of computation*, 19(90): 297–301. MR0178586. doi: <https://doi.org/10.2307/2003354>. 24
- Cressie, N. (1993). *Statistics for spatial data*. John Wiley & Sons. MR1239641. doi: <https://doi.org/10.1002/9781119115151>. 1, 4, 9, 24
- Cressie, N. and Johannesson, G. (2008). “Fixed rank kriging for very large spatial data sets.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(1): 209–226. MR2412639. doi: <https://doi.org/10.1111/j.1467-9868.2007.00633.x>. 2, 3, 24, 27, 30
- Cressie, N. and Wikle, C. K. (2015). *Statistics for spatio-temporal data*. John Wiley & Sons. MR2848400. 32
- Daniels, M. J. (1999). “A prior for the variance in hierarchical models.” *Canadian Journal of Statistics*, 27(3): 567–578. MR1745822. doi: <https://doi.org/10.2307/3316112>. 15
- Datta, A., Banerjee, S., Finley, A. O., and Gelfand, A. E. (2016). “Hierarchical nearest-neighbor Gaussian process models for large geostatistical datasets.” *Journal of the American Statistical Association*, 111(514): 800–812. MR3538706. doi: <https://doi.org/10.1080/01621459.2015.1044091>. 27, 30
- Dunn, W. L. and Shultis, J. K. (2022). *Exploring monte carlo methods*. Elsevier. MR4590414. 15
- Ekström, M. and Luna, S. S.-D. (2004). “Subsampling methods to estimate the variance of sample means based on nonstationary spatial data with varying expected

- values.” *Journal of the American Statistical Association*, 99(465): 82–95. MR2054288. doi: <https://doi.org/10.1198/016214504000000106>. 2
- Furrer, R., Genton, M. G., and Nychka, D. (2006). “Covariance tapering for interpolation of large spatial datasets.” *Journal of Computational and Graphical Statistics*, 15(3): 502–523. MR2291261. doi: <https://doi.org/10.1198/106186006X132178>. 27, 30
- Gelman, A. (2006). “Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper).” *Bayesian analysis*, 1(3): 515–534. MR2221284. doi: <https://doi.org/10.1214/06-BA117A>. 15
- Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B. (1995). *Bayesian data analysis*. Chapman and Hall/CRC. MR1385925. 15
- Gerber, F., de Jong, R., Schaepman, M. E., Schaepman-Strub, G., and Furrer, R. (2018). “Predicting missing values in spatio-temporal remote sensing data.” *IEEE Transactions on Geoscience and Remote Sensing*, 56(5): 2841–2853. 27, 30
- Gneiting, T. and Raftery, A. E. (2007). “Strictly proper scoring rules, prediction, and estimation.” *Journal of the American statistical Association*, 102(477): 359–378. MR2345548. doi: <https://doi.org/10.1198/016214506000001437>. 28
- Gneiting, T., Raftery, A. E., Westveld, A. H., and Goldman, T. (2005). “Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation.” *Monthly Weather Review*, 133(5): 1098–1118. MR2345548. doi: <https://doi.org/10.1198/016214506000001437>. 20, 28
- Gramacy, R. B. and Apley, D. W. (2015). “Local Gaussian process approximation for large computer experiments.” *Journal of Computational and Graphical Statistics*, 24(2): 561–578. MR3357395. doi: <https://doi.org/10.1080/10618600.2014.914442>. 3, 27, 30
- Guan, Y., Sherman, M., and Calvin, J. A. (2004). “A nonparametric test for spatial isotropy using subsampling.” *Journal of the American Statistical Association*, 99(467): 810–821. MR2090914. doi: <https://doi.org/10.1198/016214504000001150>. 2
- Guhaniyogi, R. and Banerjee, S. (2018). “Meta-kriging: Scalable Bayesian modeling and inference for massive spatial datasets.” *Technometrics*, 60(4): 430–444. MR3878099. doi: <https://doi.org/10.1080/00401706.2018.1437474>. 27, 30
- Guinness, J. (2019). “Spectral density estimation for random fields via periodic embeddings.” *Biometrika*, 106(2): 267–286. MR3949303. doi: <https://doi.org/10.1093/biomet/asz004>. 27, 30
- Gunawan, D., Tran, M.-N., and Kohn, R. (2017). “Fast inference for intractable likelihood problems using variational Bayes.” *arXiv preprint arXiv:1705.06679*. 2
- Handcock, M. S. and Stein, M. L. (1993). “A Bayesian analysis of kriging.” *Technometrics*, 35(4): 403–410. 1
- Hayashi, K., Imaizumi, M., and Yoshida, Y. (2020). “On random subsampling of Gaus-

- sian process regression: A graphon-based analysis.” In *International Conference on Artificial Intelligence and Statistics*, 2055–2065. PMLR. 3, 7
- Heaton, M. J., Christensen, W. F., and Terres, M. A. (2017). “Nonstationary Gaussian process models using spatial hierarchical clustering from finite differences.” *Technometrics*, 59(1): 93–101. MR3604192. doi: <https://doi.org/10.1080/00401706.2015.1102763>. 2
- 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. MR3996451. doi: <https://doi.org/10.1007/s13253-018-00348-w>. 3, 4, 24, 25, 26, 27, 29
- Hensman, J., Fusi, N., and Lawrence, N. D. (2013). “Gaussian processes for big data.” *arXiv preprint arXiv:1309.6835*. 3
- Hollander, M., Wolfe, D. A., and Chicken, E. (2013). *Nonparametric statistical methods*. John Wiley & Sons. MR3221959. 9, 14
- Hulley, G., Freepartner, R., Malakar, N., and Sarkar, S. (2016). “Moderate Resolution Imaging Spectroradiometer (MODIS) Land Surface Temperature and Emissivity Product (MxD21) User Guide.” *NASA: Washington, DC, USA*. 26
- Jurek, M. and Katzfuss, M. (2021). “Multi-resolution filters for massive spatio-temporal data.” *Journal of Computational and Graphical Statistics*, 30(4): 1095–1110. MR4356607. doi: <https://doi.org/10.1080/10618600.2021.1886938>. 27, 30
- Katzfuss, M. and Guinness, J. (2021). “A general framework for Vecchia approximations of Gaussian processes.” *Statistical Science*, 36(1): 124–141. MR4194207. doi: <https://doi.org/10.1214/19-STS755>. 2, 3
- Kazianka, H. and Pilz, J. (2012). “Objective Bayesian analysis of spatial data with uncertain nugget and range parameters.” *Canadian Journal of Statistics*, 40(2): 304–327. MR2927748. doi: <https://doi.org/10.1002/cjs.11132>. 15
- Keerthi, S. and Chu, W. (2005). “A matching pursuit approach to sparse gaussian process regression.” *Advances in neural information processing systems*, 18. 3, 7
- Kim, H.-M., Mallick, B. K., and Holmes, C. C. (2005). “Analyzing nonstationary spatial data using piecewise Gaussian processes.” *Journal of the American Statistical Association*, 100(470): 653–668. MR2160567. doi: <https://doi.org/10.1198/016214504000002014>. 2
- Knorr-Held, L. and Raßer, G. (2000). “Bayesian detection of clusters and discontinuities in disease maps.” *Biometrics*, 56(1): 13–21. 2
- Konomi, B. A., Sang, H., and Mallick, B. K. (2014). “Adaptive Bayesian nonstationary modeling for large spatial datasets using covariance approximations.” *Journal of Computational and Graphical Statistics*, 23(3): 802–829. MR3224657. doi: <https://doi.org/10.1080/10618600.2013.812872>. 2

- Lahiri, S. (1999). “Asymptotic distribution of the empirical spatial cumulative distribution function predictor and prediction bands based on a subsampling method.” *Probability Theory and Related Fields*, 114: 55–84. MR1697139. doi: <https://doi.org/10.1007/s004400050221>. 2
- Lawrence, N., Seeger, M., and Herbrich, R. (2002). “Fast sparse Gaussian process methods: The informative vector machine.” *Advances in neural information processing systems*, 15. 3, 7
- Lindgren, F., Rue, H., and Lindström, J. (2011). “An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(4): 423–498. MR2853727. doi: <https://doi.org/10.1111/j.1467-9868.2011.00777.x>. 27, 30
- Liu, H., Ong, Y.-S., Shen, X., and Cai, J. (2020). “When Gaussian process meets big data: A review of scalable GPs.” *IEEE transactions on neural networks and learning systems*, 31(11): 4405–4423. MR4169962. 3, 7
- Lohr, S. L. (2010). *Sampling: Design and Analysis, Second Edition*. Brooks/Cole, Boston. MR3057878. 2, 3
- Mallat, S. G. (1989). “A theory for multiresolution signal decomposition: the wavelet representation.” *IEEE transactions on pattern analysis and machine intelligence*, 11(7): 674–693. 24
- Mejía, J. M. and Rodríguez-Iturbe, I. (1974). “On the synthesis of random field sampling from the spectrum: An application to the generation of hydrologic spatial processes.” *Water Resources Research*, 10(4): 705–711. 24
- More, H., Halvorsen, K. B., and Berteig, V. (1989). “A Bayesian approach to kriging.” In *Geostatistics: Proceedings of the Third International Geostatistics Congress September 5–9, 1988, Avignon, France*, 109–126. Springer. 1
- Nassiuma, D. K. (2001). *Survey sampling: Theory and methods*. Nairobi University Press. 2
- Neelon, B., Gelfand, A. E., and Miranda, M. L. (2014). “A multivariate spatial mixture model for areal data: examining regional differences in standardized test scores.” *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 63(5): 737–761. MR3269410. doi: <https://doi.org/10.1111/rssc.12061>. 2
- Nordman, D. J. and Lahiri, S. N. (2004). “On optimal spatial subsample size for variance estimation.” *The Annals of Statistics*, 32(5): 1981 – 2027. URL <https://doi.org/10.1214/009053604000000779> MR2102500. doi: <https://doi.org/10.1214/009053604000000779>. 4
- Nychka, D., Bandyopadhyay, S., Hammerling, D., Lindgren, F., and Sain, S. (2015). “A multiresolution Gaussian process model for the analysis of large spatial datasets.” *Journal of Computational and Graphical Statistics*, 24(2): 579–599. MR3357396. doi: <https://doi.org/10.1080/10618600.2014.914946>. 2, 27, 30

- Nychka, D. W. (2000). “Spatial-process estimates as smoothers.” *Smoothing and regression: approaches, computation, and application*, 329: 393. 2
- Omre, H. and Halvorsen, K. B. (1989). “The Bayesian bridge between simple and universal kriging.” *Mathematical Geology*, 21: 767–786. MR1016496. doi: <https://doi.org/10.1007/BF00893321>. 1
- Paulo, R. (2005). “Default priors for Gaussian processes.” *The Annals of Statistics*, 33(2): 556 – 582. URL <https://doi.org/10.1214/009053604000001264> MR2163152. doi: <https://doi.org/10.1214/009053604000001264>. 15
- Pilz, J. and Spöck, G. (2008). “Why do we need and how should we implement Bayesian kriging methods.” *Stochastic Environmental Research and Risk Assessment*, 22(5): 621–632. MR2516843. doi: <https://doi.org/10.1007/s00477-007-0165-7>. 15
- Politis, D. N. and Romano, J. P. (1993). “On the sample variance of linear statistics derived from mixing sequences.” *Stochastic Processes and their Applications*, 45(1): 155–167. MR1204867. doi: [https://doi.org/10.1016/0304-4149\(93\)90066-D](https://doi.org/10.1016/0304-4149(93)90066-D). 3
- Possolo, A. (1991). “Subsampling a random field.” *Lecture Notes-Monograph Series*, 286–294. MR1195571. doi: <https://doi.org/10.1214/lrms/1215460508>. 2, 3
- Quiroz, M., Kohn, R., Villani, M., and Tran, M.-N. (2019). “Speeding Up MCMC by Efficient Data Subsampling.” *Journal of the American Statistical Association*, 114(526): 831–843. URL <https://doi.org/10.1080/01621459.2018.1448827> MR3963184. doi: <https://doi.org/10.1080/01621459.2018.1448827>. 2
- Ranson, K. J. (2003). “NASA’s EOS terra mission update.” In *IGARSS 2003. 2003 IEEE International Geoscience and Remote Sensing Symposium. Proceedings (IEEE Cat. No. 03CH37477)*, volume 5, 3023–3024. IEEE. 26
- Rue, H., Martino, S., and Chopin, N. (2009). “Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations.” *Journal of the royal statistical society: Series b (statistical methodology)*, 71(2): 319–392. MR2649602. doi: <https://doi.org/10.1111/j.1467-9868.2008.00700.x>. 2, 3
- Rue, H., Riebler, A., Sørbye, S. H., Illian, J. B., Simpson, D. P., and Lindgren, F. K. (2017). “Bayesian computing with INLA: a review.” *Annual Review of Statistics and Its Application*, 4: 395–421. MR3634300. doi: <https://doi.org/10.1214/16-STS576>. 2
- Rullière, D., Durrande, N., Bachoc, F., and Chevalier, C. (2018). “Nested Kriging predictions for datasets with a large number of observations.” *Statistics and Computing*, 28: 849–867. MR3766047. doi: <https://doi.org/10.1007/s11222-017-9766-2>. 4, 27, 33
- Saha, S. and Bradley, J. R. (2024). “Supplementary Materials to “Incorporating Subsampling into Bayesian Models for High-Dimensional Spatial Data”.” *Bayesian Analysis*. doi: <https://doi.org/10.1214/24-BA1426SUPP>. 5, 10, 11, 12, 13, 14, 15, 16, 18, 19, 28
- Sang, H., Jun, M., and Huang, J. Z. (2011). “Covariance approximation for large multi-

- variate spatial data sets with an application to multiple climate model errors.” *The Annals of Applied Statistics*, 2519–2548. MR2907125. doi: <https://doi.org/10.1214/11-AOAS478>. 2, 27, 30
- Seeger, M. (2003). “Bayesian Gaussian process models: PAC-Bayesian generalisation error bounds and sparse approximations.” Technical report, University of Edinburgh. 3, 7
- Sherman, M. and Carlstein, E. (1994). “Nonparametric estimation of the moments of a general statistic computed from spatial data.” *Journal of the American Statistical Association*, 89(426): 496–500. MR1294075. 2
- Stein, M. L., Chi, Z., and Welty, L. J. (2004). “Approximating likelihoods for large spatial data sets.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 66(2): 275–296. MR2062376. doi: <https://doi.org/10.1046/j.1369-7412.2003.05512.x>. 2
- Van Dyk, D. A. and Park, T. (2008). “Partially collapsed Gibbs samplers: Theory and methods.” *Journal of the American Statistical Association*, 103(482): 790–796. MR2524010. doi: <https://doi.org/10.1198/016214508000000409>. 15
- Varin, C., Reid, N., and Firth, D. (2011). “An overview of composite likelihood methods.” *Statistica Sinica*, 5–42. MR2796852. 15
- Vecchia, A. V. (1988). “Estimation and model identification for continuous spatial processes.” *Journal of the Royal Statistical Society: Series B (Methodological)*, 50(2): 297–312. MR0964183. 2
- Vigsnes, M., Kolbjørnsen, O., Hauge, V. L., Dahle, P., and Abrahamsen, P. (2017). “Fast and accurate approximation to kriging using common data neighborhoods.” *Mathematical Geosciences*, 49: 619–634. MR3660326. doi: <https://doi.org/10.1007/s11004-016-9665-7>. 2
- Wan, Z. and Dozier, J. (1996). “A generalized split-window algorithm for retrieving land-surface temperature from space.” *IEEE Transactions on geoscience and remote sensing*, 34(4): 892–905. 26
- Wan, Z. and Li, Z.-L. (1997). “A physics-based algorithm for retrieving land-surface emissivity and temperature from EOS/MODIS data.” *IEEE Transactions on Geoscience and Remote Sensing*, 35(4): 980–996. 26
- Wan, Z. et al. (2006). “MODIS land surface temperature products users’ guide.” *Institute for Computational Earth System Science, University of California: Santa Barbara, CA, USA*, 805. 26
- Xu, Y., Choi, J., Dass, S., and Maiti, T. (2013). “Efficient Bayesian spatial prediction with mobile sensor networks using Gaussian Markov random fields.” *Automatica*, 49(12): 3520–3530. MR3127559. doi: <https://doi.org/10.1016/j.automatica.2013.09.008>. 2
- Yang, H.-C. and Bradley, J. R. (2021). “Bayesian inference for big spatial data us-

- ing non-stationary spectral simulation.” *Spatial Statistics*, 43: 100507. MR4253852. doi: <https://doi.org/10.1016/j.spasta.2021.100507>. 24
- Zhang, J. and Katzfuss, M. (2022). “Multi-scale Vecchia approximations of Gaussian processes.” *Journal of Agricultural, Biological and Environmental Statistics*, 27(3): 440–460. MR4459075. doi: <https://doi.org/10.1007/s13253-022-00488-0>. 27, 29, 30