

Computationally Scalable Bayesian SPDE Modeling for Censored Spatial Responses

Indranil Sahoo¹, Suman Majumder², Arnab Hazra³, Ana G. Rappold⁴,
Dipankar Bandyopadhyay⁵

¹*Department of Statistical Sciences and Operations Research, Virginia Commonwealth University*

²*Department of Statistics, University of Missouri*

³*Department of Mathematics and Statistics, Indian Institute of Technology Kanpur*

⁴*United States Environmental Protection Agency*

⁵*Department of Biostatistics, Virginia Commonwealth University*

Abstract

Observations of groundwater pollutants, such as arsenic or Perfluorooctane sulfonate (PFOS), are riddled with left censoring. These measurements have impact on the health and lifestyle of the populace. Left censoring of these spatially correlated observations are usually addressed by applying Gaussian processes (GPs), which have theoretical advantages. However, this comes with a challenging computational complexity of $\mathcal{O}(n^3)$, which is impractical for large datasets. Additionally, a sizable proportion of the data being left-censored creates further bottlenecks, since the likelihood computation now involves an intractable high-dimensional integral of the multivariate Gaussian density. In this article, we tackle these two problems simultaneously by approximating the GP with a Gaussian Markov random field (GMRF) approach that exploits an explicit link between a GP with Matérn correlation function and a GMRF using stochastic partial differential equations (SPDEs). We introduce a GMRF-based measurement error into the model, which alleviates the likelihood computation for the censored data, drastically improving the speed of the model while maintaining admirable accuracy. Our approach demonstrates robustness and substantial computational scalability, compared to state-of-the-art methods for censored spatial responses across various simulation settings. Finally, the fit of this fully Bayesian model to the concentration of PFOS in groundwater available at 24,959 sites across California, where 46.62% responses are censored, produces prediction surface and uncertainty quantification in real time, thereby substantiating the applicability and scalability of the proposed method. Code for implementation is made available via [GitHub](#).

Keywords: Censored high-dimensional spatial data; Gaussian Markov random field; Markov chain Monte Carlo; Measurement error model; Perfluorooctane sulfonate; Stochastic partial differential equation.

1 Introduction

Analysis of censored data has been around in the statistical literature since the late 1900s regularly, with the earliest instance of a statistical analysis of censored data being in 1766 (Rosen, 1955). Measurements are often censored due to limitations of measuring instruments, physical inability to acquire data, human error, or similar.

While most analyses on censored data focus on right-censoring (Hosmer et al., 2008), some applications, such as analyzing the concentration of contaminants such as arsenic or per-and polyfluoroalkyl substances (PFAS) in groundwater call for utilizing methods involving left-censored data. This kind of censoring is prevalent in environmental monitoring and has applications in environmental and public health, epidemiology, hydrology, agriculture, and more. Typically these applications involve geostatistical data, where the measurements are censored because they fall below the minimum detection limit (MDL) of the measuring instrument. Early applications would remove censored observations or replace them with makeshift values such as MDL or MDL/2, or impute them with the mean or median of observed responses. Such ad-hoc imputations can result in biased estimates of the overall spatial variability, as demonstrated by Fridley and Dixon (2007).

Recent approaches for statistical inference of spatially distributed censored data overwhelmingly considered the Expectation-Maximization (EM) algorithm (Militino and Ugarte, 1999). Ordoñez et al. (2018) proposed an exact maximum likelihood (ML) estimation of model parameters under censoring, called ‘CensSpatial’, using the Stochastic Approximation of the Expectation Maximization (SAEM; Delyon et al., 1999) algorithm. To tackle the computational complexities arising from censored likelihoods in correlated data, Monte Carlo approximations have been employed, both within the classical framework (Stein, 1992; Rathbun, 2006), and the Bayesian paradigm (De Oliveira, 2005; Tadayon, 2017; Sa-

hoo and Hazra, 2021). For example, Schelin and Sjöstedt-de Luna (2014) introduced a semi-naive approach that utilizes an iterative algorithm and variogram estimation to determine imputed values at locations where data are censored. Finally, various data augmentation techniques have been proposed to facilitate analysis of spatially correlated censored data (Abrahamsen and Benth, 2001; Hopke et al., 2001; Fridley and Dixon, 2007; Sedda et al., 2012). However, the scalability of the suggested approaches is restricted, rendering them unsuitable for the analysis of large spatial datasets featuring censoring, a common occurrence in contemporary scientific research.

Gaussian processes (GPs; Schulz et al., 2018) are heavily used for modeling continuous spatial data due to their several theoretical and computational advantages: the likelihood involves only the first two moments, conditional independence and zeros in the underlying precision matrix are equivalent, and various linear algebraic results are well-known in the literature that are required for computing covariance matrices (Gelfand and Schliep, 2016). However, once the number of spatial sites is large and data at a large proportion of sites are censored, likelihoods based on the underlying GPs involve an intractable high-dimensional integral of a multivariate Gaussian density. This paper aims to overcome the computational challenges inherent to censored likelihoods for high-dimensional spatial settings through a combined application of two key steps:

1. We focus on a fully Bayesian method for censored point referenced data, where the underlying GP is approximated as a Matérn-like Gaussian Markov random field (GMRF, Rue and Held, 2005). The GMRF is obtained as the solution of a stochastic partial differential equation (SPDE, Lindgren et al., 2011) on a fine mesh, which yields a sparse precision matrix of the underlying basis function coefficients. This sparse spatial structure then allows for fast and scalable Bayesian computations.

2. We consider a GMRF-based measurement error model that incorporates a nugget effect in the formulation of the underlying GMRF, which expedites the imputation process for the censored observations. This inclusion effectively reduces the computational burden associated with censored likelihoods (Hazra et al., 2018; Yadav et al., 2019; Zhang et al., 2021).

We draw inferences regarding model parameters using an adaptive Markov Chain Monte Carlo (MCMC) sampling approach, where we use random walk Metropolis-Hastings (MH) steps within Gibbs sampling. Extensive simulations demonstrate the scalability and performance of the proposed methodology in comparison to the ‘CensSpatial’ algorithm, and the traditional local likelihood method (Wiens et al., 2020; Sahoo et al., 2023) applied using Vecchia’s approximation (Vecchia, 1988), across varying degrees of censoring and varying grid sizes. While the idea of a GMRF-based measurement error model has been explored in the context of spatial extremes (Hazra et al., 2021; Cisneros et al., 2023), where replications of the underlying spatial processes are available and censoring a portion of the data is artificial, per our knowledge this modeling strategy has not been explored yet for high-dimensional censored spatial data without replications. Although, the lack of temporal replications typically leads to unstable computations, the proposed stable and scalable computational framework is specifically tailored for handling censored spatial data without requiring temporal replications. Furthermore, unlike previous studies involving GMRF-based measurement error model, a novel feature of the proposed approach is the inclusion of spatial predictions.

PFAS constitute a substantial group of synthetic compounds absent in natural environments, notable for their resistance to heat, water, and oil. PFAS are persistent in the environment and can accumulate within the human body over time, and are toxic at relatively low concentrations (Wang et al., 2017). Exposure to elevated levels of PFAS can lead

to various adverse health outcomes, including developmental issues during pregnancy, cancer, liver impairment, immune system dysfunction, thyroid disruption, and alterations in cholesterol levels. Due to their chemical robustness, PFAS endure in the environment, and are resistant to degradation. Contamination of drinking water with PFAS occurs through the use or accidental spillage of products containing these substances onto land or into waterbodies (Hepburn et al., 2019). PFAS also poses a significant threat to public health, with concentrations measured as high as 82 parts per trillion (ppt) in June 2023 in the United States, above the limit set by the United States Environmental Protection Agency (EPA) at 70 ppt (Cordner et al., 2019). A recent study (Andrews and Naidenko, 2020) estimated that PFAS in publicly accessible drinking water could be affecting as many as 200 million people across the United States. Along similar lines, a robust Bayesian hierarchical approach was proposed (Smalling et al., 2023) to accommodate left-censored PFAS responses; however, the model was implemented on a limited number of sample site locations. As such, the review of existing literature highlights the necessity for further investigation into PFAS occurrences in groundwater, alongside the development of fast and efficient approaches to handle large-scale left-censored spatial data in real time. Motivated by data on PFAS concentrations collected by the Groundwater Ambient Monitoring and Assessment (GAMA) program (Moran et al., 2005) across the state of California, we develop our Bayesian scalable model for spatially-referenced left-censored PFAS responses, in an attempt to provide a more accurate quantification of the groundwater contamination within the state. These data, collected by GAMA since 2019, allow thorough quality assessments of water sources and enable the establishment of safety thresholds for select PFAS constituents. Thus, our analysis is capable of identifying possible hotspots of higher PFAS concentration, thereby providing insights for further analysis of impacts on public health.

The subsequent sections of the paper are organized as follows. In Section 2, we provide details regarding the dataset on the groundwater levels of PFAS within the state of California, along with some exploratory analyses. We outline our methodology and related computational details in Section 3, and test its scalability and predictive performance on simulated datasets in Section 4. In Section 5, we apply the proposed methodology to the PFAS dataset and report the findings. We conclude with a brief discussion in Section 6.

2 Motivating PFAS data

The groundwater PFAS data for the state of California are available online at the website GAMA Groundwater Information Systems under the label `Statewide PFOS Data`. In this paper, we focus specifically on the measurements of the chemical substance, known as Perfluorooctane sulfonate (PFOS). The dataset contains 24,959 measurements (in ng/L) of PFOS concentration and their locations (in longitudes and latitudes) as well as indicators of whether the observations are censored or not and the corresponding censoring limits within the state of California. Almost half of the measurements (46.62%) are censored observations, with varying degrees of censoring limits.

Figure 1 shows transformed PFOS concentration measurements, after transforming the raw PFOS by $g(\text{PFOS}) = \log(1 + \log(1 + \text{PFOS}))$ at the 24,959 irregularly sampled spatial locations across the state of California, prompting an approximate spatial inference model to be employed, which also accounts for the huge proportion of censored observations. Most observation sites are towards the densely populated areas on the coast. The censored observations are presented as tiny black dots in Fig 1. The censored observations are all over the spatial domain, and are not limited to one single area. Most of the measurements are below 150 ng/L, but there are observations as high as 1,330,000 ng/L, which is well-beyond

the safety limit prescribed by the EPA.

The histogram of the raw non-censored PFOS observations is presented in the left panel of Figure 2. The raw data exhibit a highly positively skewed nature, and thus a stationary Gaussian process assumption naturally becomes questionable, even after considering a spatially-smooth mean surface with covariates, like longitude and latitude. Following an exploration of different transformations of the raw data such that the histogram behaves in an approximately bell-shaped fashion, we identify that the iterated log-transformation $g(\text{PFOS}) = \log(1 + \log(1 + \text{PFOS}))$ performs reasonably well; the histogram of the transformed PFOS data is presented in the middle panel of Figure 2. We further explore the effects of the natural covariates longitude and latitude on the transformed data; following a simple linear regression, we obtain the residuals, and their histogram is presented in the right panel of Figure 2. This histogram is reasonably bell-shaped and thus we model this transformed PFOS data using a Gaussian process framework with a regression structure for the mean process where we allow longitude and latitude as covariates.

We further explore the presence of spatial correlation following a variogram analysis of the residuals (scaled by their sample standard deviation) discussed above. The sample semivariogram at distance d is defined as

$$\hat{\gamma}(d) = \frac{1}{2N(d)} \sum_{i=1}^n \sum_{j=1}^i w_{ij}(d) (R(\mathbf{s}_i) - R(\mathbf{s}_j))^2,$$

where, $R(\mathbf{s}_i)$ and $R(\mathbf{s}_j)$ are the residuals at spatial sites \mathbf{s}_i and \mathbf{s}_j , $w_{ij}(d) = 1$ if $d_{ij} \in (d - h, d + h)$ and $w_{ij} = 0$ otherwise, d_{ij} being the distance between \mathbf{s}_i and \mathbf{s}_j . Also, $N(d)$ is the number of pairs with $w_{ij}(d) = 1$. The sample semivariogram, presented in Figure 3, indicates the presence of spatial correlation, along with the presence of possible nugget effects (Bivand et al., 2008). We fit an isotropic Matérn spatial correlation function, with

PFOS concentration in California

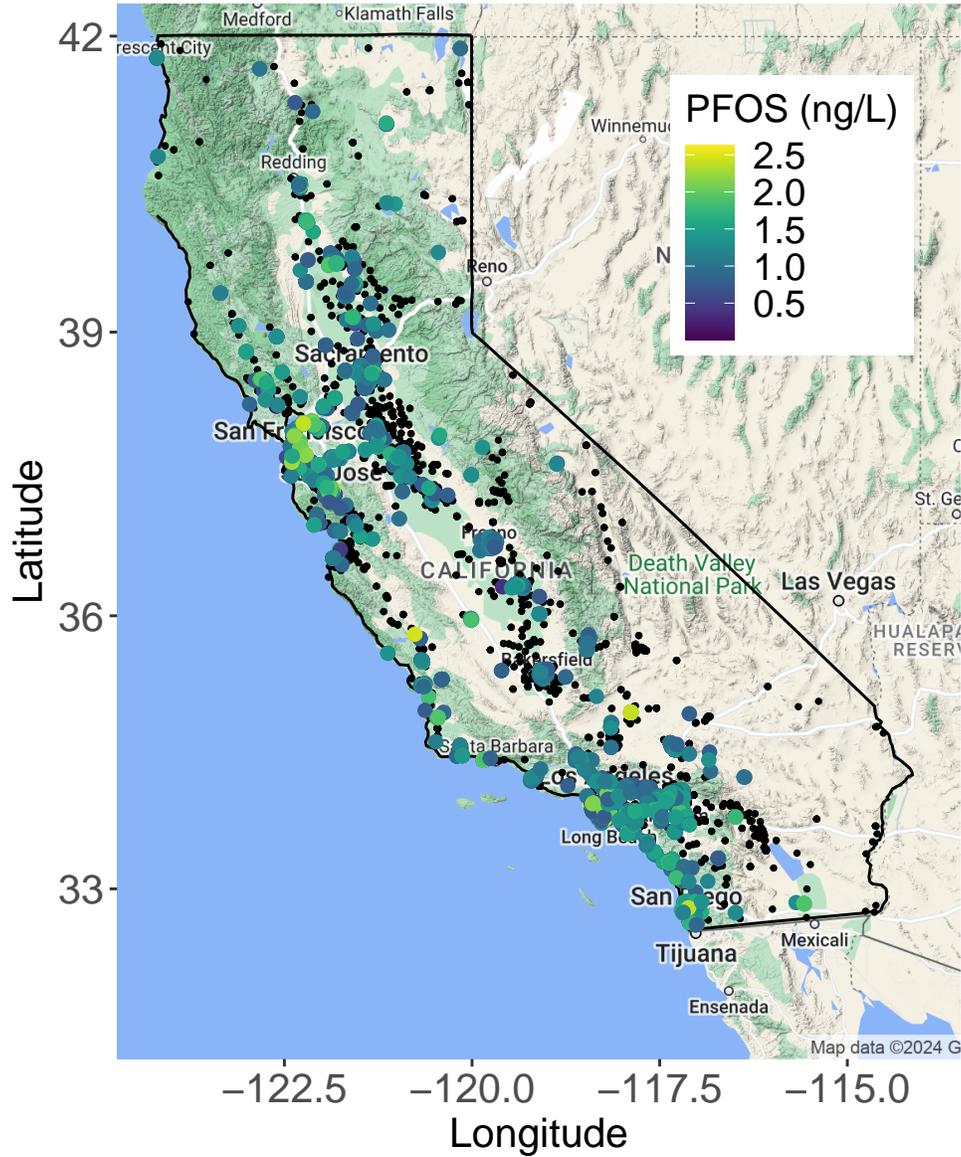


Figure 1: Concentrations of (transformed) PFOS, using the transformation $g(\text{PFOS}) = \log(1 + \log(1 + \text{PFOS}))$, measured at 24,959 irregularly-sampled spatial locations across the state of California (in ng/L). The tiny black dots indicate the sites with censored data.

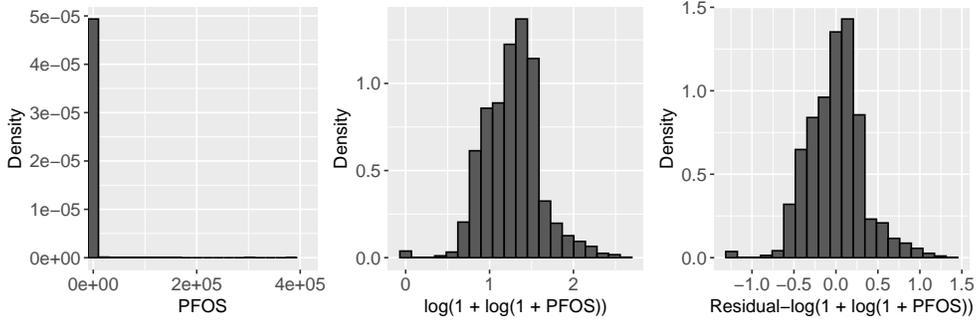


Figure 2: Pictorial representation of the raw PFOS responses. *Left panel:* Histogram of raw concentrations of PFOS across sites where the data are not censored. *Middle:* Histogram of non-censored PFOS concentrations after the transformation $g(\text{PFOS}) = \log(1 + \log(1 + \text{PFOS}))$. *Right panel:* Histogram of the residuals obtained after regressing non-censored transformed PFOS observations to longitude and latitude, via a linear model.

its smoothness parameter set to one, plus a nugget effect, given by

$$\rho(\mathbf{s}_i, \mathbf{s}_j) \equiv \rho(d) = \gamma \frac{d}{\phi} \kappa_1 \left(\frac{d}{\phi} \right) + (1 - \gamma) \mathbb{1}(\mathbf{s}_i = \mathbf{s}_j), \quad (1)$$

where, d is the Euclidean distance between locations \mathbf{s}_i and \mathbf{s}_j , $\phi > 0$ is the range parameter, $\gamma \in [0, 1]$ is the ratio of spatial to total variation, $\kappa_1(\cdot)$ is the modified Bessel function of second kind with degree 1, and $\mathbb{1}(\cdot)$ is the indicator function. The fitted population semivariance indicates a reasonable fit to the sample semivariogram. While these exploratory analyses are based on non-censored observations only, they indicate a need for proper spatial modeling after considering the censored nature of a large proportion of the data. Specifically, most of the observations near the eastern regions of the study domain are censored, and ignoring them in the spatial prediction would lead to poor spatial prediction for the nearby regions.

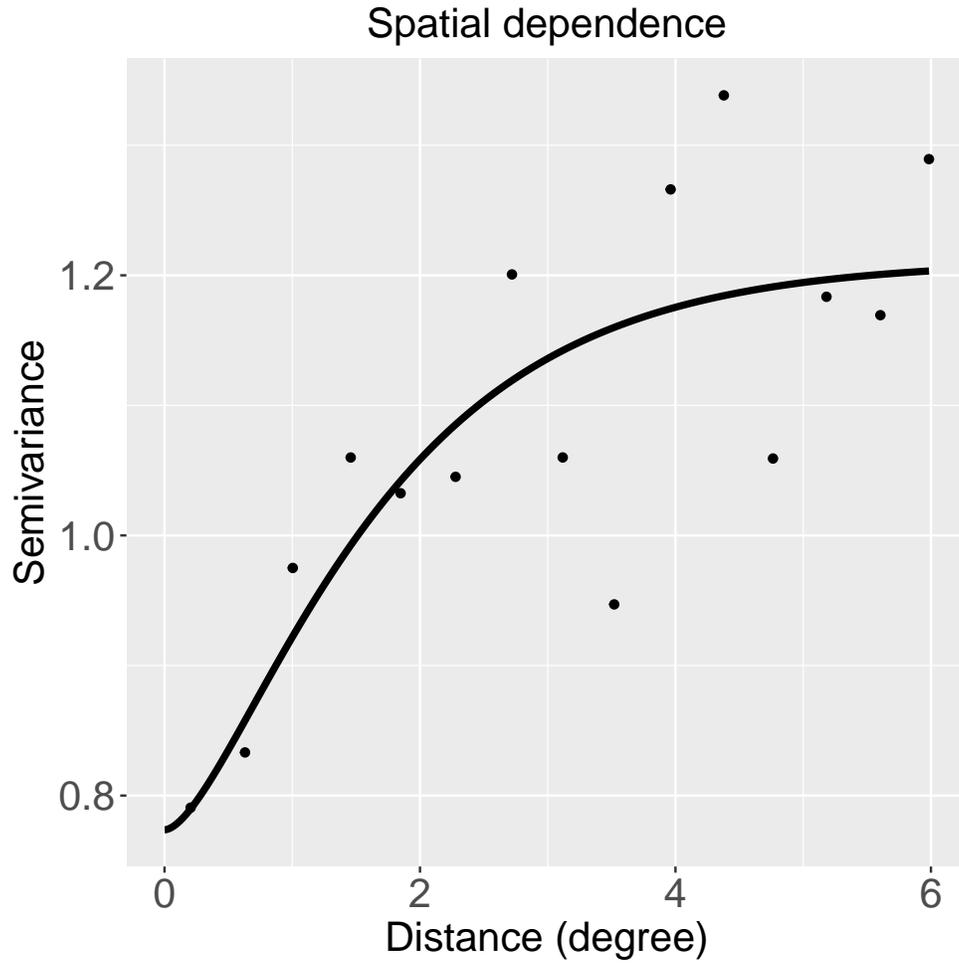


Figure 3: Sample semivariogram of residuals obtained after regressing non-censored transformed PFOS observations to longitude and latitude, as a function of distance (dots). The overlapped solid line represents the fitted population semivariance obtained from (1).

3 Methodology

Let $Y(\mathbf{s})$ represent transformed PFOS concentration at a spatial location $\mathbf{s} \in \mathcal{D} \subset \mathbb{R}^2$, where \mathcal{D} represents the spatial domain of interest, i.e., the entire state of California in our case. We model $Y(\mathbf{s})$ as

$$Y(\mathbf{s}) = \mathbf{X}(\mathbf{s})^T \boldsymbol{\beta} + \tau^{-1/2} Z(\mathbf{s}),$$

where, $X(\mathbf{s}) = [X_1(\mathbf{s}), \dots, X_p(\mathbf{s})]^T$ denotes the vector of p covariates at location \mathbf{s} , $\boldsymbol{\beta} = [\beta_1, \dots, \beta_p]^T$ is a vector of unknown regression parameters and $\tau > 0$ is a spatially-constant precision parameter. Given the absence of meaningful covariates in our data, we choose $\mathbf{X}(\mathbf{s}) = [1, \text{longitude}(\mathbf{s}), \text{latitude}(\mathbf{s})]^T$ for our analysis. We assume that $Z(\cdot)$ is a standard (mean zero and variance one at each site) GP with an isotropic Matérn spatial correlation plus a nugget effect, given by (1). While the incorporation of the nugget effect is justified by our exploratory analysis (nonzero semivariance at origin), it effectively addresses the issue of censoring in the response, thereby circumnavigating the computational burden occurring due to censored likelihoods (Hazra et al., 2018; Yadav et al., 2019; Zhang et al., 2021). Here, we fix the smoothness parameter of the Matérn correlation of the purely spatial component of (1) to one. In practice, it is difficult to estimate the smoothness parameter from the data, and hence it is generally fixed. Besides, we later build a stochastic partial differential equation-based approximation of the Matérn correlation structure, where fixing the smoothness parameter to one is a standard choice (Hazra et al., 2021; Cisneros et al., 2023).

Suppose the data are observed (either censored or non-censored) at the set of sites $\mathcal{S} = \{\mathbf{s}_1, \dots, \mathbf{s}_n\}$. In matrix notations, the spatial linear model can be written as

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \tau^{-1/2}\mathbf{Z}, \tag{2}$$

where, $\mathbf{Y}_{(n \times 1)}$ is the response vector, $\mathbf{X}_{(n \times p)}$ is the matrix of covariates, $\boldsymbol{\beta}_{(p \times 1)}$ is the vector

of regression coefficients, and $\mathbf{Z}_{(n \times 1)} \sim \text{MVN}(\mathbf{0}, \gamma \boldsymbol{\Sigma} + (1 - \gamma) \mathbf{I}_n)$, where $\boldsymbol{\Sigma}$ is the Matérn correlation matrix, and \mathbf{I}_n denotes the identity matrix of order n . Here, by construction and the PFAS dataset, $\boldsymbol{\Sigma}$ is non-singular and \mathbf{X} has full rank.

In a spatial censored linear (SCL) model, it is further assumed that $Y(\mathbf{s})$ is not fully observed at all spatial locations. Motivated by the dataset considered, we assume $Y(\cdot)$ to be left-censored at sites $\mathcal{S}^{(c)} = \{\mathbf{s}_1^{(c)}, \dots, \mathbf{s}_{n_c}^{(c)}\} \subset \mathcal{S}$ and the corresponding censoring levels be $\mathcal{U} = \{u_1, \dots, u_{n_c}\}$. However, a similar approach can be applied if the response is right or interval-censored. We define the censoring indicator $\delta(\mathbf{s})$ as

$$\delta(\mathbf{s}) = \begin{cases} 1, & \text{if } Y(\mathbf{s}) \text{ is censored at site } \mathbf{s}, \\ 0, & \text{otherwise,} \end{cases}$$

and the vector of censored observations as $\mathbf{v} = [Y(\mathbf{s}_i) : \delta(\mathbf{s}_i) = 1]^T \equiv [Y(\mathbf{s}_1^{(c)}), \dots, Y(\mathbf{s}_{n_c}^{(c)})]^T$. Then, for censored spatial data, the likelihood is given by

$$\mathcal{L}(\boldsymbol{\theta}) = \int_{\mathbf{v} \leq \mathbf{u}} f_{\text{MVN}}(\mathbf{y}; \mathbf{X}\boldsymbol{\beta}, \tau^{-1}[\gamma \boldsymbol{\Sigma} + (1 - \gamma) \mathbf{I}_n]) d\mathbf{v}, \quad (3)$$

where the integral is over the censored responses $\{\mathbf{y} : y(\mathbf{s}_i) \leq u_i \text{ if } \mathbf{s}_i \in \mathcal{S}^{(c)}\}$ and $f_{\text{MVN}}(\cdot; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes a multivariate normal density with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. A version of this likelihood has been studied in Sahoo and Hazra (2021).

3.1 Approximation of the Matérn Gaussian Process

Here, we follow the approximation strategy of the Matérn GP with nugget (Hazra et al., 2021). To ensure computational efficiency, we choose to approximate the Gaussian process $Z(\cdot)$ with a GP $\tilde{Z}(\cdot)$, constructed from a Gaussian Markov random field (GMRF) defined on a finite mesh, thereby circumventing the computational overhead associated with the dense correlation matrix inherent in the exact Matérn GP defined by (1). This strategy

capitalizes on the direct correspondence between continuous-space Matérn GP with dense covariance matrices and GMRFs with sparse precision matrices (Lindgren et al., 2011), which yields an approximate data process

$$\tilde{Y}(\mathbf{s}) = \mathbf{X}(\mathbf{s})^T \boldsymbol{\beta} + \tau^{-1/2} \tilde{Z}(\mathbf{s}), \quad \mathbf{s} \in \mathcal{D}.$$

For $\gamma = 1$, the Gaussian Matérn process $Z(\cdot)$ can be obtained as a solution to the linear Stochastic Partial Differential Equation (SPDE)

$$(\phi^{-2} - \Delta)Z(\mathbf{s}) = 4\pi\phi^{-2}\mathcal{W}(\mathbf{s}), \quad \mathbf{s} \in \mathbb{R}^2, \quad (4)$$

where $\mathcal{W}(\cdot)$ is a Gaussian white noise process, and Δ is the Laplacian. The solution $Z(\mathbf{s})$ to the SPDE can be effectively approximated through finite-element methods (Ciarlet, 2002) applied to a triangulated mesh defined within a bounded region of \mathbb{R}^2 , where the triangulation is formed through a refined Delaunay triangulation process (Borouchaki and Lo, 1995). In practical applications, the mesh can be easily constructed using the (currently depreciated) `inla.mesh.2d` function, implemented in the R package INLA (www.r-inla.org) or the `fm.mesh.2d.inla` function in the R package `fmesh` (<https://cran.r-project.org/package=fmesher>); see Lindgren and Rue (2015) for more details. The left panel of Figure 4 depicts the mesh utilized in the data application discussed in Section 5.

Let $\mathcal{S}^* = \{\mathbf{s}_1^*, \dots, \mathbf{s}_N^*\}$ denote the set of mesh nodes. We construct a finite-element solution by writing $\tilde{Z}(\mathbf{s}) = \sum_{j=1}^N \zeta_j(\mathbf{s})Z_j^*$, and plugging it in (4) in place of $Z(\cdot)$. Here, $\{\zeta_j(\cdot)\}$ are piecewise linear and compactly-supported “hat” basis functions defined over the mesh, and $\{Z_j^*\}$ are normally distributed weights defined for each basis function (that is, one for each mesh node in \mathcal{S}^*). Then, $\mathbf{Z}^* = [Z_1^*, \dots, Z_N^*]^T \sim \text{MVN}(\mathbf{0}, \mathbf{Q}_\phi^{-1})$, where the $(N \times N)$ -dimensional precision matrix \mathbf{Q}_ϕ can be written as

$$\mathbf{Q}_\phi = \frac{\phi^2}{4\pi} \left[\frac{1}{\phi^4} \mathbf{D} + \frac{2}{\phi^2} \mathbf{G}_1 + \mathbf{G}_2 \right], \quad (5)$$

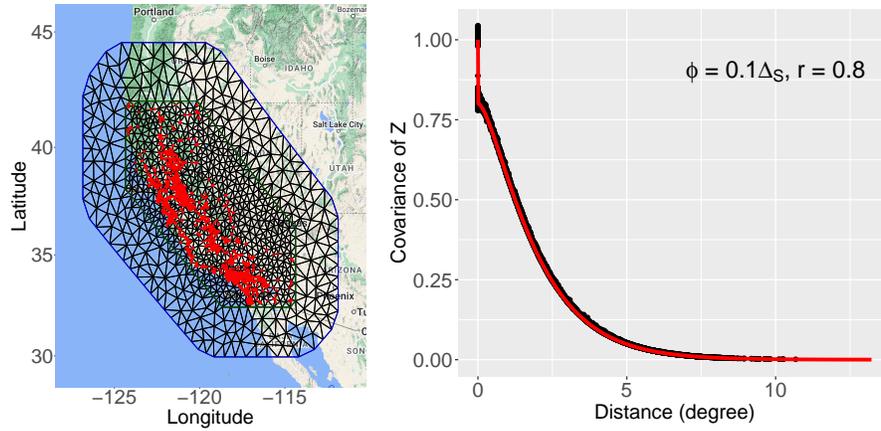


Figure 4: *Left:* Triangulated mesh over California, that is used to approximate the spatial SPDE process $Z(\cdot)$. *Right:* Comparison of the true Matérn correlation (solid line) and the pairwise covariances between two spatial locations obtained by the SPDE approximation (points), as a function of distance. The parameters are set to $\phi = 0.1\Delta_S$, with Δ_S being the maximum spatial distance between two locations in the domain, and $r = 0.8$.

where \mathbf{D} , \mathbf{G}_1 , and \mathbf{G}_2 are sparse $(N \times N)$ -dimensional finite-element matrices that can be obtained as follows. The matrix \mathbf{D} is diagonal with its j^{th} diagonal entry $D_{j,j} = \langle \zeta_j(\cdot), 1 \rangle$, where $\langle f, g \rangle = \int f(\mathbf{s})g(\mathbf{s})d\mathbf{s}$ denotes an inner product. Similarly, \mathbf{G}_1 has the elements $G_{1;j,k} = \langle \nabla \zeta_j(\cdot), \nabla \zeta_k(\cdot) \rangle$ and $\mathbf{G}_2 = \mathbf{G}_1 \mathbf{D}^{-1} \mathbf{G}_1$. Efficient computation of these sparse matrices is implemented using the function `inla.mesh.fem` from the R package INLA. For further theoretical details, see Bakka et al. (2018) and Lindgren et al. (2022).

In order to map the spatial random effects \mathbf{Z}^* (defined across mesh nodes) back to the observation locations \mathcal{S} , we use an $(n \times N)$ -dimensional projection matrix \mathbf{A} . The $(i, j)^{\text{th}}$ element of this matrix corresponds to $\zeta_j(\mathbf{s}_i)$ for every spatial location $\mathbf{s}_i \in \mathcal{S}$ and mesh node $\mathbf{s}_j^* \in \mathcal{S}^*$, allowing us to compute \mathbf{AZ}^* , the projection of \mathbf{Z}^* at the data locations. The generation of the matrix \mathbf{A} is carried out through the function `inla.spde.make.A` within the R package INLA.

In presence of a nugget effect $\gamma \in [0, 1]$ in (1), the model for $\tilde{\mathbf{Z}} = [\tilde{Z}(\mathbf{s}_1), \dots, \tilde{Z}(\mathbf{s}_n)]^T \equiv [\tilde{Z}_1, \dots, \tilde{Z}_n]^T$ becomes

$$\tilde{\mathbf{Z}} = \sqrt{\gamma} \mathbf{AZ}^* + \sqrt{1 - \gamma} \boldsymbol{\eta},$$

where $\boldsymbol{\eta} = [\eta_1, \dots, \eta_n]^T$ with $\eta_i \stackrel{IID}{\sim} \text{N}(0, 1), i = 1, \dots, n$. Then, the covariance matrix of $\tilde{\mathbf{Z}}$ can be written as

$$\boldsymbol{\Sigma}_{\tilde{\mathbf{Z}}} = \gamma \mathbf{A} \mathbf{Q}_{\phi}^{-1} \mathbf{A}^T + (1 - \gamma) \mathbf{I}_n. \quad (6)$$

The SPDE approach yields a precise approximation of the true correlation structure; see the right panel of Figure 4. Leveraging the sparsity of the matrix \mathbf{Q}_{ϕ}^{-1} , we facilitate rapid Bayesian computations. Furthermore, we exploit the conditional independence structure of $\tilde{\mathbf{Z}} | \mathbf{Z}^* \sim \text{MVN}(\sqrt{\gamma} \mathbf{AZ}^*, (1 - \gamma) \mathbf{I}_n)$ for swift imputation of left-censored observations.

3.2 Final hierarchical model

Suppose we write the vector of the final approximate data process, $\tilde{Y}(\cdot)$, evaluated at \mathcal{S} by $\tilde{\mathbf{Y}} = [\tilde{Y}(\mathbf{s}_1), \dots, \tilde{Y}(\mathbf{s}_n)]^T$. We define a rescaled random effects vector defined at mesh nodes by $\tilde{\mathbf{Z}}^* = \sqrt{\gamma/\tau} \mathbf{Z}^*$. Then, similar to (2), the model for $\tilde{\mathbf{Y}}$ can be written hierarchically as

$$\begin{aligned} \tilde{\mathbf{Y}}|\tilde{\mathbf{Z}}^* &\sim \text{MVN}\left(\mathbf{X}\boldsymbol{\beta} + \mathbf{A}\tilde{\mathbf{Z}}^*, \tau^{-1}(1 - \gamma)\mathbf{I}_n\right), \\ \tilde{\mathbf{Z}}^* &\sim \text{MVN}(\mathbf{0}, \gamma\tau^{-1}\mathbf{Q}_\phi^{-1}), \\ \{\boldsymbol{\beta}, \tau, \phi, \gamma\} &\sim \pi(\boldsymbol{\beta}) \times \pi(\tau) \times \pi(\phi) \times \pi(\gamma). \end{aligned} \quad (7)$$

Here the last layer of the model indicates prior choices for the model parameters that we discuss in Section 3.4. Instead of the likelihood based on (2), we fit the approximate data process (7) to the actual observation process.

3.3 Prediction

Let $\mathcal{S}^{(0)} = \{\mathbf{s}_1^{(0)}, \dots, \mathbf{s}_{n_0}^{(0)}\} \subset \mathcal{D}$ denote a set of n_0 prediction sites, and define $\tilde{\mathbf{Y}}^{(0)} = [\tilde{Y}(\mathbf{s}_1^{(0)}), \dots, \tilde{Y}(\mathbf{s}_{n_0}^{(0)})]^T$. Also, let $\mathbf{X}^{(0)}$ denote the $(n_0 \times p)$ -dimensional design matrix, with its i^{th} row $\mathbf{X}(\mathbf{s}_i^{(0)})$, $i = 1, \dots, n_0$ denoting the vector of covariates at prediction location $\mathbf{s}_i^{(0)}$. For mapping the (scaled) spatial random effects $\tilde{\mathbf{Z}}^*$ (defined across mesh nodes) to the prediction locations $\mathcal{S}^{(0)}$, we use an $(n_0 \times N)$ -dimensional projection matrix $\mathbf{A}^{(0)}$. The $(i, j)^{\text{th}}$ element of this matrix corresponds to $\zeta_j(\mathbf{s}_i^{(0)})$ for every spatial location $\mathbf{s}_i^{(0)} \in \mathcal{S}^{(0)}$ and mesh node $\mathbf{s}_j^* \in \mathcal{S}^*$, allowing us to compute $\mathbf{A}^{(0)}\tilde{\mathbf{Z}}^*$, the projection of $\tilde{\mathbf{Z}}^*$ at $\mathcal{S}^{(0)}$. Then, given $\tilde{\mathbf{Z}}^*$, the conditional distribution of $\tilde{\mathbf{Y}}^{(0)}$ is

$$\tilde{\mathbf{Y}}^{(0)}|\tilde{\mathbf{Z}}^* \sim \text{MVN}\left(\mathbf{X}^{(0)}\boldsymbol{\beta} + \mathbf{A}^{(0)}\tilde{\mathbf{Z}}^*, \tau^{-1}(1 - \gamma)\mathbf{I}_n\right) \quad (8)$$

3.4 Computational details

Inference concerning the model parameters is conducted through Markov chain Monte Carlo (MCMC) sampling, implemented in R. Given the computational dependence on prior selections for the model parameters, we first specify these priors. Whenever feasible, we opt for conjugate priors and employ Gibbs sampling to update them iteratively. When prior conjugacy is unavailable, we resort to random walk Metropolis-Hastings (MH) steps for parameter updates. During the burn-in period, we adjust the candidate distributions within the MH steps to ensure that the acceptance rate throughout the post-burn-in period remains within the range of 0.3 to 0.5.

Here we draw samples from the full posterior

$$\pi(\boldsymbol{\beta}, \tau, \phi, \gamma, \tilde{\mathbf{Z}}^*, \tilde{\mathbf{Y}}^{(c)} | \tilde{\mathbf{Y}}^{(nc)}),$$

where $\tilde{\mathbf{Y}}^{(c)}$ is the vector of censored data vector and $\tilde{\mathbf{Y}}^{(nc)}$ is the vector of non-censored data vector. For the vector of regression coefficients $\boldsymbol{\beta}$, we consider weakly-informative conjugate prior $\boldsymbol{\beta} \sim \text{MVN}(\mathbf{0}, 100^2 \mathbf{I}_p)$. The full conditional posterior of $\boldsymbol{\beta}$ is then multivariate normal and hence updated using direct sampling within Gibbs steps. Due to the strong posterior correlation between $\boldsymbol{\beta}$ and $\tilde{\mathbf{Z}}^*$, they are updated jointly within each Gibbs sampling step. For the hyperparameters involved in the correlation function in (1), we consider the non-informative priors as well. Specifically, we choose $\phi \sim \text{Uniform}(0, 0.5\Delta_S)$ for the spatial range parameter, where Δ_S is the largest Euclidean distance between two data locations, and $\gamma \sim \text{Uniform}(0, 1)$ for the nugget effect γ . We further designate a non-informative conjugate prior for the spatially-constant precision parameter τ in the process model, namely $\tau \sim \text{Gamma}(0.1, 0.1)$. The full conditional posterior distribution of $\tilde{\mathbf{Y}}^{(c)}$ is $\text{MVN}\left(\mathbf{X}^{(c)}\boldsymbol{\beta} + \mathbf{A}^{(c)}\tilde{\mathbf{Z}}^*, \tau^{-1}(1 - \gamma)\mathbf{I}_n\right)$, where $\mathbf{X}^{(c)}$ and $\mathbf{A}^{(c)}$ are design matrix and SPDE projection matrix (from mesh nodes), respectively, corresponding to the locations with cen-

sored data, i.e., they comprise of the rows of \mathbf{X} and \mathbf{A} that correspond to the censored entries of \mathbf{Y} .

3.5 Software

We have developed an open-source R package, called **CensSpBayes**, which implements the proposed approximate Matérn GP model for large left-censored spatial data. Implementation code, along with details of execution using simulated data are made available at <https://github.com/SumanM47/CensSpBayes>.

4 Simulation Study

In this section, we conduct simulation studies using synthetic data to assess the efficacy of our proposed scalable modeling framework in terms of spatial prediction, while imputing censored values. We simulate 100 datasets over grids $\mathcal{D}^* = \{(i, j) : i, j \in \{1/K, 2/K, \dots, 1\}\}$ of varying sizes within a spatial domain $[0, 1]^2$. We consider $K \times K$ grids with $K = 20, 50, 100$, and 200 , to demonstrate the computational power and scalability inherent to our proposed methodology.

For simulating the datasets, we consider an intercept term only and no other covariates, and we assume the true value of the regression coefficient to be $\beta^{\text{true}} = 5$. The true value of the range parameter of the spatial Matérn correlation is chosen to be $\phi^{\text{true}} = 0.15 \times \Delta^*$, where $\Delta^* = \sqrt{2}$, the maximum spatial distance between two locations in $[0, 1]^2$. The smoothness parameter is set to one and not estimated while fitting our proposed model. The true ratio of partial sill to total variation is chosen to be $\gamma^{\text{true}} = 0.9$, and the true precision parameter is chosen to be $\tau^{\text{true}} = 1/5$. Exact simulations are conducted to generate datasets from a GP with Matérn correlation, as given in (1).

Once the datasets are generated, we divide each dataset randomly into 80% training and 20% test datasets. Within each training set, we consider two different levels of censoring (denoted by L1 and L2) for the response by setting different values of the minimum detection limit (MDL):

L1 Low censoring: The MDL is at the 15th percentile point of observations and thus 15% data are censored.

L2 High censoring: The MDL is at the 45th percentile point of observations and thus 45% data are censored.

For each of the two levels of censoring, we implement our proposed approximate Matérn GP model under three different settings, denoted by S1, S2, and S3:

S1 We selectively exclude spatial locations where observations are censored and apply the spatial model approximated via SPDE, as elaborated in Section 3.1, exclusively to the observed locations. This does not require any imputation of the censored observations.

S2 We set the censored observations at the MDL and employ the SPDE-approximated Matérn GP model, as detailed in Section 3.1. This once again circumvents the need for imputing the censored observations.

S3 We fit the full proposed model, treating the observations below MDL as censored observations, and implement the SPDE-approximated Matérn GP model, as in Section 3.1, while simultaneously performing imputations for the censored observations.

In each of the three scenarios, the approximated spatial process using SPDE has been fitted to assess the effects of removing or considering ad hoc imputations of censored observations, in terms of mean squared prediction error (MSPE), in contrast to treating them

as genuinely censored. The prior distributions for β and γ as described in Section 3.4 remain unchanged in the simulation study. However, for the range parameter, we assume $\phi \sim \text{Uniform}(0, 0.25\Delta^*)$.

For comparison, we consider two additional models, S4 and S5:

S4 We use ML estimation to locally fit the full Matérn Gaussian censored likelihood, using Vecchia’s approximation with $m = 30$ nearest neighbors, thereby estimating the model parameters under censoring. Here, we fix the Matérn smoothness parameter to 1 in keeping with the data generation mechanism.

S5 We use the ‘CensSpatial’ algorithm to perform an exact ML estimation of model parameters, which implements the SAEM algorithm of Ordoñez et al. (2018).

Table 1 presents the average (across 100 simulated datasets) mean squared prediction errors (MSPE), along with corresponding average standard errors, obtained from fitting the models S1-S5 to data in test sets that varies according to censoring levels and grid sizes. Notably, under low levels of censoring (L1), the proposed model in scenario S3 yields comparable performance to situations where censored observations are excluded from analysis. However, ignoring spatial locations with censored observations entirely leads to unreliable estimates, particularly for the covariance parameters. Conversely, in instances of high data censoring (L2), the final model along with the imputation of censored observations (S3) outperforms all other models. It is noteworthy that the ‘CensSpatial’ method also demonstrates relatively favorable performance for a grid size of 20×20 , when the data-generating model is Matérn GP; however, its computational inefficiency and inadequate scaling impeded our ability to apply the method to the higher-dimensional simulated datasets. In fact, Ordoñez et al. (2018) showcased the efficacy of the ‘CensSpatial’ algorithm through

Table 1: Average mean squared prediction error (MSPE) corresponding to model fitting under the five different settings (S1-S5) to simulated test data from a Matérn GP that varies with censoring levels L1 (low-censoring; 15%) and L2 (high-censoring; 45%), and grid sizes. The values in parenthesis represent the corresponding average prediction standard errors. The lowest MSPE in each row is in bold. Since model S5 (‘CensSpatial’) was infeasible for larger grids, table entries (MSPE and standard errors) appear as ‘-’.

Censoring Level	Grid size	Mean Squared Prediction Error				
		S1	S2	S3	S4	S5
L1	20 × 20	1.13(0.82)	2.10(0.86)	0.79 (0.85)	99.15(6.0E18)	0.89(0.97)
	50 × 50	0.87(0.74)	2.03(0.80)	0.61 (0.77)	136.55(3.3E84)	-
	100 × 100	0.80(0.70)	1.88(0.76)	0.55 (0.74)	189.94(3.5E124)	-
	200 × 200	0.79(0.69)	2.00(0.77)	0.54 (0.73)	242.44(4.1E122)	-
L2	20 × 20	2.95(0.82)	4.92(0.63)	0.90 (0.90)	52.77(2.4E34)	1.51(1.08)
	50 × 50	2.52(0.74)	5.22(0.60)	0.69 (0.82)	103.02(4.3E109)	-
	100 × 100	2.29(0.69)	4.89(0.57)	0.61 (0.77)	137.25(2.8E119)	-
	200 × 200	2.29(0.66)	4.91(0.55)	0.58 (0.75)	176.24(1.2E124)	-

simulations involving only 50 and 200 spatial locations, clearly indicating its inadequacy regarding scalability.

Table 2 presents the median computation time corresponding to fitting the models S1-S5 to data in test sets that varies according to censoring levels and grid size. The computation times for S1, S2, and S3 are comparable, as they all employ the SPDE-approximated Matérn GP, with runtime approximately proportional to the size of the INLA mesh used for process approximation (between 557 and 673 nodes for different datasets of different

Table 2: Median computation time (in minutes) corresponding to model fitting under the five different settings (S1-S5) to simulated test data from a Matérn GP that varies with censoring levels L1 (low-censoring; 15%) and L2 (high-censoring; 45%), and grid sizes. The values in parenthesis represent the median absolute deviation for the corresponding computing times. Since model S5 (‘CensSpatial’) was infeasible for larger grids, table entries appear as ‘-’.

Censoring	Grid	Computation Time (in minutes)				
		S1	S2	S3	S4	S5
L1	20 × 20	27.91(0.79)	28.64(0.46)	29.94(0.67)	0.97(0.30)	42.70(2.23)
	50 × 50	22.32(0.30)	22.61(0.48)	22.39(0.51)	7.41(0.90)	-
	100 × 100	31.22(1.13)	31.58(1.18)	31.47(0.80)	31.75(2.22)	-
	200 × 200	25.40(0.33)	25.46(0.34)	25.98(0.31)	129.11(8.60)	-
L2	20 × 20	26.47(0.50)	28.82(0.48)	29.68(0.92)	3.99(0.26)	80.40(3.53)
	50 × 50	21.67(0.45)	22.70(0.56)	22.78(0.22)	25.82(0.92)	-
	100 × 100	29.28(1.05)	30.18(1.41)	31.06(1.14)	104.14(2.69)	-
	200 × 200	24.74(0.51)	25.47(0.27)	26.56(0.36)	423.26(8.90)	-

sizes). As anticipated, the runtime for the local likelihood approach utilizing Vecchia’s approximation increases with larger grid sizes. As discussed earlier, the ‘CensSpatial’ algorithm was infeasible for larger grids. The initial four methods, S1 - S4, were executed on SLURM clusters with one core per job and 8 GB RAM allocation. However, due to the current version of ‘CensSpatial’ on CRAN being incompatible with the UNIX system, the algorithm was implemented on a personal Dell 7210 computer featuring 16 GB RAM, an Intel Core i5 dual-core processor, and a Windows 11 Enterprise 64-bit operating system.

5 Application: California PFAS Data

5.1 Analysis Plan and Hyperparameters

We use the iterated log-transformed data (as described in Section 2) as our input to the proposed method. Since we have no covariates in this dataset, we use the coordinates of the locations (longitude, latitude) as covariates. The hyperparameters for the priors are the same as mentioned in Section 3.4. We fit a variogram model on the non-censored observations to obtain an initial set of parameter estimates for β , τ , ϕ , and r . We run three chains with different starting values that are all close to the initial parameter estimates obtained by variogram fitting to allow for checking convergence and increasing the reliability of the model output. Each of the chains was run for 25,000 iterations with the first 15,000 samples discarded as burn-in. We thinned the post-burn-in samples by 5 to obtain 2,000 samples from the posterior distribution of the parameters.

5.2 Results

For the observed data comprising 24,959 locations and the prediction grid set at $0.1^\circ \times 0.1^\circ$, resulting in 405,893 prediction locations across California, each of the three chains completed within approximately 62 minutes, with computations conducted on a SLURM cluster with an 8GB RAM allocation. We observe a reasonable well-mixing of the three chains. Each of the estimated covariate effects turned out to be significant (see Table 3). The estimated spatial range is low ($\sim 7\text{km}$). The prediction surface is smooth at places (left panel of 5) with higher detailing around the regions with observed data. The prediction standard deviation is low towards the western parts where we have more observed data, but has an intriguing pattern on the east-southeastern parts (right panel of Figure 5). We

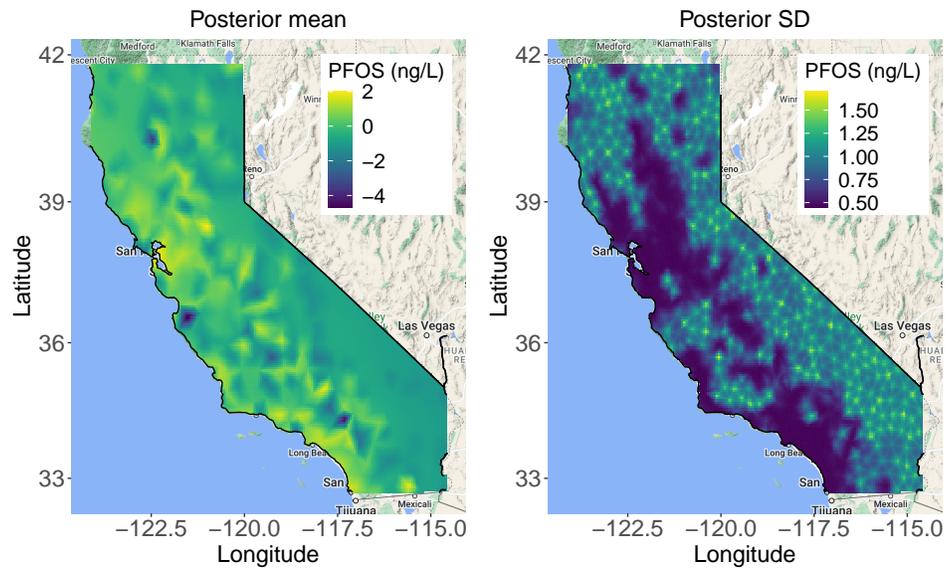


Figure 5: *Left*: The predicted surface map for $g(\text{PFOS}) = \log(1 + \log(1 + \text{PFOS}))$ concentration in the state of California. *Right*: The corresponding uncertainty estimates associated with the predictions for $g(\text{PFOS})$ across the pixels.

predict high values of PFOS concentration along the western part of the state, surpassing the EPA safety limit of approximately 0.45 in the transformed scale. Particularly notable are the elevated PFOS levels observed in and around Sonoma, Napa, Solano, Contra Costa, Alameda, San Francisco, and Santa Clara counties in central-west California, as well as in Los Angeles, Orange, and San Diego counties in southwest California.

Table 3: Table of estimates and standard deviations corresponding to the parameters $\beta_0 - \beta_2$ denoting the intercept and the 2 covariates, respectively, ϕ (the spatial range), τ (the precision) and r (ratio of partial sill to total variance).

	Estimates	Standard Deviation
β_0	-22.04	7.60
β_1	-0.23	0.08
β_2	-0.14	0.07
ϕ	0.07	0.02
τ	0.22	0.08
r	0.95	0.02

6 Discussion

We present a novel method to address the problem of modeling censored outcomes that are spatially correlated in big data settings. We observe that the proposed model scales nicely with an increased number of total observations and performs better than all other competing methods, even when nearly half of the observed data are censored. Despite being a fully Bayesian model, the runtime is moderate and better than the competing methods, highlighting its scalability, which combined with its demonstrated accuracy and precision

makes this method an optimally efficient approximate method for modeling large spatial data in the presence of (left-) censoring. The real data analysis demonstrated this further as the model achieved satisfactory mixing of three chains for a large dataset in only an hour, producing sensible prediction surfaces and uncertainty quantification.

However, the data presents specific challenges during modeling, which may also be considered as limitations of the proposed method. The predicted surface in the left panel of Figure 5 is very smooth towards the east-southeast end of California. This is expected, as we have very few observations around that area to inform our spatial process. This, while being non-desirable, makes sense and is in line with what one would expect to happen for such a dataset. We can not hope to manufacture information in the absence of observations and we do not, reflecting the consistency of statistical principles being adhered to here in our analysis. The map of prediction uncertainty (right panel of Figure 5) also reflects this. We have nearly zero uncertainty for the bulk of the region, where we observe numerous instances and have higher uncertainty whenever we move far from observations. Interestingly, we also notice a quilting pattern in the right panel of Figure 5. This is a byproduct of the computation mesh considered here along with the lack of observations in the east-southeast region.

Further consideration are therefore needed for choosing the mesh and smoothness parameters for fitting the model. We explain our choice of mesh in Section 2. But this process is ad-hoc in nature, and a concrete workflow for selecting a mesh would be greatly beneficial for users. We consider this as a plausible future research direction. Another development on both the software and methodological fronts would be to include fractional smoothness parameters in the model, which is currently restricted to integer smoothness (we use $\nu = 1$ for all our analyses). One possible approach would be to use the fractional rational approximations to the SPDE model (Bolin et al., 2024; Bolin and Kirchner, 2020). Combining

the theory for fractional approximations with the software should render additional model flexibility, and more well-suited to real data applications. Further developments for multivariate extensions of the model that can model multiple spatial processes, both having a mix of censored and uncensored observations, simultaneously, are underway.

Other than proposing a novel, scalable spatial model in its own right, the implications of this study extend beyond academic interest. By elucidating the spatial distribution of PFAS/PFOS contamination and its associated factors, we can inform targeted interventions, policy recommendations, and resource allocation to mitigate the impact of PFAS exposure on public health. Additionally, our research provides a framework, which can be adapted to analyze censored data in other environmental contexts, fostering a deeper understanding of complex contamination scenarios and enabling evidence-based decision-making.

Acknowledgement

Arnab Hazra is partially supported by the Indian Institute of Technology Kanpur and Rice University collaborative research grant under Award No. DOIR/2023246. Dipankar Bandyopadhyay acknowledges funding support from grant R21DE031879 from the US National Institutes of Health.

References

- Abrahamsen, P. and Benth, F. E. (2001). Kriging with inequality constraints. *Mathematical Geology*, 33(6):719–744.
- Andrews, D. Q. and Naidenko, O. V. (2020). Population-wide exposure to per-and polyflu-

- oroalkyl substances from drinking water in the United States. *Environmental Science & Technology Letters*, 7(12):931–936.
- Bakka, H., Rue, H., Fuglstad, G.-A., Riebler, A., Bolin, D., Illian, J., Krainski, E., Simpson, D., and Lindgren, F. (2018). Spatial modeling with r-inla: A review. *Wiley Interdisciplinary Reviews: Computational Statistics*, 10(6):e1443.
- Bivand, R. S., Pebesma, E. J., Gómez-Rubio, V., and Pebesma, E. J. (2008). *Applied spatial data analysis with R*, volume 747248717. Springer.
- Bolin, D. and Kirchner, K. (2020). The rational spde approach for gaussian random fields with general smoothness. *Journal of Computational and Graphical Statistics*, 29(2):274–285.
- Bolin, D., Simas, A. B., and Xiong, Z. (2024). Covariance-based rational approximations of fractional spdes for computationally efficient bayesian inference. *Journal of Computational and Graphical Statistics*, 33(1):64–74.
- Borouchaki, H. and Lo, S. (1995). Fast Delaunay triangulation in three dimensions. *Computer Methods in Applied Mechanics and Engineering*, 128(1-2):153–167.
- Ciarlet, P. G. (2002). *The finite element method for elliptic problems*. SIAM.
- Cisneros, D., Gong, Y., Yadav, R., Hazra, A., and Huser, R. (2023). A combined statistical and machine learning approach for spatial prediction of extreme wildfire frequencies and sizes. *Extremes*, 26(2):301–330.
- Cordner, A., De La Rosa, V. Y., Schaidler, L. A., Rudel, R. A., Richter, L., and Brown, P. (2019). Guideline levels for pfoa and pfos in drinking water: the role of scientific

- uncertainty, risk assessment decisions, and social factors. *Journal of exposure science & environmental epidemiology*, 29(2):157–171.
- De Oliveira, V. (2005). Bayesian inference and prediction of Gaussian random fields based on censored data. *Journal of Computational and Graphical Statistics*, 14(1):95–115.
- Delyon, B., Lavielle, M., and Moulines, E. (1999). Convergence of a stochastic approximation version of the EM algorithm. *Annals of statistics*, 27(1):94–128.
- Fridley, B. L. and Dixon, P. (2007). Data augmentation for a Bayesian spatial model involving censored observations. *Environmetrics: The official journal of the International Environmetrics Society*, 18(2):107–123.
- Gelfand, A. E. and Schliep, E. M. (2016). Spatial statistics and gaussian processes: A beautiful marriage. *Spatial Statistics*, 18:86–104.
- Hazra, A., Huser, R., and Bolin, D. (2021). Realistic and fast modeling of spatial extremes over large geographical domains. *arXiv preprint arXiv:2112.10248*.
- Hazra, A., Reich, B. J., Shaby, B. A., and Staicu, A.-M. (2018). A semiparametric spatiotemporal bayesian model for the bulk and extremes of the fosberg fire weather index. *arXiv preprint arXiv:1812.11699*.
- Hepburn, E., Madden, C., Szabo, D., Coggan, T. L., Clarke, B., and Currell, M. (2019). Contamination of groundwater with per-and polyfluoroalkyl substances (pfas) from legacy landfills in an urban re-development precinct. *Environmental Pollution*, 248:101–113.
- Hopke, P. K., Liu, C., and Rubin, D. B. (2001). Multiple imputation for multivariate data

- with missing and below-threshold measurements: time-series concentrations of pollutants in the arctic. *Biometrics*, 57(1):22–33.
- Hosmer, D. W., Lemeshow, S., and May, S. (2008). Applied survival analysis. *Wiley Series in Probability and Statistics*, page 60.
- Lindgren, F., Bolin, D., and Rue, H. (2022). The spde approach for gaussian and non-gaussian fields: 10 years and still running. *Spatial Statistics*, 50:100599.
- Lindgren, F. and Rue, H. (2015). Bayesian spatial modelling with r-inla. *Journal of statistical software*, 63(19).
- 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.
- Militino, A. F. and Ugarte, M. D. (1999). Analyzing censored spatial data. *Mathematical Geology*, 31(5):551–561.
- Moran, J. E., Hudson, G. B., Eaton, G. F., and Leif, R. (2005). California GAMA Program: Groundwater Ambient Monitoring and Assessment Results for the Sacramento Valley and Volcanic Provinces of Northern California. Technical report, Lawrence Livermore National Lab.(LLNL), Livermore, CA (United States).
- Ordoñez, J. A., Bandyopadhyay, D., Lachos, V. H., and Cabral, C. R. (2018). Geostatistical estimation and prediction for censored responses. *Spatial statistics*, 23:109–123.
- Rathbun, S. L. (2006). Spatial prediction with left-censored observations. *Journal of agricultural, biological, and environmental statistics*, 11(3):317–336.

- Rosen, G. (1955). Problems in the application of statistical analysis to questions of health: 1700-1880. *Bulletin of the History of Medicine*, 29(1):27–45.
- Rue, H. and Held, L. (2005). *Gaussian Markov random fields: theory and applications*. CRC press.
- Sahoo, I., Guinness, J., and Reich, B. J. (2023). Estimating atmospheric motion winds from satellite image data using space-time drift models. *Environmetrics*, page e2818.
- Sahoo, I. and Hazra, A. (2021). Contamination mapping in bangladesh using a multivariate spatial bayesian model for left-censored data. *arXiv preprint arXiv:2106.15730*.
- Schelin, L. and Sjöstedt-de Luna, S. (2014). Spatial prediction in the presence of left-censoring. *Computational Statistics & Data Analysis*, 74:125–141.
- Schulz, E., Speekenbrink, M., and Krause, A. (2018). A tutorial on Gaussian process regression: Modelling, exploring, and exploiting functions. *Journal of Mathematical Psychology*, 85:1–16.
- Sedda, L., Atkinson, P. M., Barca, E., and Passarella, G. (2012). Imputing censored data with desirable spatial covariance function properties using simulated annealing. *Journal of geographical systems*, 14(3):265–282.
- Smalling, K. L., Romanok, K. M., Bradley, P. M., Morriss, M. C., Gray, J. L., Kanagy, L. K., Gordon, S. E., Williams, B. M., Breitmeyer, S. E., Jones, D. K., et al. (2023). Per-and polyfluoroalkyl substances (PFAS) in United States tapwater: Comparison of underserved private-well and public-supply exposures and associated health implications. *Environment International*, page 108033.

- Stein, M. L. (1992). Prediction and inference for truncated spatial data. *Journal of Computational and Graphical Statistics*, 1(1):91–110.
- Tadayon, V. (2017). Bayesian analysis of censored spatial data based on a non-Gaussian model. *arXiv preprint arXiv:1706.05717*.
- Vecchia, A. V. (1988). Estimation and model identification for continuous spatial processes. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 50(2):297–312.
- Wang, Z., DeWitt, J. C., Higgins, C. P., and Cousins, I. T. (2017). A never-ending story of per-and polyfluoroalkyl substances (pfass)?
- Wiens, A., Nychka, D., and Kleiber, W. (2020). Modeling spatial data using local likelihood estimation and a matérn to spatial autoregressive translation. *Environmetrics*, 31(6):e2652.
- Yadav, R., Huser, R., and Opitz, T. (2019). Spatial hierarchical modeling of threshold exceedances using rate mixtures. *Environmetrics*, page e2662.
- Zhang, L., Shaby, B. A., and Wadsworth, J. L. (2021). Hierarchical transformed scale mixtures for flexible modeling of spatial extremes on datasets with many locations. *Journal of the American Statistical Association*, pages 1–13.