

Robust Distributed Learning of Functional Data From Expensive Simulators through Data Sketching

R. Jacob Andros

Department of Statistics, Texas A&M University

Rajarshi Guhaniyogi

Department of Statistics, Texas A&M University

Devin Francom

Los Alamos National Laboratories

Donatella Paqualini

Los Alamos National Laboratories

May 1, 2024

Abstract

In environmental studies, realistic simulations are essential for understanding complex systems. Statistical emulation with Gaussian processes (GPs) in functional data models have become a standard tool for this purpose. Traditional centralized processing of such models requires substantial computational and storage resources, leading to emerging distributed Bayesian learning algorithms that partition data into shards for distributed computations. However, concerns about the sensitivity of distributed inference to shard selection arise. Instead of using data shards, our approach employs multiple random matrices to create random linear projections, or sketches, of the dataset. Posterior inference on functional data models is conducted using random data sketches on various machines in parallel. These individual inferences are combined across machines at a central server. The aggregation of inference across random matrices makes our approach resilient to the selection of data sketches, resulting in *robust distributed Bayesian learning*. An important advantage is its ability to maintain the privacy of sampling units, as random sketches prevent the recovery of raw data. We highlight the significance of our approach through simulation examples and showcase the performance of our approach as an emulator us-

ing surrogates of the Sea, Lake, and Overland Surges from Hurricanes (SLOSH) simulator—an important simulator for government agencies.

Keywords: Data Sketching; Distributed Inference; Gaussian process; Low-rank models; Parallel computing; SLOSH Emulator.

1 Introduction

In environmental applications, scientific analysis often relies on high-resolution spatiotemporal physics-based simulations. For instance, various physics-based simulators for climate and weather systems generate detailed spatial simulations based on input attributes [Petersen et al., 2019, Borge et al., 2008]. This manuscript specifically centers on the Sea, Lake, and Overland Surges from Hurricanes (SLOSH) simulator [Jelesnianski et al., 1992], developed by the National Weather Service for operational hurricane monitoring and response. Configured with a specific area of interest, SLOSH takes a hurricane’s track and a vector of hurricane-related attributes as input, and produces a spatio-temporal grid of storm surge, or a measure of flood depth above normal levels (e.g., water depth above ground level for land spatial locations). In our study, we focus on determining the maximum flood depth at each spatial location for a given hurricane, with emphasis on flooding of electrical substations. Our study area encompasses a part of the Delaware Bay, which separates the southern end of New Jersey from the northern side of Delaware (see Figure 1).

To characterize the spatial variation of maximum flood depth and its relationships with input attributes, as well as to address input uncertainty, one approach is to treat the simulator as a black box and explore it by running it with various settings [Sacks and Welch, 1989]. However, this approach involves conducting numerous runs of the simulator with various settings, and each of these runs comes with substantial computational costs. As a result, the process of simulating multiple runs is significantly expensive and often infeasible in terms of computational resources. A solution to these challenges often involves creating a statistical surrogate model, or emulator, for the simulator. Once trained, these emulators provide rapid predictions for new input settings, facilitating studies of model response and parameter uncertainty. Hutchings et al. [2023] compare four emulation methods for SLOSH, though a significant numbers of model runs were necessary for satisfactory performance of the framework. In contrast, the focus of this work is to build a SLOSH emulator with *very few model runs* while taking special care to treat the spatial association of the

water depth level and allowing for part of data stored in different centers or servers locally.

Considering flood depth as functional data across spatial locations, Gaussian processes (GPs) can offer a suitable framework for accounting for the correlation between outputs at different locations. GPs have gained prominence as effective predictive tools in the domain of computer experiments, as they can significantly reduce the computational burden associated with running simulations while maintaining flexibility and allowing for comprehensive uncertainty analysis encompassing both parameter and code uncertainties [Kennedy and O’Hagan, 2001, Salter and Williamson, 2016]. In the specific context of the SLOSH simulator, corresponding to an input vector $\mathbf{z}_s \in \mathbb{R}^p$ in the s th simulation ($s = 1, \dots, S$), we observe noisy outputs $y_s(\mathbf{u}_1), \dots, y_s(\mathbf{u}_n) \in \mathbb{R}$ from the simulator at d -dimensional index points $\mathbf{u}_1, \dots, \mathbf{u}_n \in \mathcal{U} \subseteq \mathcal{R}^d$, respectively. These input attributes remains same corresponding to the output at all index points, and is referred to as *global attributes* throughout the article. We observe an additional q -dimensional covariates $\mathbf{x}(\mathbf{u}_1), \dots, \mathbf{x}(\mathbf{u}_n)$ at each index point, referred to as *local attributes*. With global and local attributes, one can build a GP emulator described by the following model,

$$y_s(\mathbf{u}_i) = \mathbf{z}_s^T \boldsymbol{\gamma} + \mathbf{x}(\mathbf{u}_i)^T \boldsymbol{\beta} + w_s(\mathbf{u}_i) + \epsilon_s(\mathbf{u}_i), \quad i = 1, \dots, n, \quad s = 1, \dots, S, \quad (1)$$

where $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$ are p and q dimensional coefficients corresponding to the global and local attributes, respectively. The term $\epsilon_s(\mathbf{u}_i)$ denotes idiosyncratic errors, assumed to be independently and identically distributed with a Normal distribution with mean 0 and variance τ^2 , for simplicity. The unknown function $w_s(\cdot)$, accounting for variability across the domain \mathcal{U} is considered a realization from a GP with a mean of 0 and a covariance kernel $\kappa(\cdot, \cdot)$, independently over the simulations $s = 1, \dots, S$.

Performing Bayesian inference for the model (1) with a Gaussian process prior on $w_s(\mathbf{u})$ becomes computationally impractical for a large sample size due to the $O(n^3)$ computational cost and the $O(n^2)$ storage cost associated with estimating $w_s(\mathbf{u})$. The domain of modeling high-dimensional dependent functional data has seen substantial growth in the last decade, which has been largely adapted and built upon scalable spatial models. While the extensive literature in this area cannot be fully covered here, comprehensive reviews can be found in Heaton et al. [2019].

This literature has largely operated within a centralized data processing framework, where all

data is stored and processed at a central location. However, an alternative approach involves distributed Bayesian learning for functional data [Guhaniyogi and Banerjee, 2018, Guhaniyogi et al., 2023, Guhaniyogi and Banerjee, 2019, Guhaniyogi et al., 2022]. This approach extends the scalability of existing well-established algorithms to estimate (1), without necessitating the development of new algorithms or software. In essence, this methodology partitions the data into a multitude of shards, wherein a suitable functional data model such as (1) is fitted to each shard to draw parallel posterior inferences. These individual inferences are then combined at a central server, reducing computation and storage requirements on the server itself. Distributed inference eliminates the need for extensive raw data exchange and communication between the central server and processors during statistical analysis. Consequently, this leads to lower latency and reduced communication traffic. Furthermore, as each processor analyzes a smaller data shard, the computation and storage becomes more efficient, contributing to overall faster inference. Significantly, numerous state-of-the-art methods in the analysis of large functional data come from spatial statistics research, though they crucially use smaller dimensionality of \mathcal{U} in their design, as exemplified by methods that utilize nearest neighbors Datta et al. [2016]. These methods may not be entirely suitable when \mathcal{U} is higher-dimensional since all points in higher dimensions are approximately equidistant. In contrast, distributed approaches offer scalability without the need to leverage the reduced dimensionality of \mathcal{U} , thus making them suitable for the seamless application to functional data models with higher dimensional \mathcal{U} . Nevertheless, it is important to note that inference in distributed Bayesian learning can become sensitive to the choice of data subsets, as highlighted by simulation studies in Guhaniyogi et al. [2023].

We will address this issue by introducing a distributed Bayesian inferential framework for functional data modeling that employs Bayesian data sketching techniques [Vempala, 2005, Halko et al., 2011, Mahoney et al., 2011, Woodruff et al., 2014, Guhaniyogi and Dunson, 2015, 2016]. The concept of data sketching involves compressing the complete dataset before applying a model, facilitating more efficient computation and storage. While data sketching methods have gained popularity in high-dimensional and penalized regression contexts with large datasets [Zhang et al., 2013, Chen et al., 2015, Dobriban and Liu, 2018, Drineas et al., 2011, Ahfock et al., 2017, Huang, 2018], their application to resolving computational challenges in Bayesian high-dimensional and functional data regression remains limited, with a few notable exceptions [Guhaniyogi and Scheffler, 2021].

The overarching outline of this framework is as follows. First, we generate H random matrices, each with dimensions $m \times n$, where m is significantly smaller than n , following the literature on data sketching, and construct H random linear sketches of both the response vector and the predictor matrix. Secondly, we undertake parallel posterior computations on these k data sketches using the functional data model (1). In this process, the likelihood in the model is raised to the power of n/m for each sketch. This manipulation ensures that the variance of each sketch’s posterior distribution is of the same order of n as that of the full data posterior distribution. This set of pseudo posterior distributions, constructed for each data sketch, is referred to as the “sketched posterior.” Thirdly, we calculate the Wasserstein mean of these k sketched posterior distributions, yielding a single probability distribution termed the “collaborative sketched posterior” distribution. This collaborative sketched posterior serves as a computationally tractable approximation to the full posterior distribution.

The proposed approach addresses a number of important issues on distributed Bayesian learning simultaneously. First, unlike current distributed Bayesian learning methods that perform inference conditional on a fixed choice of data shards, our proposed approach takes a distinct route. Instead of being confined to a single set of data shards, we compute the Wasserstein mean of sketched posteriors obtained using different random matrices. This innovative technique mitigates the sensitivity of inference stemming from the selection of specific random matrix in the construction of data sketches. As a result, we introduce a novel category of distributed learning paradigm, referred to as the “robust distributed learning,” that is resilient to the influences of data shard choices. Second, a crucial aspect of the proposed methodology is its ability to uphold the privacy of sampling units throughout the inference. This is achieved by revealing only lower-dimensional sketches to the analysts which are designed in such a way that the mutual information between them and the full data converges to zero as sample size becomes large, rendering the recovery of the full data from data sketches practically impossible [Zhou et al., 2008]. Importantly, the framework allows constructing a sketch for the full data from sketches of subsets of data stored privately in different research centers.

The rest of this article progresses as follows. Section 2 describes the data simulated from SLOSH simulator. Section 3 describes the robust distributed learning approach in detail. Implementation of the approach on our own simulated data, followed by the SLOSH simulator analysis, are presented

Predictor	Lower	Upper	Units
Heading	204.0349	384.0244	degrees
Velocity	0	40	knots
Latitude	38.32527	39.26811	degrees
Pressure	930	980	millibars
Sea level rise (2100)	-20	350	cm

Table 1: Parameters varied in SLOSH simulations.

in Sections 4 and 5, respectively. Finally, Section 6 concludes the paper with an eye towards future work.

2 SLOSH Emulator Data

Storm surge models simulate floodwater depth resulting from hurricanes, and are used for emergency response, planning, and research. The Sea, Lake, and Overland Surges from Hurricanes (SLOSH) simulator (Jelesnianski et al., 1992) is one such model developed by the National Weather Service. Various other storm surge models exist, with more complex models requiring extreme computation time to run. While SLOSH is not prohibitively expensive, our purpose in this paper is to develop emulator for expensive models often encountered in national laboratories that can only be run at a handful of times. Therefore, our dataset comprises an ensemble of only 10 simulations executed using the SLOSH simulator, representing 10 distinct simulated storms. Each storm within this ensemble is characterized by a unique combination of five input parameters or global attributes. Four of these attributes describe hurricane characteristics when the hurricane makes landfall, including the heading of the eye, the velocity of the eye, the latitude of the eye, and the minimum air pressure experienced. The fifth attribute is the projected sea level rise for the year 2100, which of course is uncertain. General ranges for these parameters are given in Table 1, from which 10 simulations were generated.

Our SLOSH simulations predict hurricane-induced flooding in the southern tip of New Jersey (see Figure 1). Since the major focus is study the damage to electrical power stations, and majority of the power stations are far enough in-land, we only focus on the functional data over $n = 49719$ spatial locations. A map of of output from one SLOSH model run is shown in Figure 1.

In addition, elevation data is available for each spatial location of the electrical power stations. Recognizing the significance of this attribute in assessing the maximum flood height at a given

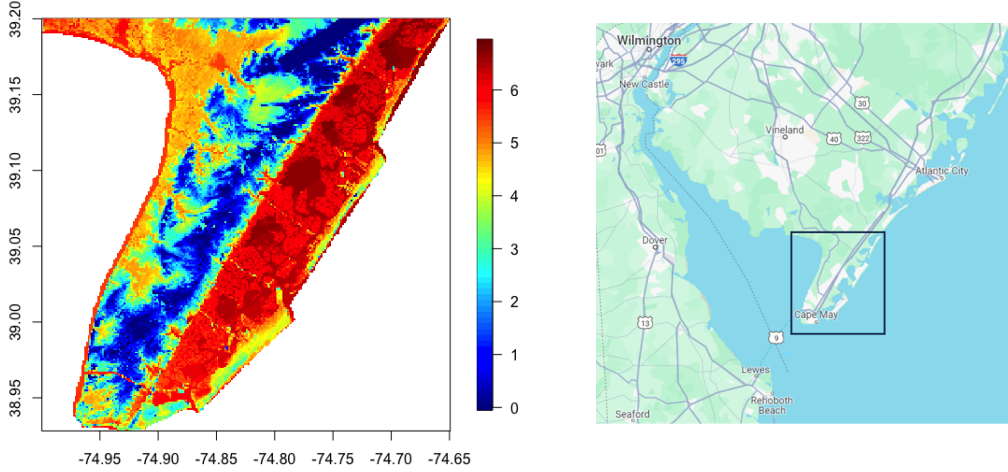


Figure 1: Output from one SLOSH model run, side-by-side with the location of the peninsula in Google Maps. Here, x and y axes represent longitude and latitude, respectively.

location, we incorporate elevation as a local attribute in our analysis.

Power stations within this region are typically designed to withstand flooding up to a level of *four feet*, with higher levels leading to catastrophic damage. Consequently, our focus is on evaluating the emulator’s capability to reliably predict whether a storm surge has exceeded the four-foot mark. This predictive information is crucial for making informed decisions about potential interventions, such as the need to initiate a station shutdown in anticipation of an approaching storm.

3 Proposed Approach: A Collaborative Sketching Framework

Let $\mathbf{y}_s = (y_s(\mathbf{u}_1), \dots, y_s(\mathbf{u}_n))^T$ and $\boldsymbol{\epsilon}_s = (\epsilon_s(\mathbf{u}_1), \dots, \epsilon_s(\mathbf{u}_n))^T$ be the vector of responses and errors collected over all index points for the s th simulation. Further, assume \mathbf{X} is an $n \times q$ matrix with its i th row given by $\mathbf{x}(\mathbf{u}_i)^T$ and $\mathbf{w}_s = (w_s(\mathbf{u}_1), \dots, w_s(\mathbf{u}_n))^T$ is the vector of functional random effects at all the index points corresponding to the s th simulation. The model (1) yields the Gaussian linear mixed model

$$\mathbf{y}_s = (\mathbf{1}_n \otimes \mathbf{z}_s^T)\boldsymbol{\gamma} + \mathbf{X}\boldsymbol{\beta} + \mathbf{w}_s + \boldsymbol{\epsilon}_s, \quad \boldsymbol{\epsilon}_s \sim N(0, \tau^2 \mathbf{I}_n), \quad s = 1, \dots, S, \quad (2)$$

where $\mathbf{1}_n$ refers to the n -dimensional vector of ones.

Specifying Gaussian process prior on the unknown function $w_s(\mathbf{u}) \sim GP(0, \sigma^2 \kappa(\cdot, \cdot; \theta))$ independently over $s = 1, \dots, S$ leads to $\mathbf{w}_s \stackrel{i.i.d.}{\sim} N(\mathbf{0}, \sigma^2 \mathbf{K}(\theta))$, where $\mathbf{K}(\theta)$ is an $n \times n$ matrix with its

(i, j) th entry given by $\kappa(\mathbf{u}_i, \mathbf{u}_j; \theta)$. Here θ and σ^2 are referred to the length-scale parameter and variance parameter, respectively, for the functional regression. Assuming response over simulators are independent, a customary Bayesian hierarchical model is formulated from (2) as follows

$$p(\boldsymbol{\gamma}, \boldsymbol{\beta}, \tau^2, \sigma^2) \times \prod_{s=1}^S N(\mathbf{y}_s | (\mathbf{1}_n \otimes \mathbf{z}_s^T) \boldsymbol{\gamma} + \mathbf{X} \boldsymbol{\beta}, \sigma^2 \mathbf{K}(\theta) + \tau^2 \mathbf{I}_n), \quad (3)$$

where $p(\boldsymbol{\gamma}, \boldsymbol{\beta}, \tau^2, \sigma^2)$ is the prior distribution on the parameters $(\boldsymbol{\gamma}, \boldsymbol{\beta}, \tau^2, \sigma^2)$. The hierarchical model (3) fixes the length-scale parameter θ . While the data can inform about this parameter, it is inconsistently estimable for the general Matern class of correlation functions [Zhang, 2004] often resulting in poorer convergence. Hence, recent studies proposed an approach where inference is drawn using hierarchical models like (3) by fitting the model across several fixed values of θ and subsequently combining the inferences tailored to draw achieve a specific inferential goal [Zhang et al., 2023]. We will adopt a similar strategy of fixing θ during model fitting with every data sketch. Even with a fixed θ , the Bayesian computation of (3) involves inverting the $n \times n$ matrix $\sigma^2 \mathbf{K}(\theta) + \tau^2 \mathbf{I}_n$ which is infeasible for the simulator data in Section 2.

To circumvent the computational issues posed by $\mathbf{K}(\theta)$ induced from full GP prior on w_s , variants of GP prior are proposed on w_s , leading to efficient computation of $\mathbf{K}(\theta)^{-1}$. This includes low-rank GP priors and sparse GP priors, among others. This article employs modified predictive process (MPP) prior [Finley et al., 2009] and nearest neighbor GP (NNGP) prior [Datta et al., 2016] as the representative low-rank and sparse GP priors, respectively. Let $\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_{n_{knot}}$ be a set of ‘‘knot’’ points randomly selected from \mathcal{U} , with $n_{knot} \ll n$. Let $\tilde{\mathbf{K}}(\theta)$ and $\tilde{\tilde{\mathbf{K}}}(\theta)$ be $n_{knot} \times n_{knot}$ and $n \times n_{knot}$ matrix with the (i, j) th entries given by $\kappa(\tilde{\mathbf{u}}_i, \tilde{\mathbf{u}}_j; \theta)$ and $\kappa(\mathbf{u}_i, \tilde{\mathbf{u}}_j; \theta)$, respectively. The MPP specifies $\mathbf{K}(\theta)$ as $\mathbf{K}(\theta) = \tilde{\tilde{\mathbf{K}}}(\theta) \tilde{\mathbf{K}}(\theta)^{-1} \tilde{\tilde{\mathbf{K}}}(\theta) + \mathbf{D}(\theta)$, where $\mathbf{D}(\theta)$ is an $n \times n$ diagonal matrix with the i th diagonal entry $1 - \tilde{\boldsymbol{\kappa}}(\mathbf{u}_i; \theta)^T \tilde{\mathbf{K}}(\theta)^{-1} \tilde{\boldsymbol{\kappa}}(\mathbf{u}_i; \theta)$, $\tilde{\boldsymbol{\kappa}}(\mathbf{u}_i; \theta) = (\kappa(\mathbf{u}_i, \tilde{\mathbf{u}}_1; \theta), \dots, \kappa(\mathbf{u}_i, \tilde{\mathbf{u}}_{n_{knot}}; \theta))^T$. In contrast, NNGP sparsifies $\mathbf{K}(\theta)^{-1}$ by replacing $w_s(\mathbf{u}_i) | w_s(\mathbf{u}_1), \dots, w_s(\mathbf{u}_{i-1})$ with $w_s(\mathbf{u}_i) | w_s(\mathbf{U}_{i,nn})$, where $\mathbf{U}_{i,nn}$ is the few nearest neighbors of \mathbf{u}_i [Datta et al., 2016]. While MPP and Sparse GP offers model based efficient alternatives to the full GP, this article considers an approach orthogonal to them. Specifically, it considers a two-stage distributed Bayesian learning approach, referred to as the collaborative sketching framework, which enables scaling up full-GP, as well as its computationally efficient alternatives (e.g., MPP and Sparse-GP), as elaborated in the subsequent sections.

3.1 Construction of Sketched Posteriors

At the first stage, we construct multiple low-dimensional random linear mapping or “sketches” of the full data and fit model (3) with these data sketches in parallel. To elaborate on this, let Φ_1, \dots, Φ_H are $m \times n$ dimensional sketching matrices with random entries encoding random linear mapping of the data to lower dimensions, with $m \ll n$. For $h = 1, \dots, H$, the sketching matrix Φ_h is applied to \mathbf{y}_s , \mathbf{X} and $\mathbf{Z}_s = \mathbf{1}_n \otimes \mathbf{z}_s^T$ to construct the $m \times 1$ sketched response vector $\mathbf{y}_{s, \Phi_h} = \Phi_h \mathbf{y}_s$ and sketched predictor matrices $\mathbf{X}_{\Phi_h} = \Phi_h \mathbf{X}$ and $\mathbf{Z}_{s, \Phi_h} = \Phi_h \mathbf{Z}_s$. We will return to the specification of Φ_h , which, of course, will be crucial for relating the inference from the compressed data with the full model. For now assuming that we have fixed Φ_h and a fixed value of the length-scale parameter θ_h , we define the sketched posterior distribution of the model parameters by

$$\Pi_h(\boldsymbol{\gamma}, \boldsymbol{\beta}, \tau^2, \sigma^2 | \Phi_h, \theta_h) = \frac{p(\boldsymbol{\gamma}, \boldsymbol{\beta}, \tau^2, \sigma^2) \times \prod_{s=1}^S L_{s,h}(\sigma^2, \tau^2, \boldsymbol{\gamma}, \boldsymbol{\beta})^{n/m}}{\int p(\boldsymbol{\gamma}, \boldsymbol{\beta}, \tau^2, \sigma^2) \times \prod_{s=1}^S L_{s,h}(\sigma^2, \tau^2, \boldsymbol{\gamma}, \boldsymbol{\beta})^{n/m} d\sigma^2 d\tau^2 d\boldsymbol{\gamma} d\boldsymbol{\beta}}, \quad (4)$$

where $L_{s,h}(\sigma^2, \tau^2, \boldsymbol{\gamma}, \boldsymbol{\beta}) = N(\mathbf{y}_{s, \Phi_h} | \mathbf{Z}_{s, \Phi_h} \boldsymbol{\gamma} + \mathbf{X}_{\Phi_h} \boldsymbol{\beta}, \sigma^2 \Phi_h \mathbf{K}(\theta_h) \Phi_h^T + \tau^2 \mathbf{I}_m)$ denotes the likelihood of data sketch from the s th simulator run, after fixing $\{\Phi, \theta\}$ at $\{\Phi_h, \theta_h\}$. Here $\mathbf{K}(\theta_h)$ can be induced from full-GP or its computationally efficient variants as discussed before. The likelihood $L_{s,h}$ differs from the one obtained by applying Φ_h to the likelihood in (2) because the error distribution in $L_{s,h}$ is retained as the usual noise distribution without any influence of Φ_h . Thus, the likelihood $L_{s,h}$ corresponds to the likelihood derived from a model similar to (2), but it is applied to the sketched dataset $\mathbf{y}_{s, \Phi_h}, \mathbf{X}_{\Phi_h}, \mathbf{Z}_{s, \Phi_h}$. Employing a Φ_h -transformed model on (2), where noise distribution is transformed by $\Phi_h \epsilon_s$ will not deliver the computational benefits.

The modification of likelihood to yield the sketched posterior density in (4) is referred to as stochastic approximation [Guhaniyogi et al., 2023, 2022]. Conceptually, raising the likelihood to the power of n/m can be viewed as replicating the sketched data n/m times, so stochastic approximation accounts for the fact that the h th sketched posterior distribution Π_h conditions on an m -dimensional random sketch of the full data and ensures that its variance aligns in the same order (as a function of n) as that of the full data posterior distribution.

3.1.1 Computation of sketched posteriors

Assume an $IG(a_\tau, b_\tau)$, $IG(a_\sigma, b_\sigma)$ and $N(\mathbf{0}, \mathbf{I})$ prior distributions on τ^2, σ^2 and $\boldsymbol{\lambda} = (\boldsymbol{\beta}^T, \boldsymbol{\gamma}^T)^T$. With the proposed stochastic approximation, the posterior computation of sketched posterior Π_h with the h th data sketch involves Markov Chain Monte Carlo (MCMC) which cycles through drawing samples from the following distributions:

- Sample $(\boldsymbol{\beta}^T, \boldsymbol{\gamma}^T)^T | - \sim N(\boldsymbol{\mu}_{\boldsymbol{\beta}, \boldsymbol{\gamma}}, \boldsymbol{\Sigma}_{\boldsymbol{\beta}, \boldsymbol{\gamma}})$, where $\boldsymbol{\Sigma}_{\boldsymbol{\beta}, \boldsymbol{\gamma}} = \left\{ (m/n) \sum_{s=1}^S \mathbf{A}_s^T \boldsymbol{\Sigma}^{-1} \mathbf{A}_s + \mathbf{I}_p \right\}^{-1}$, $\boldsymbol{\mu}_{\boldsymbol{\beta}, \boldsymbol{\gamma}} = \boldsymbol{\Sigma}_{\boldsymbol{\beta}, \boldsymbol{\gamma}} \left\{ (m/n) \sum_{s=1}^S \mathbf{A}_s^T \boldsymbol{\Sigma}^{-1} \mathbf{y}_{s, \Phi_h} \right\}$. Here $\mathbf{A}_s = [\mathbf{X}_{\Phi_h} : \mathbf{Z}_{s, \Phi_h}]$ is an $m \times n$ matrix and $\boldsymbol{\Sigma} = (\boldsymbol{\Phi}_h \mathbf{K}(\theta_h) \boldsymbol{\Phi}_h^T + \tau^2 \mathbf{I}_m)$ is an $m \times m$ matrix. This step incurs a computation complexity of $O(m^3)$, since $\boldsymbol{\Sigma}$ is an $m \times m$ covariance matrix that needs to be inverted.
- Sample σ^2 and τ^2 through Metropolis-Hasting sampler. Since θ_h is kept fixed throughout the analysis of Π_h , we need to compute $\boldsymbol{\Phi}^T \mathbf{K}(\theta_h) \boldsymbol{\Phi}$ only once which leads to substantial computational benefit. Additionally, the strategy needs storage of the $m \times m$ matrix $\boldsymbol{\Phi}^T \mathbf{K}(\theta_h) \boldsymbol{\Phi}$ instead of the full data covariance matrix, which reduces the storage cost from $O(n^2)$ to $O(m^2)$.

Next we focus on drawing predictive inference at n^* index points $\mathcal{U} = \{\mathbf{u}_1^*, \dots, \mathbf{u}_{n^*}^*\}$ for a new simulator run. To this end, let \mathbf{y}^* and \mathbf{X}^* be $n^* \times 1$ dimensional response vector and $n^* \times q$ dimensional matrix of local attributes at n^* index points, respectively. Assume \mathbf{z}^* is the input attributes corresponding to the new simulator run, $\mathbf{Z}^* = \mathbf{1}_{n^*} \otimes \mathbf{z}^{*T}$, and $\mathbf{w}_s^* = (w_s(\mathbf{u}_1^*), \dots, w_s(\mathbf{u}_{n^*}^*))^T$. The posterior predictive density of \mathbf{y}^* is given by

$$\begin{aligned}
 p(\mathbf{y}^* | \mathbf{X}^*, \mathbf{Z}^*, \mathbf{X}_{\Phi_h}, \mathbf{Z}_{s, \Phi_h}, \mathbf{y}_{s, \Phi_h}) &= \int f(\mathbf{y}^* | \mathbf{X}^*, \mathbf{Z}^*, \boldsymbol{\gamma}, \boldsymbol{\beta}, \sigma^2, \tau^2) \Pi_h(\boldsymbol{\gamma}, \boldsymbol{\beta}, \sigma^2, \tau^2 | \boldsymbol{\Phi}_h, \theta_h) \\
 f(\mathbf{y}^* | \mathbf{X}^*, \mathbf{Z}^*, \boldsymbol{\gamma}, \boldsymbol{\beta}, \sigma^2, \tau^2) &= N(\mathbf{y}^* | \boldsymbol{\mu}^* + \boldsymbol{\Sigma}_{21} \boldsymbol{\Sigma}_{11}^{-1} (\mathbf{y} - \boldsymbol{\mu}), \boldsymbol{\Sigma}_{22} - \boldsymbol{\Sigma}_{21} \boldsymbol{\Sigma}_{11}^{-1} \boldsymbol{\Sigma}_{12}) \\
 \boldsymbol{\mu}^* &= (\mathbf{1}_{n^*} \otimes \mathbf{z}^*) \boldsymbol{\gamma} + \mathbf{X}^* \boldsymbol{\beta}, \quad \boldsymbol{\mu} = (\mathbf{1}_n \otimes \mathbf{z}) \boldsymbol{\gamma} + \mathbf{X} \boldsymbol{\beta} \\
 \boldsymbol{\Sigma}_{22} &= \text{Var}(\mathbf{w}_s^*), \quad \boldsymbol{\Sigma}_{11} = \text{Var}(\mathbf{w}_s), \quad \boldsymbol{\Sigma}_{12} = \boldsymbol{\Sigma}_{21}^T = \text{Cov}(\mathbf{w}_s, \mathbf{w}_s^*).
 \end{aligned} \tag{5}$$

We utilize composition sampling to generate MCMC samples from the posterior predictive density (5) of \mathbf{y}^* . The entire analysis of Π_1, \dots, Π_H will be distributed across H separate CPUs, enabling parallel processing for enhanced efficiency.

3.1.2 Choice of Sketching Matrices

To specify the matrix Φ_h , we adopt the concept of data oblivious Gaussian sketching as proposed in Sarlos [2006]. This involves independently selecting elements of Φ_h from a normal distribution with mean 0 and variance $1/n$, and then keeping them fixed. The dominant computational effort required for generating the sketched data using Gaussian sketches follows a time complexity of $O(mn^2)$. While there are alternative data oblivious methods available for efficiently sketching Φ_h , like the Hadamard sketch [Ailon and Chazelle, 2009] and the Clarkson-Woodruff sketch [Clarkson and Woodruff, 2017], these options hold less relevance in Bayesian contexts. This is due to the fact that the computational time needed for the operation in (4) is significantly greater than that needed for constructing the sketching matrix itself. The sketched data serves as a surrogate for the Bayesian estimation of model coefficients. Given that the count of sketched records is much smaller than the total records in the complete data matrix, the process of fitting the model becomes both computationally efficient and resource-friendly. This efficiency extends to storage requirements as well as the count of floating-point operations (flops) needed, as outlined in Section 3.1.1.

3.1.3 Data Privacy

Crucially, even when Φ_h is known, the linear systems $\Phi_h \mathbf{y}_s$, $\Phi_h \mathbf{X}$, and $\Phi_h \mathbf{Z}_s$ are significantly under-determined due to the substantial difference in the dimensions m and n . This results in a safeguarding of the privacy of the data samples. To evaluate the privacy implications in terms of information theory, an upper bound on the average mutual information per unit can be used, denoted as $I(\mathbf{X}_{\Phi_h}, \mathbf{X})/nq$. It can be shown that the supremum of $I(\mathbf{X}_{\Phi_h}, \mathbf{X})/nq$ is bounded by $O(m/n)$ [Zhou et al., 2008], where the supremum is taken across all possible distributions of \mathbf{X} . As m grows at a much slower rate than n , and $n \rightarrow \infty$, the supremum of the average mutual information approaches 0. This implies that, intuitively, the compressed data reveals no more information about the original data than what could be derived from an independent sample. It's important to note that this bound is derived under the assumption that Φ_h is known. In practical scenarios, only \mathbf{X}_{Φ_h} (along with \mathbf{y}_{s, Φ_h} and \mathbf{Z}_{s, Φ_h}) would be disclosed to the analyst, without revealing the actual matrix Φ_h itself. Consequently, the privacy imposed through sketching turns out to be more stringent than what is implicated by this theoretical outcome. It is worth emphasizing that these privacy implications stem from the information theory and are distinct from the broader notion of

data privacy explored through concepts like ϵ -differential privacy.

3.1.4 Distributed Storage in Research Centers

An essential benefit of the proposed framework lies in its capacity to analyze data collectively from research centers that are restricted from sharing their data with one another. Consider a scenario with J centers denoted as $Center_1, \dots, Center_J$, where the j th center holds data for n_j indexing points, $n = \sum_{j=1}^J n_j$. In this context, the j th center privately stores an n_j -dimensional response vector $\mathbf{y}_s^{(j)}$ and an $n_j \times p$ -dimensional matrix $\mathbf{X}^{(j)}$ for local attributes, such that $\mathbf{y}_s = (\mathbf{y}_s^{(1)}, \dots, \mathbf{y}_s^{(J)})^T$ and $\mathbf{X} = [\mathbf{X}^{(1)T} : \dots : \mathbf{X}^{(J)T}]^T$. Let $\Phi_h = [\Phi_h^{(1)T} : \dots : \Phi_h^{(J)T}]^T$, where $\Phi_h^{(j)}$ represents the $m \times n_j$ sketching matrix. The sketching matrix $\Phi_h^{(j)}$ can be constructed within each research center, and local computations of $\Phi_h^{(j)} \mathbf{y}_s^{(j)}$ and $\Phi_h^{(j)} \mathbf{X}^{(j)}$ can be performed within the j th center before releasing them to calculate the h th sketched posterior. The privacy guarantee, as discussed in Section 3.1.3, ensures that the original data cannot be reconstructed from the sketches obtained from centers. By aggregating these local sketches $\Phi_h \mathbf{y}_s = \sum_{j=1}^J \Phi_h^{(j)} \mathbf{y}_s^{(j)}$ and $\Phi_h \mathbf{X} = \sum_{j=1}^J \Phi_h^{(j)} \mathbf{X}^{(j)}$, the framework allows the computation of complete data sketches. The schematic representation in the Figure 2 illustrates the strategy for securely computing sketched posteriors while preserving privacy.

3.2 Construction of the Collaborative Sketched Posterior

The approach described below for combining sketched posteriors constructed in Section 3.1 is in the same spirit as the combination of “subset posteriors” outlined in the divide-and-conquer strategies in Bayesian inference [Guhaniyogi et al., 2023]. The salient feature of the combination technique is that it is agnostic to the model- or data-specific assumptions, such as assuming independence among samples in the training data.

This process of combining sketched posteriors involves leveraging the concept of the Wasserstein barycenter, a technique extensively used in the realm of scalable Bayesian methods for distributed inference, as highlighted in works such as Guhaniyogi et al. [2023] and Srivastava et al. [2018]. Let (Ω, ρ) be a complete separable metric space and $\mathcal{P}_2(\Omega)$ denotes the probability distributions on (Ω, ρ) with finite second moments. Let $\tilde{\Pi}_1, \tilde{\Pi}_2$ be two probability measures in $\mathcal{P}_2(\Omega)$. Assume $\Delta(\tilde{\Pi}_1, \tilde{\Pi}_2)$ is the set of all probability measures on $\Omega \times \Omega$ with marginals $\tilde{\Pi}_1, \tilde{\Pi}_2 \in \mathcal{P}_2(\Omega)$. Then the Wasserstein distance of order 2, denoted as W_2 , between $\tilde{\Pi}_1, \tilde{\Pi}_2$ is defined as $W_2(\tilde{\Pi}_1, \tilde{\Pi}_2) =$




Figure 2: Distributed storage of data for J separate storage centers.

$\left\{ \inf_{\pi \in \Delta(\bar{\Pi}_1, \bar{\Pi}_2)} \int_{\Omega \times \Omega} \rho^2(x, y) d\pi(x, y) \right\}^{1/2}$. In the context of Section 3.1, if Π_1, \dots, Π_H all have finite second moments, then the Wasserstein barycenter of Π_1, \dots, Π_H is defined as

$$\bar{\Pi} = \underset{\Pi \in \mathcal{P}_2(\Omega)}{\operatorname{argmin}} \frac{1}{H} \sum_{h=1}^H W_2^2(\Pi, \Pi_h). \quad (6)$$

It is known that $\bar{\Pi}$ exists and is unique [Agueh and Carlier, 2011]. The Wasserstein barycenter $\bar{\Pi}$ is referred to as the collaborative sketched posterior for the parameters and replaces full data posterior as its computationally efficient approximation for inference on (2). The two-step procedure for the construction of collaborative sketched posterior can be visualized in Figure 3.

When focus lies on drawing inference on one-dimensional functional of the parameters, denoted as α , the sketched pseudo posterior distribution of α can be easily derived by averaging the empirical sketched posterior quantiles [Guhaniyogi et al., 2023]. This is facilitated by the fact that the Wasserstein distance (W_2 distance) between two univariate distributions corresponds to the L_2 distance between their respective quantile functions (as demonstrated in Lemma 8.2 of Bickel and Freedman [1981]). More precisely, for a given quantile level $\xi \in (0, 1)$, let $\hat{\alpha}_{\xi, h}$ represent the ξ th

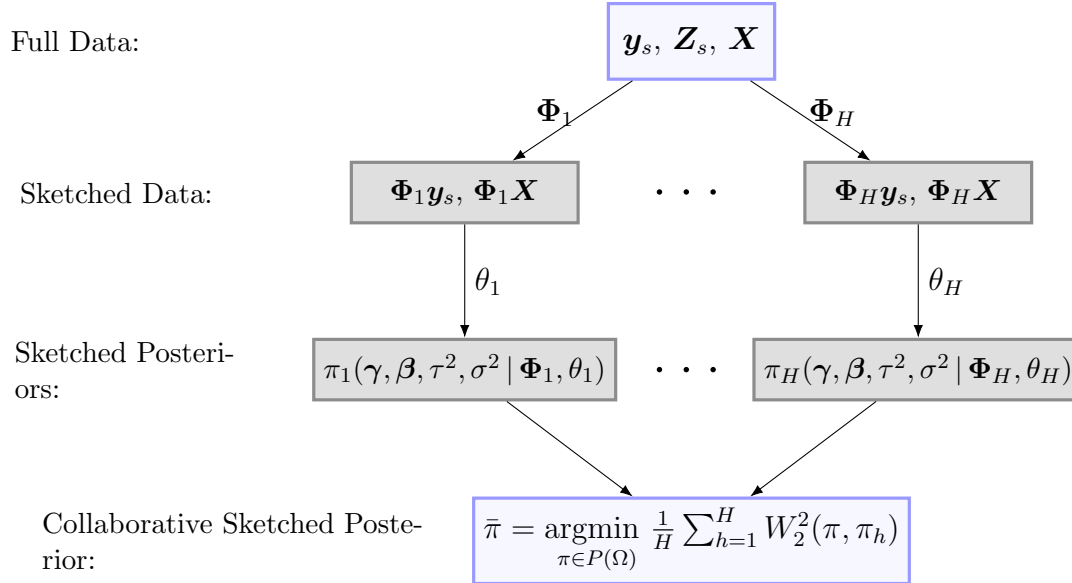


Figure 3: Visual overview of distributed learning algorithm for H parallel nodes.

empirical quantile of α based on MCMC samples drawn from the marginal posterior distribution of α from h th subset. With $\hat{\alpha}_\xi$ as the ξ th empirical quantile of α following the collaborative sketched posterior, $\hat{\alpha}_\xi$ can be expressed as,

$$\hat{\alpha}_\xi = \frac{1}{H} \sum_{h=1}^H \hat{\alpha}_{\xi,h}, \quad (7)$$

where the value of ξ is varied across a fine grid within the interval $(0,1)$. By utilizing a sufficiently fine ξ -grid in the equation (7), MCMC samples from the marginal collaborative sketched posterior distribution of α can be obtained by inverting the empirical distribution function supported on the estimated quantiles.

In real-world applications, the primary focus often centers on the marginal distributions of model parameters and predicted values, i.e., on one-dimensional functionals of parameters. Consequently, the process of obtaining the univariate Wasserstein barycenter by averaging quantiles, as described in equation (7), accomplishes this objective with remarkable versatility and a user-friendly implementation. Due to its general applicability and ease of use, in the subsequent sections, we concentrate exclusively on scenarios where α is one-dimensional and utilize equation (7) to calculate the collaborative sketched posterior by utilizing its empirical quantiles. It is worth noting that there are existing approaches which allow combination of joint posteriors instead of marginal

posteriors, see the recent work by Guhaniyogi et al. [2022] and the references therein. However, these methods are computationally more demanding and yield only marginal improvements over the univariate quantile combination approach mentioned in equation (7).

While our proposed approach shares similarities with the evolving field of distributed Bayesian inference for correlated data [Guhaniyogi et al., 2023, 2022, Guhaniyogi and Banerjee, 2018], it offers a distinct advantage over these approaches. Notably, all the existing distributed Bayesian learning approaches involve dividing the data into shards based on a user-defined partitioning scheme. However, a crucial observation from simulation studies in Guhaniyogi et al. [2023] is that the effectiveness of distributed inference is somewhat influenced by the specific choice of data shards. This necessitates application-specific judgments when partitioning data into these shards. In contrast, our approach sidesteps the need to create data shards and instead focuses on constructing data sketches with multiple random sketching matrices. The fundamental distinction lies in the fact that while the Wasserstein barycenter of shard posteriors in traditional distributed learning methods is contingent upon the chosen data partitioning, our approach incorporate data sketching as an intrinsic facet of the model development process, and mitigates the impact of variations due to the choice of sketching matrices by computing the Wasserstein barycenter of sketched posteriors, each constructed using a different random sketching matrix. This innovation leads to the novel concept of *robust distributed Bayesian learning*, where the resulting inference is more resilient to the uncertainties introduced by random sketching.

4 Simulation Studies

One salient feature of the proposed approach is that it is agnostic to the dimension d of the index points. However, we empirically validate the efficacy of the collaborative sketched posterior for the case of $d = 2$, where the index points represent locations in space, since $d = 2$ for the SLOSH emulator data. To comprehensively evaluate the performance of the collaborative sketched posterior, it is constructed with $w_s(\mathbf{u})$ modeled using each of the three strategies: (i) the full-rank GP, (ii) the low-rank Modified predictive process [Finley et al., 2009] and (iii) NNGP [Datta et al., 2016]. The inferential and predictive performance of the collaborative sketched posterior under (i)-(iii) is compared with existing distributed Bayesian approaches utilizing different data subsetting techniques. To simulate the data, we use $n+n_0$ spatial locations $\mathbf{u}_1, \dots, \mathbf{u}_{n+n_0}$ drawn uniformly over

the domain $\mathcal{D} = [0, 10] \times [0, 10]$. The data generation procedure involves $p = 1$ global attribute and $q = 2$ local attributes. The global attribute values are generated from the $N(0, 1)$ distribution for $S + S_0$ simulations, i.e., $z_1, \dots, z_{S+S_0} \stackrel{i.i.d.}{\sim} N(0, 1)$. Similarly, the local attributes $\mathbf{x}(\mathbf{u}_1), \dots, \mathbf{x}(\mathbf{u}_{n+n_0})$ are simulated independently from $N(\mathbf{0}, \mathbf{I})$. For each $s = 1, \dots, S + S_0$ and $i = 1, \dots, n + n_0$, the response $y_s(\mathbf{u}_i)$ is drawn independently from $N(z_s \gamma_0 + \mathbf{x}(\mathbf{u}_i)^T \boldsymbol{\beta}_0 + w_{s0}(\mathbf{u}_i), \tau_0^2)$ following (1), where the noise variance τ_0^2 is set to be 0.2.

The true coefficient γ_0 for the global attribute is set to be 5, where as the true coefficient $\boldsymbol{\beta}_0$ for local attributes is kept at $(2, -1)$. The true function $w_{s0}(\mathbf{u})$ s are generated from a Gaussian process with mean 0 and covariance kernel $\sigma_0^2 \kappa(\cdot, \cdot; \theta_0)$ independently over s , i.e., $(w_{s0}(\mathbf{u}_1), \dots, w_{s0}(\mathbf{u}_{n+n_0}))^T \stackrel{ind.}{\sim} N(\mathbf{0}, \sigma_0^2 \mathbf{K}(\theta_0))$ for $s = 1, \dots, S + S_0$, where $\mathbf{K}(\theta_0)$ is an $(n + n_0) \times (n + n_0)$ matrix with the (j, j') th element $\kappa(\mathbf{u}_j, \mathbf{u}_{j'}; \theta_0)$. We set the covariance kernel $\kappa(\cdot, \cdot; \theta_0)$ to be the exponential correlation function given by $\kappa(\mathbf{u}_j, \mathbf{u}_{j'}; \theta_0) = \exp(-\theta_0 \|\mathbf{u}_j - \mathbf{u}_{j'}\|)$, with the true values σ_0^2 and θ_0 set to 2 and 3, respectively.

Out of $S + S_0 = 15$ simulations, we randomly choose $S = 10$ simulations for model fitting, and the remaining $S_0 = 5$ simulations are used for testing the predictive performance of the model. For the $S = 10$ training simulations, the model is fitted on $n = 10,000$ spatial locations, whereas prediction is performed on $n_0 = 1000$ different locations for the remaining $S_0 = 5$ simulations. In other words, we assess predictive accuracy not only on different locations in space but also on entirely different simulations.

Throughout all simulations, the collaborative sketched posterior is calculated using a sketching dimension of $m = 500$. Moreover, when fitting sketched posteriors with the MPP for each $w_s(\mathbf{u})$, we employ 500 knots, and for fitting NNGP for each $w_s(\mathbf{u})$, we consider 10 nearest neighbors.

4.1 Competitors

For all simulations, we compare the collaborative sketching framework proposed here with the Distributed Kriging (DISK) approach [Guhaniyogi et al., 2023], a divide-and-conquer Bayesian approach that computes subset posteriors with user-defined data subsets and combines marginal distributions of these subset posteriors by averaging their quantiles, similar to the strategy in (7). To demonstrate sensitivity in inference due to the choice of different data subsetting, DISK is fitted with two different strategies: (a) subdomain and (b) stratification. The subdomain strategy

divides the domain into rectangular subdomains with each subset constituting all data points from a subdomain, while the stratification strategy constructs each subset with representative samples from each subdomain. We refer to them as DISK-subdomain and DISK-stratified, respectively, and they both maintain the number of samples per subset to be around the sketching dimension $m = 500$ to ensure a fair comparison with our collaborative sketching approach.

4.2 Metrics of Comparison

Let $\mathbf{Y}^{(l)}$ denote the l th post burn-in MCMC sample from the posterior predictive distribution of $\mathbf{Y} = (y_s(\mathbf{u}_i) : i = n + 1, \dots, n + n_0; s = S + 1, \dots, S + S_0)^T$, for $l = 1, \dots, L$. The point prediction for competitors is assessed using mean squared prediction error (MSPE), defined as the discrepancy between the true and the predicted responses $\|\bar{\mathbf{Y}} - \mathbf{Y}_t\|^2 / (n_0 S_0)$, where $\bar{\mathbf{Y}} = \frac{1}{L} \sum_{l=1}^L \mathbf{Y}^{(l)}$ and \mathbf{Y}_t denotes the true value of \mathbf{Y} . The predictive uncertainty is determined with the coverage probability, interval score and energy score [Gneiting and Raftery, 2007] for 95% predictive intervals (PIs) for all the competing methods over n_0 out of sample locations in S_0 out of sample simulations. Interval scores favors model with the smallest possible intervals that still contain the data. On the other hand, energy score is a multivariate extension to Continuous Rank Probability Score (CRPS) and is calculated as $\frac{1}{L} \sum_{l=1}^L \|\mathbf{Y}^{(l)} - \mathbf{Y}_t\| - \frac{1}{2L^2} \sum_{l=1}^L \sum_{l'=1}^L \|\mathbf{Y}^{(l)} - \mathbf{Y}^{(l')}\|$. Energy score takes into account not only the predictive accuracy of each sample from the posterior predictive distribution, but also the level of uncertainty in the distribution. For this reason, energy score has gained interest in recent literature as a model ranking mechanism [Heaton et al., 2019]. Finally, we compare performance of all distributed Bayesian methods for parameter estimation using the posterior medians or point estimates and the 95% credible intervals (CIs) for $\boldsymbol{\beta} = (\beta_1, \beta_2), \gamma, \sigma^2, \tau^2$.

4.3 Results of Simulated Data Analysis

The simulation results presented in Table 2 highlight significant sensitivity in parameter estimation, contingent upon the nature of data partitioning in the DISK approach. For instance, when a full GP or NNGP is fitted to each data subset, the 95% CIs for σ^2 in DISK-Subdomain and DISK-Stratified do not overlap. Additionally, the 95% CI for β_1 exhibits a much larger confidence interval for DISK-subdomain than DISK-stratified. Furthermore, under the MPP fitting, DISK-Stratified underestimates τ^2 , whereas the estimation remains precise when DISK-Subdomain is fitted with MPP. However, DISK-subdomain underestimates σ^2 more egregiously than DISK-

Model	Competitors	σ^2	τ^2	β_1	β_2	γ
	Truth	2.00	0.20	2.00	-1.00	5.00
Full GP	DISK-Subdomain	2.13 (2.09, 2.16)	0.21 (0.21, 0.22)	1.77 (1.39, 2.14)	-1.00 (-1.01, -0.98)	5.00 (4.99, 5.01)
	DISK-Stratified	1.94 (1.91, 1.96)	0.20 (0.19, 0.21)	1.99 (1.89, 2.08)	-1.02 (-1.02, -0.98)	5.00 (5.00, 5.00)
	Collaborative Sketching	2.07 (2.03, 2.10)	0.21 (0.19, 0.22)	2.00 (1.98, 2.01)	-0.99 (-1.00, -0.97)	5.00 (5.00, 5.00)
MPP	DISK-Subdomain	1.62 (1.59, 1.65)	0.18 (0.17, 0.20)	1.79 (1.55, 2.03)	-1.00 (-1.02, -0.98)	5.00 (4.99, 5.00)
	DISK-Stratified	1.80 (1.79, 1.82)	0.05 (0.04, 0.05)	1.94 (1.86, 2.01)	-0.99 (-1.02, -0.97)	5.00 (5.00, 5.00)
	Collaborative Sketching	2.07 (1.99, 2.16)	0.11 (0.10, 0.13)	2.01 (2.00, 2.03)	-0.95 (-0.96, -0.95)	5.00 (5.00, 5.00)
NNGP	DISK-Subdomain	2.14 (2.10, 2.18)	0.21 (0.21, 0.22)	1.76 (1.37, 2.14)	-1.00 (-1.01, -0.98)	5.00 (5.00, 5.01)
	DISK-Stratified	1.99 (1.97, 2.01)	0.21 (0.20, 0.21)	1.99 (1.90, 2.08)	-1.00 (-1.01, -0.98)	5.00 (4.99, 5.00)
	Collaborative Sketching	2.04 (2.00, 2.07)	0.19 (0.18, 0.20)	1.97 (1.95, 1.99)	-0.98 (-0.99, -0.97)	5.00 (5.00, 5.00)

Table 2: We calculate the posterior median with 95% confidence intervals for all model parameters for all the distributed Bayesian competitors, fitting a full-GP, low-rank modified predictive process (MPP), and NNGP. The MPP utilizes 500 knots for model implementation and NNGP uses 10 nearest neighbors. We set both the sketching dimension for our approach and the size of each subset for the DISK approach to be $m = 500$ to ensure comparability.

Model	Competitors	MSPE	Coverage	Interval Score	Energy Score
Full GP	DISK-Subdomain	2.44	0.95	7.02	0.89
	DISK-Stratified	2.30	0.94	6.89	0.86
	Collaborative Sketching	2.27	0.95	6.87	0.86
MPP	DISK-Subdomain	2.50	0.91	7.24	0.90
	DISK-Stratified	2.32	0.92	7.05	0.89
	Collaborative Sketching	2.30	0.95	6.87	0.86
NNGP	DISK-Subdomains	2.43	0.95	7.00	0.89
	DISK-Stratified	2.29	0.94	6.88	0.86
	Collaborative Sketching	2.28	0.95	6.88	0.86

Table 3: MSPE, coverage, interval score and energy score for all competing methods.

stratified. Consequently, there is uncertainty regarding which data partitioning scheme to rely on while providing inference with DISK on model parameters. In contrast, the collaborative sketching approach aggregates inference over the data sketching mechanism, resulting in robust and accurate point estimation, as well as 95% credible intervals that cover the true parameters.

The MSPE values presented in Table 3 exhibit some variation between DISK-Subdomain and DISK-Stratified, although it is less pronounced compared to the parameter estimation. DISK-Subdomain performs inferior to DISK-Stratified across full-GP, MPP, and NNGP fittings on each subset. The coverage of both approaches is approximately equal, with DISK-Stratified showing slightly smaller interval score and energy score. This finding aligns with Guhaniyogi et al. [2023], which demonstrates DISK-Stratified to be the best strategy among alternatives (e.g., DISK-Subdomain) for partitioning data subsets, as it includes representative samples from every subdomain in each subset, resulting in better model fitting using each subset posterior. Notably, the collaborative sketching strategy yields a slightly smaller MSPE than DISK-Stratified. While offering




Figure 4: MSPE and computation time (in hours) for model fit and prediction in our simulation data, as a function of the compressed dimension size m .

similar coverage, interval score, and energy score with DISK-Stratified for full-GP and NNGP, collaborative sketching provides better coverage and smaller interval and energy scores when sketched posteriors are fitted with MPP. In summary, the empirical investigation demonstrates the excellent performance of the collaborative sketching framework in terms of both inference and prediction when fitting each sketched posterior with a full-GP and its most popular computationally efficient alternatives.

4.4 Choice of the Rank m of Sketching Matrices

We explore the selection of an appropriate compression matrix size, denoted as m . The theoretical consideration suggests that m should be on the order of $n/\log(n)$. However, in practical applications, it is feasible to achieve robust and accurate inference with smaller values of m . To illustrate this with simulated data having a sample size of $n = 10,000$, we conducted model runs for various values of m . As expected, the MSPE decreases as m increases, albeit with a diminishing rate of decline (refer to Figure 4) until m reaches 1% of the full sample size n . Beyond this point, the MSPE stabilizes, a trend observed in our empirical experiments. Further increasing m may result in marginal reductions in MSPE but comes at the expense of longer computation times (see Figure 4). Therefore, for the real data analysis, and in particular, for SLOSH emulator data, we recommend utilizing m as 1% of the total sample size in the subsequent sections.

5 SLOSH Emulator Data Analysis

We fit the collaborative sketching approach with full Gaussian process on the SLOSH emulator data described in Section 2. As noted in Sections 1 and 2, this article offers the first principled functional analysis of SLOSH emulator data incorporating global and local attributes. Out of 10 simulated storms each on $n = 49719$ locations, 5 storms are randomly selected for model fitting and rest are employed to assess predictive performance of the proposed approach. Following the discussion in Section 4.4, m is set at 1% of n , i.e., $m = 498$.

We compare our approach with Bayesian adaptive spline surfaces (BASS) [Francom and Sansó, 2020], a commonly used approach in national laboratories for analyzing functional data from multiple simulations. BASS is not able to incorporate local attributes (i.e., elevation in the SLOSH data) and is susceptible to high variance of estimation when the number of simulations is as low. However, given its popularity among researchers in national laboratories for the functional data analysis, we include it as a competitor and implement it using the R package `BASS`. Note that in this article, our primary focus centers on distributed Bayesian approaches. However, for benchmarking purposes, we also introduce a non-distributed method. Spatial statistics methods typically rely on a single simulation, making it challenging to compare our approach with an appropriate non-distributed state-of-the-art spatial techniques. In addressing this challenge, we employ the non-distributed NNGP on the complete dataset Datta et al. [2016] as a comparative method, albeit with a slight modification. Due to constraints in the NNGP package that allow only independent simulation runs, we perform predictive inference on NNGP by averaging results across multiple simulation runs. Despite the potential suboptimal performance of NNGP under this approach, we designate it as a benchmark to evaluate the non-distributed spatial model’s performance across the entire dataset, referring to it as NNGP-ind.

Figure 5 presents densities of the posterior distribution of model parameters with posterior median and 95% CIs marked within each figure. All global and local attributes turn out to be significantly associated with the storm surge with none of their 95% CIs includes zero. Understandably, the storm surge is positively associated with the sea level rise, minimum velocity of the eye of the storm, and negatively associated with the elevation, direction of the wind and minimum air pressure of the eye of the storm. The estimated posterior median of spatial variance σ^2 is

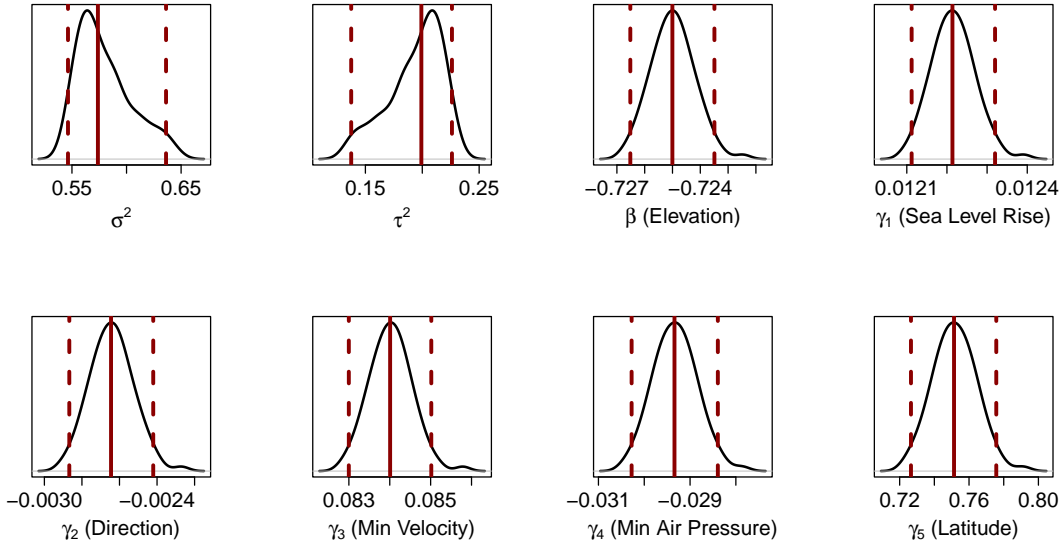


Figure 5: Densities for each parameter under our sketching model. Solid maroon lines indicate the posterior means, and dotted maroon lines indicate the upper and lower 95% credible interval bounds.

	MSPE	Error %	Coverage	Score	Energy score
Sketching	1.07	0.06	0.88	5.60	0.63
NNGP-ind	0.83	0.09	0.18	20.45	1.52
BASS	1.45	0.10	0.59	11.23	1.02

Table 4: Predictive diagnostics for the storm surge analysis. For interval score and Energy score, lower values indicate better scores.

dominant over the error variance τ^2 , justifying the spatial analysis.

Table 4 presents predictive inferences for all the competing approaches. The actual storm surge and the predicted storm surge in a randomly selected test storm in Figure 6 demonstrate that the collaborative sketching approach, along with its competitors, accurately captures the local features of storm surge over space. Moreover, the results from Table 4 suggest that NNGP-ind might outperform other methods in terms of point predictions, as evidenced by its lowest MSPE. However, when considering the energy score, which utilizes predictive samples rather than the mean, the sketching approach is substantially favored over the NNGP-ind model. Therefore, while NNGP-ind provides excellent mean predictions, the sketching approach offers superior uncertainty quantification in prediction. This observation is further supported by the coverage and interval scores, where sketching exhibits significantly higher predictive coverage than its competitors, coupled with

a much smaller interval score. Understandably, BASS underperforms both in terms of point estimation and uncertainty quantification due to the small number $S = 5$ simulations for model fitting.

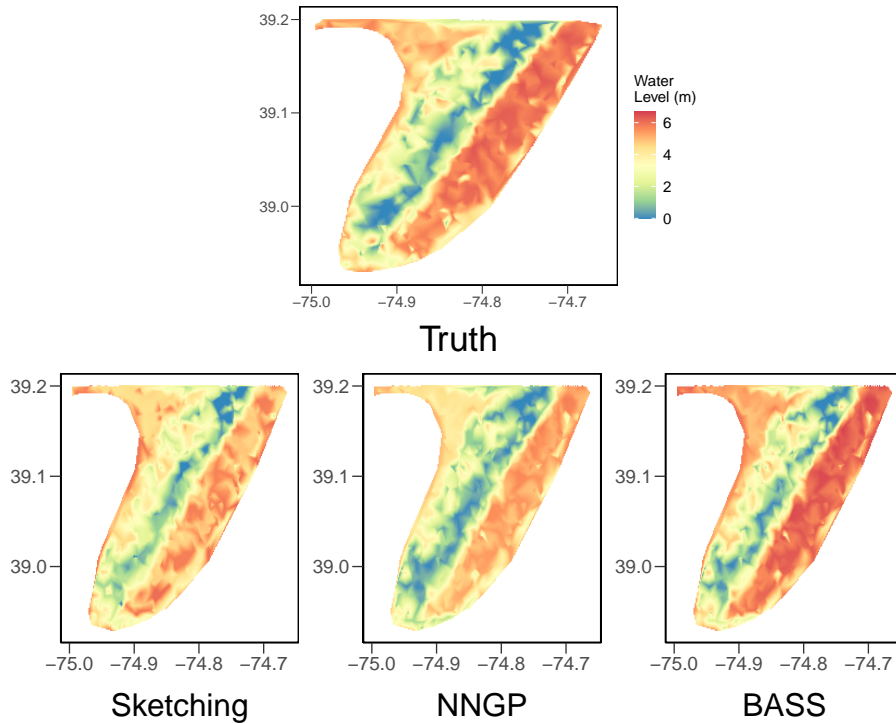


Figure 6: Actual water level (in meters) for a randomly selected storm at each coordinate in the testing dataset, along with the predicted water level under each model.

Section 2 highlights the significance of a flooding threshold set at four feet, considering its real-world implications, particularly in the context of power stations designed to withstand this level of floodwater. Given its importance, it is crucial for an emulator to accurately predict flood levels above four feet. To evaluate emulators in terms of this feature, we will examine the percentage of predictions that mis-classify the need for intervention, i.e., the percentage of predictions in the test simulations where the true SLOSH output exceeds four feet but the predicted output is less than four feet, or, the predicted output is above four feet and the true SLOSH output is less than four feet. Error percentage in Table 4 presents this metric for all the competitors. Notably, all competitors show small mis-classification rate with the rate being minimum for the collaborative sketching approach.

6 Conclusion and Future Work

This article introduces a novel distributed Bayesian inferential framework that utilizes the theory of data sketching through random compression matrices. The article proposes fitting a powerful functional data model on multiple random sketches of the full data constructed using multiple random sketching matrices in parallel, followed by combining these inferences in a central server. By aggregating inference across diverse random sketches, our approach proves resilient to the selection of data sketches, leading to the development of novel robust distributed Bayesian learning approach. The proposed framework allows joint analysis of data stored within different research centers without leaking the privacy of samples. The proposed framework offers accurate inference on the association between water surge with storm-specific characteristics in the SLOSH simulator data.

Our framework’s ample generality enables its application in scaling complex data models. An immediate avenue for future work involves extending the framework to facilitate robust distributed inference with multivariate functional data models. Additionally, we plan to broaden the framework to support distributed Bayesian inference with Gaussian Cox process models, specifically tailored for large point process data.

References

- M. Agueh and G. Carlier. Barycenters in the wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2):904–924, 2011.
- D. Ahfock, W. J. Astle, and S. Richardson. Statistical properties of sketching algorithms. *arXiv preprint arXiv:1706.03665*, 2017.
- N. Ailon and B. Chazelle. The fast johnson–lindenstrauss transform and approximate nearest neighbors. *SIAM Journal on computing*, 39(1):302–322, 2009.
- P. J. Bickel and D. A. Freedman. Some asymptotic theory for the bootstrap. *The annals of statistics*, 9(6):1196–1217, 1981.
- R. Borge, V. Alexandrov, J. J. Del Vas, J. Lumbreras, and E. Rodríguez. A comprehensive sensitiv-

- ity analysis of the wrf model for air quality applications over the iberian peninsula. *Atmospheric Environment*, 42(37):8560–8574, 2008.
- S. Chen, Y. Liu, M. R. Lyu, I. King, and S. Zhang. Fast relative-error approximation algorithm for ridge regression. In *UAI*, pages 201–210, 2015.
- K. L. Clarkson and D. P. Woodruff. Low-rank approximation and regression in input sparsity time. *Journal of the ACM (JACM)*, 63(6):1–45, 2017.
- A. Datta, S. Banerjee, A. O. Finley, and A. E. Gelfand. Hierarchical nearest-neighbor gaussian process models for large geostatistical datasets. *Journal of the American Statistical Association*, 111(514):800–812, 2016.
- E. Dobriban and S. Liu. A new theory for sketching in linear regression. *arXiv preprint arXiv:1810.06089*, 2018.
- P. Drineas, M. W. Mahoney, S. Muthukrishnan, and T. Sarlós. Faster least squares approximation. *Numerische mathematik*, 117(2):219–249, 2011.
- A. O. Finley, S. Banerjee, P. Waldmann, and T. Ericsson. Hierarchical spatial modeling of additive and dominance genetic variance for large spatial trial datasets. *Biometrics*, 65(2):441–451, 2009.
- D. Francom and B. Sansó. BASS: An R package for fitting and performing sensitivity analysis of Bayesian adaptive spline surfaces. *Journal of Statistical Software*, 94(LA-UR-20-23587), 2020.
- T. Gneiting and A. E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, 102(477):359–378, 2007.
- R. Guhaniyogi and S. Banerjee. Meta-kriging: Scalable bayesian modeling and inference for massive spatial datasets. *Technometrics*, 60(4):430–444, 2018.
- R. Guhaniyogi and S. Banerjee. Multivariate spatial meta kriging. *Statistics & probability letters*, 144:3–8, 2019.
- R. Guhaniyogi and D. B. Dunson. Bayesian compressed regression. *Journal of the American Statistical Association*, 110(512):1500–1514, 2015.

- R. Guhaniyogi and D. B. Dunson. Compressed gaussian process for manifold regression. *The Journal of Machine Learning Research*, 17(1):2472–2497, 2016.
- R. Guhaniyogi and A. Scheffler. Sketching in bayesian high dimensional regression with big data using gaussian scale mixture priors. *arXiv preprint arXiv:2105.04795*, 2021.
- R. Guhaniyogi, C. Li, T. D. Savitsky, and S. Srivastava. Distributed bayesian varying coefficient modeling using a gaussian process prior. *The Journal of Machine Learning Research*, 23(1):3642–3700, 2022.
- R. Guhaniyogi, C. Li, T. Savitsky, and S. Srivastava. Distributed bayesian inference in massive spatial data. *Statistical science*, 38(2):262–284, 2023.
- N. Halko, P.-G. Martinsson, and J. A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.
- M. J. Heaton, A. Datta, A. O. Finley, R. Furrer, J. Guinness, R. Guhaniyogi, F. Gerber, R. B. Gramacy, D. Hammerling, M. Katzfuss, et al. A case study competition among methods for analyzing large spatial data. *Journal of Agricultural, Biological and Environmental Statistics*, 24:398–425, 2019.
- Z. Huang. Near optimal frequent directions for sketching dense and sparse matrices. In *International Conference on Machine Learning*, pages 2048–2057. PMLR, 2018.
- G. Hutchings, B. Sansó, J. Gattiker, D. Francom, and D. Pasqualini. Comparing emulation methods for a high-resolution storm surge model. *Environmetrics*, 34(3):e2796, 2023.
- C. P. Jelesnianski, J. Chen, and W. A. Shaffer. Slosh : Sea, lake, and overland surges from hurricanes. *NOAA Technical Report NWS*, 48, 1992. URL <https://repository.library.noaa.gov/view/noaa/7235>.
- M. C. Kennedy and A. O’Hagan. Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(3):425–464, 2001.

- M. W. Mahoney et al. Randomized algorithms for matrices and data. *Foundations and Trends® in Machine Learning*, 3(2):123–224, 2011.
- M. R. Petersen, X. S. Asay-Davis, A. S. Berres, Q. Chen, N. Feige, M. J. Hoffman, D. W. Jacobsen, P. W. Jones, M. E. Maltrud, S. F. Price, et al. An evaluation of the ocean and sea ice climate of e3sm using mpas and interannual core-ii forcing. *Journal of Advances in Modeling Earth Systems*, 11(5):1438–1458, 2019.
- W. Sacks and T. Welch. Mitchell, and hp wynn, “design and analysis of computer experiments,”. *Statistical Science*, 4(4):409–453, 1989.
- J. M. Salter and D. Williamson. A comparison of statistical emulation methodologies for multi-wave calibration of environmental models. *Environmetrics*, 27(8):507–523, 2016.
- T. Sarlos. Improved approximation algorithms for large matrices via random projections. In *2006 47th annual IEEE symposium on foundations of computer science (FOCS’06)*, pages 143–152. IEEE, 2006.
- S. Srivastava, C. Li, and D. B. Dunson. Scalable bayes via barycenter in wasserstein space. *The Journal of Machine Learning Research*, 19(1):312–346, 2018.
- S. S. Vempala. *The random projection method*, volume 65. American Mathematical Soc., 2005.
- D. P. Woodruff et al. Sketching as a tool for numerical linear algebra. *Foundations and Trends® in Theoretical Computer Science*, 10(1–2):1–157, 2014.
- H. Zhang. Inconsistent estimation and asymptotically equal interpolations in model-based geostatistics. *Journal of the American Statistical Association*, 99(465):250–261, 2004.
- L. Zhang, M. Mahdavi, R. Jin, T. Yang, and S. Zhu. Recovering the optimal solution by dual random projection. In *Conference on Learning Theory*, pages 135–157, 2013.
- L. Zhang, W. Tang, and S. Banerjee. Exact bayesian geostatistics using predictive stacking. *arXiv preprint arXiv:2304.12414*, 2023.
- S. Zhou, L. Wasserman, and J. D. Lafferty. Compressed regression. In *Advances in Neural Information Processing Systems*, pages 1713–1720, 2008.