



# Bayesian Nonparametric Generative Modeling of Large Multivariate Non-Gaussian Spatial Fields

Paul F. V. WIEMANN and Matthias KATZFUSS

Multivariate spatial fields are of interest in many applications, including climate model emulation. Not only can the marginal spatial fields be subject to nonstationarity, but the dependence structure among the marginal fields and between the fields might also differ substantially. Extending a recently proposed Bayesian approach to describe the distribution of a nonstationary univariate spatial field using a triangular transport map, we cast the inference problem for a multivariate spatial field for a small number of replicates into a series of independent Gaussian process (GP) regression tasks with Gaussian errors. Due to the potential nonlinearity in the conditional means, the joint distribution modeled can be non-Gaussian. The resulting nonparametric Bayesian methodology scales well to high-dimensional spatial fields. It is especially useful when only a few training samples are available, because it employs regularization priors and quantifies uncertainty. Inference is conducted in an empirical Bayes setting by a highly scalable stochastic gradient approach. The implementation benefits from mini-batching and could be accelerated with parallel computing. We illustrate the extended transport-map model by studying hydrological variables from non-Gaussian climate-model output.

**Key Words:** Climate-model emulation; Gaussian process; Generative modeling; Multivariate spatial field; Non-stationarity.

## 1. INTRODUCTION

Multivariate spatial fields play a significant role in various scientific disciplines, including environmental modeling and climate science, where multiple spatially referenced variables are observed. To highlight the need for statistical models to effectively capture the intricate relationships among multiple variables in spatial fields, consider measures for apparent temperature. The heat index, for instance, relies on temperature and humidity to assess the apparent temperature. More comprehensive measures of apparent temperature, such as the wet-bulb globe temperature, which incorporates the additional factors of wind speed and

---

P. F. V. Wiemann · M. Katzfuss (✉) Department of Statistics, Texas A & M University, College Station, TX, USA (E-mail: [katzfuss@gmail.com](mailto:katzfuss@gmail.com)).

radiation, have been proposed to better assess the stress of exposure to high temperatures on the human body. Consequently, in the context of statistical climate-model emulators, capturing the inter-variable dependence is crucial for reliable predictions, conditional predictions, and accurate uncertainty quantification. Accurately inferring the joint distribution and understanding the conditional relationships among these variables is challenging, especially when dealing with complex dependencies and non-Gaussian characteristics.

Most existing methods for univariate or multivariate spatial analysis were developed for inference based on a single training sample and assume Gaussian processes (GPs) with simple parametric covariance functions (e.g., Cressie 1993; Banerjee et al. 2004). Extensions to nonparametric covariances (e.g., Huang et al. 2011; Choi et al. 2013; Porcu et al. 2021) or multivariate fields (e.g., Genton and Kleiber 2015) typically still rely on implicit or explicit assumptions of Gaussianity. For emulation of univariate spatial climate-model output, one can combine locally fit anisotropic Matérn covariances into a global Gaussian model (Nychka et al. 2018; Wiens et al. 2020). Generative machine-learning approaches (e.g., Kobyzhev et al. 2020; Goodfellow et al. 2016; Kovachki et al. 2020) often require many training samples and may be sensitive to tuning-parameter and network-architecture choices (e.g., Arjovsky and Bottou 2017; Hestness et al. 2017; Mescheder et al. 2018).

Spatial-temporal dependencies of multivariate global fields were captured by Jun (2011). Nonstationarities in latitude are supported by Castruccio and Stein (2013) for the univariate field describing annually averaged surface temperature. Edwards et al. (2019) extend this method to multivariate fields relying on a parametric approach and assuming a Gaussian distribution. Relying on a stationary assumption while allowing for arbitrary missingness patterns, Guinness (2022) estimates the spectral form of gridded multivariate spatio-temporal data. In contrast to the method we propose, these multivariate models required all variables to be observed at the same spatial locations.

Triangular transport maps (e.g., Marzouk et al. 2016) can be used to characterize continuous multivariate distributions. A transport map transforms the target distribution into a reference distribution, such as the standard Gaussian. Non-Gaussian target distributions can be obtained by introducing nonlinearities into the map. With an invertible transport map, one can sample from the target distribution and its conditionals or convert non-Gaussian data to the reference space, where linear regression or interpolation can be applied. Transport maps are often estimated from training data by iteratively expanding a finite-dimensional parameterization of the transport map (e.g., El Moselhy and Marzouk 2012; Bigoni et al. 2016; Marzouk et al. 2016; Parno et al. 2016).

Katzfuss and Schäfer (2023) instead proposed a Bayesian nonparametric approach, in which the components of the transport map are modeled as GPs. This results in closed-form inference that quantifies uncertainty and avoids under- and over-fitting even when the number of training samples is small. For target distributions corresponding to spatial fields, Katzfuss and Schäfer (2023) proposed specific priors that exploit the screening effect via suitable conditional-independence assumptions that guarantee computational scalability for very large datasets. The resulting sparse nonlinear transport maps can be seen as a nonparametric and non-Gaussian generalization of Vecchia approximations (e.g., Vecchia 1988; Stein et al. 2004; Datta et al. 2016; Katzfuss and Guinness 2021; Schäfer et al. 2021a),

which implicitly utilize linear transport map given by a sparse inverse Cholesky factor. [Kidd and Katzfuss \(2022\)](#) proposed a Bayesian nonparametric inference on the Cholesky factor.

Our contribution is a novel extension of the scalable Bayesian transport map (BTM) approach developed by [Katzfuss and Schäfer \(2023\)](#), tailored explicitly for learning the distribution of multivariate spatial fields from a few replicates. The essential contribution lies in the introduction of an augmented input space that incorporates both the spatial locations and latent locations referencing the component from the multivariate response (see [Apanasovich and Genton 2010](#), for a similar concept used with a stationary parametric covariance function). By leveraging the augmented input space, we expand the scope of BTM to encompass multivariate spatial fields without fundamentally altering the core principles and estimation algorithms. Consequently, the extension benefits from the good approximation properties, making the method scale well to very large spatial data sets while being trainable from a small number of replicates.

The remainder of the paper is organized as follows. In Sect. 2, we provide an overview of the methodology, including a review of Bayesian transport maps, the proposed extension estimation procedures, and computational considerations. Section 3.1 presents numerical comparisons in a simulation study. In Sect. 3.2, we discuss the results of an application to climate-model output. Finally, we conclude in Sect. 4, summarizing the contributions, discussing future directions, and emphasizing the importance of our findings. An implementation of the method described in this manuscript can be found on GitHub (<https://github.com/katzfuss-group/batram>).

## 2. METHODOLOGY

Consider a  $P$ -variate spatial field, where the locations at which the observations are available might differ among the marginal fields. Let  $y(p, s)$  denote the random variable associated with the  $p$ -th marginal field at location  $s \in \mathcal{S}$ , where  $\mathcal{S}$  is the spatial domain and  $p = 1, \dots, P$ . Moreover, let  $\mathbf{y}_p = (y(p, s_1^{(p)}), \dots, y(p, s_{N_p}^{(p)}))^T$  be the random vector associated with the  $p$ -th spatial field observed at locations  $s_1^{(p)}, \dots, s_{N_p}^{(p)} \in \mathcal{S}$ . Finally,  $\mathbf{y} = (\mathbf{y}_1^T, \dots, \mathbf{y}_P^T)^T$  combines the marginal fields by concatenating  $\mathbf{y}_1, \dots, \mathbf{y}_P$ .

We would like to infer the distribution of  $\mathbf{y}$  from  $R$  independent replicates denoted as  $\mathbf{Y} = \mathbf{y}^{(1)}, \dots, \mathbf{y}^{(R)}$ , drawn from the same distribution as  $\mathbf{y}$ . The scalable Bayesian transport map method, proposed by [Katzfuss and Schäfer \(2023\)](#), provides an approach to infer the distribution of a univariate field. The method focuses on modeling the dependence structure using a nonparametric approach assuming  $\mathbf{y}$  to be a field with zero expectation. In this section, we review the existing methodology and propose an extension to multivariate fields.

### 2.1. REVIEW OF THE SCALABLE BAYESIAN TRANSPORT MAP

Throughout this subsection, we assume  $P = 1$  and drop the corresponding index for simplicity. A transport map  $\mathcal{T} : \mathbb{R}^N \rightarrow \mathbb{R}^N$  is a map that characterizes the distribution of  $\mathbf{y}$  by providing a transformation of  $\mathbf{y}$  to a simple reference distribution, e.g.,  $\mathcal{T}(\mathbf{y}) \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_N)$ . Without loss of generality, the map can have a lower triangular form ([Rosenblatt](#)

1952; Carlier et al. 2009) such that

$$\mathcal{T}(\mathbf{y}) = \begin{pmatrix} \mathcal{T}_1(y_1) \\ \mathcal{T}_2(y_1, y_2) \\ \vdots \\ \mathcal{T}_N(y_1, \dots, y_N) \end{pmatrix}$$

with  $\mathcal{T}_n$  being strictly monotone in the  $n$ -th argument. [Katzfuss and Schäfer \(2023\)](#) model  $\mathcal{T}_n(\mathbf{y}_{1:n}) = (y_n - f_n(\mathbf{y}_{1:n-1}))d_n^{-1}$ , where  $\mathbf{y}_{i:j} = (y_i, \dots, y_j)^T$  and  $\mathbf{d} = (d_1, \dots, d_N)^T$ , and  $\mathbf{f} = (f_1, \dots, f_N)^T$  are random vectors. This facilitates the factorization of the joint distribution of  $\mathbf{y}$  as

$$p(\mathbf{y}) = p(y_1) \prod_{n=2}^N p(y_n | y_1, \dots, y_{n-1}) = \prod_{n=1}^N \int \mathcal{N}(y_n | f_n(\mathbf{y}_{1:n-1}), d_n^2) p(f_n, d_n) dd_n df_n, \quad (1)$$

to which we refer as the integrated likelihood.

### 2.1.1. Priors

An independent Gaussian-process inverse-Gamma prior is placed on each pair  $(f_n, d_n^2)$  for  $n = 1, \dots, N$ :

$$\begin{aligned} d_n^2 &\sim \mathcal{IG}(\alpha_n, \beta_n), & \text{with } \alpha_n > 1, \beta_n > 0 \\ f_n | d_n &\sim \mathcal{GP}(0, d_n^2 K_n) & \text{with covariance function } K_n. \end{aligned}$$

The priors' parameters, including parameters determining  $K_n$  here referred to as  $\kappa_n$ ,  $\psi = \{\alpha_1, \dots, \alpha_N, \beta_1, \dots, \beta_N, \kappa_1, \dots, \kappa_N\}$  may depend on a hyperparameter vector  $\theta$ . In the spatial case described by [Katzfuss and Schäfer \(2023\)](#) as well as in our multivariate extension,  $\psi$  deterministically depends on  $\theta$ . Moreover, the dimensionality of  $\theta$  is much smaller than that of  $\psi$ , as the number of hyperparameters does not increase with  $N$ .

Given the substantial number of random variables involved, it becomes imperative to select the prior parameters judiciously. These priors are formulated based on the principles of shrinkage and theoretical considerations pertinent to Gaussian processes with specific covariance functions.

To provide more clarity, we introduce some additional notations. We denote  $o(n)$  as a sequence that orders the values preceding  $y_n$  by increasing the distance of their associated locations to  $s_n$ . Formally,  $o(n)$  is the sequence such that for  $1 \leq i < j < n$ , it holds that  $\|s_n - s_{o(n)i}\| \leq \|s_n - s_{o(n)j}\|$ . Additionally, we define  $\ell_n$  as the minimum distance from  $s_n$  to any of its preceding neighbors, specifically as  $\ell_n = \min_{i \in 1, \dots, n-1} \|s_n - s_i\| = \|s_n - s_{o(n)1}\|$ .

Based on the reasoning of [Schäfer et al. \(2021b\)](#), [Katzfuss and Schäfer \(2023\)](#) observe a roughly power-functional decay of the conditional variance  $\text{Var}[y_n | f_n, \mathbf{y}_{1:n}] = d_n^2$  with respect to the distance to the nearest neighbor  $\ell_n$ . The prior on  $d_n^2$  captures this relationship by incorporating hyperparameter values  $\alpha_n$  and  $\beta_n$  derived from the condition that the prior expectation and prior standard deviation of  $d_n^2$  should satisfy  $E[d_n^2] = \exp(\theta_1^d + \exp(\theta_2^d) \log(\ell_n))$  and  $\text{SD}[d_n^2] = gE[d_n^2]$  with constant  $g > 0$ ,

respectively. A relatively weak prior is constructed by setting  $g = 4$ . The number of hyperparameters in the priors of  $d_n^2$  that need to be estimated reduces from  $2N$  to 2.

The inference of  $f_n$  becomes feasible through the application of two key principles. Firstly, as the distance to the nearest neighbor, denoted as  $\ell_n$ , decreases, the prior exerts a stronger tendency to push  $f_n$  toward linearity. Secondly, drawing motivation from the “screening effect” (Stein 2011) and the work by Schäfer et al. (2021b), the inputs to  $f_n$  are regularized such that the relevance of each input diminishes exponentially according to its position in the ordered input sequence, namely,  $\mathbf{y}_{o(n)} = (y_{o(n)_1}, \dots, y_{o(n)_{n-1}})^T$ . These principles lead to the covariance function

$$K_n(\mathbf{y}_{o(n)}, \mathbf{y}'_{o(n)}) = (\mathbb{E}[d_n^2])^{-1} \left( \mathbf{y}_{o(n)}^T \mathbf{Q}_n \mathbf{y}'_{o(n)} + \sigma_n^2 \rho \left( \frac{\sqrt{(\mathbf{y}_{o(n)} - \mathbf{y}'_{o(n)})^T \mathbf{Q}_n (\mathbf{y}_{o(n)} - \mathbf{y}'_{o(n)})}}{\gamma} \right) \right).$$

Here,  $\mathbf{Q}_n = \text{diag}_j(\exp(-j \exp(\theta^q)))$  encodes the decreasing relevance of more distant inputs,  $\rho$  represents the Matérn correlation function with three-halves smoothness, and  $\gamma = \exp(\theta^\gamma)$  serves as a range parameter. The parameter  $\sigma_n$  governs the nonlinearity of  $f_n$  and, a priori, decays in a manner similar to  $\mathbb{E}(d_n^2)$ , specifically as  $\sigma_n^2 = \exp(\theta_1^\sigma + \exp(\theta_2^\sigma) \log(\ell_n))$ . The hyperparameters that determine the values of  $\psi$  are encompassed within the hyperparameter vector  $\theta = (\theta^q, \theta^\gamma, \theta_1^d, \theta_2^d, \theta_1^\sigma, \theta_2^\sigma)^T$ .

### 2.1.2. Vecchia Approximation and Hyperparameter Estimation

To enable inference for small training sizes  $R$  and to ensure scalability to large spatial fields  $N$ ,  $f_n$  is assumed to depend only on a conditioning set of restricted size. Concretely,  $f_n(\mathbf{y}_{1:n-1})$  is replaced by  $f_n(\mathbf{y}_{c(n)})$ , where  $c(n) \subseteq \{1, \dots, n-1\}$  with  $|c(n)| = \min(m, n-1)$  and  $\mathbf{y}_{c(n)}$  is the subvector of  $\mathbf{y}_{1:n-1}$  with the indices found in  $c(n)$ . The good approximation properties of the reduced conditioning set are achieved by ordering the vector  $\mathbf{y}$  according to the maximum-minimum (maxmin) ordering and then selecting the nearest neighbors for the reduced conditioning set  $c(n)$ . For a comprehensive discussion of different orderings, refer to Guinness (2018). In the case of the scalable Bayesian transport map, the computational complexity to determine the  $n$ -th element of the posterior map reduces from  $\mathcal{O}(R^3 + nR^2)$  to  $\mathcal{O}(R^3 + mR^2)$ .

The size of the conditioning set is driven by the relevance decay incorporated into the prior for  $f_n$ . This decay follows an exponential pattern based on the position within the conditioning set. Consequently, we limit the conditioning set size such that the conditioning set only comprises neighbors whose contributions are still considered relevant, e.g.,  $m = \max\{j \geq 1 : \exp(-j \exp(\theta^q)) \geq \epsilon\}$ . In our analysis, we set  $\epsilon = 0.01$  and estimate the value of  $\theta^q$ . Thus, the size of the conditioning set is automatically determined.

Katzfuss and Schäfer (2023) suggest employing an empirical Bayes (EB) approach for inference. Due to conjugacy, the integrated likelihood (with  $f$  and  $\mathbf{d}$  integrated out) in Eq. (1) is available in closed form and can be maximized using numerical methods, leading

to  $\hat{\boldsymbol{\theta}} = \arg \max_{\boldsymbol{\theta}} p(\mathbf{Y}|\boldsymbol{\theta})$ . For the spatial case, the authors discuss details that motivate the prior choice, the relationship between  $\boldsymbol{\theta}$  and  $\boldsymbol{\psi}$ , as well as closed-form expressions for the integrated likelihood and predictive posterior distribution of  $\mathbf{y}|\mathbf{Y}, \hat{\boldsymbol{\theta}}$ .

## 2.2. EXTENSION TO MULTIVARIATE SPATIAL FIELDS

As in the univariate case, we aim to estimate the joint distribution of  $\mathbf{y}$ , allowing for non-Gaussian dependence. We propose employing the BTM approach in a higher-dimensional input space  $\tilde{\mathcal{S}}$ . We construct the new input space by combining the spatial domain  $\mathcal{S}$  with a latent process space  $\tilde{\mathcal{S}} = \mathbb{R}^{P-1}$  in which the marginal spatial fields are positioned relative to each other. These process positions  $\tilde{\mathbf{s}}_p \in \tilde{\mathcal{S}}, p = 1, \dots, P$ , serve to position the univariate spatial fields relative to one another. In the latent process space, smaller distances between  $\tilde{\mathbf{s}}_p$  and  $\tilde{\mathbf{s}}_{p'}$  indicate stronger associations between fields  $\mathbf{y}_p$  and  $\mathbf{y}_{p'}$ .

We combine the process space with the spatial domain, resulting in the augmented input space  $\tilde{\mathcal{S}} = \mathcal{S} \times \tilde{\mathcal{S}}$ . Now, each element  $y_n$  of  $\mathbf{y}$  is associated with a location in the augmented input space, represented by  $\tilde{\mathbf{s}}_n \in \tilde{\mathcal{S}}$ . Figure 1 provides a visual representation of this concept using a toy example consisting of a two-variate field on a one-dimensional spatial domain. The figure displays the spatial domain along the  $x$ -axis and the process space along the  $y$ -axis. Each point in the plot represents a location in the augmented input space, combining the spatial domain and the process space. The top row of the figure depicts less-correlated processes with a greater distance between them, while the lower row illustrates strongly dependent processes with a smaller distance. In the figure, two situations for the position in the ordering are considered, one early ( $n = 7$ ) and one late ( $n = 23$ ) in the ordering. The locations are ordered according to maxmin ordering and the reduced conditioning sets are indicated by the red circles, showcasing how reduced conditioning sets early in the ordering consist of observations from both processes while later in the ordering observations from the other process enter the conditioning sets only if the processes are close. The latter part is explained with more detail below.

### 2.2.1. Vecchia Approximation in the Augmented Input Space

The transport map formulation casts the problem of inferring the distribution of  $\mathbf{y}$  into the task of solving  $N$  independent regression problems of the form

$$y_n | f_n, d_n, \mathbf{y}_{1:n-1} \sim \mathcal{N}(f_n(\mathbf{y}_{1:n-1}), d_n^2) \quad (2)$$

$$f_n | d_n, K_n \sim \mathcal{GP}(0, d_n^2 K_n) \quad (3)$$

$$d_n^2 | \alpha_n, \beta_n \sim \mathcal{IG}(\alpha_n, \beta_n). \quad (4)$$

As described above, the estimation is computationally infeasible for larger dataset sets. Similar to [Katzfuss and Schäfer \(2023\)](#), we address the problem by conditioning  $y_n$  only on the sub-vector  $\mathbf{y}_{c(n)}$ ,  $c(n) \subseteq \{1, \dots, n-1\}$  of  $\mathbf{y}_{1:n-1}$ , rendering  $y_n$  independent of  $\mathbf{y}_{\{1, \dots, n-1\} \setminus c(n)}$  given  $\mathbf{y}_{c(n)}$ ,  $f_n$ , and  $d_n$ .

Our method uses the Euclidean distance within the augmented input space for computing the maxmin ordering and the conditioning sets. Following the maxmin ordering, we define

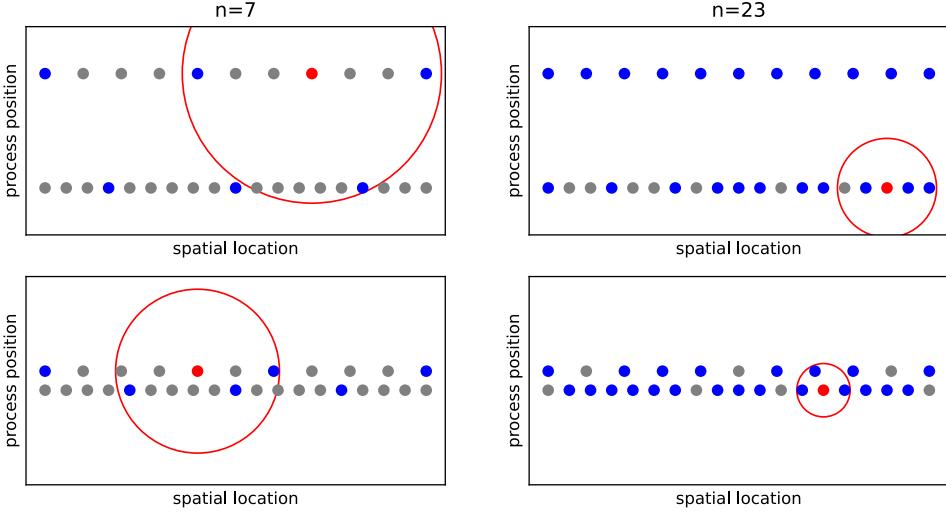


Figure 1. Stylized illustration of the augmented input space for two one-dimensional spatial fields ( $P = 2$  and  $\dim(\mathcal{S}) = 1$ ). The top row corresponds to weakly associated fields, while the bottom row represents strongly associated fields. The dots in the figure represent locations within the augmented input space. The colors of the dot correspond to the position relative to  $n$  in the maxmin ordering. The  $n$ -th ordered location is colored red, while the locations previously ordered are colored blue and the locations subsequent in the ordering are colored gray. Here, the situation for  $n$  equals 7 (left column) and  $n$  equals 23 (right column) are considered. The circles in the figure depict the radius of the largest distance from the  $n$ -th location to locations included in the reduced conditioning set  $c(n)$  of size  $m = 3$ . The conditioning sets are constructed using the three nearest neighbors. It is important to note that only observations preceding the  $n$ -th location in the maxmin ordering, these are the blue colored dots, can be included in the conditioning set. Consequently, blue dots within the circle form the conditioning set for the red dot. Analyzing the illustration, we can observe that the conditioning set comprises observations from both fields in the early stages of the ordering. However, as the index progresses in the ordering, the conditioning set primarily includes spatial neighbors from the same field. This effect is more pronounced for weakly associated fields.

the  $n$ -th index in the sequence  $\Omega$  that orders  $\mathbf{y}$  such that it maximizes the minimum distance to all previously ordered locations. In mathematical terms,

$$\Omega_n = \arg \max_{i \in \{1, \dots, N\} \setminus \Omega_{1:n-1}} \min_{j \in \Omega_{1:n-1}} \|\tilde{s}_i - \tilde{s}_j\|.$$

Notably, the first index in  $\Omega$  can be selected arbitrarily, and in our implementation, we opt for the most central point as the initial choice. This ordering concept exhibits parallels with the space-filling maxmin-distance design (Johnson et al. 1990; Pronzato and Müller 2012), although it operates within the spatial constraints of available locations. Furthermore, it applies the maxmin criterion sequentially in each step, as opposed to global optimization. We view this ordering as space-filling on different resolutions, initially filling the space on a coarse scale and subsequently decreasing the distance to the nearest location.

Following the literature, we define reduced conditioning sets  $c(n)$  as the  $\min(n-1, m)$  nearest (with respect to Euclidean distance) previously ordered neighbors. More precisely, the  $i$ -th element is given as  $c(n)_i = \arg \min_{j \in \{1, \dots, n-1\} \setminus \Omega_{1:i-1}} \|\tilde{s}_n - \tilde{s}_j\|$ . This choice is motivated by the so-called screening effect (Stein 2011), which states that for many popular covariance functions, such as the Matérn covariance function, random variables in a spatial

field are (almost) independent of distant random variables conditioned on the values in-between.

In a stylized example shown in Fig. 1, we demonstrate this approach. Consider two relatively smooth spatial fields with a weak association between them. In this scenario, the distance to the nearest spatial neighbor is smaller than the distance to the next process. As we employ the maxmin ordering, locations chosen almost alternate between processes early in the ordering. Consequently, the first conditioning sets are likely to contain information from both processes. However, as the ordering progresses, the conditioning sets will predominantly include information about the spatial neighbors. On the other hand, when dealing with fields that exhibit a strong association, the conditioning sets will encompass a larger number of values from the other process even later in the ordering.

Using this Vecchia-type approximation introduces dependence of  $\mathbf{y}$ 's modeled distribution on its order, and the choice of the conditioning sets  $c_{(n)}$ . The combination of employing the maxmin ordering and selecting nearest neighbors for the conditioning set has been widely recognized in the literature as an effective approach (Datta et al. 2016; Heaton et al. 2019; Huang et al. 2021; Katzfuss and Guinness 2021).

### 2.2.2. Distance Metric

Besides their direct involvement in the prior construction, the input locations, or, more strictly speaking, the distances, also determine the ordering and are used to find the nearest neighbors. Both are essential to the quality of the Vecchia approximations. The involvement of distances in the posterior density allows gradient computations for the distances and, consequently, the process positions, thus, enabling a gradient-based update of the process positions.

Conversely, choosing an adequate distance metric is crucial for the model's quality. Using the Euclidean distance may be a natural choice when applied to spatial locations. In other scenarios, including the multivariate case discussed in this article, the Euclidean distance is not necessarily meaningful, and other distance measures can be used. For example, Kang and Katzfuss (2023) explore a correlation-based distance metric for GP regression. In contrast, we propose employing the Euclidean distance in the augmented input space. We argue that the Euclidean distance is meaningful in the higher-dimensional augmented input space since the distance between each input pair  $\tilde{s}_p, \tilde{s}_{p'}$  is composed of both the distance between the spatial locations and the distance between the process locations. By appropriately scaling the latent process space, the distances between the processes can be interpreted similarly to distances in the spatial domain.

### 2.2.3. Parameterization of the Process Positions

We aim to include the process positions in the hyperparameter vector  $\boldsymbol{\theta}$ . However, the model depends only on the relative process positions, because the model depends on the input locations only through the distances. For any two indices  $n$  and  $n'$  within the range of  $1, \dots, N$ , the squared Euclidean distance between the augmented input locations  $\tilde{s}_n$  and  $\tilde{s}_{n'}$  can be decomposed as  $\|\tilde{s}_n - \tilde{s}_{n'}\|^2 = \|s_n - s_{n'}\|^2 + \|\check{s}_n - \check{s}_{n'}\|^2$ . Here,  $s_n$  represents the spatial location and  $\check{s}_n$  represents the process position associated with  $y_n$ ; analogous for

symbols with index  $n'$ . As the model depends on the process positions only through their distances, we can fix rotation and the first location. For this purpose, we define  $\check{s}_1 = \mathbf{0}$  and decrease the degrees of freedom for each subsequent process position from  $P - 1$  to 1 by one at a time. Jointly, the process positions  $\check{\mathbf{S}}$  are parameterized as

$$\check{\mathbf{S}} = \begin{pmatrix} \check{s}_1^T \\ \vdots \\ \check{s}_P^T \end{pmatrix} = \begin{pmatrix} \mathbf{0}^T \\ \mathbf{Q} \mathbf{R} \end{pmatrix}, \quad (5)$$

where the columns of  $\mathbf{Q}$  form an orthonormal basis of  $\mathbb{R}^{P-1}$  and  $\mathbf{R}$  is an upper triangular matrix with positive entries on the diagonal. Only the  $P(P - 1)/2$  nonzero entries of  $\mathbf{R}$  are included as hyperparameters in the model.

### 2.3. ESTIMATION

The estimation procedure consists of two stages. In the first stage, we fit a parametric and separable GP model to obtain initial values for the process positions  $\check{\mathbf{S}}$ . These initial values serve as a starting point for the subsequent estimation of the transport map hyperparameters. For the estimation of the hyperparameters, we explore three different approaches, which are detailed below. Finally, using the estimated hyperparameters, the transport map can be employed as a generative model to simulate new samples that capture the spatial dependencies as learned from the data. The transport map can also be used for uncertainty quantification. The multivariate extension allows us to study the conditional distribution of the spatial fields corresponding to one or more variables given the observed spatial fields of other variables.

#### 2.3.1. Obtaining the Initial Process Positions

To determine the initial process positions, specifically  $\check{s}_j \in \mathcal{S}_p$  referring to the second part of the augmented input locations, we fit a separable parametric model. In this model, we evaluate a separable covariance function parameterized with  $\zeta$  at the locations  $s_1, \dots, s_N$  to obtain the covariance matrix  $\mathbf{K}_\zeta$ . For the two elements  $y_n, y_{n'}$  with the associated spatial locations  $s_n, s_{n'}$  and processes  $p, p'$ , the covariance between  $y_n$  and  $y_{n'}$  is given by

$$\text{Cov}[y_n, y_{n'} | \zeta] = \tau_\zeta^2 C_1(\|s_n - s_{n'}\| | \zeta) C_2(p, p' | \zeta) + \sigma_\zeta^2 \mathbb{1}_{s_n = s_{n'} \wedge p = p'}.$$

Here,  $\tau_\zeta^2, \sigma_\zeta^2$  are variance parameters,  $C_1$  corresponds to a parametric isotropic correlation function, and  $C_2$  corresponds to an unstructured  $P \times P$  correlation matrix  $\check{\mathbf{K}}_\zeta$ . Moreover,  $\mathbb{1}$  is the indicator function. To map  $\binom{P}{2}$  unrestricted elements to  $\check{\mathbf{K}}_\zeta$ , we employ the mapping described in Stan Development Team (2023, Section 10.12).

To estimate  $\zeta$ , we maximize the likelihood of  $\mathbf{y} \sim \mathcal{N}(\mathbf{0}, \mathbf{K}_\zeta)$  by considering all available samples of  $\mathbf{y}$ . Considering the availability of multiple observations of  $\mathbf{y}$ , we have found that thinning the spatial density can effectively manage computational time constraints. By randomly choosing a subset of locations, we can alleviate the computational burden asso-

ciated with factorizing the covariance matrix  $\mathbf{K}_\zeta$ . This approach strikes a balance between computational efficiency and ensuring reliable parameter estimation for  $\zeta$ .

**Retrieving the hyperparameter values** Using the estimated parameter  $\hat{\zeta}$ , a distance between two marginal spatial fields can be computed by applying the inverse of the correlation function  $C_1$  on the value of  $C_2$ . For instance, the estimated distance between the  $p$ -th and  $p'$ -th marginal spatial fields is  $\hat{d}_{pp'} = C_1^{-1}(|C_2(p, p')|\hat{\zeta})$ . Arranging the estimated distances in a  $P \times P$  distance matrix  $\hat{\mathbf{D}} = (\hat{d}_{ij})_{1 \leq i \leq P, 1 \leq j \leq P}$  allows us to easily compute a set of process positions that give rise to  $\hat{\mathbf{D}}$  (Young and Householder 1938; Torgerson 1952).

Given a distance matrix  $\mathbf{D} = (d_{ij})_{1 \leq i \leq P, 1 \leq j \leq P}$ , it is possible to determine a set of locations that generate this matrix up to Euclidean transformations. Let the elements of the matrix  $\mathbf{E} = (e_{ij})_{1 \leq i \leq P, 1 \leq j \leq P}$  be  $e_{ij} = d_{1,j}^2 + d_{i0}^2 - d_{ij}^2$ . Then, the matrix representing the process locations is  $\check{\mathbf{S}} = \mathbf{u} \Lambda^{1/2}$  where  $\mathbf{u}$  and  $\Lambda$  are obtained by performing the eigenvalue decomposition  $\mathbf{E} = \mathbf{u} \Lambda \mathbf{u}^{-1}$  and defining  $\Lambda^{1/2}$  as the element-wise square root. Subsequently, a QR decomposition can be performed on the bottom  $P - 1$  rows of  $\check{\mathbf{S}}$  to calculate  $\mathbf{Q}$  and  $\mathbf{R}$ . (see Equation (5)). Should the values on the diagonal of  $\mathbf{R}$  not be positive, use  $\mathbf{Q}\Upsilon$  and  $\Upsilon\mathbf{R}$  instead of  $\mathbf{Q}$  and  $\mathbf{R}$  where  $\Upsilon$  is a square diagonal matrix with the entries 1 or -1 corresponding to the sign of the diagonal elements of  $\mathbf{R}$ .

### 2.3.2. Estimation of the Multivariate Transport Map

We extend the empirical Bayes approach used by Katzfuss and Schäfer (2023) to estimate the hyperparameters  $\theta$ . The nonzero elements of  $\mathbf{R}$  are included in  $\theta$ , and suitable transformations are applied to ensure that  $\theta$  remains unrestricted.

To estimate the hyperparameters  $\theta$ , we employ a gradient-based optimization algorithm to maximize the integrated likelihood

$$p(\mathbf{y}|\theta) = \prod_{n=1}^N \int p(y_n|\mathbf{y}_{c(n)}, \mathbf{d}_n, \mathbf{f}_n) p(\mathbf{f}_n, \mathbf{d}_n|\theta) d\mathbf{d}_n d\mathbf{f}_n. \quad (6)$$

The presented method focuses on inferring the distribution of  $\mathbf{y}$  rather than on conducting inference for  $\mathbf{f}$  and  $\mathbf{d}$ , providing a significant computational advantage, especially since the formulation with  $\mathbf{f}$  and  $\mathbf{d}$  integrated out permits the utilization of mini-batching, allowing for efficient computations on data subsets. We utilize the Adam optimizer (Kingma and Ba 2014), which adapts the learning rate during optimization, and implement early stopping as a regularization technique. The early stopping criterion is based on monitoring the improvement of the integrated log-likelihood on a separate validation dataset. We define a patience parameter, typically set to 5% or 10% of the maximum number of iterations. If no improvement in the integrated log-likelihood is observed within the last patience steps, the optimization process is terminated. Upon completion of the optimization, we select the parameter set that achieved the highest integrated log-likelihood on the test data. This parameter set represents the optimal configuration based on the performance of the model on unseen data.

However, in terms of estimating the process positions, we consider three strategies:

- (i) Constant process positions (CPP): In this strategy, the process positions, estimated using the parametric model, are assumed to be known and remain fixed during the hyperparameter estimation process. No updates to the process positions are made.
- (ii) Frozen Ordering (FO): With this approach, the process positions are updated during the hyperparameter estimation via a gradient-based update. Recall, the process positions enter the integrated likelihood through the integration of the distance to the nearest neighbor  $\ell_n$  in the prior on  $d_n^2$ . The ordering and conditioning sets, however, are determined using the initial process positions and remain unchanged throughout the optimization.
- (iii) Occasional Re-ordering (OR): Here, the process positions are updated during the hyperparameter estimation via a gradient-based update. Following the idea of [Kang and Katzfuss \(2023\)](#), we recompute the ordering and conditioning sets after a pre-specified number of iterations (e.g., after 4, 8, 16, 32, . . . iterations). As the integrated likelihood is based on a Vecchia approximation (using Equation (6) instead of Equation (1)), reordering may produce distinct likelihood values. Furthermore, the reordering procedure does not necessarily lead to an improvement in the integrated likelihood. Hence, we reset the patience counter and update the best encountered integrated likelihood to the current value. This ensures that the algorithm will continue in searching for a better parameter value for at least the specified number of epochs and considers only the integrated likelihood values encountered in the current ordering. Additionally, determining these updates can be computationally expensive as one has to consider the entire data set. However, [Schäfer et al. \(2021b\)](#) presented an algorithm to compute ordering and conditioning sets in quasi-linear time complexity, i.e.,  $\mathcal{O}(N \log^2(N))$ .

For ease of comprehension, we provide the estimation algorithm in pseudo-code in Appendix 4.

### 3. NUMERICAL COMPARISON

#### 3.1. SIMULATION STUDY

Our study focuses on two objectives: (1) learning the distribution of  $\mathbf{y}$  and (2) learning the conditional distribution of one variable given the others, i.e., the distribution of  $\mathbf{y}^1|\mathbf{y}^2, \dots, \mathbf{y}^P$ . Additionally, we aim to compare the performance of the different estimation strategies: Constant process positions (CPP), Frozen Ordering (FO), and Occasional Re-ordering (OR). To assess the of the MVTM, we compare it with a parametric model. We refrain from a comparison to VAE ([Kingma and Welling 2014](#)) and a GAN designed for climate-model output ([Besombes et al. 2021](#)) as those deep-learning methods have turned out as not competitive in an application similar to ours (see the supplementary materials in [Katzfuss and Schäfer 2023](#) for details).

To evaluate the performance, we assess the average log-density of the learned distribution at 20 test samples. This evaluation metric provides an approximation, up to an additive

constant, of the negative Kullback–Leibler (KL) divergence  $D(p||\hat{p})$  between the true distribution  $p$  and the estimated distribution  $\hat{p}$ . The log-density serves as a positively oriented measure, allowing us to compare the accuracy and goodness-of-fit of the learned distribution to the true distribution.

### 3.1.1. Experimental Setup

We vary two key factors in our simulation study: the training size and the number of variables. Specifically, we consider training sizes  $R \in \{10, 20, 30, 40, 60, 80\}$ , and explore scenarios with  $P \in \{2, 3, 4, 5\}$  variables. By varying these factors, we aim to examine the behavior and performance of the estimation procedures under different data settings.

For the data-generating process, we adopt Scenario NR900 in [Katzfuss and Schäfer \(2023\)](#). This data-generating process can be characterized by a transport map. Using the hierarchical formulation in Sect. 2.2.1, we specify the functions  $f_i$  as additively composed from a linear and nonlinear part

$$f_i(\mathbf{y}_{1:i-1}) = \mathbf{b}_i^T \mathbf{y}_{c(i)} + 2 \sin(4(b_{i,1} y_{c(i)1} + b_{i,2} y_{c(i)2}))$$

with  $\mathbf{b}_i = (b_{i,1}, \dots, b_{i,|c(i)|})^T$ , where the  $b_{i,k}$  are based on the exponential covariance function with range 0.3 and the distance between  $\check{\mathbf{s}}_i$  and  $\check{\mathbf{s}}_{c(i)k}$ . The  $P$ -variate spatial field is observed on a regular grid of size  $32 \times 32$  on the unit-square and we use the augmented input space with the process positions  $\check{\mathbf{s}}_1 = \mathbf{0}$ ,  $\check{\mathbf{s}}_2 = (0.2, 0, 0)^T$ ,  $\check{\mathbf{s}}_3 = (0.0, 0.3, 0)^T$ ,  $\check{\mathbf{s}}_4 = (0.0, 0, 0.4)^T$ , and  $\check{\mathbf{s}}_5 = (0.3, 0.3, 0)^T$ . Thus, the  $P$ -variate spatial field is observed in 1,024 locations per process, which gives 5,120 locations for a five-variate spatial field in the augmented input space. In total, roughly 400,000 datapoints must be considered in the scenario with 80 replicates of the five-variate field.

For fast estimation of the  $6 + \binom{P}{2}$  hyperparameters  $\boldsymbol{\theta}$ , we utilize mini-batching with a batch size of 256, resulting in  $4P$  gradient updates per epoch. Convergence monitoring is conducted by evaluating the integrated log-likelihood on an independently generated validation dataset with 20 replicates. We employ early stopping with a patience of 25 and a maximum of 500 iterations. Remarkably, all estimations terminate early, indicating successful convergence. To optimize the model, we utilize the Adam optimizer with an initial learning rate of 0.01 and apply cosine annealing to mitigate the variance induced by mini-batching. For estimating the initial process positions, we randomly select 256 of the 1,024 spatial locations. Importantly, we observe no adverse effects on the estimated parameters due to the subsampling. In our evaluation, we also include as a competitor procedure a mean zero Gaussian model with isotropic Matérn covariance function combined with an unstructured  $P \times P$  correlation matrix as described in Sect. 2.3.1, whose  $3 + \binom{P}{2}$  hyperparameters are estimated via maximum likelihood. We refer to this model as the parametric model.

### 3.1.2. Results and Analysis

The results indicate that the transport map outperforms the parametric model when the ensemble size exceeds approximately 20–25 (refer to Fig. 2). Interestingly, incorporating the process positions into the hyperparameter estimation of the transport map (OR) did not

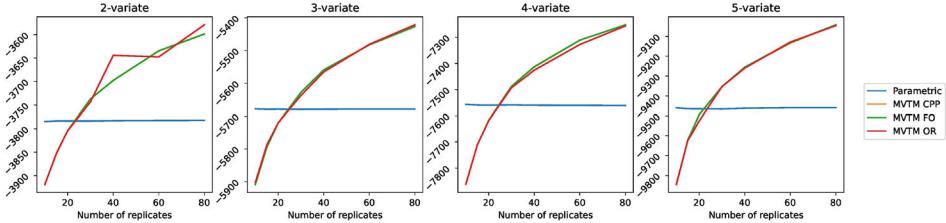


Figure 2. Mean of the estimated log-density of test data sets for different ensemble sizes  $R$  (x-axis) and varying dimension  $P$  of the multivariate spatial field (panels). The “Parametric” approach refers to the parametric GP, while approaches labeled with “MVTM CPP”, “MVTM FO”, and “MVTM OR” correspond to the multivariate Bayesian transport map with the estimation strategies in Sect. 2.3.2. Note that the orange line is almost completely covered by the green line (Color figure online).

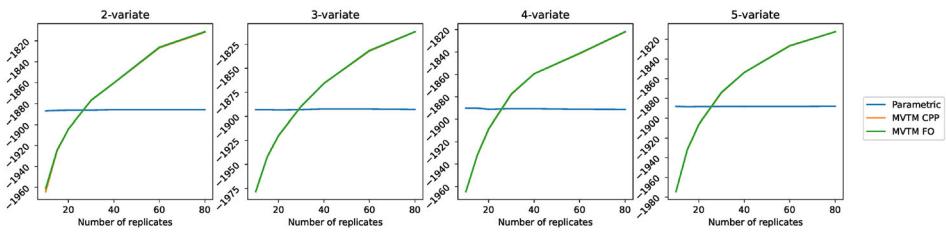


Figure 3. Log-density of  $y_1$  conditioned on  $y_2, \dots, y_P$ . The plots show the mean log-density evaluated on test data sets for different ensemble sizes  $R$  (x-axis) and different dimensions  $P$  of the multivariate spatial field. “Parametric” refers to the parametric approach, “MVTM CPP” and “MVTM FO” refer to the Bayesian transport map with the estimation strategies described in Sect. 2.3.2.

lead to an improvement in model fit. However, it is worth noting that this might be attributed to the limited nonlinearity in the data-generating process.

In addition to estimating the joint distribution of  $y$ , we also investigate the conditional distribution of  $y_1$  given  $y_2, \dots, y_P$ . To account for the conditional distribution, we use a modified maxmin-ordering in which we order the indices corresponding to  $y_1$  subsequent to all other indices. Based on the findings from the initial study, we exclude the OR estimation strategy from this analysis as it did not yield improved results. Again, the approximated KL-divergences suggest that the MVTM is superior given enough training data. See Fig. 3 for a visual presentation of the estimated log densities.

In summary, the findings from our simulation study indicate that the MVTM approach performs better when a sufficient amount of training data is available. We observe that the MVTM approach outperforms the alternative methods in capturing the underlying spatial and inter-process dependencies given enough available training data. Interestingly, the choice of estimation strategies seems to have no significant impact on the performance. This lack of impact could be due to the limited nonlinearity present in the data generating process. Specifically, only the two nearest neighbors in the conditioning set, which belong to the same process for a large portion of the data, contribute to the conditional expected value in a nonlinear manner. Consequently, the parametric approach is able to effectively capture the dependence between the spatial fields. We also explored an alternative configuration, involving the scaling of weights  $b_{i,1}$  and  $b_{i,2}$  based on the process affiliations of both  $y_i$

and its corresponding neighbors  $y_{c(i)_1}$  and  $y_{c(i)_2}$ . This leads to a scaling of the nonlinear functions based on process affiliations. We omit a detailed presentation, as the results are comparable, with MVTM outperforming the parametric GP with only  $R = 20$  training replicates. Note that our study does not explore spatial data sparsity or density, as previous results by [Katzfuss and Schäfer \(2023\)](#) suggests that the performance differences among the methods are similar under such conditions. Therefore, in the context of our simulation study, the focus is primarily on the performance of the MVTM approach in relation to the availability of training data.

### 3.2. APPLICATION: CLIMATE MODEL OUTPUT

Climate models serve as computational tools essential for simulating and comprehending the Earth's climate system, playing a vital role in climate research and policy-making. In essence, climate models are computer programs that describe the Earth's climate system through sets of differential equations. Developing as well as running these models demands a lot of resources, time, and specialized high-performance computers. Each run takes a considerable amount of time and consumes significant energy. For instance, 17 million core hours were spent on the computation of large ensemble of the Community Earth System Model (CESM) using the Yellowstone supercomputer ([Kay et al. 2015](#)). It takes approximately three weeks to produce each ensemble member.

Nevertheless, conducting multiple runs is typically imperative due to the potential impact of even slight perturbations in the initial conditions, which can result in larger variations at the end of a model run. Therefore, it is crucial to perform multiple runs of climate models to account for the uncertainty arising from these initial conditions. Due to this nature, a climate model can be interpreted as encoding a distribution of climate rather than predicting an exact Earth-system state, rendering the analysis of climate model even more challenging. Statistical emulators, which replicate this distribution, can be employed to summarize the distribution and generate additional samples at much lower computational costs.

Climate models produce numerous variables with high temporal and spatial resolution, often exhibiting significant non-stationarity. In this study, we specifically examine data obtained from the large ensemble project (LENS) of the CESM developed by the National Center for Atmospheric Research. The LENS consists of 42 ensemble members with variations due to slightly different initial atmospheric state.

Our focus centers on four hydrological variables generated by the land-surface-model (LSM) component.<sup>1</sup> We restrict our focus further to the conterminous United States and consider the yearly average of the variables *SNOW*, *RAIN*, *SOILWATER\_10CM*, and *QRUNOFF* in the year 2001. The data is available on a roughly  $1^\circ$  longitude-latitude grid yielding 690 values per variable and ensemble member. In a prepossessing step, we combine the first two variables to a new variable representing the combined precipitation (*PRECIP*). As we aim to model anomalies, all raw observations are transformed by subtracting the pixel-wise mean and dividing by the pixel-wise standard deviation.

---

<sup>1</sup>Data is available at <https://www.earthsystemgrid.org/dataset/ucar.cgd.cesm4.cesmLE.html>.

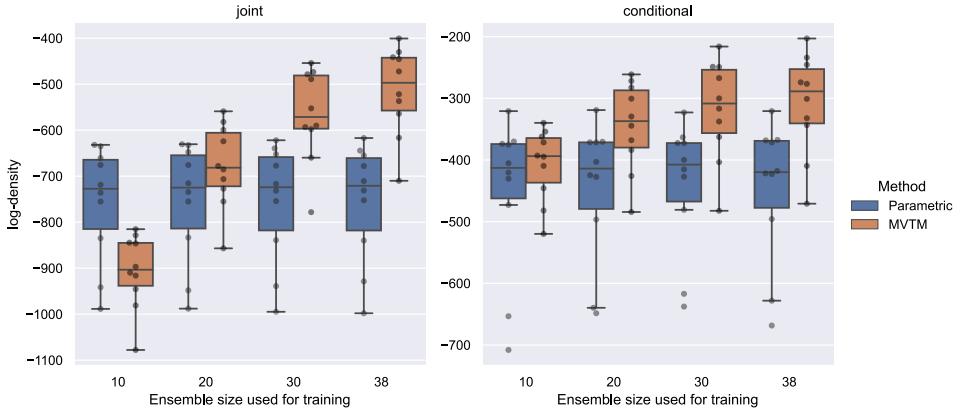


Figure 4. The box plots illustrate the estimated log-density of the holdout data, with the points representing the underlying data. The left panel concerns the log-density of the joint model while the right panel depicts the log-density of *QRUNOFF* conditioned on the remaining variables. The x-axis indicates the size of the ensemble used for training .

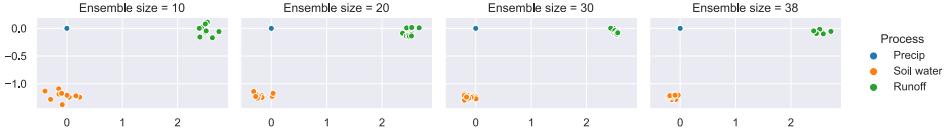


Figure 5. Scatterplots of the latent positions estimated in the joined model for 10 CV splits, stratified by ensemble size used for training. Each point represents the estimated position of process in one replication. The location of the first process (PRECIP) is by construction 0 .

We fit the MVTM using estimation strategy CPP and compare it to the parametric model employing the log-density evaluated for holdout data. Since the other estimation strategies produce similar results, we omit their presentation. In a cross-validation (CV) setting, we select 4 ensemble members as holdout data. From the remaining 38 members, we use 10, 20, 30 and 38 for training to investigate the effect of sample size. The process is repeated until all but the last two data sets are once used as holdout, giving us ten values for the estimated log-density. The left panel in Fig. 4 visualizes the results, strengthening the impression from the simulation study. The MVTM's performance improves with increasing samples size. In comparison, it adapts better to the underlying dependence in the anomalies. The estimated latent locations are illustrated in Fig. 5. Notably, these estimated locations exhibit relatively little variation across CV splits as the training size increases. Nevertheless, it is worth noting a discernible trend of increased point concentration among replications as the sample size grows.

On an Apple M1 Pro equipped with eight performance cores and 32GB of RAM, the estimation process takes an average of 91 seconds of CPU time for ten replications and 194 seconds for 38 replications. This estimation time is distributed between two tasks: estimating the initial latent locations and the hyperparameters of the transport map. Specifically, estimating the initial latent locations consumes 12, 22, 33, and 44 seconds for 10, 20, 30, and 38 replications, respectively. Meanwhile, the estimation of hyperparameters requires between 79, 98, 122, and 150 seconds for 10, 20, 30, and 38 replications, respectively.

The standard deviation of these estimations is roughly 1%, indicating a relatively stable performance across CV-replications.

In addition to studying the joint distribution of the variables, we also investigate the conditional distribution of  $QRUNOFF$  given the other two variables. This analysis is particularly valuable as runoff in an area cannot be directly measured but must be modeled. Having a statistical model that describes the conditional distribution of runoff can be immensely beneficial for practitioners. For example, it enables the study of the conditional distribution of catchment runoff within the drainage basin of a river, which represents the water from sources like rain, snow, and soil moisture flowing into the river. Thus, a model that characterizes the conditional distribution of runoff in an area based on observable quantities can assist in estimating the amount of water a river system needs to handle and evaluating whether infrastructure, such as dams, is adequately designed. From the right panel in Fig. 4, we see that the MVTM describes the conditional distribution already for the smallest training size better than the parametric model. Judging from visual impression of samples drawn from the estimated distribution, the parametric model cannot capture the anisotropy present in the data and seems overall too smooth (see Fig. 6).

## 4. CONCLUSIONS

We have presented a Bayesian approach for learning the distribution of multivariate spatial fields based on a relatively small number of training samples by estimating a transport map. Our method can capture the potential nonlinearity in the conditional dependencies, enabling the learning of non-Gaussian distributions and mapping them to the standard normal. The scalability of the approach is enhanced by the mini-batching capability of the estimation algorithm, allowing for efficient analysis of large datasets. From our numerical demonstrations, we can confidently recommend the multivariate transport map approach, particularly when dealing with non-stationarity in the data and an ample number of replications.

We do not account for uncertainty in the estimation of the hyperparameters  $\theta$ . While fully Bayesian approaches like Markov chain Monte Carlo (MCMC) can in principle address this, we opt for using empirical Bayes (EB) estimation due to computational constraints. This decision is supported by insights from [Katzfuss and Schäfer \(2023\)](#), which indicate minimal impact on the estimated posterior distribution of  $y$  when contrasting EB results with a full Bayesian approach. MCMC's requirement for full dataset processing in each iteration poses scalability challenges for large datasets. Stochastic gradient MCMC methods ([Nemeth and Fearnhead 2021](#)) show promise in bypassing this limitation, offering potential for enhanced hyperparameter uncertainty quantification while maintaining scalability.

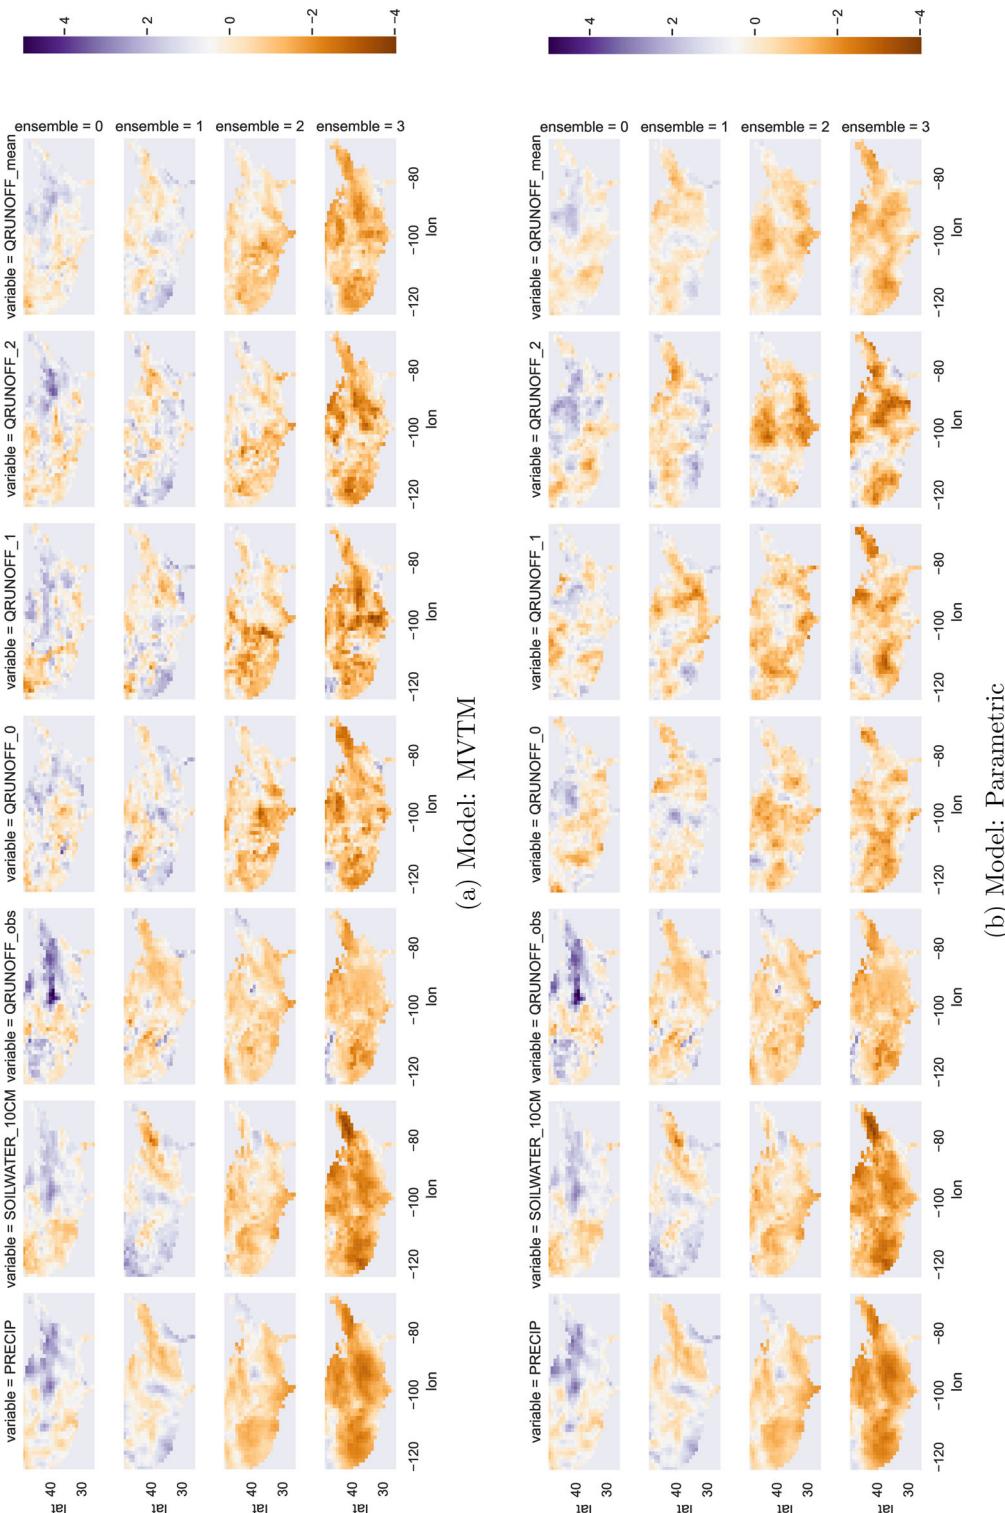


Figure 6. Plots of the anomalies of the variables PRECIP, SOILWATER\_10CM, and QRUNOFF. The figure is divided into two panels: (a) Model: MVTM and (b) Model: Parametric. Each panel contains a 5x5 grid of maps. The columns are labeled 'variable = PRECIP', 'variable = SOILWATER\_10CM', 'variable = QRUNOFF\_obs', 'variable = QRUNOFF\_0', 'variable = QRUNOFF\_1', 'variable = QRUNOFF\_2', and 'variable = QRUNOFF\_mean'. The rows are labeled 'ensemble = 0', 'ensemble = 1', 'ensemble = 2', 'ensemble = 3', and 'ensemble = 4'. Each map shows anomalies on a geographic grid with longitude (lon) from -120 to 30 and latitude (lat) from 30 to 40. Color bars at the top of each column indicate the scale of the anomalies, ranging from -4 (dark blue) to -2 (light blue) to 0 (white) to 2 (light orange) to 4 (dark orange). Panel (a) is titled '(a) Model: MVTM' and panel (b) is titled '(b) Model: Parametric'.

Moving forward, there are several potential extensions and avenues for future research. One promising direction is the incorporation of a temporal component, enhancing the proposed methodology's utility in climate model emulation and capturing temporal variations. Another useful direction to enhance the flexibility and applicability of the MVTM approach is the inclusion of covariates. Here, the distribution of  $y$  is related to covariate values. In the context of climate models, this would allow to interpolation between emission scenarios.

A continuation of the extension presented is to make the model highly multivariate. While providing most flexibility, the current formulation requires relating  $\binom{P}{2}$  entries in the hyperparameter to the process positions. This approach is suitable for relatively small  $P$ . However, it seems prohibitive when considering all 1168 variables in the CESM.

Furthermore, an important extension to consider is relaxing the assumption of conditional normality at each location. By allowing for more flexible modeling, such as accounting for skewed or heavy-tailed data, we can extend the Bayesian transport map to multivariate spatial fields where one field may represent extreme values.

In summary, the MVTM approach offers a powerful tool for learning the distribution of multivariate spatial fields. Its ability to handle nonlinearity, scalability, and potential for future enhancements make it a promising methodology in spatial statistics.

## ACKNOWLEDGEMENTS

PW and MK were supported by NASA's Advanced Information Systems Technology Program (AIST-21). MK was also partially supported by National Science Foundation (NSF) Grant DMS-1953005. We would like to thank Jonathan Hobbs for helpful comments and discussions as well as Daniel Drennan for spotting typos in the manuscript suggestions for improvements.

### Declarations

**Conflicts of interest** The authors report no conflict of interests.

*[Received July 2023. Revised September 2023. Accepted October 2023. Published Online November 2023.]*

## A ALGORITHMS

**Data:** Training data  $Y$ , validation data  $V$ , initial hyper parameters  $\theta$ , batchsize  $bs$ , learning rate  $lr$ , max epochs  $E$ , patience, estimation mode

**Result:** Estimated hyperparameter  $\hat{\theta}$

```

best_loglik ← −∞
 $\hat{\theta} \leftarrow \theta$ 
patience_counter ← 0
/* epoch loop */  

for epoch ← 0 to E do
    batches ← create batches of  $Y$  of size  $bs$ 
    /* for each batch, perform gradient based update */  

    for batch in batches do
        |  $\theta = \theta + lr \nabla_{\theta} \log(\text{int\_lik}(\text{batch}, \theta))$ 
    end
    /* update patience counter; break if patience is exhausted */  

    if  $\log(\text{int\_lik}(V, \theta)) > \text{best\_loglik}$  then
        |  $\text{best\_loglik} \leftarrow \log(\text{int\_lik}(V, \theta))$ 
        |  $\hat{\theta} \leftarrow \theta$ 
        | patience_counter ← 0
    else
        | patience_counter ← patience_counter + 1
    end
    if patience_counter > patience then
        | break loop
    end
    /* check if reordering is required and perform if so */  

    if mode is OR and epoch indicates reordering then
        |  $Y \leftarrow$  update ordering and conditioning sets of  $Y$  using current positions  $\theta$ 
        |  $V \leftarrow$  update ordering and conditioning sets of  $V$  using current positions  $\theta$ 
        | /* reset patience counter */  

        |  $\text{best\_loglik} \leftarrow \log(\text{int\_lik}(V, \theta))$ 
        |  $\hat{\theta} \leftarrow \theta$ 
        | patience_counter ← 0
    end
end
return  $\hat{\theta}$ 

```

**Algorithm 1:** Pseudocode of the algorithm estimating the hyperparameters. The training data  $Y$  and validation data  $V$  contains the ordering and conditioning sets. When the estimation mode is CPP, the elements of  $\theta$  associated with the latent positions are not updated. The epochs after which a reordering happens in estimation mode OR are predefined.

## REFERENCES

Apanasovich TV, Genton MG (2010) Cross-covariance functions for multivariate random fields based on latent dimensions. *Biometrika* 97(1):15–30

Arjovsky M, Bottou L (2017) Towards principled methods for training generative adversarial networks. In: International conference on learning representations

Banerjee S, Carlin BP, Gelfand AE (2004) Hierarchical modeling and analysis for spatial data. Chapman & Hall, London

Besombes C, Pannekoucke O, Lapeyre C, Sanderson B, Thual O (2021) Producing realistic climate data with generative adversarial networks. *Nonlinear Process Geophys* 28(3):347–370

Bigoni D, Spantini A, Marzouk YM (2016) Adaptive construction of measure transports for Bayesian inference. In: NIPS 2016 workshop on advances in approximate Bayesian inference

Carlier G, Galichon A, Santambrogio F (2009) From Knothe's transport to Brenier's map and a continuation method for optimal transport. *SIAM J Math Anal* 41(6):2554–2576

Castruccio S, Stein ML (2013) Global space–time models for climate ensembles. *Ann Appl Stat* 7(3):1593–1611

Choi IK, Li B, Wang X (2013) Nonparametric estimation of spatial and space–time covariance function. *J Agric Biol Environ Stat* 18(4):611–630

Cressie N (1993) Statistics for spatial data, revised. Wiley, New York

Datta A, Banerjee S, Finley AO, Gelfand AE (2016) Hierarchical nearest-neighbor Gaussian process models for large geostatistical datasets. *J Am Stat Assoc* 111(514):800–812

Edwards M, Castruccio S, Hammerling D (2019) A multivariate global spatiotemporal stochastic generator for climate ensembles. *J Agric Biol Environ Stat* 24(3):464–483

El Moselhy TA, Marzouk YM (2012) Bayesian inference with optimal maps. *J Comput Phys* 231(23):7815–7850

Genton MG, Kleiber W (2015) Cross-covariance functions for multivariate geostatistics. *Stat Sci* 30(2):147–163

Goodfellow I, Bengio Y, Courville A (2016) Deep learning. MIT Press, Cambridge

Guinness J (2018) Permutation and grouping methods for sharpening Gaussian process approximations. *Technometrics* 60(4):415–429

Guinness J (2022) Nonparametric spectral methods for multivariate spatial and spatial–temporal data. *J Multivar Anal* 187:104823

Heaton MJ, Datta A, Finley AO, Furrer R, Guinness J, Guhaniyogi R, Gerber F, Gramacy RB, Hammerling DM, Katzfuss M, Lindgren F, Nychka DW, Sun F, Zammit-Mangion A (2019) A case study competition among methods for analyzing large spatial data. *J Agric Biol Environ Stat* 24(3):398–425

Hestness J, Narang S, Ardalani N, Diamos G, Jun H, Kianinejad H, Patwary MMA, Yang Y, Zhou Y (2017) Deep learning scaling is predictable, empirically. [arXiv:1712.00409](https://arxiv.org/abs/1712.00409)

Huang C, Hsing T, Cressie N (2011) Nonparametric estimation of the variogram and its spectrum. *Biometrika* 98(4):775–789

Huang H, Blake LR, Katzfuss M, Hammerling DM (2021) Nonstationary spatial modeling of massive global satellite data. [arXiv:2111.13428](https://arxiv.org/abs/2111.13428)

Johnson ME, Moore LM, Ylvisaker D (1990) Minimax and maximin distance designs. *J Stat Plan Inference* 26(2):131–148

Jun M (2011) Non-stationary cross-covariance models for multivariate processes on a globe. *Scand J Stat* 38(4):726–747

Kang M, Katzfuss M (2023) Correlation-based sparse inverse Cholesky factorization for fast Gaussian-process inference. *Stat Comput* 33(56):1–17

Katzfuss M, Guinness J (2021) A general framework for Vecchia approximations of Gaussian processes. *Stat Sci* 36(1):124–141

Katzfuss M, Schäfer F (2023) Scalable Bayesian transport maps for high-dimensional non-Gaussian spatial fields. *J the Am Stat Assoc*. <https://doi.org/10.1080/01621459.2023.2197158>

Kay JE, Deser C, Phillips A, Mai A, Hannay C, Strand G, Arblaster JM, Bates SC, Danabasoglu G, Edwards J, Holland M, Kushner P, Lamarque J-F, Lawrence D, Lindsay K, Middleton A, Munoz E, Neale R, Oleson K, Polvani L, Vertenstein M (2015) The community earth system model (CESM) large ensemble project: a community resource for studying climate change in the presence of internal climate variability. *Bull Am Meteor Soc* 96(8):1333–1349

Kidd B, Katzfuss M (2022) Bayesian nonstationary and nonparametric covariance estimation for large spatial data (with discussion). *Bayesian Anal* 17(1):291–351

Kingma DP, Ba J (2014) ADAM: a method for stochastic optimization. arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980)

Kingma DP, Welling M (2014) Auto-encoding variational Bayes. In: International conference on learning representations

Kobyzev I, Prince S, Brubaker M (2020) Normalizing flows: an introduction and review of current methods. *IEEE Trans Pattern Anal Mach Intell* 43:3964–3979

Kovachki NB, Hosseini B, Baptista R, Marzouk YM (2020) Conditional sampling with monotone GANs. [arXiv:2006.06755](https://arxiv.org/abs/2006.06755)

Marzouk YM, Moseley T, Parno M, Spantini A (2016) Sampling via measure transport: an introduction. In: Ghanem R, Higdon D, Owhadi H (eds) *Handbook of uncertainty quantification*. Springer, Berlin

Mescheder L, Geiger A, Nowozin S (2018) Which training methods for GANs do actually converge? In: International conference on machine learning, pp 3481–3490

Nemeth C, Fearnhead P (2021) Stochastic gradient Markov chain Monte Carlo. *J Am Stat Assoc* 116(533):433–450

Nychka DW, Hammerling DM, Krock M, Wiens A (2018) Modeling and emulation of nonstationary Gaussian fields. *Spat Stat* 28:21–38

Parno M, Moseley T, Marzouk Y (2016) A multiscale strategy for Bayesian inference using transport maps. *SIAM/ASA J Uncertain Quantif* 4(1):1160–1190

Porcu, E, Pier GB, Felipe T, Rubén S, Quintana FA (2021) Nonparametric bayesian modeling and estimation of spatial correlation functions for global data. *Bayesian Anal* 16(3):845–873. <https://doi.org/10.1214/20-BA1228>

Pronzato L, Müller WG (2012) Design of computer experiments: space filling and beyond. *Stat Comput* 22(3):681–701

Rosenblatt M (1952) Remarks on a multivariate transformation. *Ann Math Stat* 23(3):470–472

Schäfer F, Katzfuss M, Owhadi H (2021) Sparse Cholesky factorization by Kullback–Leibler minimization. *SIAM J Sci Comput* 43(3):A2019–A2046

Schäfer F, Sullivan TJ, Owhadi H (2021) Compression, inversion, and approximate PCA of dense kernel matrices at near-linear computational complexity. *Multiscale Model Simul* 19(2):688–730

Stan Development Team (2023) Stan reference manual v2.32

Stein ML (2011) When does the screening effect hold? *Ann Stat* 39(6):2795–2819

Stein ML, Chi Z, Welty L (2004) Approximating likelihoods for large spatial data sets. *J R Stat Soc B* 66(2):275–296

Torgerson W (1952) Multidimensional scaling: I. Theory and method. *Psychometrika* 17:401–419

Vecchia A (1988) Estimation and model identification for continuous spatial processes. *J R Stat Soc B* 50(2):297–312

Wiens A, Nychka DW, Kleiber W (2020) Modeling spatial data using local likelihood estimation and a Matérn to spatial autoregressive translation. *Environmetrics* 31(6):1–15

Young G, Householder AS (1938) Discussion of a set of points in terms of their mutual distances. *Psychometrika* 3(1):19–22

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.