



Scalable Bayesian Transport Maps for High-Dimensional Non-Gaussian Spatial Fields

Matthias Katzfuss & Florian Schäfer

To cite this article: Matthias Katzfuss & Florian Schäfer (2023): Scalable Bayesian Transport Maps for High-Dimensional Non-Gaussian Spatial Fields, *Journal of the American Statistical Association*, DOI: [10.1080/01621459.2023.2197158](https://doi.org/10.1080/01621459.2023.2197158)

To link to this article: <https://doi.org/10.1080/01621459.2023.2197158>



© 2023 The Author(s). Published with
license by Taylor & Francis Group, LLC



[View supplementary material](#)



Published online: 02 May 2023.



[Submit your article to this journal](#)



Article views: 2105



[View related articles](#)



[View Crossmark data](#)

Scalable Bayesian Transport Maps for High-Dimensional Non-Gaussian Spatial Fields

Matthias Katzfuss^a and Florian Schäfer^b

^aDepartment of Statistics, Texas A&M University, College Station, TX; ^bSchool of Computational Science and Engineering, Georgia Institute of Technology, Atlanta, GA

ABSTRACT

A multivariate distribution can be described by a triangular transport map from the target distribution to a simple reference distribution. We propose Bayesian nonparametric inference on the transport map by modeling its components using Gaussian processes. This enables regularization and uncertainty quantification of the map estimation, while resulting in a closed-form and invertible posterior map. We then focus on inferring the distribution of a nonstationary spatial field from a small number of replicates. We develop specific transport-map priors that are highly flexible and are motivated by the behavior of a large class of stochastic processes. Our approach is scalable to high-dimensional distributions due to data-dependent sparsity and parallel computations. We also discuss extensions, including Dirichlet process mixtures for flexible marginals. We present numerical results to demonstrate the accuracy, scalability, and usefulness of our methods, including statistical emulation of non-Gaussian climate-model output. Supplementary materials for this article are available online.

ARTICLE HISTORY

Received March 2022

Accepted March 2023

KEYWORDS

Climate-model emulation; Dirichlet process mixture; Gaussian process; Generative modeling; Maximin ordering; Nonstationarity

1. Introduction

Inference on a high-dimensional joint distribution based on a relatively small number of replicates is important in many applications. For example, generative modeling of nonstationary and non-Gaussian spatial distributions is crucial for statistical climate-model emulation (e.g., Castruccio et al. 2014; Nychka et al. 2018; Haugen et al. 2019), in ensemble-based data assimilation (e.g., Houtekamer and Zhang 2016; Katzfuss, Stroud, and Wikle 2016), and design studies for new satellite observing systems at NASA using observing system simulation experiments (Errico et al. 2013).

Continuous multivariate distributions can be characterized via triangular transport maps (see Marzouk et al. 2016, for a review) that transform the target distribution to a reference distribution (e.g., standard Gaussian), as illustrated in Figure 1. For Gaussian target distributions, such a map is linear and given by the Cholesky factor of the precision matrix; non-Gaussian distributions can be obtained by allowing nonlinearities in the map. Given an invertible transport map, it is straightforward to sample from the target distribution and some of its conditionals, or to transform the non-Gaussian data to the reference space, in which simple linear operations such as regression or interpolation can be applied. Typically, the map is estimated based on training data, often by iteratively expanding a finite-dimensional parameterization of the transport map (e.g., El Mouselhy and Marzouk 2012; Bigoni, Spantini, and Marzouk 2016; Marzouk et al. 2016; Parno, Mouselhy, and Marzouk 2016; Baptista, Zahm, and Marzouk 2020); subsequent inference is then carried out assuming that the map is known.

We propose an approach for Bayesian inference on a transport map that describes a multivariate continuous distribution

and is learned from a limited number of samples from the distribution. We model the map components using nonparametric, conjugate Gaussian-process priors, which probabilistically regularize the map and shrink toward linearity. The resulting generative model is flexible, naturally quantifies uncertainty, and adjusts to the amount of complexity that is discernible from the training data, thus, avoiding both over- and under-fitting. The conjugacy results in simple, closed-form inference. Instead of assuming Gaussianity for the multivariate target distribution, our approach is equivalent to a series of conditional GP regression problems that together characterize a non-Gaussian target distribution.

We then focus on learning or emulating structured target distributions corresponding to spatial fields observed at a finite but large number of locations, based on a relatively small number of training replicates. In this setting, our Bayesian transport maps impose sparsity and regularization motivated by the behavior of diffusion-type processes that are encountered in many environmental applications. After applying a so-called maximin ordering of the spatial locations, determining the triangular transport map essentially consists of conditional spatial-prediction problems on an increasingly fine scale. We discuss how this scale decay results in conditional near-Gaussianity for a large class of non-Gaussian stochastic processes associated with quasilinear partial differential equations. Hence, our prior distributions are motivated by the behavior of Gaussian fields with Matérn-type covariance, for which the so-called screening effect leads to a decay of influence that motivates sparse transport maps that only consider nearby observations in the spatial prediction problems, corresponding to assumptions of conditional independence. The degree of shrinkage and sparsity

CONTACT Matthias Katzfuss  katzfuss@gmail.com  Department of Statistics, Texas A&M University, College Station, TX.

 Supplementary materials for this article are available online. Please go to www.tandfonline.com/r/JASA.

© 2023 The Author(s). Published with license by Taylor & Francis Group, LLC.

This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

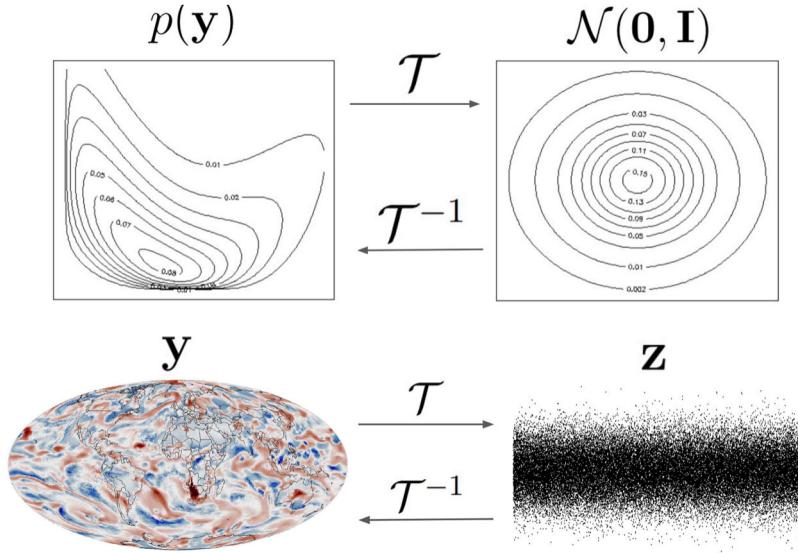


Figure 1. Top panel: Illustration of a transport map T transforming a (bivariate) non-Gaussian distribution $p(\mathbf{y})$ to a standard Gaussian distribution $N(\mathbf{0}, \mathbf{I})$. Bottom: Equivalently, T converts a realization (here, a spatial field) $\mathbf{y} \triangleq p(\mathbf{y})$ to standard Gaussian coefficients $\mathbf{z} = T(\mathbf{y}) \triangleq N(\mathbf{0}, \mathbf{I})$. Under maximin ordering (Figure 2), \mathbf{z} can be viewed as scores corresponding to a nonlinear version of principal components, and they decrease in importance and in corresponding spatial scale from left to right. The spatial field is output from a climate model on a grid of size $N = 288 \times 192 = 55,296$; we want to learn T characterizing the N -dimensional distribution based on an ensemble of $n < 100$ training samples (see Section 6).

are determined by hyperparameters that are inferred from data. The resulting Bayesian methods require little user input, scale near-linearly in the number of spatial locations, and the main computations are trivially parallel.

We further increase the flexibility in the (continuous) marginal distributions by modeling the GP-regression error terms using Dirichlet process mixtures, which can be fit using a Gibbs sampler. The resulting method lets the data decide the degrees of nonlinearity, nonstationarity, and non-Gaussianity, without manual tuning or model-selection. We also discuss an extension for settings in which Euclidean distance between the locations is not meaningful or in which variables are not identified by spatial locations (e.g., multivariate spatial processes).

Most existing methods for spatial inference are in principle applicable in our emulation setting, but they are often geared toward spatial prediction based on a single training replicate and assume Gaussian processes (GPs) with simple parametric covariance functions (e.g., Cressie 1993; Banerjee et al. 2004), whereas our methodology is not designed for spatial prediction at unobserved locations. Many extensions to nonstationary (e.g., as reviewed by Risser 2016) or nonparametric covariances (e.g., Huang, Hsing, and Cressie 2011; Choi, Li, and Wang 2013; Porcu et al. 2021) have been proposed, but these typically still rely on implicit or explicit assumptions of Gaussianity. This includes locally parametric methods specifically developed for climate-model emulation (Nychka et al. 2018; Wiens, Nychka, and Kleiber 2020; Wiens 2021) that locally fit anisotropic Matérn covariances in small windows and then combine the local fits into a global model. For non-Gaussian spatial data, GPs can be transformed or used as latent building blocks (see, e.g., Gelfand and Schliep 2016; Xu and Genton 2017 and references therein), but relying on a GP’s covariance function limits the types of dependence that can be captured. Parametric non-Gaussian Matérn fields can be constructed using

stochastic partial differential equations driven by non-Gaussian noise (Wallin and Bolin 2015; Bolin and Wallin 2020). Models for non-Gaussian spatial data can also be built using copulas; for example, Gräler (2014) proposed vine copulas for spatial fields with extremal behavior, and the factor copula approach of Krupskii, Huser, and Genton (2018) assumes all locations in a homogeneous spatial region to be affected by a common latent factor. Many existing non-Gaussian spatial methods are not scalable to large datasets.

A popular way to achieve scalability for Gaussian spatial fields with parametric covariances is via the Vecchia approximation (e.g., Vecchia 1988; Stein, Chi, and Welty 2004; Datta et al. 2016; Katzfuss and Guinness 2021; Schäfer, Katzfuss, and Owhadi 2021), which implicitly uses a linear transport map given by a sparse inverse Cholesky factor. Kidd and Katzfuss (2022) proposed a Bayesian approach to infer the Cholesky factor nonparametrically. Our (sparse) nonlinear transport maps can be viewed as a Bayesian, nonparametric, and non-Gaussian generalization of Vecchia approximations.

A close relative of transport maps in machine learning are normalizing flows (see Kobyzhev, Prince, and Brubaker 2020 for a review), where triangular layers ensure easy evaluation and inversion of likelihood objectives. Normalizing flows have been used to model point-process intensity functions over the sphere (Ng and Zammit-Mangion 2022, 2023) and random fields in cosmology (Rouhiainen, Giri, and Münchmeyer 2021). Variational autoencoders (VAEs) and generative adversarial networks (GANs) relying on deep neural networks (e.g., Goodfellow, Bengio, and Courville 2016) can be highly expressive and have been employed for climate-model emulation (e.g., Ayala et al. 2021; Besombes et al. 2021). Kovachki et al. (2020) designed GANs with triangular generators that allow conditional sampling. Our approach can be viewed as a Bayesian shallow autoencoder, with the posterior transport map and its inverse acting as the encoder and decoder, respectively. In

contrast to our method, deep-learning approaches typically require massive training data, can be expensive to train, and are often highly sensitive to tuning-parameter and network-architecture choices (e.g., Arjovsky and Bottou 2017; Hestness et al. 2017; Mescheder, Geiger, and Nowozin 2018). Hence, in many low-data applications such approaches are only useful when paired with laborious and application-specific techniques, such as data augmentation, transfer learning, or advances in physics-informed machine learning (e.g., Kashinath et al. 2021).

The remainder of this article is organized as follows. In Section 2, we develop Bayesian transport maps. In Section 3, we consider the special case of high-dimensional spatial distributions. In Section 4, we discuss extensions to non-Gaussian errors using Dirichlet process mixtures. Sections 5 and 6 provide comparisons and applications to simulated data and climate-model output, respectively. Section 7 concludes and discusses future work. Appendices A–G in the supplementary materials contain proofs and further details. Fully automated implementations of our methods, along with code to reproduce all results, are available at <https://github.com/katzfuss-group/BaTraMaSpa>.

2. Bayesian Transport Maps

2.1. Transport Maps and Regression

Consider a continuous random vector $\mathbf{y} = (y_1, \dots, y_N)^\top$, for example describing a spatial field at N locations as in Figure 8. For simplicity, assume that \mathbf{y} has been centered to have mean zero.

For a multivariate Gaussian distribution, $\mathbf{y} \in N_N(\mathbf{0}, \mathbf{G})$ with $\mathbf{G}^{-1} = \mathbf{L}^\top \mathbf{L}$, the (lower-triangular) Cholesky factor \mathbf{L} represents a transformation to a standard normal: $\mathbf{z} = \mathbf{L}\mathbf{y} \in N_N(\mathbf{0}, \mathbf{I}_N)$. As a natural extension, we can characterize any continuous N -variate distribution $p(\mathbf{y})$ by a potentially nonlinear transport map $\mathcal{T} : \mathbb{R}^N \rightarrow \mathbb{R}^N$ (Villani 2009), such that $\mathbf{z} = \mathcal{T}(\mathbf{y}) \in N_N(\mathbf{0}, \mathbf{I}_N)$ for $\mathbf{y} \in p(\mathbf{y})$. Like \mathbf{L} , we can assume without loss of generality that the transport map \mathcal{T} is lower-triangular (Rosenblatt 1952; Carlier, Galichon, and Santambrogio 2009),

$$\mathcal{T}(\mathbf{y}) = \begin{bmatrix} T_1(y_1) \\ T_2(y_1, y_2) \\ \vdots \\ T_N(y_1, y_2, \dots, y_N) \end{bmatrix}, \quad (1)$$

where each $T_i(y_{1:i})$ with $\mathbf{y}_{1:i} = (y_1, \dots, y_i)^\top$ is an increasing function of its i th argument to ensure that \mathcal{T} is invertible and implies a proper density $p(\mathbf{y})$. Letting $N(x|\mu, \sigma^2)$ denote a Gaussian density with parameters μ and σ^2 evaluated at x , we then have

$$p(\mathbf{y}) = p_{\mathbf{z}} \prod_{i=1}^N |T_i(y_{1:i})| \det \mathcal{T} = \prod_{i=1}^N N(T_i(y_{1:i}) | 0, 1) \prod_{i=1}^N \frac{\partial T_i(y_{1:i})}{\partial y_i}, \quad (2)$$

as the triangular \mathcal{T} also implies a triangular $\mathcal{T}^{-1} = (\frac{\partial T_i(y_{1:i})}{\partial y_j})_{i,j=1,\dots,N}$.

Throughout, we assume each T_i to be linearly additive in its i th argument,

$$T_i(y_{1:i}) = (y_i - f_i(\mathbf{y}_{1:i-1}))/d_i, \quad i = 1, \dots, N, \quad (3)$$

for some $d_i \in \mathbb{R}^+$, $f_i : \mathbb{R}^{i-1} \rightarrow \mathbb{R}$ for $i = 2, \dots, N$, and $f_i(\mathbf{y}_{1:i-1}) \equiv 0$ for $i = 1$. Then, $\partial_i T_i(\mathbf{y}_{1:i}) = 1/d_i > 0$, as required. Using (2), it is easy to show that

$$p(\mathbf{y}) = \prod_{i=1}^N \exp\left(-\frac{1}{2d_i^2}(y_i - f_i(\mathbf{y}_{1:i-1}))^2\right) \frac{1}{d_i} \prod_{i=1}^N N(y_i | f_i(\mathbf{y}_{1:i-1}), d_i^2). \quad (4)$$

Thus, the transport-map approach has turned the difficult problem of inferring the N -variate distribution of \mathbf{y} into N independent regressions of y_i on $\mathbf{y}_{1:i-1}$ of the form

$$y_i = f_i(\mathbf{y}_{1:i-1}) + \varepsilon_i, \quad \varepsilon_i \in N(0, d_i^2), \quad i = 1, \dots, N. \quad (5)$$

Sparsity in the map \mathcal{T} corresponds to conditional independence in the joint distribution $p(\mathbf{y})$ (see Spantini, Bigoni, and Marzouk 2018). Specifically, if we assume $f_i(\mathbf{y}_{1:i-1}) = f_i(\mathbf{y}_{c_i})$ for a subset $c_i \subseteq \{1, \dots, i-1\}$, then \mathcal{T} is sparse in that T_i only depends on y_j if $j \in c_i$ (or if $j = i$). Making such a sparsity assumption for $i = 2, \dots, N$ (and setting $\mathbf{y}_c = \emptyset$), we have from (4) that $p(\mathbf{y}) = \prod_{i=1}^N p(y_i | \mathbf{y}_c)$, meaning that y_i is independent of $\{y_j : j \in c_i, j < i\}$ conditional on \mathbf{y}_c . We will exploit this sparsity for computational gain for inferring large non-Gaussian spatial fields in Section 3.

2.2. Modeling the Map Functions Using Gaussian Processes

In the existing transport-map literature (e.g., Marzouk et al. 2016), $f_i : \mathbb{R}^{i-1} \rightarrow \mathbb{R}$ and $d_i \in \mathbb{R}^+$ in (3)–(5) are often assumed to have parametric form, whose parameters are estimated and then assumed known. Instead, we here assume a flexible, non-parametric prior on the map \mathcal{T} by specifying independent conjugate Gaussian-process-inverse-Gamma priors for the f_i and d_i^2 . These prior assumptions induce prior distributions on the map components T_i in (3), and thus on the entire map \mathcal{T} in (1).

Specifically, for the “noise” variances d_i^2 , we assume inverse-Gamma distributions,

$$d_i^2 \stackrel{\text{ind}}{\sim} IG(\alpha_i, \beta_i), \quad \text{with } \alpha_i > 1, \beta_i > 0, \quad i = 1, \dots, N. \quad (6)$$

Conditional on d_i^2 , each function f_i is modeled as a Gaussian process (GP) with inputs $\mathbf{y}_{1:i-1}$,

$$f_i | d_i \stackrel{\text{ind}}{\sim} GP(0, d_i^2 K_i), \quad i = 1, \dots, N, \quad (7)$$

where $K_i(\cdot, \cdot) = C_i(\cdot, \cdot)/E(d_i^2)$, $E(d_i^2) = \beta_i/(\alpha_i - 1)$,

$$C_i(\mathbf{y}_{1:i-1}, \mathbf{y}_{1:i-1}^\top)$$

$$= \mathbf{y}_{1:i-1}^\top \mathbf{Q}_i \mathbf{y}_{1:i-1} + \sigma_i^2 \rho_i(\mathbf{y}_{1:i-1}, \mathbf{y}_{1:i-1}^\top), \quad i = 1, \dots, N, \quad (8)$$

$\sigma_i \in \mathbb{R}_0^+$, and ρ_i is a positive-definite correlation function such that $\rho_i(\mathbf{y}_{1:i-1}, \mathbf{y}_{1:i-1}) = 1$. This prior on f_i is motivated by considering $f_i | \mathbf{b}_i \stackrel{\text{ind}}{\sim} GP(\mathbf{b}_i^\top(\cdot), \sigma_i^2 \rho_i(\cdot, \cdot))$ with inputs $\mathbf{y}_{1:i-1}$, where

$\mathbf{b}_i \in N(\mathbf{0}, \mathbf{Q}_i)$. Integrating out \mathbf{b}_i , we obtain $f_i \stackrel{\text{ind}}{\sim} GP(0, C_i)$ with C_i as in (8), and hence $f_i = (d_i^2 / E(d_i^2)) \hat{f}_i$ as in (7). The degree of nonlinearity of f_i is determined by σ_i^2 ; if $\sigma_i^2 = 0$, then f_i is a linear function of $\mathbf{y}_{1:i-1}$. The prior distributions (i.e., α_i, β_i, C_i) may depend on hyperparameters θ ; see Section 2.4 for more details.

2.3. The Posterior Map

Now assume that we have observed n independent training samples $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(n)}$ from the distribution in Section 2.1 conditional on $\mathbf{f} = (f_1, \dots, f_N)$ and $\mathbf{d} = (d_1, \dots, d_N)$, such that $\mathbf{y}^{(j)} \stackrel{\text{iid}}{\sim} p(\mathbf{y}^{(j)} | \mathbf{f}, \mathbf{d})$ with $T(\mathbf{y}^{(j)}) | \mathbf{f}, \mathbf{d} \sim N_N(\mathbf{0}, \mathbf{I}_N)$, $j = 1, \dots, n$. We combine the samples into an $n \times N$ data matrix \mathbf{Y} whose j th row is given by $\mathbf{y}^{(j)}$. Then, for the regression in (5), the responses \mathbf{y}_i and the covariates $\mathbf{Y}_{1:i-1}$ are given by the i th and the first $i-1$ columns of \mathbf{Y} , respectively. Below, let $\mathbf{y}^?$ denote a new observation sampled from the same distribution, $\mathbf{y}^? \sim p(\mathbf{y} | \mathbf{f}, \mathbf{d})$, independently of \mathbf{Y} .

Based on the prior distribution for \mathbf{f} and \mathbf{d} in Section 2.2, we can now determine the posterior map $\mathbf{\Phi}$ learned from the training data \mathbf{Y} , with \mathbf{f} and \mathbf{d} integrated out. This map is available in closed form and invertible:

Proposition 1. The transport map $\mathbf{\Phi}$ from $\mathbf{y}^? \sim p(\mathbf{y} | \mathbf{Y})$ to $\mathbf{z}^? = \mathbf{\Phi}(\mathbf{y}^?) \sim N_N(\mathbf{0}, \mathbf{I}_N)$ is a triangular map with components

$$\begin{aligned} z_i^? &= \Phi_i(y_1^?, \dots, y_i^?) \\ &= 8^{-1} \hat{F}_{2\tilde{\alpha}_i}^{-1} \hat{d}_i^{-1} (v_i(y_{1:i-1}^?) + 1)^{-1/2} (y_i^? - \hat{f}_i(y_{1:i-1}^?)) \quad (9) \end{aligned}$$

$i = 1, \dots, N$,

where $\tilde{\alpha}_i = \alpha_i + n/2$, $\beta_i = \beta_i + y_i^? \mathbf{G}_i^{-1} \mathbf{y}_i / 2$, $\hat{d}_i^2 = \beta_i^2 / \tilde{\alpha}_i$, $\mathbf{G}_i = \mathbf{K}_i + \mathbf{I}_n$, $\mathbf{K}_i = K_i(\mathbf{Y}_{1:i-1}, \mathbf{Y}_{1:i-1}) = K_i(y_{1:i-1}^?, y_{1:i-1}^?)_{j,l=1,\dots,n}$,

$$\hat{f}_i(y_{1:i-1}^?) = K_i(y_{1:i-1}^?, \mathbf{Y}_{1:i-1}) \mathbf{G}_i^{-1} \mathbf{y}_i, \quad (10)$$

$$\begin{aligned} v_i(y_{1:i-1}^?) &= K_i(y_{1:i-1}^?, y_{1:i-1}^?) \\ &\quad - K_i(y_{1:i-1}^?, \mathbf{Y}_{1:i-1}) \mathbf{G}_i^{-1} K_i(\mathbf{Y}_{1:i-1}, y_{1:i-1}^?), \quad (11) \end{aligned}$$

for $i = 2, \dots, N$, $\hat{f}_1 = v_1 = 0$ for $i = 1$, and 8 and F_κ denote the cumulative distribution functions of the standard normal and the t distribution with κ degrees of freedom, respectively. The inverse map $\mathbf{\Phi}^{-1}$ can be evaluated at a given $\mathbf{z}^?$ by solving the nonlinear triangular system $\mathbf{\Phi}(\mathbf{y}^?) = \mathbf{z}^?$ for $\mathbf{y}^?$; because $\mathbf{\Phi}$ is triangular, the solution can be expressed recursively as

$$y_i^? = f_i(y_{1:i-1}^?) + F_{2\tilde{\alpha}_i}^{-1}(8(z_i^?)) \hat{d}_i (v_i(y_{1:i-1}^?) + 1)^{1/2}, \quad i = 1, \dots, N. \quad (12)$$

All proofs are provided in Appendix A, [supplementary materials](#). We can write the prior map in a similar form, but this is only useful in the case of highly informative priors.

Determining $\mathbf{\Phi}_i$ requires $O(n^3 + in^2)$ time, mostly for computing and decomposing the $n \times n$ matrix \mathbf{G}_i , for each $i = 1, \dots, N$. However, note that the N rows or components of $\mathbf{\Phi}$ can be computed completely in parallel, as in the optimization-based transport-map estimation reviewed in Marzouk et al. (2016). Each application of the transport map or its inverse then consists of the GP prediction in (10)–(11) and only requires $O(n^2 + in)$ time for $i = 1, \dots, N$, but the inverse map is evaluated recursively (i.e., not in parallel).

In contrast to existing transport-map approaches, our approach is Bayesian and naturally quantifies uncertainty in the nonlinear transport functions. The GP priors on the f_i automatically adapt to the amount of information available, only resulting in strongly nonlinear function estimates when supplied the requisite evidence by the data. If n is increasing,

then $\tilde{\alpha}_i$ increases, $F_{2\tilde{\alpha}_i}$ converges to 8, and $v_i(y_{1:i-1}^?)$ typically converges to zero, and so the map components simplify to

$$\Phi_i(y_1^?, \dots, y_i^?) = (y_i^? - f_i(y_{1:i-1}^?)) / \hat{d}_i \quad \text{and} \quad y_i^? = \hat{f}_i(y_{1:i-1}^?) + z_i \hat{d}_i. \quad (13)$$

When employed for finite n , this simplified version of the map ignores posterior uncertainty in \mathbf{f} and \mathbf{d} and instead relies on the point estimates $\hat{f}_i(y_{1:i-1})$ and \hat{d}_i^2 . If we further assume that $\sigma_i = 0$ in (8) for all $i = 1, \dots, N$, then all f_i and all Φ_i become linear functions; we can think of the resulting linear map $\mathbf{\Phi}(\mathbf{y}^?) = \mathbf{L}^? \mathbf{y}^?$ as an inverse Cholesky factor, in the sense that $\mathbf{y}^? | \mathbf{Y} \sim N(\mathbf{0}, \mathbf{3}^{-1})$ with $\mathbf{3} = \mathbf{L} \mathbf{L}^?$.

Transport maps can be used for a variety of purposes. For example, we can obtain new samples $\mathbf{y}^?$ from the posterior predictive distribution $p(\mathbf{y} | \mathbf{Y})$ by sampling $\mathbf{z}^? \sim N_N(\mathbf{0}, \mathbf{I}_N)$ and computing $\mathbf{y}^? = \mathbf{\Phi}^{-1}(\mathbf{z}^?)$ using (12). The map $\mathbf{\Phi}$ in (9) provides a transformation from a non-Gaussian vector $\mathbf{y}^?$ to the standard Gaussian $\mathbf{z}^? = \mathbf{\Phi}(\mathbf{y}^?)$; we call $\mathbf{z}^? = (z_1^?, \dots, z_N^?)$ the map coefficients corresponding to $\mathbf{y}^?$ (see Figure 1 for an illustration). Because the nonlinear dependencies have been removed, many operations are more meaningful on $\mathbf{z}^?$ than on $\mathbf{y}^?$, including linear regressions, translations using linear shifts, and quantifying similarity using inner products. We can also detect inadequacies of the map $\mathbf{\Phi}$ for describing the target distribution by examining the degree of non-Gaussianity and dependence in $\mathbf{z}^?$. These uses of transport maps will be considered further in Section 3.5.

2.4. Hyperparameters

The prior distributions on the f_i and d_i in Section 2.2 may depend on unknown hyperparameters θ . For example, by making inference on hyperparameters in the σ_i in (8), we can let the data decide the degree of nonlinearity in the map and thus the non-Gaussianity in the resulting joint target distribution. We can write in closed form the integrated likelihood $p(\mathbf{Y})$, where \mathbf{f} and \mathbf{d} have been integrated out.

Proposition 2. The integrated likelihood is

$$p(\mathbf{Y}) \stackrel{Y^N}{\sim} \prod_{i=1}^N |G_i|^{-1/2} \times (\beta_i^{\alpha_i} / \beta_i^{\tilde{\alpha}_i}) \times \theta(\tilde{\alpha}_i) / \theta(\alpha_i), \quad (14)$$

where $\theta(\cdot)$ denotes the gamma function, and $\tilde{\alpha}_i$, β_i , \mathbf{G}_i are defined in Proposition 1.

Now denote by $p_\theta(\mathbf{Y})$ the integrated likelihood $p(\mathbf{Y})$ computed based on a particular value θ of the hyperparameters. There are two main possibilities for inference on θ . First, an empirical Bayesian approach consists of estimating θ by the value that maximizes $\log p_\theta(\mathbf{Y})$, and then regarding θ as fixed and known. As $\log p_\theta(\mathbf{Y})$ is a sum of N simple terms, it is straightforward to optimize this function using stochastic gradient ascent based on automatic differentiation. Second, we can carry out fully Bayesian inference by specifying a prior $p(\theta)$, and sampling θ from its posterior distribution $p(\theta | \mathbf{Y}) \propto p_\theta(\mathbf{Y})p(\theta)$ using Metropolis-Hastings; subsequent inference then relies on these posterior draws.

For our numerical results, we employed the empirical Bayesian approach, because it is faster and preserves the closed-form map properties in Section 2.3. In exploratory numerical

experiments, we observed no significant decrease in inferential accuracy relative to the fully Bayesian approach, likely due to working with a small number of hyperparameters in θ .

3. Bayesian Transport Maps for Large Spatial Fields

Now assume that $\mathbf{y} = (y_1, \dots, y_N)^\top$ consists of spatial observations or computer-model output at spatial locations $\mathbf{s}_1, \dots, \mathbf{s}_N$ in a region or domain $D \subseteq \mathbb{R}^{\text{dim}}$. We assume Bayesian transport maps as in Section 2.1, with regressions of the form (5) in $(i-1)$ -dimensional space for $i = 1, \dots, N$. As N is very large in many relevant applications, we will specify priors distributions of the form described in Section 2.2 that induce substantial regularization and sparsity, as a function of hyperparameters $\theta = (\theta_{\sigma,1}, \theta_{\sigma,2}, \theta_{d,1}, \theta_{d,2}, \theta_\gamma, \theta_q)$ to be introduced in Sections 3.2–3.4.

3.1. Maximin Ordering and Nearest Neighbors

A triangular map $T(\mathbf{y})$ as in (1) depends on the ordering of the variables y_1, \dots, y_N . We assume a maximum-minimum-distance (maximin) ordering of the corresponding locations $\mathbf{s}_1, \dots, \mathbf{s}_N$ (see Figure 2), in which we sequentially choose each location to maximize the minimum distance to all previously ordered locations. Specifically, the first index i_1 is chosen arbitrarily (e.g., $i_1 = 1$), and then the subsequent indices are selected as $i_j = \arg \max_{i \in I_j} \min_{k \in I_j} \|\mathbf{s}_i - \mathbf{s}_k\|$ for $j = 2, \dots, N$, where $I_j = \{i_1, \dots, i_{j-1}\}$. For notational simplicity, we assume throughout that $\mathbf{y} = (y_1, \dots, y_N)$ follows maximin ordering (i.e., $y_j = y_{i_j}$). Define $c_i(k)$ as the index of the k th nearest (previously ordered) neighbor of the i th location (and so $\mathbf{s}_{c_i(1)}, \dots, \mathbf{s}_{c_i(4)}$ are indicated by \mathbf{x} in Figure 2).

The maximin ordering can be interpreted as a multiresolution decomposition into coarse scales early in the ordering and fine scales later in the ordering. In particular, the minimal pairwise distance $\gamma_i = \|\mathbf{s}_i - \mathbf{s}_{c_i(1)}\|$ among the first i locations of the ordering decays roughly as $\gamma_i \approx i^{-1/\text{dim}}$, where dim here is the dimension of the spatial domain (see Figure 2(d)). As a result of the maximin ordering, the i th regression in (5) can be viewed as a spatial prediction at location \mathbf{s}_i based on data at locations $\mathbf{s}_1, \dots, \mathbf{s}_{i-1}$ that lie roughly on a regular grid with distance (i.e., scale) γ_i .

When the variables y_1, \dots, y_N are not associated with spatial locations or when Euclidean distance between the locations is not meaningful (e.g., nonstationary, multivariate, spatio-temporal, or functional data), the maximin and neighbor ordering can be carried out based on other distance metrics, such as $(1 - |\text{correlation}|)^2$ based on some guess or estimate of the correlation between variables (Kang and Katzfuss 2023; Kidd and Katzfuss 2022).

3.2. Priors on the Conditional Non-Gaussianity σ_i^2

In (8), σ_i^2 determines the degree of nonlinearity in f_i ; hence, $\sigma_1^2, \dots, \sigma_N^2$ together determine the conditional non-Gaussianity in the distribution of $\mathbf{y}_{i:N}$ given $\mathbf{y}_{1:i-1}$. A priori, we assume that the degree of nonlinearity decays polynomially with length scale γ_i , namely $\sigma_i^2 = e^{\theta_{\sigma,1} \theta_{\sigma,2} i}$, which allows the conditional distributions of $\mathbf{y}_{i:N}$ given $\mathbf{y}_{1:i-1}$ to be increasingly Gaussian as i increases, as a function of hyperparameters $\theta_{\sigma,1}, \theta_{\sigma,2}$.

This prior assumption is motivated by the behavior of stochastic processes with quasiquadratic loglikelihoods. A quasiquadratic loglikelihood of order r is the sum of a quadratic leading-order term that depends on the r th derivatives of the process, and a nonquadratic term that may only depend on derivatives up to order $r-1$. Gaussian smoothness priors (with quadratic loglikelihoods) such as the Matérn model (Whittle 1954, 1963) are closely related to linear elliptic PDEs. They can formally be thought of as having log-densities $-\text{hu}, A\text{u}/2 - \text{hu}, b\text{u}$ that are maximized by solutions of the linear equation $A\text{u} = b$. Similarly, the maximizers of quasiquadratic log-densities $-\text{hu}, L(D^r\text{u})/2i - V(D^{r-1}\text{u}, \dots, \text{u})$ are solutions of quasilinear PDEs $L(D^r\text{u}) = -\frac{d}{du}V(D^{r-1}\text{u}, \dots, \text{u})$. A wide range of physical phenomena is governed by quasilinear PDEs. For instance, the Cahn-Hilliard (Cahn and Hilliard 1958) and Allen-Cahn (Allen and Cahn 1972) equations describe phase separation in multi-component systems, the Navier-Stokes equation describes the dynamics of fluids, and the Föppl-von Kármán equations describe the large deformations of thin elastic plates. If the order of a data-generating quasilinear PDE is known, a Matérn model with matching regularity is a sensible choice to model the leading-order behavior. However, most of the time, the observations will not arise from a known PDE model, and so the above arguments primarily motivate us to

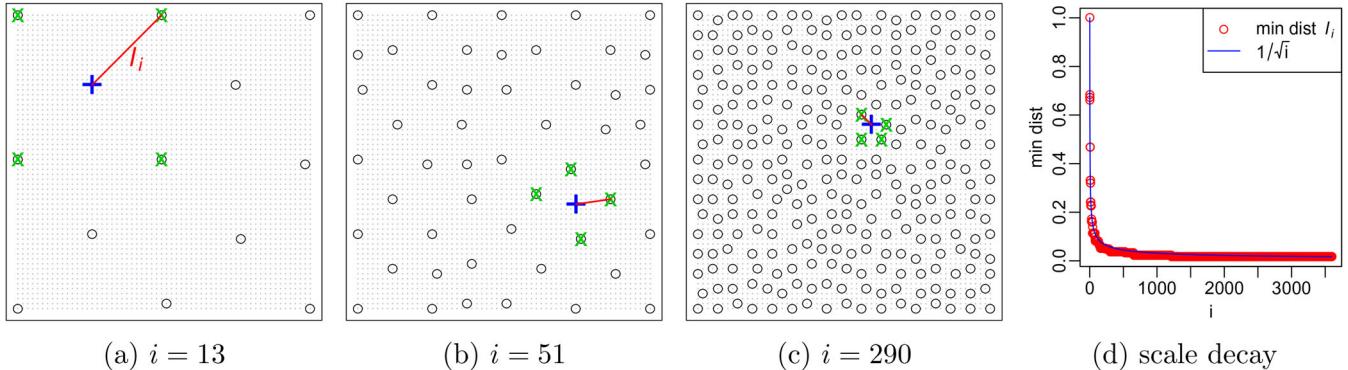


Figure 2. Maximin ordering (Section 3.1) for locations on a grid (small gray points) of size $N = 60 \times 60 = 3600$ on a unit square, $[0, 1]^{\text{dim}}$ with $\text{dim} = 2$. (a)–(c): The i th ordered location (+), the previous $i-1$ locations (○), including the nearest $m = 4$ neighbors (x) and the distance γ_i to the nearest neighbor (—). (d): For $i = 1, \dots, N$, the length scales (i.e., minimum distances) decay as $\gamma_i = i^{-1/\text{dim}}$.

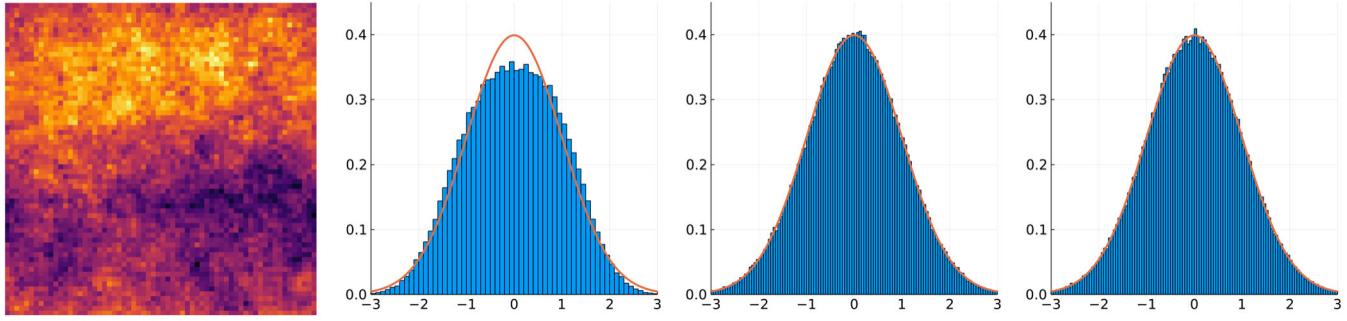


Figure 3. Samples from the non-Gaussian process in (21) in Appendix B, [supplementary materials](#) feature regions of negative (dark) and positive (light) values (first panel). The distribution at a given location is a mixture of these two possibilities and thus non-Gaussian (second panel). By contrast, after conditioning on averages over regions of size $\epsilon = 2^{-1}$ (third panel) or $\epsilon = 2^{-5}$ (fourth panel), the conditional distribution is close to Gaussian, as these averages determine with high probability whether the location is in a positive or negative region. (See Appendix B, [supplementary materials](#) for details.)

expect screening and power laws, without quantifying their effects.

In its simplest form, the mechanism relating local conditioning and approximate Gaussianity is captured by the classical Poincaré inequality (Adams and Fournier 2003).

Lemma 1 (Poincaré inequality). Let $\tilde{A} \subseteq \mathbb{R}^d$ be a Lipschitz-bounded domain with diameter ϵ , let u and its first derivative be square-integrable, and let $u_{\tilde{A}} = \frac{1}{|\tilde{A}|} \int_{\tilde{A}} u(x) dx$ be the mean

of u over \tilde{A} . Then, we have

$$\|u - u_{\tilde{A}}\|_{L^2(\tilde{A})} \leq \epsilon \|u\|_{L^2(\tilde{A})}. \quad (15)$$

The Poincaré inequality directly implies the following corollary.

Corollary 1. Let \tilde{A} be a Lipschitz-bounded domain. Let τ be a partition of \tilde{A} into Lipschitz-bounded subdomains with diameter upper-bounded by ϵ , and assume that u, v and their first derivatives are square-integrable and satisfy $\int_{\tilde{A}} (u - v) dx = 0$ for all $t \in \tau$. Then,

$$\|u - v\|_{L^2(\tilde{A})} \leq \epsilon \|u - v\|_{L^2(\tilde{A})}. \quad (16)$$

This means that after conditioning a stochastic process on averages of diameter $\epsilon \leq 1$, even a minor perturbation in u results in a large change of u . For a quasiquadratic likelihood of order $r = 1$ with lower-bounded curvature of the quadratic part, a perturbation that effects even a minor change in the nonlinear part, depending only on u , must effect a major change in the leading-order, quadratic term, assuming the curvature of the latter is bounded from below. Under suitable growth conditions on the nonlinear term, this means that the conditional density of a quasiquadratic likelihood of order $r = 1$ is dominated by the leading-order quadratic term as ϵ approaches zero. As a result, the conditional stochastic process is approximately Gaussian. This is illustrated in a numerical example in Figure 3. Additional details are provided in Appendix B, [supplementary materials](#). Using generalizations of the Poincaré inequality to $r > 1$ and point-wise measurements (e.g., Schäfer, Sullivan, and Owhadi 2021, thm. 5.9(2)), the above argument can be extended to the setting of $r > 1$ and conditioning on point sets with distance ϵ (instead of local averages).

3.3. Priors on the Conditional Variances d_i^2

As we have argued in Section 3.2, even non-Gaussian stochastic processes with quasiquadratic loglikelihoods exhibit conditional near-Gaussianity on fine scales. Thus, we will now describe prior assumptions for the d_i^2 and f_i in (5) that are motivated by the behavior of a transport map T for a Gaussian target distribution with Matérn covariance (see Figure 4), which is a highly popular assumption in spatial statistics. The Matérn covariance function is also the Green's function of an elliptic PDE (Whittle 1954, 1963).

Schäfer, Sullivan, and Owhadi (2021, thm. 2.3) show that Gaussian processes with covariance functions given by the Green's function of elliptic PDEs of order r have conditional variance of order ϵ^{2r} when conditioned on the first i elements of the maximin ordering (see Figure 4(b)).

Hence, for the noise or conditional variances $d_i^2 \sim IG(\alpha_i, \beta_i)$ as in Section 2.2, we set $E(d_i^2) = \beta_i / (\alpha_i - 1) = e^{\theta_{d,1}} \epsilon^{d,2}$. Assuming the prior standard deviation of d_i^2 to be equal to g times the mean, we obtain $\alpha_i = 2 + 1/g^2$ and $\beta_i = e^{\theta_{d,1}} \epsilon^{d,2} (1 + 1/g^2)$. For our numerical experiments, we chose $g = 4$ to obtain a relatively vague prior for the d_i .²

3.4. Priors on the Regression Functions f_i

The regression functions $f_i : \mathbb{R}^{i-1} \rightarrow \mathbb{R}$ in (5) were specified to be GPs in $(i-1)$ -dimensional space in Section 2.2. For the covariance function in (8), we assume that $\rho_i(\mathbf{y}_{1:i-1}, \mathbf{y}_{1:i-1}) = \rho h_i(\mathbf{y}_{1:i-1}, \mathbf{y}_{1:i-1}) / \gamma$, where $h_i(\mathbf{y}_{1:i-1}, \mathbf{y}_{1:i-1}) = (\mathbf{y}_{1:i-1} - \mathbf{y}_{1:i-1}^0)^\top \mathbf{Q}_i(\mathbf{y}_{1:i-1} - \mathbf{y}_{1:i-1}^0)$, $\gamma = \exp(\theta_\gamma)$ is a range parameter, and ρ is an isotropic correlation function, taken to be Matérn with smoothness 1.5 for our numerical experiments.

To make this potentially high-dimensional regression feasible, we again use the example of a spatial GP with Matérn covariance to motivate regularization and sparsity via the relevance matrix $\mathbf{Q}_i = \text{diag}(q_{i,1}^2, \dots, q_{i,i-1}^2)$. We assume that the relevance of the k th neighbor (see Section 3.1) decays exponentially as a function of k , such that $q_{i,c_i(k)}$ decays as $\exp(-\theta_q k)$. This type of behavior, often referred to as the screening effect (e.g., Stein 2011), is illustrated in Figure 4(c), and it has been exploited for covariance estimation of a Gaussian spatial field by Kidd and Katzfuss (2022). Recently, Schäfer, Sullivan, and Owhadi (2021) proved exponential rates of screening for Gaussian

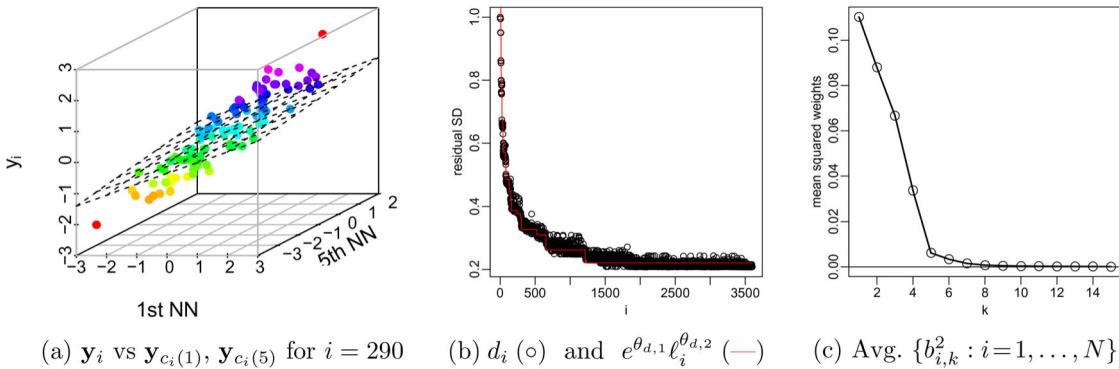


Figure 4. For a Gaussian process with exponential covariance on the grid and with the ordering from Figure 2, expressing the joint distribution $p(\mathbf{y})$ using a transport map as in (1)–(3) results in a series of regressions as in (5) with linear predictors, $f_i(\mathbf{y}_{1:i-1}) = \sum_{k=1}^{i-1} y_{c_i(k)} b_{i,k}$, where $c_i(k)$ indicates the k th nearest (previously ordered) neighbor of the i th location. (For non-Gaussian $p(\mathbf{y})$, the functions f are nonlinear.) (a): For $n = 100$ simulations, the values of \mathbf{y} and its first and fifth nearest neighbor (NN) lie on a low-dimensional manifold; the regression plane (assuming all other variables to be fixed) indicates a stronger influence of the first NN (see the slope of the intersection of the regression plane with the front of the box) than of the fifth NN. (b): The conditional standard deviations decay as a function of the length scale γ (see Figure 2(d)). (c) The squared regression coefficients decay rapidly as a function of neighbor number k .

processes derived from elliptic boundary-value problems; following the discussion in Section 3.2, we expect similar conditional-independence phenomena to hold on the fine scales of processes with quasiquadratic loglikelihoods. As shown in Figure 9(b), we also observed this behavior for climate data.

Given this exponential decay as a function of the neighbor number k , the relevance will be essentially zero for sufficiently large k , and so we achieve sparsity by setting

$$q_{i,c_i(k)} = \begin{cases} \exp(\theta_q k), & k \leq m, \\ 0, & k > m, \end{cases} \quad (17)$$

where the sparsity parameter $m = \max\{k : \exp(\theta_q k) \geq \varepsilon\}$ is determined by the data through the hyperparameter θ_q . We used $\varepsilon = 0.01$ for our numerical examples, which produced highly accurate inference and usually resulted in $m < 10$. Assumption (17) induces a sparse transport map, in that f_i (and thus T_i) depend on $\mathbf{y}_{1:i-1}$ only through the m nearest neighbors $y_{c_i(1)}, \dots, y_{c_i(m)}$, where ρ_i is isotropic as a function of the scaled inputs $y_{c_i(k)}/q_{i,c_i(k)}$. Sparsity in the transport map is equivalent to an assumption of ordered conditional independence. Similar ordered-conditional-independence assumptions are also popular for Vecchia approximations of Gaussian fields with parametric covariance functions.

Identifying the regression functions f_i in m -dimensional space is further aided by the data approximately concentrating on a lower-dimensional manifold due to the strong dependence between most $y_{c_i(k)}$ and $y_{c_i(l)}$ for small $k, l \leq m$ (e.g., see Figure 4(a)).

3.5. Inference

Based on the prior distributions in Sections 3.2–3.4, we can carry out inference and compute the transport map as in Section 2.3. The prior distributions depend on a vector of hyperparameters, $\theta = (\theta_{\sigma,1}, \theta_{\sigma,2}, \theta_{d,1}, \theta_{d,2}, \theta_{\gamma}, \theta_q)$. When making inference on θ as described in Section 2.4, we effectively let the training data \mathbf{Y} decide the degree of sparsity (through θ_q via m) and the degree of nonlinearity (through $\theta_{\sigma,1}, \theta_{\sigma,2}$ via σ_i). Algorithm 1 summarizes the inference procedure. Figure 5 illustrates estimation of transport-map components in a simulated example.

Algorithm 1: Inference for the spatial transport map

- 1: Order y_1, \dots, y_N in maximin ordering and compute scales γ_i and nearest-neighbor indices $c_i(1), \dots, c_i(m_{\max})$ (e.g., $m_{\max} = 30$) for each $i = 1, \dots, N$ (see Section 3.1)
- 2: Compute $\theta = \arg \max_{\theta} \log p(\mathbf{Y})$ via stochastic gradient ascent, where $p(\mathbf{Y}) \triangleq \prod_{i=1}^N |\mathbf{G}_i|^{-1/2} \times (\beta_i^{\alpha_i} / \beta_i^{\tilde{\alpha}_i}) \times \theta(\tilde{\alpha}_i) / \theta(\alpha_i)$, with $\theta = (\theta_{\sigma,1}, \theta_{\sigma,2}, \theta_{d,1}, \theta_{d,2}, \theta_{\gamma}, \theta_q)$, $\tilde{\alpha}_i = \alpha_i + n/2$, $\beta_i = \beta_i + \gamma_i^2 \mathbf{G}_i^{-1} \mathbf{y}_i / 2$, $\alpha_i = 2 + 1/g_i^2$, $\beta_i = e^{\theta_{d,1}} \gamma_i^{\theta_{d,2}} (1 + 1/g_i^2)$, $g_i = 4$, $\mathbf{G}_i = (C_i(\mathbf{y}^{(j)}, \mathbf{y}^{(l)}))_{j,l=1,\dots,n}/(e^{\theta_{d,1}} \gamma_i^{\theta_{d,2}}) + \mathbf{I}_n$, $C_i(\mathbf{y}^{(j)}, \mathbf{y}^{(l)}) = \sum_{k=1}^m \tilde{y}_{c_i(k)}^{(j)} \tilde{y}_{c_i(k)}^{(l)} + \sigma_i^2 \rho \sum_{k=1}^m (\tilde{y}_{c_i(k)}^{(j)} - \tilde{y}_{c_i(k)}^{(l)})^2 / \gamma_i^2$, $\tilde{y}_{c_i(k)}^{(j)} = y_{c_i(k)}^{(j)} e^{\theta_{\gamma} k}$, $m = \max\{k : e^{\theta_q k} \geq 0.01\}$, $\gamma_i^2 = e^{\theta_{\sigma,1}} \gamma_i^{\theta_{\sigma,2}}$, $\gamma = e^{\theta_{\gamma}}$, $\rho(x) = (1 + x^2)^{-3} \exp(-x^2/3)$
- 3: Use fitted map as desired. For example, generate a new sample $\mathbf{y}^* = \mathbf{F}^{-1}(\mathbf{z}^*)$ using (12) based on $\mathbf{z}^* \sim N_N(\mathbf{0}, \mathbf{I}_N)$.

Due to the sparsity assumption in (17), the computational complexity is lower than in Section 2.3; specifically, determining \mathbf{F}_i now only requires $O(n^3 + mn^2)$ time, again in parallel for each $i = 1, \dots, N$. Each application of the transport map or its inverse then requires $O(N(n^2 + mn))$ time. The maximin ordering and nearest neighbors can also be computed in quasilinear time in N (Schäfer, Katzfuss, and Owhadi 2021, Alg. 7).

In Section 2.3, we discussed using \mathbf{F} in (9) to transform the non-Gaussian \mathbf{y}^* to standard Gaussian map coefficients $\mathbf{z}^* = \mathbf{T}(\mathbf{y}^*)$. This concept, which is illustrated in Figure 6, is especially interesting in our spatial setting. Due to the maximin ordering (Figure 2), the scales γ_i are arranged in decreasing order, and in our prior the d_i^2 also follow a decreasing stochastic order with $E(d_i^2) = e^{\theta_{d,1}} \gamma_i^{\theta_{d,2}}$ (see, e.g., Figure 4(b)). Thus, we can view the map components as a form of nonlinear principal components (NPCs), with the map coefficients as the corresponding component scores. For Gaussian processes with covariance functions

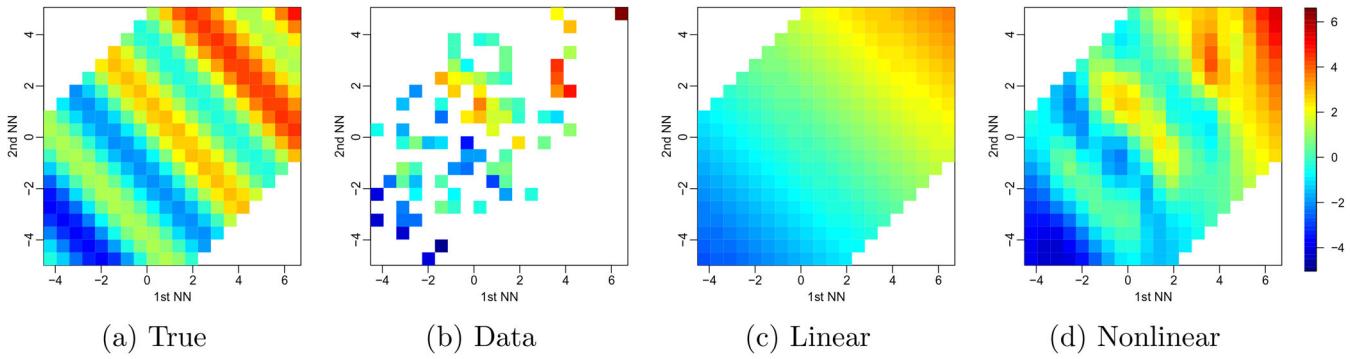


Figure 5. Simulation from a nonlinear map with sine structure in f_i , described as NR900 in Section 5. For $n = 100$ and $i = 80$, y_i versus its first and second nearest neighbor (NN): true f_i (a), observations y_i (b), together with linear (c) and nonlinear (d) fit (i.e., posterior means) of f_i , with further variables in $\mathbf{y}_{1:i-1}$ held at their mean levels. The linear map in (c) is estimated under the restriction $\sigma_i = 0$. In (d), we have a nonlinear regression in 79-dimensional space, with $m = 5$ active variables in the estimated (via θ) nonlinear model.

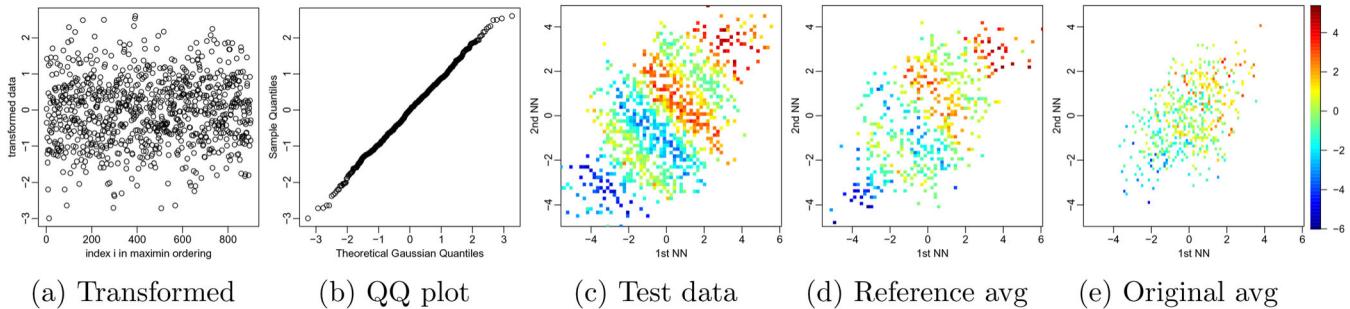


Figure 6. Illustration of map coefficients $\mathbf{z}^2 = T\mathbf{y}^2$ (see Sections 2.3 and 3.5) for the simulated NR900 data using T inferred from $n = 100$ training data. (a)-(b): The $N = 900$ map coefficients corresponding to one test sample are roughly iid Gaussian. (c): For 1000 test samples, y_i versus $y_{1:i-1}$ (d) When averaging pairs of two map-coefficient vectors in reference space and transforming back to the original space using (12), the sinusoidal relationship between y_{80}^2 and its NNs is preserved in the resulting 500 averages. (e) When averaging test samples directly in original space, the nonlinear structure is lost.

given by the Green's function of elliptic PDEs, similar to the Matérn family, it can be shown that these principal components based on the maximin ordering are approximately optimal (Schäfer, Sullivan, and Owhadi 2021). For example, as illustrated in Appendix F, [supplementary materials](#), these NPCs can be used for dimension reduction by only storing or modeling the first k , say, map coefficients $\mathbf{z}_{1:k}^2 = (z_1^2, \dots, z_k^2)^2$. Note that if we set $\mathbf{z}_{k+1:N}^2 = \mathbf{0}$, we assume $y_i^2 = f_i(\mathbf{y}_{1:i-1}^2)$ for $i > k$, which overestimates dependence and underestimates variability; hence, it is preferable to draw $\mathbf{z}_{k+1:N}^2 \sim N(\mathbf{0}, \mathbf{I})$. In addition to reducing storage, we can also use this approach for conditional simulation (Marzouk et al. 2016, Lemma 1), in which we fix the large-scale features of an observed field by fixing the first k map coefficients (see Figure 12 for an illustration). To model a time series of spatial fields, we could assume a linear vector autoregressive model for the NPCs, such that the map coefficients at time $t+1$, say $\mathbf{z}_{1:k}^{(t+1)}$, linearly depend on $\mathbf{z}_{1:k}^{(t)}$. When it is of interest to regress some response on a spatial field, one could also use the first k map coefficients of the field as the covariates, similar to the use of function principal component scores in regression.

4. Non-Gaussian Errors

So far, we have focused on nonlinear, non-Gaussian dependence structures. The model described in Sections 2 and 3 assumes Gaussian errors in the regressions (5), which implies a marginal Gaussian distribution for y_1 , the first variable in the maximin

ordering. If this does not hold at least approximately, extensions based on additional marginal (i.e., pointwise) transformations, especially of the first few variables in the ordering, are straightforward. For example, assume that the model from Sections 2 and 3 holds for \mathbf{y} , but that we actually observe $\tilde{\mathbf{y}} = G(\mathbf{y})$ such that $\tilde{y}_i = g_i(y_i)$. If the g_i are one-to-one differentiable functions, the resulting posterior map is a simple extension of that in Proposition 1. The g_i can be pre-determined (see Section 6 for an example with a log transform) or may depend on θ and thus be inferred based on a minor modification of the integrated likelihood in Proposition 2.

To increase flexibility of the marginal distributions, the GP errors $z_i^{(j)}$ can be modeled using Bayesian nonparametrics for all $i = 1, \dots, N$. More precisely, we will use Dirichlet process mixtures (DPMs). In (5), we now assume that $f_i(\cdot) \sim GP(0, C_i)$, and the $z_i^{(j)}$ are distributed according to a DPM for $j = 1, \dots, n$:

$$z_i^{(j)} | \mu_i^{(j)}, d_i^{(j)} \sim N(\mu_i^{(j)}, (d_i^{(j)})^2), \quad (\mu_i^{(j)}, (d_i^{(j)})^2) | F_i \sim F_i, \\ F_i \sim DP(NIG(\zeta_i, \eta_i, \alpha_i, \beta_i), \zeta_i), \quad (18)$$

where ζ_i is the concentration parameter, and the base measure $NIG(\zeta_i, \eta_i, \alpha_i, \beta_i)$ is a normal-inverse-Gamma distribution with density $p(x, y) = \eta^{1/2} (2\pi y)^{-1/2} \beta^{\alpha_i} / \Gamma(\alpha_i) y^{-\alpha_i - 1} \exp(-(2\beta_i + \eta_i(x - \zeta_i)^2)/(2y))$, where we assume $\zeta_i = 0$. The degree of non-Gaussianity allowed for the $z_i^{(j)}$ is determined by η_i and ζ_i . A small value of ζ_i concentrates the Dirichlet process near the NIG base measure, for which a large value of η_i shrinks

the $\mu_i^{(j)}$ toward zero. Thus, in the limit as $\zeta_i \rightarrow 0$ and $\eta_i \rightarrow \infty$, we obtain a model similar to that in Section 2.2 (except that here the $d_i^{(j)}$ do not appear in the variance of the GP f_i). Conversely, for large ζ_i (or large n), the posterior of $z_i^{(j)}$ will be a Gaussian mixture that may differ substantially from the posterior implied by the model in Section 2.2.

For the spatial setting with maximin ordering of Section 3, we can again find a sparse parameterization in terms of hyperparameters $\theta = (\theta_{\sigma,1}, \theta_{\sigma,2}, \theta_{d,1}, \theta_{d,2}, \theta_{\gamma}, \theta_q, \theta_{\zeta,1}, \theta_{\zeta,2}, \theta_{\eta,1}, \theta_{\eta,2})$. We parameterize the α_i, β_i, C_i in terms of the first six hyperparameters as in Sections 3.2–3.4. For the concentration parameter $\zeta_i = e^{\theta_{\zeta,1} - \theta_{\zeta,2}}$, we allow increasing shrinkage toward Gaussianity for increasing i . We similarly set $\eta_i = e^{\theta_{\eta,1} - \theta_{\eta,2}}$. For this DPM model, we take a fully Bayesian perspective and assume an improper uniform prior for θ over \mathbb{R}^{10} .

The resulting model is fully nonparametric with the exception of the additivity assumption in (3). Specifically, due to the nonparametric nature of the DPM, the universal approximation property of GPs (Micchelli, Xu, and Zhang 2006), and nonzero prior probability for the dense (non-sparse) transport map, the posterior distribution obtained using this model contracts (for $n \rightarrow \infty$ and fixed N) to the Kullback-Leibler (KL) projection of the actual distribution of \mathbf{y} onto the space of distributions that can be described by a transport map whose components are additive in the i th argument as in (3), due to the KL optimality of the Knothe-Rosenblatt map (Marzouk et al. 2016, sec. 4.1). In other words, as the number of replicates increases, the learned distribution gets as close as possible to the truth under the additivity restriction.

Inference for our DPM model cannot be carried out in closed form anymore and instead relies on a Metropolis-within-Gibbs Markov chain Monte Carlo (MCMC) sampler. We can also compute and draw samples from the posterior predictive distribution

$$p(\mathbf{y}^2 | \mathbf{Y}) = p(y_i^2 | \mathbf{y}_{1:i-1}^2, \mathbf{Y}), \quad i=1$$

for which each $p(y_i^2 | \mathbf{y}_{1:i-1}^2, \mathbf{Y})$ is approximated as a Gaussian mixture based on the MCMC output. Details for the MCMC procedure and the posterior predictive distribution are given in Appendix C, [supplementary materials](#).

In the spatial setting with sparsity parameter m , each MCMC iteration still has time complexity $O(N(n^3 + n^2m))$ and the computations within each iteration are highly parallel; however, the actual computational cost for this sampler is much higher (typically, roughly two orders of magnitude higher) than for the empirical Bayes approach in Section 2.4 due to the large number of MCMC iterations required. Because of this larger computational expense and the loss of a closed-form transport map for the DPM model, we recommend the empirical Bayes approach (potentially after a pre-transformation G as described above) as the first option in most large-scale applications; the DPM model is most useful for settings in which its computational expense is not crucial, the training size n is sufficiently large to discern non-Gaussian error structure, and only posterior sampling (as opposed to other functions that transport maps can provide) is of interest.

5. Simulation Study

We compared the following methods:

nonlin: Our method with Bayesian uncertainty quantification described in Section 3.

S-nonlin: Simplified version of nonlin ignoring uncertainty in the f_i and d_i , as in (13).

linear: Same as nonlin, but forcing $\theta_{\sigma,1} = -\infty$ and hence linear f_i .

S-linear: Simplified version of linear ignoring uncertainty in the f_i and d_i as in (13), which results in a joint Gaussian posterior predictive distribution and is similar to the approach proposed and used in numerical comparisons in Kidd and Katzfuss (2022).

DPM: The model with Dirichlet process mixture residuals described in Section 4.

MatCov: Gaussian with zero mean and isotropic Matérn covariance, whose three hyperparameters are inferred via maximum likelihood estimation.

tapSamp: Gaussian with a covariance matrix given by the sample covariance tapered (i.e., element-wise multiplied) by an exponential correlation matrix with range equal to the maximum pairwise distance among the locations.

autoFRK: resolution-adaptive automatic fixed rank kriging (Tzeng and Huang 2018; Tzeng et al. 2021) with approximately \sqrt{N} basis functions.

local: a locally parametric method for climate data (Wiens 2021) that fits anisotropic Matérn covariances in local windows and combines the local fits into a global model.

We also compared to a VAE (Kingma and Welling 2014) and a GAN designed for climate-model output (Besombes et al. 2021), but these deep-learning methods were not competitive in our simulation settings or for the climate data in Section 6 (see Appendix G, [supplementary materials](#)).

We considered four simulation scenarios, for which samples are illustrated in the top row of Figure 7, consisting of a Gaussian distribution with an exponential covariance and three non-Gaussian extensions thereof. All scenarios can be characterized via transport maps as in Section 2.1, with d_i as given by a Gaussian with exponential covariance in the form (3):

LR900: Linear map (i.e., a Gaussian distribution) with components $f_i^L(\mathbf{y}_{1:i-1}) = \sum_{k=1}^{i-1} b_{i,k} y_{c_i(k)}$, where the $b_{i,k}$ are based on an exponential covariance with unit variance and range parameter 0.3 on a Regular grid of size $N = 30 \times 30 = 900$ on the unit square.

NR900: Nonlinear extension of LR900 by a sine function of a weighted sum of the nearest two neighbors: $f_i^{NL}(\mathbf{y}_{1:i-1}) = f_i^L(\mathbf{y}_{1:i-1}) + 2 \sin(4(b_{i,1} y_{c_i(1)} + b_{i,2} y_{c_i(2)}))$ (see Figure 5(a))

NI3600: Same as NR900, but at $N = 3600$ Irregularly spaced locations sampled uniformly at random

NR900B: Same as NR900, but with a Bimodal distribution for the z_i in (5): $z_i | \mu_i, d_i \sim N(\mu_i, d_i^2)$ with μ_i sampled from $\{-3.5d_i, 3.5d_i\}$ with equal probability

For computational simplicity, each (true) f_i was assumed to only depend on the nearest 30 previously ordered neighbors, but this gives a highly accurate approximation of a “full” exponential covariance in the LR900 case, as the true fields exhibit strong

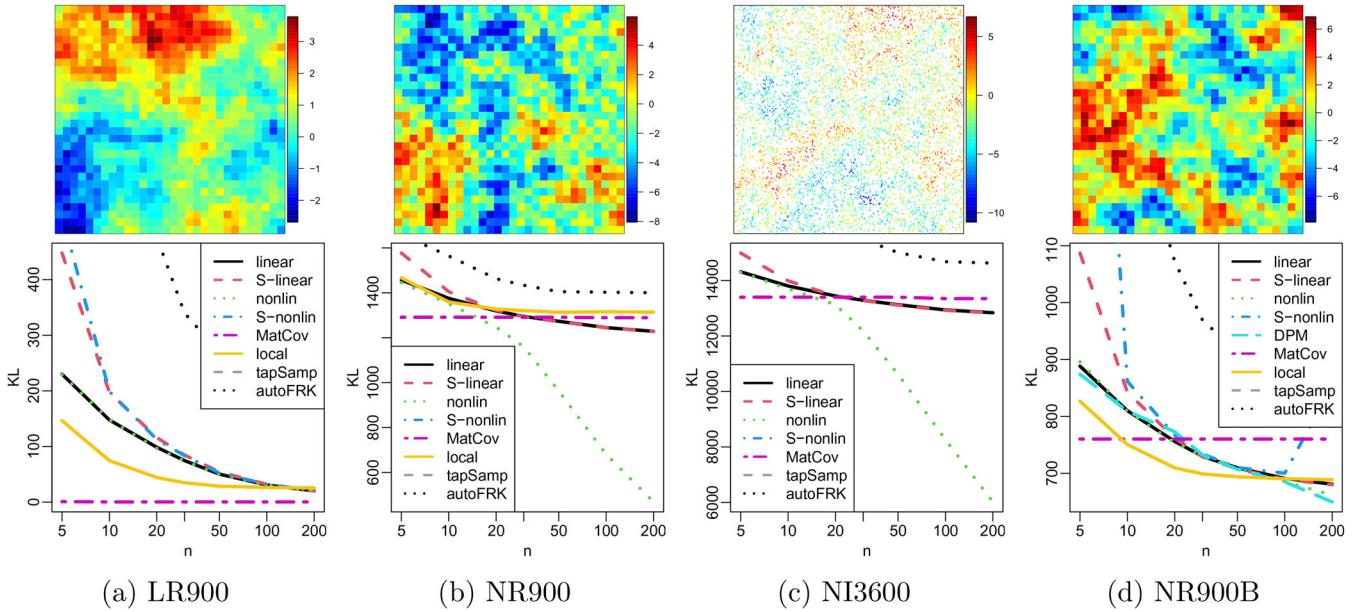


Figure 7. Top row: Simulated spatial fields for four simulation scenarios described in Section 5. Bottom row: Corresponding comparisons of KL divergence as a function of ensemble size n (on a log scale) for different methods. The KL divergences for `tapSamp` in (a)–(d) and for `S-nonlin` in (b)–(c) were too high and are not visible. DPM is only included in (d), while `local` is omitted from (c) because it was created for regular grids.

screening due to being based on the same maximin ordering as our methods. A further ordering-invariant simulation scenario is considered in Appendix D, [supplementary materials](#).

We compared the accuracy of the methods via the Kullback-Leibler (KL) divergence,

$$E(\log p_0(\mathbf{y})) - E(\log p(\mathbf{y}|\mathbf{Y})), \quad (19)$$

between the true distribution $p_0(\mathbf{y})$ and the inferred distribution $p(\mathbf{y}|\mathbf{Y})$ implied by the posterior map (see (18) in Appendix A, [supplementary materials](#)), where the expectations are taken with respect to the true distribution. We approximated the expectations by averaging over 50 simulated test fields \mathbf{y}^* , and so the resulting KL divergence is the difference of the log-scores (e.g., Gneiting and Katzfuss 2014) of the true and inferred distributions.

The results are shown in Figure 7. Whenever nonlinear structure was not discernible from the data (because the true map was linear or because the ensemble size n was too small), `nonlin` performed similarly to `linear` and hence did not suffer due to its over-flexibility. For larger ensemble size and nonlinear truths, `nonlin` at times far outperformed `linear`. `S-linear` and `S-nonlin` were generally less accurate than their counterparts with uncertainty quantification; in the linear LR900 setting, this was only an issue for small ensemble size, but `S-nonlin` performed extremely poorly when the nonlinear structure was clearly apparent in the data, likely due to overfitting without accounting for uncertainty. `tapSamp` and `autoFRK` performed uniformly worst. As `MatCov` (with smoothness 0.5) is the true model for LR900, it was almost exact in that scenario. The other three scenarios are extensions of a Matérn GP, and so `MatCov` also performed well for $n < 20$ or so. The `local` Matérn method was less accurate than `MatCov` for LR900 and NR900 but performed well for NR900B. For simulation scenarios that deviate more strongly from a Matérn

GP, `nonlin` was uniformly more accurate than `MatCov` and `local` (see Appendix D, [Supplementary Materials](#)).

Estimating θ via stochastic gradient ascent with three epochs and fitting the map based on $n = 20$ samples took less than 7 sec for the scenarios with $N = 900$ and less than 44 sec for the larger NI3600 scenario for `nonlin`, `linear`, `S-linear`, and `S-nonlin` on a single core on a laptop (2.5 GHz Intel Core i7 with 16GB RAM); DPM required a total of around 16 min for 500 MCMC iterations for NR900B.

6. Climate-Data Application

An important application of our methods is the analysis and emulation of output from climate models. Climate models are essentially large sets of computer code describing the behavior of the Earth system (e.g., the atmosphere) via systems of differential equations. Much time and resources have been spent on developing these models, and enormous computational power is required to produce ensembles (i.e., solve the differential equations for different starting conditions) on fine latitude-longitude grids for various scenarios of greenhouse-gas emissions. Of the large amount of data and output that have been generated, only a small fraction has been fully explored or analyzed (e.g., Benestad et al. 2017). Stochastic weather generators infer the distribution of one or more variables, so that relevant summaries or additional samples can be computed more cheaply than via more runs of the computer model.

We considered log-transformed total precipitation rate (in m/s) on a roughly 1° longitude-latitude global grid of size $N = 288 \times 192 = 55,296$ in the middle of the Northern summer (July 1) in 98 consecutive years (the number of years contained in one NetCDF data file), starting in the year 402, from the Community Earth System Model (CESM) Large Ensemble Project (Kay et al. 2015). We obtained precipitation anomalies by standardizing the data at each grid location to mean zero

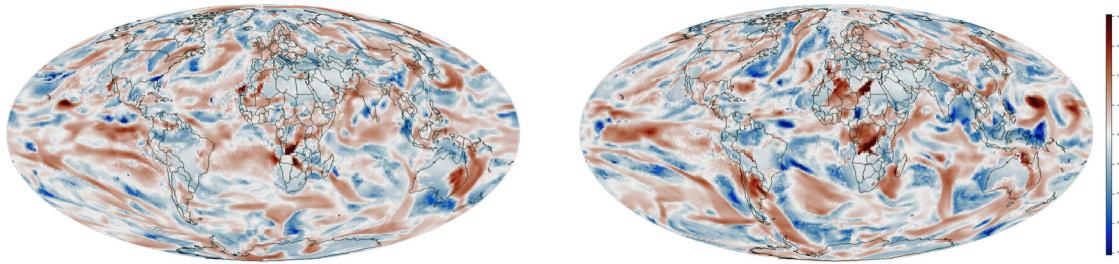


Figure 8. Two members of an ensemble of log-transformed precipitation anomalies produced by a climate model, on a global grid of size $N = 288 \times 192 = 55,296$. We want to infer the underlying N -dimensional distribution based on an ensemble of $n < 100$ training samples.

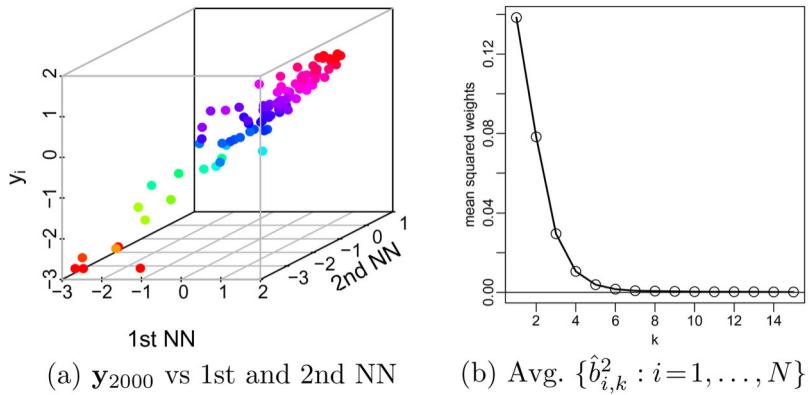


Figure 9. The precipitation anomalies (in the Americas subregion) have similar properties as the Gaussian distribution with exponential covariance in Figure 4: (a) Our approach can be viewed as N regressions as in (5) of each y_i on ordered nearest neighbors (NNs), with the regression data on low-dimensional manifolds. (b) For linear regressions with $f_i(\mathbf{y}_1: i-1) = \sum_{k=1}^{i-1} y_{c_i(k)} b_{i,k}$ fitted via Lasso, the squared (estimated) regression coefficients decay rapidly as a function of neighbor number k .

and variance one, shown in Figure 8. For our methods, we used chordal distance to compute the maximin ordering and nearest neighbors.

For ease of comparison and illustration, we first considered a smaller grid of size $N = 37 \times 74 = 2738$ in a subregion containing large parts of the Americas (45°S to 45°N and 130°W to 30°W) containing ocean, land, and mountains. As shown in Figure 9, the precipitation anomalies exhibited similar features as our simulated data in Figure 4, with regression data concentrating on lower-dimensional manifolds and weights decaying rapidly as a function of neighbor number.

For comparing the methods from Section 5 on the precipitation anomalies, computing the KL divergence as in (19) was not possible, as the true distribution $p_0(\mathbf{y})$ was unknown. Hence, we compared the methods using various training data sizes n in terms of log-scores, which approximate the KL divergence up to an additive constant; specifically, these log-scores consist of the second part of (19), $-E(\log p(\mathbf{y}|\mathbf{Y}))$, with the expectation approximated by averaging over 18 test replicates and over five random training/test splits.

The comparison for the Americas subregion is shown in Figure 10(a). (A prediction comparison for partially observed test data provided in Appendix E, [supplementary materials](#) produced similar results.) `nonlin` outperformed `linear`, and `DPM` was even more accurate than `nonlin` for large n , indicating that the precipitation anomalies exhibit joint and marginal non-Gaussian features. As in Section 5, `S-linear` and `S-nonlin` performed poorly due to ignoring uncertainty in the estimated map. `local` performed similarly to `linear` but was less accurate than `nonlin` and `DPM` for all n . `VAE`,

`MatCov`, `tapSamp`, and `autoFRK` were not competitive for this dataset.

We also considered the map coefficients $\mathbf{z}^2 = \mathbf{f}(\mathbf{y}^2)$ discussed in Sections 2.3 and 3.5, using the map obtained by fitting `nonlinear` to the first $n = 97$ replicates as training data. In Figure 11(a), the map coefficients for a held-out test field appeared roughly iid standard Gaussian, with sample autocorrelations near zero (not shown). Figure 11 illustrates that the map coefficients offer similar properties for non-Gaussian fields as principal-component scores do for Gaussian settings. For example, the medians of the posterior distributions of the d_i (see (17) in Appendix A, [supplementary materials](#)) decreased rapidly as a function of i , which means that the map coefficients early in the maximin ordering captured much more (nonlinear) variation than later-ordered coefficients (see, e.g., (12) and (13)). Further, we computed the map coefficients for all 98 replicates for July 2–30 (still based on the posterior map trained on July 1 data), and the lag-1 autocorrelation over time between map coefficients also decreased with i . Specifically, while most of the first 100 were greater than 0.2, many later autocorrelations were negligible; this indicates that a spatio-temporal analysis could proceed by fitting a simple (linear) autoregressive model over time to only the first k , say, map coefficients, while treating the remaining $N - k$ coefficients as independent over time. As shown in Appendix F, [supplementary materials](#), the `nonlinear` map coefficients strongly outperformed standard linear principal components in terms of dimension reduction and reconstruction of the Americas climate fields.

To demonstrate scalability to large datasets, we compared `linear` and `nonlinear` on the entire global precipitation

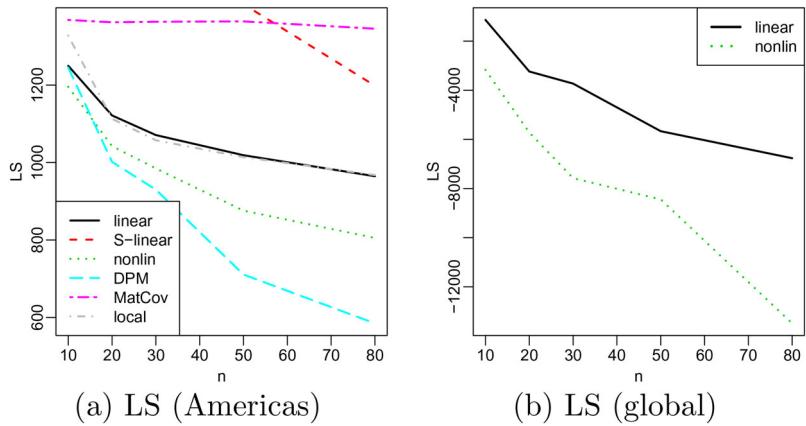


Figure 10. For precipitation anomalies, comparison of log-score (LS; equal to KL divergence up to an additive constant) for estimated joint distribution as a function of ensemble size n : (a) Americas subregion; S-nonlin, tapSamp, and autoFRK are not shown because their LS were too high. (b) LS for linear and nonlin for precipitation anomalies on the global grid.

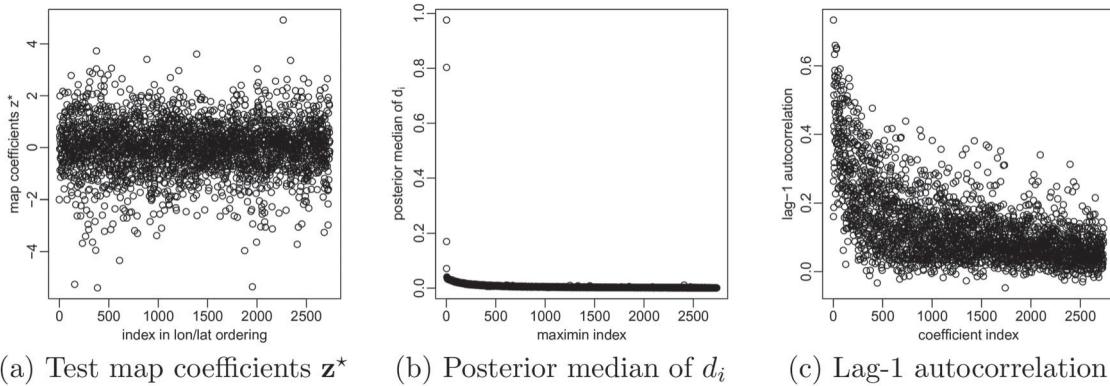


Figure 11. Properties of the map coefficients $z^2 = \mathbb{E}(\mathbf{y}^2)$ for the precipitation anomalies on the grid of size $N = 2738$ in the Americas subregion. (a): The map coefficients corresponding to the test field in Figure 12(a) in the original data ordering (first by longitude, then latitude) appeared roughly iid standard Gaussian, aside from slightly heavy tails. (b): The posterior medians of the d_i decreased rapidly as a function of i , meaning that the first few map coefficients captured much more variation than later-ordered coefficients. (c): The autocorrelation between consecutive days also decreased with i ; while most were greater than 0.2 for $i < 100$, many autocorrelations for $i > 100$ were negligible.

anomaly fields of size $N = 288 \times 192 = 55,296$. As shown in Figure 10(b), nonlin outperformed linear even more decisively than for the Americas subregion. Even in the largest and most accurate setting ($n = 80$), the estimated θ for nonlin implied $m = 9$, meaning that the corresponding transport maps were extremely sparse and hence computationally efficient; estimating θ (4 epochs) and fitting the map for nonlin took only around 6 min on a single core on a laptop (2.5 GHz Intel Core i7 with 16GB RAM) for $n = 10$. In contrast, MatCov and local (which already took about two hours for the much smaller Americas region) were too computationally demanding for the global data. A Vecchia approximation of MatCov resulted in a log-score above +77,000 and was thus not competitive. Also, for nonlin all but 113 of the $N = 55,296$ posterior medians of the d_i were more than 20 times smaller than the largest posterior median (i.e., that of d_1), indicating that our approach could be used for massive dimension reduction without losing too much information.

Finally, the fitted map (or rather, its inverse \mathbf{F}^{-1}) can also be viewed as a stochastic emulator of the climate model. Specifically, we can produce a new precipitation-anomaly sample by drawing $\mathbf{z}^{\mathbb{B}} \sim N(\mathbf{0}, \mathbf{I}_N)$ and then computing $\mathbf{y}^{\mathbb{B}} = \mathbf{F}^{-1}(\mathbf{z}^{\mathbb{B}})$. One such sample (for the full global grid) is shown in Figure 12(d) and appears qualitatively similar to the model

output in Figure 8; while producing the latter requires a supercomputer, the former can be generated in a few seconds on a laptop. Further, our approach can also be used to draw conditional samples, in which we fix the first i , say, map coefficients, for example at the values corresponding to a given spatial field. Such draws, which maintain the large-scale features in the held-out (98th) test field but allow for newly sampled fine-scale features, are shown in Figure 12. This is related to the supervised conditional sampling ideas in Kovachki et al. (2020), with their inputs given by our first i ordered test observations.

7. Conclusions

We have developed a Bayesian approach to inferring a non-Gaussian target distribution via a transport map from the target to a standard normal distribution. The components of the map are modeled using Gaussian processes. For the distribution of spatial fields, we have developed specific prior assumptions that result in sparse maps and thus scalability to high dimensions. Instead of manually or iteratively expanding a finite-dimensional parameterization of the transport map, our Bayesian approach probabilistically regularizes the map; the resulting approach is flexible and nonparametric, but guards against overfitting and quantifies uncertainty in the estimation

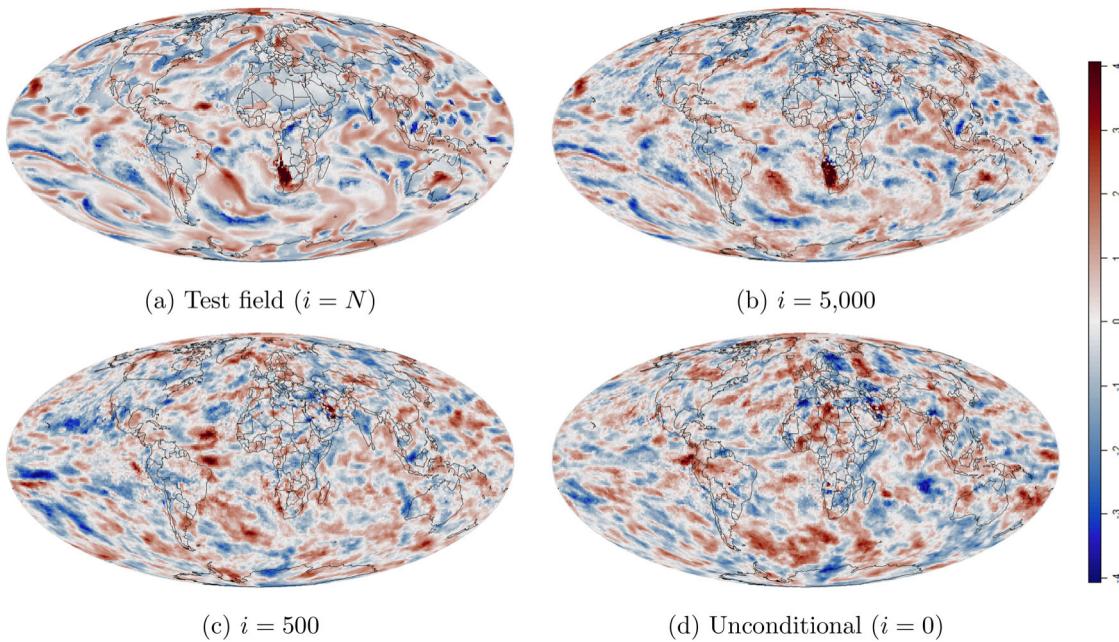


Figure 12. For the global climate data ($N = 55,296$), we fitted a stochastic emulator using `nonlin` based on $n = 97$ training replicates. Given a held-out test field \mathbf{y}^{H} in (a), we show conditional simulations based on fixing the first i map coefficients in $\mathbf{z} = \mathbf{T}^{\text{H}}(\mathbf{y}^{\text{H}})$. (b): Only differs in some fine-scale features from (a). (c): Some large-scale features from (a) are preserved. (d): Unconditional simulation (i.e., independent from (a)).

of the map. Because our method can be fitted rapidly, is fully automated, and was highly accurate in our numerical comparisons, we recommend it for most spatial emulation tasks, except for applications in which very few replicates are available or for which exploratory analyses have shown that a (Gaussian) parametric approach can provide a good fit. In addition, due to conjugate priors and the resulting closed-form expressions for the posterior map and its inverse, our approach also allows us to convert non-Gaussian data into iid Gaussian map coefficients, which can be thought of as a nonlinear extension of principal components.

As our approach essentially turns estimation of a high-dimensional joint distribution into a series of GP regressions, it is straightforward to include additional covariates and examine their nonlinear, non-Gaussian effect on the distribution. Shrinkage toward a joint Gaussian distribution with a parametric covariance function could be achieved by assuming the mean for the GP regressions to be the one implied by a Vecchia approximation of that covariance function (Kidd and Katzfuss 2022); this could enable meaningful predictions at unobserved spatial locations (see Appendix E, [supplementary materials](#)). Extensions to more complicated input domains (e.g., space-time) could be obtained using correlation-based ordering ([Section 3.1](#)). Another major avenue of future work would be to use the inferred distribution as the prior of a latent field, which we then update to obtain a posterior given noisy observations; among numerous other applications, this would enable the use of our technique to infer the forecast distribution and account for uncertainty in ensemble-based data assimilation (Boyles and Katzfuss 2021), leading to nonlinear updates for non-Gaussian applications. We are currently pursuing multiple extensions and applications of our methods to climate science, including climate-change detection and attribution, climate-model calibration, and climate-model emulation and interpolation in covariate space (e.g., as a function of CO₂ emissions).

Supplementary Materials

Appendices A–G contain proofs, a discussion of conditional near-Gaussianity for quasiquadratic loglikelihoods, details on the Gibbs sampler for Dirichlet process mixture model, and additional numerical results and comparisons.

Acknowledgments

We would like to thank Joe Guinness and several reviewers for helpful comments. We are especially grateful to Jian Cao, who wrote a Python implementation, produced timing results, and obtained GAN and VAE results, and to Trevor Harris, who created the VAE implementation for our numerical comparisons.

Disclosure Statement

The authors report there are no competing interests to declare.

Funding

Katzfuss was partially supported by National Science Foundation (NSF) grants DMS–1654083, DMS–1953005, and CCF–1934904, and by NASA’s Advanced Information Systems Technology Program (AIST-21). Schäfer gratefully acknowledges support by the Air Force Office of Scientific Research under award number FA9550-18-1-0271, and the Office of Naval Research under award N00014-18-1-2363.

References

- Adams, R. A., and Fournier, J. J. F. (2003), *Sobolev Spaces*, volume 140 of Pure and Applied Mathematics (2nd ed.), Amsterdam: Elsevier/Academic Press. [6]
- Allen, S. M., and Cahn, J. W. (1972), “Ground State Structures in Ordered Binary Alloys with Second Neighbor Interactions,” *Acta Metallurgica*, 20, 423–433. [5]

Arjovsky, M., and Bottou, L. (2017), "Towards Principled Methods for Training Generative Adversarial Networks," in *International Conference on Learning Representations*. [3]

Ayala, A., Drazic, C., Hutchinson, B., Kravitz, B., and Tebaldi, C. (2021), "Loosely Conditioned Emulation of Global Climate Models with Generative Adversarial Networks," arXiv:2105.06386. [2]

Banerjee, S., Carlin, B. P., and Gelfand, A. E. (2004), *Hierarchical Modeling and Analysis for Spatial Data*, London: Chapman & Hall. [2]

Baptista, R., Zahm, O., and Marzouk, Y. (2020), "An Adaptive Transport Framework for Joint and Conditional Density Estimation," arXiv preprint arXiv:2009.10303. [1]

Benestad, R., Sillmann, J., Thorarinsdottir, T. L., Guttorp, P., Mesquita, M. D., Tye, M. R., Uotila, P., Maule, C. F., Thejll, P., Drews, M., and Pardig, K. M. (2017), "New Vigour Involving Statisticians to Overcome Ensemble Fatigue," *Nature Climate Change*, 7, 697–703. [10]

Besombes, C., Pannekoucke, O., Lapeyre, C., Sanderson, B., and Thual, O. (2021), "Producing Realistic Climate Data with Generative Adversarial Networks," *Nonlinear Processes in Geophysics*, 28, 347–370. [2,9]

Bigoni, D., Spantini, A., and Marzouk, Y. M. (2016), "Adaptive Construction of Measure Transports for Bayesian Inference," in *NIPS 2016 Workshop on Advances in Approximate Bayesian Inference*. [1]

Bolin, D., and Wallin, J. (2020), "Multivariate Type G Matérn Stochastic Partial Differential Equation Random Fields," *Journal of the Royal Statistical Society, Series B*, 82, 215–239. [2]

Boyles, W., and Katzfuss, M. (2021), "Ensemble Kalman Filter Updates Based On Regularized Sparse Inverse Cholesky Factors," *Monthly Weather Review*, 149, 2231–2238. [13]

Cahn, J. W., and Hilliard, J. E. (1958), "Free Energy of a Nonuniform System. I. Interfacial Free Energy," *The Journal of Chemical Physics*, 28, 258–267. [5]

Carlier, G., Galichon, A., and Santambrogio, F. (2009), "From Knothe's Transport to Brenier's Map and a Continuation Method for Optimal Transport," *SIAM Journal on Mathematical Analysis*, 41, 2554–2576. [3]

Castruccio, S., McInerney, D. J., Stein, M. L., Crouch, F. L., Jacob, R. L., and Moyer, E. J. (2014), "Statistical Emulation of Climate Model Projections based on Precomputed GCM Runs," *Journal of Climate*, 27, 1829–1844. [1]

Choi, I. K., Li, B., and Wang, X. (2013), "Nonparametric Estimation of Spatial and Space-Time Covariance Function," *Journal of Agricultural, Biological, and Environmental Statistics*, 18, 611–630. [2]

Cressie, N. (1993), *Statistics for Spatial Data* (rev. ed.), New York: Wiley. [2]

Datta, A., Banerjee, S., Finley, A. O., and Gelfand, A. E. (2016), "Hierarchical Nearest-Neighbor Gaussian Process Models for Large Geostatistical Datasets," *Journal of the American Statistical Association*, 111, 800–812. [2]

El Mouselhy, T. A., and Marzouk, Y. M. (2012), "Bayesian Inference with Optimal Maps," *Journal of Computational Physics*, 231, 7815–7850. [1]

Errico, R. M., Yang, R., Privé, N. C., Tai, K. S., Todling, R., Sienkiewicz, M. E., and Guo, J. (2013), "Development and Validation of Observing-System Simulation Experiments at NASA's Global Modeling and Assimilation Office," *Quarterly Journal of the Royal Meteorological Society*, 139, 1162–1178. [1]

Gelfand, A. E., and Schliep, E. M. (2016), "Spatial Statistics and Gaussian Processes: A Beautiful Marriage," *Spatial Statistics*, 18, 86–104. [2]

Gneiting, T., and Katzfuss, M. (2014), "Probabilistic Forecasting," *Annual Review of Statistics and Its Application*, 1, 125–151. [10]

Goodfellow, I., Bengio, Y., and Courville, A. (2016), *Deep Learning*, Cambridge, MA: MIT Press. [2]

Gräler, B. (2014), "Modelling Skewed Spatial Random Fields through the Spatial Vine Copula," *Spatial Statistics*, 10, 87–102. [2]

Haugen, M. A., Stein, M. L., Srivastava, R. L., and Moyer, E. J. (2019), "Future Climate Emulations Using Quantile Regressions on Large Ensembles," *Advances in Statistical Climatology, Meteorology, and Oceanography*, 5, 37–55. [1]

Hestness, J., Narang, S., Ardalani, N., Diamos, G., Jun, H., Kianinejad, H., Patwary, M. M. A., Yang, Y., and Zhou, Y. (2017), "Deep Learning Scaling is Predictable, Empirically," arXiv:1712.00409. [3]

Houtekamer, P. L., and Zhang, F. (2016), "Review of the Ensemble Kalman Filter for Atmospheric Data Assimilation," *Monthly Weather Review*, 144, 4489–4532. [1]

Huang, C., Hsing, T., and Cressie, N. (2011), "Nonparametric Estimation of the Variogram and its Spectrum," *Biometrika*, 98, 775–789. [2]

Kang, M., and Katzfuss, M. (2023), "Correlation-based Sparse Inverse Cholesky Factorization for Fast Gaussian-Process Inference," *Statistics and Computing*, 33, 1–17. [5]

Kashinath, K., Mustafa, M., Albert, A., Wu, J. L., Jiang, C., Esmaeilzadeh, S., Azizzadenesheli, K., Wang, R., Chattopadhyay, A., Singh, A., Manepalli, A., Chirila, D., Yu, R., Walters, R., White, B., Xiao, H., Tchelapi, H. A., Marcus, P., Anandkumar, A., Hassanzadeh, P., and Prabhat (2021), "Physics-Informed Machine Learning: Case Studies for Weather and Climate Modelling," *Philosophical Transactions of the Royal Society A*, 379, 20200093. [3]

Katzfuss, M., and Guinness, J. (2021), "A General Framework for Vecchia Approximations of Gaussian Processes," *Statistical Science*, 36, 124–141. [2]

Katzfuss, M., Stroud, J. R., and Wikle, C. K. (2016), "Understanding the Ensemble Kalman Filter," *The American Statistician*, 70, 350–357. [1]

Kay, J. E., Deser, C., Phillips, A., Mai, A., Hannay, C., Strand, G., Arblaster, J. M., Bates, S. C., 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., and 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," *Bulletin of the American Meteorological Society*, 96, 1333–1349. [10]

Kidd, B., and Katzfuss, M. (2022), "Bayesian Nonstationary and Nonparametric Covariance Estimation for Large Spatial Data," (with Discussion), *Bayesian Analysis*, 17, 291–351. [2,5,6,9,13]

Kingma, D. P., and Welling, M. (2014), "Auto-Encoding Variational Bayes," in *2nd International Conference on Learning Representations, ICLR 2014 - Conference Track Proceedings*. [9]

Kobyzev, I., Prince, S., and Brubaker, M. (2020), "Normalizing Flows: An Introduction and Review of Current Methods," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43, 3964–3979. [2]

Kovachki, N. B., Hosseini, B., Baptista, R., and Marzouk, Y. M. (2020), "Conditional Sampling with Monotone GANs," arXiv:2006.06755. [2,12]

Krupskii, P., Huser, R., and Genton, M. G. (2018), "Factor Copula Models for Replicated Spatial Data," *Journal of the American Statistical Association*, 113, 467–479. [2]

Marzouk, Y. M., Moseley, T., Parno, M., and Spantini, A. (2016), "Sampling via Measure Transport: An Introduction," in *Handbook of Uncertainty Quantification*, eds. R. Ghanem, D. Higdon, and H. Owhadi, pp. 785–783, Cham: Springer. [1,3,4,8,9]

Mescheder, L., Geiger, A., and Nowozin, S. (2018), "Which Training Methods for GANs do Actually Converge?" in *International Conference on Machine Learning*, pp. 3481–3490. [3]

Micchelli, C. A., Xu, Y., and Zhang, H. (2006), "Universal Kernels," *Journal of Machine Learning Research*, 7, 2651–2667. [9]

Ng, T. L. J., and Zammit-Mangion, A. (2022), "Spherical Poisson Point Process Intensity Function Modeling and Estimation with Measure Transport," *Spatial Statistics*, 50, 100629. [2]

——— (2023), "Mixture Modeling with Normalizing Flows for Spherical Density Estimation," arXiv:2301.06404. [2]

Nychka, D. W., Hammerling, D. M., Krock, M., and Wiens, A. (2018), "Modeling and Emulation of Nonstationary Gaussian Fields," *Spatial Statistics*, 28, 21–38. [1,2]

Parno, M., Moseley, T., and Marzouk, Y. (2016), "A Multiscale Strategy for Bayesian Inference Using Transport Maps," *SIAM/ASA Journal on Uncertainty Quantification*, 4, 1160–1190. [1]

Porcu, E., Bissiri, P. G., Tagle, F., and Quintana, F. (2021), "Nonparametric Bayesian Modeling and Estimation of Spatial Correlation Functions for Global Data," *Bayesian Analysis*, 16, 845–873. [2]

Risser, M. D. (2016), "Review: Nonstationary Spatial Modeling, with Emphasis on Process Convolution and Covariate-Driven Approaches," arXiv:1610.02447. [2]

Rosenblatt, M. (1952), "Remarks on a Multivariate Transformation," *The Annals of Mathematical Statistics*, 23, 470–472. [3]

Rouhainen, A., Giri, U., and Münchmeyer, M. (2021), "Normalizing Flows for Random Fields in Cosmology," arXiv:2105.12024. [2]

Schäfer, F., Katzfuss, M., and Owhadi, H. (2021), “Sparse Cholesky Factorization by Kullback-Leibler Minimization,” *SIAM Journal on Scientific Computing*, 43, A2019–A2046. [2,7]

Schäfer, F., Sullivan, T. J., and Owhadi, H. (2021), “Compression, Inversion, and Approximate PCA of Dense Kernel Matrices at Near-Linear Computational Complexity,” *Multiscale Modeling & Simulation*, 19, 688–730. [6,8]

Spantini, A., Bigoni, D., and Marzouk, Y. M. (2018), “Inference via Low-Dimensional Couplings,” *Journal of Machine Learning Research*, 19, 1–71. [3]

Stein, M. L. (2011), “2010 Rietz lecture: When Does the Screening Effect Hold? *The Annals of Statistics*, 39, 2795–2819. [6]

Stein, M. L., Chi, Z., and Welty, L. (2004), “Approximating Likelihoods for Large Spatial Data Sets,” *Journal of the Royal Statistical Society, Series B*, 66, 275–296. [2]

Tzeng, S. L., and Huang, H.-C. (2018), “Resolution Adaptive Fixed Rank Kriging,” *Technometrics*, 60, 198–208. [9]

Tzeng, S. L., Huang, H.-C., Wang, W.-T., Nychka, D. W., and Gillespie, C. (2021), “autoFRK: Automatic Fixed Rank Kriging.” [9]

Vecchia, A. (1988), “Estimation and Model Identification for Continuous Spatial Processes,” *Journal of the Royal Statistical Society, Series B*, 50, 297–312. [2]

Villani, C. (2009), *Optimal Transport: Old and New*, Berlin: Springer. [3]

Wallin, J., and Bolin, D. (2015), “Geostatistical Modelling Using non-Gaussian Matérn Fields,” *Scandinavian Journal of Statistics*, 42, 872–890. [2]

Whittle, P. (1954), “On Stationary Processes in the Plane,” *Biometrika*, 41, 434–449. [5,6]

— (1963), “Stochastic Processes in Several Dimensions,” *Bulletin of the International Statistical Institute*, 40, 974–994. [5,6]

Wiens, A. (2021), “Nonstationary Covariance Modeling for Gaussian Processes and Gaussian Markov Random Fields,” available at <https://github.com/ashtonwiens/nonstationary>. [2,9]

Wiens, A., Nychka, D. W., and Kleiber, W. (2020), “Modeling Spatial Data Using Local Likelihood Estimation and a Matérn to Spatial Autoregressive Translation,” *Environmetrics*, 31, 1–15. [2]

Xu, G., and Genton, M. G. (2017), “Tukey g-and-h Random Fields,” *Journal of the American Statistical Association*, 112, 1236–1249. [2]