

NIASRA

NATIONAL INSTITUTE FOR APPLIED
STATISTICS RESEARCH AUSTRALIA



***National Institute for Applied Statistics Research
Australia***

University of Wollongong, Australia

Working Paper

04-21

**Modeling Nonstationary and Asymmetric Multivariate
Spatial Covariances via Deformations**

Quan Vu, Andrew Zammit-Mangion, and Noel Cressie

*Copyright © 2021 by the National Institute for Applied Statistics Research Australia, UOW.
Work in progress, no part of this paper may be reproduced without permission from the Institute.*

National Institute for Applied Statistics Research Australia, University of Wollongong,
Wollongong NSW 2522, Australia Phone +61 2 4221 5435, Fax +61 2 4221 4998.

Email: karink@uow.edu.au

Modeling Nonstationary and Asymmetric Multivariate Spatial Covariances via Deformations

Quan Vu^{1, *}, Andrew Zammit-Mangion¹, and Noel Cressie¹

¹School of Mathematics and Applied Statistics, University of Wollongong, Australia

*Corresponding author: quanv@uow.edu.au

Abstract

Multivariate spatial-statistical models are often used when modeling environmental and socio-demographic processes. The most commonly used models for multivariate spatial covariances assume both stationarity and symmetry for the cross-covariances, but these assumptions are rarely tenable in practice. In this article we introduce a new and highly flexible class of nonstationary and asymmetric multivariate spatial covariance models that are constructed by modeling the simpler and more familiar stationary and symmetric multivariate covariances on a warped domain. Inspired by recent developments in the univariate case, we propose modeling the warping function as a composition of a number of simple injective warping functions in a deep-learning framework. Importantly, covariance-model validity is guaranteed by construction. We establish the types of warpings that allow for cross-covariance symmetry and asymmetry, and we use likelihood-based methods for inference that are computationally efficient. The utility of this new class of models is shown through two data illustrations: a simulation study on nonstationary data and an application on ocean temperatures at two different depths.

Keywords: Cross-Covariance, Deep Learning, Gaussian Process, Spatial Statistics, Warping

1 Introduction

Multivariate spatial-statistical models are used to jointly model two or more variables that are spatially indexed. They find widespread use in several application domains, such as the environmental sciences and the social sciences, where spatial processes interact. The utility of multivariate models lies in the concept of “borrowing strength,” where in this setting, information on one process (obtained, for example, through observation) imparts information on the other processes being jointly modeled, but that may or may not be directly observed. Multivariate spatial models need to adequately model both the marginal behavior of the spatial processes and the joint dependence between the processes. Often, the central object of interest when constructing a multivariate spatial model is the cross-covariance matrix function.

This function encodes the marginal covariances and cross-covariances of the spatial processes, and its use improves inference over one-at-a-time univariate analyses of each process.

Typically, the two simple assumptions of stationarity and symmetry are made when modeling the marginal behavior of, and the joint dependence between, multiple processes. For example, the popular linear model of coregionalization (LMC) (e.g., Goulard and Voltz, 1992; Wackernagel, 2003; Gu and Shen, 2020) assumes both of these properties, as does a recent multivariate model constructed using spectral representations (Qadir and Sun, 2020) and the multivariate Matérn model (Gneiting et al., 2010; Apanasovich et al., 2012). In the multivariate Matérn model, the elements of the cross-covariance matrix function are all Matérn covariance functions. However, although this has proven to be considerably more flexible than the traditional LMC, covariance nonstationarity and asymmetry are present in many scientific applications and should be modeled. For example, the rates of ice loss in Antarctica are clearly nonstationary, because more ice loss occurs in regions of high ice-stream velocity, which are at the boundary of the continent (Zammit-Mangion et al., 2015b). In addition, in the inversion of a trace-gas, the cross-covariance between the flux field and the mole-fraction field is asymmetric because of atmospheric transport (Zammit-Mangion et al., 2015a). Asymmetry of the cross-covariances is clearly present in the ocean temperature data in Section 4.2, due to diffusive and advective oceanographic processes.

Nonstationarity in a univariate setting has been addressed using spatial deformations (e.g., Sampson and Guttorp, 1992; Damian et al., 2001; Schmidt and O’Hagan, 2003; Fouedjio et al., 2015; Porcu et al., 2020); basis functions (e.g., Cressie and Johannesson, 2008); stochastic partial differential equations (SPDEs) (e.g., Lindgren et al., 2011; Fuglstad et al., 2015); and process convolution with a spatially varying kernel, which leads to spatially varying model parameters (e.g., Higdon et al., 1999; Paciorek and Schervish, 2006). Addressing nonstationarity in a multivariate setting is more problematic, because one needs to ensure validity, that is, nonnegative-definiteness, of all possible covariance matrices constructed through a proposed cross-covariance matrix function. Valid spatial multivariate modeling approaches that account for nonstationarity include those of Gelfand et al. (2004) and Kleiber and Nychka (2012), who extended the LMC and multivariate Matérn model, respectively, to contain spatially varying parameters. Other approaches consider basis functions (Nguyen et al., 2017) or systems of SPDEs (Hu and Steinsland, 2016; Hildeman et al., 2019).

Some multivariate models also model asymmetry. For example, Li and Zhang (2011) built on the bivariate example given by Ver Hoef and Cressie (1993) and extended the general stationary multivariate model to the asymmetric case. Apanasovich and Genton (2010) used latent dimensions to model asymmetric cross-covariances, and Cressie and Zammit-Mangion (2016) used a noncentered kernel to introduce asymmetry in the joint dependence structure; see also Ver Hoef and Barry (1998) and Majumdar and Gelfand (2007).

An attractive way to introduce nonstationarity and asymmetry in multivariate spatial-statistical models is through a generalization of the univariate deformation approach of Sampson and Guttorp (1992). It

is well known that nonstationarity can be modeled by deforming space; specifically, a stationary process on a warped domain can induce a highly nonstationary process on the original (geographic) domain. In the multivariate case, one may apply a common deformation to all of the processes, or separate deformations to each process. As we demonstrate here, using a common deformation function enforces symmetry and constrains the nonstationary behavior (i.e., the local anisotropies and scales) to be common for each process. However, process-specific deformation functions allow for distinct nonstationary behavior and practically guarantee asymmetry. Multivariate models built using spatial deformations bring with them several advantages over some of the other models mentioned above. In particular, they are able to capture complex nonstationary and asymmetric behavior. Their cross-covariance functions are valid by construction. Furthermore, if deep compositional warping functions are used, they are computationally efficient to fit and predict with.

The remainder of this paper is organized as follows. In Section 2, we provide some background on multivariate spatial models and univariate deep compositional spatial models (DCSMs). In Section 3, we introduce a multivariate generalization of the univariate DCSM, and show how asymmetry can be induced in a simple manner through what we call “aligning functions.” In Section 4, we show results from two data illustrations. First, we show the utility of multivariate DCSMs for modeling symmetric nonstationary simulated data. Then, we show its utility in an application on modeling Atlantic Ocean temperatures at two different depths. In these illustrations of our methodology, we show, through cross-validation and visualization, that spatial predictions from the multivariate DCSMs are generally superior to those from conventional multivariate spatial models. In Section 5, we summarize our conclusions. Additional material is provided in the Supplementary Material.

2 Background

The multivariate DCSM that we construct in Section 3 requires the specification of a conventional symmetric, stationary, possibly isotropic, multivariate covariance model and a deep warping function. In this section, we briefly review these two building blocks.

2.1 Multivariate Spatial Covariance Models

Consider a p -variate spatial process $\mathbf{Y}(\mathbf{s}) \equiv (Y_1(\mathbf{s}), \dots, Y_p(\mathbf{s}))'$, $\mathbf{s} \in G$, where we refer to $G \subset \mathbb{R}^d$ as the *geographic* domain in d -dimensional Euclidean space. We assume that $\text{var}(Y_i(\mathbf{s})) < \infty$, for all $\mathbf{s} \in G$ and all $i = 1, \dots, p$. Therefore, this multivariate process has finite expectation $\boldsymbol{\mu}(\cdot) \equiv (\mu_1(\cdot), \dots, \mu_p(\cdot))'$, and a valid cross-covariance matrix function $\mathbf{C}_G(\cdot, \cdot) \equiv (C_{ij,G}(\cdot, \cdot) : i, j = 1, \dots, p)$, where $C_{ij,G}(\mathbf{s}, \mathbf{u}) = \text{cov}(Y_i(\mathbf{s}), Y_j(\mathbf{u}))$; $\mathbf{s}, \mathbf{u} \in G$. For $i, j = 1, \dots, p$, the covariance function $C_{ii,G}(\cdot, \cdot)$ is the covariance function of the process $Y_i(\cdot)$ and, for $i \neq j$, $C_{ij,G}(\cdot, \cdot)$ is the cross-covariance function of $(Y_i(\cdot), Y_j(\cdot))'$.

In some cases, the cross-covariance matrix function depends only on $\mathbf{h} \equiv \mathbf{s} - \mathbf{u}$. That is, for $i, j =$

$1, \dots, p$, $C_{ij,G}(\mathbf{s}, \mathbf{u}) \equiv C_{ij,G}^o(\mathbf{h}); \mathbf{s}, \mathbf{u} \in G$, where now each $C_{ij,G}^o(\cdot)$ is a function of displacement. In this case, we say that $\mathbf{C}_G^o(\cdot) = (C_{ij,G}^o(\cdot) : i, j = 1, \dots, p)$ is a stationary cross-covariance matrix function. A cross-covariance matrix function is said to be symmetric if, for $i, j = 1, \dots, p$, $C_{ij,G}(\mathbf{s}, \mathbf{u}) = C_{ji,G}(\mathbf{s}, \mathbf{u})$, for $\mathbf{s}, \mathbf{u} \in G$. In the stationary case, symmetry is given by $C_{ij,G}^o(\mathbf{h}) = C_{ji,G}^o(\mathbf{h})$, for $\mathbf{h} = \mathbf{s} - \mathbf{u}$ and $\mathbf{s}, \mathbf{u} \in G$.

Stationarity and symmetry are strong assumptions in practice, but they remain popular because they facilitate the construction of valid cross-covariance matrix functions with a relatively small set of parameters. Among the most popular stationary symmetric multivariate covariance models is the multivariate Matérn model (Gneiting et al., 2010), where the marginal covariance functions and cross-covariance functions are all Matérn covariance functions. In this model, every process may have a different degree of smoothness, thus circumventing a key limitation of the LMC, where, for $i = 1, \dots, p$, the smoothness of each $Y_i(\cdot)$ is generally the same by construction.

The isotropic Matérn *correlation* function is given by $\mathcal{M}(\mathbf{h}|\nu, a) = \frac{2^{1-\nu}}{\Gamma(\nu)} (a\|\mathbf{h}\|)^\nu K_\nu(a\|\mathbf{h}\|)$, where ν is the smoothness parameter, $K_\nu(\cdot)$ is the modified Bessel function of the second kind of order ν , a is the scale parameter, and $\Gamma(\cdot)$ is the gamma function. A multivariate spatial-statistical process $\mathbf{Y}(\cdot)$ has a multivariate Matérn cross-covariance matrix function if, for $i, j = 1, \dots, p$, and $\mathbf{s}, \mathbf{u} \in G$,

$$\begin{aligned} C_{ii,G}(\mathbf{s}, \mathbf{u}) &\equiv \text{cov}(Y_i(\mathbf{s}), Y_i(\mathbf{u})) = \sigma_i^2 \mathcal{M}(\mathbf{h}|\nu_{ii}, a_{ii}), \\ C_{ij,G}(\mathbf{s}, \mathbf{u}) &\equiv \text{cov}(Y_i(\mathbf{s}), Y_j(\mathbf{u})) = \rho_{ij} \sigma_i \sigma_j \mathcal{M}(\mathbf{h}|\nu_{ij}, a_{ij}), \quad \text{for } i \neq j, \end{aligned} \tag{2.1}$$

where $\{a_{ij}\}$ are scale parameters, $\{\nu_{ij}\}$ are smoothness parameters, $\{\sigma_i^2\}$ are variance parameters, and $\{\rho_{ij}\}$ are cross-correlation parameters. From (2.1), we can see that multivariate Matérn cross-covariance functions are stationary, symmetric, and isotropic cross-covariance functions. In order to ensure validity, some constraints must be placed on the parameters of the multivariate Matérn covariance models shown in (2.1). The parsimonious Matérn covariance models have even stricter constraints than the more general multivariate Matérn models do (see also Apanasovich et al., 2012), but they have been shown to be flexible enough to model several environmental processes of interest (see Gneiting et al., 2010, for more details). We shall use the multivariate parsimonious Matérn covariance model in Section 4 to construct multivariate DCSMs for two bivariate spatial data sets.

2.2 Deep Compositional Spatial Models

The univariate deep compositional spatial modeling approach of Zammit-Mangion et al. (2021) uses injective warpings to construct nonstationary covariance models from simple covariance models. The idea of using deformations (or warpings) to modify the properties of a process stems from the work of Sampson and Guttorp (1992); see also Meiring et al. (1997), Sampson et al. (2001), Schmidt and O’Hagan (2003), Calandra et al. (2016), and the references therein. In this article, we extend the univariate deep compositional approach to the important multivariate case.

Consider, for the moment, a univariate process $Y(\cdot)$ with $\text{var}(Y(\mathbf{s})) < \infty$, $\mathbf{s} \in G$, and with nonstationary covariance function $C_G(\cdot, \cdot)$. After warping the space G , suppose that $C_G(\cdot, \cdot)$ can be expressed as a simpler stationary covariance function, $C_D(\cdot, \cdot)$, on a deformed space D , using a warping function $\mathbf{f} : G \rightarrow D$. Specifically, $C_G(\mathbf{s}, \mathbf{u}) \equiv C_D(\mathbf{f}(\mathbf{s}), \mathbf{f}(\mathbf{u}))$, for $\mathbf{s}, \mathbf{u} \in G$, where $C_D(\cdot, \cdot)$ is a familiar (stationary) covariance function. In DCSMs, the warping function \mathbf{f} is constrained to be smooth and injective in order to preclude the possibility of space-folding; see also Perrin and Monestiez (1999). In particular, it is expressed as the composition $\mathbf{f}(\cdot) \equiv \mathbf{f}_{[L]} \circ \mathbf{f}_{[L-1]} \circ \cdots \circ \mathbf{f}_{[1]}(\cdot)$, where $\mathbf{f}_{[1]}(\cdot), \dots, \mathbf{f}_{[L]}(\cdot)$ are simple elemental injective functions, and L is the number of warpings (or layers). This compositional construction is very flexible in that it can model highly nonstationary spatial processes, yet it is simple enough to facilitate parameter estimation from relatively sparse data. Zammit-Mangion et al. (2021) call the functions $\mathbf{f}_{[1]}(\cdot), \dots, \mathbf{f}_{[L]}(\cdot)$ warping units, and propose three types: axial warping units, radial basis function units, and Möbius transformation units. In this article, we also use these three types of warping units; see Table S1 in the Supplementary Material for more details. For example, in Section 4.1, the warping function $\mathbf{f}(\cdot)$ is a composition of $L = 4$ warping units, where $\mathbf{f}_{[1]}(\cdot)$ and $\mathbf{f}_{[2]}(\cdot)$ are two axial warping units (one for each spatial dimension), $\mathbf{f}_{[3]}(\cdot)$ is a radial basis function unit, and $\mathbf{f}_{[4]}(\cdot)$ is a Möbius transformation unit.

Zammit-Mangion et al. (2021) modeled a low-rank univariate process that was approximately stationary on the warped domain. Here, in a multivariate setting, we construct valid flexible models for covariances and cross-covariances on a geographic domain G by considering a stationary and symmetric cross-covariance matrix function on the warped domain D . We use the warping functions to model the nonstationary and asymmetric behavior of the multivariate spatial process.

3 Multivariate DCSMs

3.1 Model

We now construct multivariate DCSMs by extending the univariate construction of Zammit-Mangion et al. (2021). In the univariate case, one warping function is all that is required; however, in the p -variate case, we could use p different warping functions, one for each process.

We start off with the special case where a single warping function is used for all of the p processes. In this case, for $i, j = 1, \dots, p$, we have

$$C_{ij,G}(\mathbf{s}, \mathbf{u}) = C_{ij,D}(\mathbf{f}(\mathbf{s}), \mathbf{f}(\mathbf{u})) = C_{ij,D}^o(\mathbf{f}(\mathbf{s}) - \mathbf{f}(\mathbf{u})); \quad \mathbf{s}, \mathbf{u} \in G, \quad (3.1)$$

where $\mathbf{C}_D^o(\cdot) \equiv (C_{ij,D}^o(\cdot) : i, j = 1, \dots, p)$ is a stationary nonnegative-definite cross-covariance matrix function.

Proposition 1. *If $\mathbf{C}_D^o(\cdot)$ is symmetric, then the cross-covariance matrix function $\mathbf{C}_G(\cdot, \cdot)$ defined in (3.1), is symmetric.*

Proof: See Section S1.1 of the Supplementary Material.

Consider now the case where p warpings, one for each process, are used to construct the cross-covariance matrix function of the p -variate process. In this case, for $i, j = 1, \dots, p$,

$$C_{ij,G}(\mathbf{s}, \mathbf{u}) = C_{ij,D}(\mathbf{f}_i(\mathbf{s}), \mathbf{f}_j(\mathbf{u})) = C_{ij,D}^o(\mathbf{f}_i(\mathbf{s}) - \mathbf{f}_j(\mathbf{u})); \quad \mathbf{s}, \mathbf{u} \in G, \quad (3.2)$$

where $\{\mathbf{f}_i(\cdot) : i = 1, \dots, p\}$ are process-specific warping functions and, as in Proposition 1, the valid stationary cross-covariance matrix function $\mathbf{C}_D^o(\cdot)$ is symmetric.

Proposition 2. *If $\mathbf{C}_D^o(\cdot)$ is symmetric, then the cross-covariance matrix function $\mathbf{C}_G(\cdot, \cdot)$ defined in (3.2) is not necessarily symmetric.*

Proof: See Section S1.2 of the Supplementary Material.

The validity of the cross-covariance matrix function constructed using warping functions is established through the following proposition.

Proposition 3. *Assume that $\mathbf{C}_D^o(\cdot)$ is a valid, stationary cross-covariance matrix function. Consider the spatial locations $\{\mathbf{s}_{i1}, \dots, \mathbf{s}_{in_i}\}$, where $n_i > 0$, for $i = 1, \dots, p$. Let $\Sigma_G = (\Sigma_{ij,G} : i, j = 1, \dots, p)$, where $\Sigma_{ij,G} = (C_{ij,G}(\mathbf{s}_{ik}, \mathbf{s}_{jl}) : k = 1, \dots, n_i, l = 1, \dots, n_j)$, $N = \sum_{i=1}^p n_i$, and $C_{ij,G}(\cdot, \cdot)$ is given by (3.2). Then, Σ_G is nonnegative-definite.*

Proof: See Section S1.3 of the Supplementary Material.

Summarizing the results of Proposition 1–3, we see that if $\mathbf{C}_D^o(\cdot)$ is a valid stationary cross-covariance matrix function, then the cross-covariance matrix function $\mathbf{C}_G(\cdot, \cdot)$ constructed through (3.2) is valid (i.e., nonnegative-definite). Furthermore, if $\mathbf{f}_i(\cdot) \neq \mathbf{f}_j(\cdot)$, for any $i, j = 1, \dots, p$, then the cross-covariance matrix function is not necessarily symmetric.

Using p general warpings, as in (3.2), will yield a highly flexible parameterized model, but one that may be prone to over-fitting. In practice, any asymmetry present is likely to be simple and dominated by global shifts and rotations. Hence, to model asymmetry, we propose expressing each $\mathbf{f}_i(\cdot)$ as a composition of a shared warping function $\mathbf{f}(\cdot)$ and a process-specific “aligning function” $\mathbf{g}_i(\cdot)$, for $i = 1, \dots, p$. That is, for $i, j = 1, \dots, p$, we let

$$\begin{aligned} C_{ij,G}(\mathbf{s}, \mathbf{u}) &= C_{ij,D}(\mathbf{f} \circ \mathbf{g}_i(\mathbf{s}), \mathbf{f} \circ \mathbf{g}_j(\mathbf{u})) \\ &= C_{ij,D}^o(\mathbf{f} \circ \mathbf{g}_i(\mathbf{s}) - \mathbf{f} \circ \mathbf{g}_j(\mathbf{u})); \quad \mathbf{s}, \mathbf{u} \in G, \end{aligned} \quad (3.3)$$

where $\mathbf{g}_i(\cdot)$, for $i = 1, \dots, p$, are simple transformations commonly used to align spatial fields and that can include translations and rotations (e.g., Wiens et al., 2020). It can be seen from Propositions 2 and 3 that the effect of the aligning functions $\{\mathbf{g}_i\}$ is to introduce asymmetry, while preserving nonnegative-definiteness. Note that a common frame of reference for the aligning functions needs to be chosen when aligning fields in this way. Therefore, without loss of generality, we fix $\mathbf{g}_1(\cdot)$ to be the identity map, in

which case, $\mathbf{f}_1(\cdot) = \mathbf{f} \circ \mathbf{g}_1(\cdot)$ is simply the shared warping function $\mathbf{f}(\cdot)$. The cross-covariance model in (3.3) is a generalization of the asymmetric cross-covariance model of Li and Zhang (2011), where the shared warping function $\mathbf{f}(\cdot)$ is the identity map and the aligning functions are translations, that is, where $\mathbf{g}_i(\mathbf{s}) = \mathbf{s} + \mathbf{d}_i$, for some $\mathbf{d}_i \in \mathbb{R}^2$, $i = 2, \dots, p$.

Under the cross-covariance-matrix model (3.3), nonstationarity can be introduced through both the shared warping function $\mathbf{f}(\cdot)$ and the aligning functions $\{\mathbf{g}_i(\cdot) : i = 1, \dots, p\}$. Perhaps not surprisingly, the aligning functions can induce nonstationarity in the cross-covariance functions even when $\mathbf{f}(\cdot)$ is the identity map, as we demonstrate in the following proposition.

Proposition 4. *Consider the p -variate cross-covariance matrix model (3.3), where $\mathbf{f}(\cdot)$ is the identity map; one of the aligning functions $\mathbf{g}_k(\cdot)$, for some $k \in \{2, \dots, p\}$, is an affine transformation, and $\{\mathbf{g}_i(\cdot), i \neq k\}$ are identity maps. Then, $C_{ik,G}(\cdot, \cdot)$, for $i \neq k$, is not necessarily stationary.*

Proof: See Section S1.4 of the Supplementary Material.

Proposition 4 represents one simple way to introduce nonstationarity. More generally, when one has p warping functions $\{\mathbf{f}_i(\cdot) : i = 1, \dots, p\}$, nonstationarity of $\mathbf{C}_{ij,G}(\cdot, \cdot)$ is obtained by choosing $\mathbf{f}_i(\cdot)$ and $\mathbf{f}_j(\cdot)$ such that $\mathbf{f}_i(\mathbf{s}) - \mathbf{f}_j(\mathbf{u})$ is not a function of $\mathbf{s} - \mathbf{u}$, for $\mathbf{s}, \mathbf{u} \in G$.

3.2 Parameter Estimation

Assume now that we have observations $\{Z_{ik} : k = 1, \dots, n_i; i = 1, \dots, p\}$ of a p -variate Gaussian process $\tilde{\mathbf{Y}}(\cdot)$, where

$$Z_{ik} = \tilde{Y}_i(\mathbf{s}_{ik}) + \epsilon_{ik}; \quad k = 1, \dots, n_i, \quad i = 1, \dots, p. \quad (3.4)$$

In (3.4), $\{\epsilon_{ik} : k = 1, \dots, n_i; i = 1, \dots, p\}$ are independent Gaussian measurement errors that satisfy $\epsilon_{ik} \sim \text{Gau}(0, \tau_i^2)$ for $k = 1, \dots, n_i$ and $i = 1, \dots, p$, and $\tau_1^2, \dots, \tau_p^2$ are measurement-error variances that are assumed to be process-specific, and hence all potentially different. We model the p -variate Gaussian process $\tilde{\mathbf{Y}}(\cdot) = (\tilde{Y}_1(\cdot), \dots, \tilde{Y}_p(\cdot))'$ to have first moment that is linear in covariates $\mathbf{x}(\cdot) \equiv (x_1(\cdot), x_2(\cdot), \dots, x_q(\cdot))'$. That is,

$$\tilde{Y}_i(\cdot) = \mathbf{x}(\cdot)' \boldsymbol{\beta}_i + Y_i(\cdot), \quad i = 1, \dots, p, \quad (3.5)$$

where $\boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_p \in \mathbb{R}^q$ are vectors of unknown coefficients that need to be estimated, and now $\mathbf{Y}(\cdot) \equiv (Y_1(\cdot), \dots, Y_p(\cdot))'$ is a zero-mean second-order nonstationary multivariate Gaussian process on the geographic domain G .

Let $\tilde{\mathbf{Y}}_i \equiv (\tilde{Y}_i(\mathbf{s}_{i1}), \dots, \tilde{Y}_i(\mathbf{s}_{in_i}))'$, $\tilde{\mathbf{Y}} \equiv (\tilde{\mathbf{Y}}_1', \dots, \tilde{\mathbf{Y}}_p')$, $\mathbf{Y}_i \equiv (Y_i(\mathbf{s}_{i1}), \dots, Y_i(\mathbf{s}_{in_i}))'$, $\mathbf{Y} \equiv (\mathbf{Y}_1', \dots, \mathbf{Y}_p')$, $\mathbf{X}_i \equiv (\mathbf{x}(\mathbf{s}_{i1}), \dots, \mathbf{x}(\mathbf{s}_{in_i}))'$, $\mathbf{X} = \text{bdiag}(\mathbf{X}_1, \dots, \mathbf{X}_p)$, where $\text{bdiag}(\cdot)$ returns a block diagonal matrix from its arguments, and $\boldsymbol{\beta} = (\boldsymbol{\beta}_1', \dots, \boldsymbol{\beta}_p')$. Then, (3.5) can be written compactly as

$$\tilde{\mathbf{Y}} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Y}. \quad (3.6)$$

The covariance matrix of \mathbf{Y} , $\Sigma_G \equiv \text{cov}(\mathbf{Y})$, is given by $(\Sigma_{ij,G} : i, j = 1, \dots, p)$, where $\Sigma_{ij,G} \equiv (C_{ij,G}(\mathbf{s}_{ik}, \mathbf{s}_{jl}) : k = 1, \dots, n_i; l = 1, \dots, n_j)$. Furthermore, $\text{cov}(\tilde{\mathbf{Y}}) = \text{cov}(\mathbf{Y})$.

Let $\mathbf{Z}_i \equiv (Z_{i1}, \dots, Z_{in_i})'$, $\mathbf{Z} \equiv (\mathbf{Z}'_1, \dots, \mathbf{Z}'_p)'$, $\boldsymbol{\epsilon}_i \equiv (\epsilon_{i1}, \dots, \epsilon_{in_i})'$, and $\boldsymbol{\epsilon} \equiv (\boldsymbol{\epsilon}'_1, \dots, \boldsymbol{\epsilon}'_p)'$. Then, from (3.4) and (3.6), we have, $\mathbf{Z} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Y} + \boldsymbol{\epsilon}$, where the covariance matrix of $\boldsymbol{\epsilon}$, $\mathbf{V} \equiv \text{cov}(\boldsymbol{\epsilon})$, is diagonal. The model for the observations, \mathbf{Z} , is therefore

$$\mathbf{Z} \sim \text{Gau}(\mathbf{X}\boldsymbol{\beta}, \Sigma_Z), \quad (3.7)$$

where $\Sigma_Z = \Sigma_G + \mathbf{V}$.

Likelihood-based inference can be used to estimate the parameters (including the warping parameters) in (3.7) (Zammit-Mangion et al., 2021). Here, we use restricted maximum likelihood (REML) in which the parameters in $\boldsymbol{\beta}$ have been marginalized out to estimate the parameters in Σ_Z . REML is known to provide less biased estimators of variance-component parameters (Cressie and Lahiri, 1996), and can help avoid degenerate predictions (e.g., Gu et al., 2018). Let $\boldsymbol{\theta}$ be the vector containing all parameters appearing in the covariance matrix Σ_Z , which includes the unknown parameters appearing in the warping function $\mathbf{f}(\cdot)$ (and the aligning functions $\{\mathbf{g}_i(\cdot)\}$, if present), the parameters in the cross-covariance matrix function of the process on D (i.e., the scale parameter a , smoothness parameters $\{\nu_{ij}\}$, variance parameters $\{\sigma_i^2\}$, and cross-correlation parameters $\{\rho_{ij}\}$ for the parsimonious Matérn covariance function), and the measurement-error variances. The restricted maximum likelihood estimate $\hat{\boldsymbol{\theta}}$ of $\boldsymbol{\theta}$ is found by maximizing the log-restricted likelihood $\mathcal{L}(\boldsymbol{\theta}; \mathbf{Z})$ with respect to $\boldsymbol{\theta}$, for some given \mathbf{Z} . After obtaining $\hat{\boldsymbol{\theta}}$, the associated estimate $\hat{\boldsymbol{\beta}}$ of $\boldsymbol{\beta}$ is found through generalized least squares. The REML parameter estimates are “plugged in” and allow spatial predictions of the hidden processes at an unobserved location \mathbf{s}^* . For more details on fitting and prediction, see Section S2 of the Supplementary Material.

The computational time complexity of evaluating the restricted likelihood is the sum of that of evaluating the deformation function and that of factorizing the full joint covariance matrix on the deformed space. The computational complexity of evaluating the aligning functions $\{\mathbf{g}_i(\cdot)\}$, when these are affine transformations, is $O(\sum_{i=2}^p n_i)$. The computational complexity of evaluating the shared warping-function layer $\mathbf{f}_{[l]}(\cdot)$ is $O(Nr_l)$, where $N = \sum_{i=1}^p n_i$ is the total number of observations for all processes, and r_l is the number of basis functions in $\mathbf{f}_{[l]}(\cdot)$. Hence, the total complexity for evaluating the deformation function is $O(\sum_{i=2}^p n_i + N \sum_{l=1}^L r_l)$. The complexity of factorizing the covariance matrix on the deformed space is $O(N^3)$. Usually, we choose $r_l \ll N$, for $l = 1, \dots, L$, so that the computational complexity is dominated by the factorization of the covariance matrix on the deformed space. The actual runtime also changes with the number of iterations used to optimize the parameter estimates, which needs to be larger when the model is more complex. Parameter estimation was done using gradient-based optimization via the R package `tensorflow` (Allaire and Tang, 2019), which computes gradients using automatic differentiation, and can be run on a graphics processing unit (GPU).

3.3 Fixing the Frame of Reference

While setting $\mathbf{g}_1(\cdot)$ to be the identity map establishes a common frame of reference for the aligning functions, the shared warping function $\mathbf{f}(\cdot)$ and any scale parameters appearing in the cross-covariance matrix function are themselves non-identifiable unless this common frame of reference is fixed. Non-identifiability occurs when there exist at least two distinct parameters, $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2$ say, for which $\mathcal{L}(\boldsymbol{\theta}_1; \cdot) = \mathcal{L}(\boldsymbol{\theta}_2; \cdot)$ (Kadane, 1974). Note that this notion of non-identifiability of parameters differs from the identifiability problem of consistent estimation under infill asymptotics described by Zhang (2004). In our case, if we use a stationary symmetric cross-covariance matrix function on the warped domain that is also isotropic, the likelihood is invariant to translation, rotation, or reflection of $\mathbf{f}(\cdot)$. Because we also allow $\mathbf{f}(\cdot)$ to stretch and contract the geographic domain, any scale parameter a_{ij} associated with the cross-covariance function $C_{ij,D}^o(\cdot)$ is also non-identifiable (see Anderes and Stein, 2008; Anderes and Chatterjee, 2009, for more details). While this invariance and lack of identifiability does not pose a problem for prediction, it does mean that we cannot make inference on certain properties of the warping function, such as stretches/contractions and rotations, without further assumptions. As we discuss next, it helps to use a function $\mathbf{b}_0(\cdot)$, which we call a *homogenizing function*, to place the estimates of $\mathbf{f}(\cdot)$ in a fixed frame of reference and to obtain transformations of the scale parameters that are identifiable.

We illustrate our methodology on the two-dimensional Euclidean space with $d = 2$. Specifically, we establish a fixed frame of reference (which can be easily generalized for $d > 2$) when we assume isotropy on D , as follows. Consider three locations, say $\mathbf{s}_k, \mathbf{s}_l$, and \mathbf{s}_m , in $G \subset \mathbb{R}^2$ such that $\mathbf{f}(\mathbf{s}_k), \mathbf{f}(\mathbf{s}_l)$, and $\mathbf{f}(\mathbf{s}_m)$ are not colinear. Then, we use the homogenizing function to shift, scale, rotate, and reflect the warped domain D such that $\mathbf{b}_0 \circ \mathbf{f}(\mathbf{s}_k) = (0, 0)'$, $\mathbf{b}_0 \circ \mathbf{f}(\mathbf{s}_l) = (1, 0)'$, and $b_{0,2} \circ \mathbf{f}(\mathbf{s}_m) > 0$, where $b_{0,2}(\cdot)$ refers to the second element of $\mathbf{b}_0(\cdot)$. A homogenizing function that accomplishes these transformations is given by

$$\mathbf{b}_0(\cdot) \equiv \mathbf{b}_3 \circ \mathbf{b}_2 \circ \mathbf{b}_1(\cdot), \quad (3.8)$$

where $\mathbf{b}_1(\cdot)$ shifts and scales, $\mathbf{b}_2(\cdot)$ rotates around the origin, and $\mathbf{b}_3(\cdot)$ reflects around the horizontal axis. Figure S1 in the Supplementary Material illustrates the effect of the homogenizing function $\mathbf{b}_0(\cdot)$ on points in D .

Denote $\tilde{\mathbf{s}}_k \equiv \mathbf{f}(\mathbf{s}_k)$, $\tilde{\mathbf{s}}_l \equiv \mathbf{f}(\mathbf{s}_l)$, and $\tilde{\mathbf{s}}_m \equiv \mathbf{f}(\mathbf{s}_m)$. The shifting and scaling is done through the function $\mathbf{b}_1(\mathbf{s}) \equiv \frac{1}{\|\tilde{\mathbf{s}}_l - \tilde{\mathbf{s}}_k\|}(\mathbf{s} - \tilde{\mathbf{s}}_k)$; $\mathbf{s} \in D$. Denote the scaled and shifted domain as D_1 , where $D_1 \equiv \{\mathbf{b}_1(\mathbf{s}) : \mathbf{s} \in D\}$. Note that the distance between $\mathbf{b}_1(\tilde{\mathbf{s}}_l)$ and $\mathbf{b}_1(\tilde{\mathbf{s}}_k)$ is fixed to one in D_1 . The operation that rotates $\mathbf{b}_1(\tilde{\mathbf{s}}_l)$ to the point $(1, 0)'$ is given by

$$\mathbf{b}_2(\mathbf{s}) \equiv \begin{pmatrix} \cos \psi_l & \sin \psi_l \\ -\sin \psi_l & \cos \psi_l \end{pmatrix} \mathbf{s}; \quad \mathbf{s} \in D_1,$$

where $\psi_l = \text{atan2}(b_{1,2}(\tilde{\mathbf{s}}_l), b_{1,1}(\tilde{\mathbf{s}}_l))$ is the angle of $\mathbf{b}_1(\tilde{\mathbf{s}}_l)$, and $b_{1,i}(\cdot)$ refers to the i th element of $\mathbf{b}_1(\cdot)$. Denote the scaled, shifted, and rotated domain as D_2 , where $D_2 \equiv \{\mathbf{b}_2(\mathbf{s}) : \mathbf{s} \in D_1\}$. Finally, the

reflection operation that ensures that $b_{0,2}(\tilde{\mathbf{s}}_m) > 0$ is given by

$$\mathbf{b}_3(\mathbf{s}) \equiv \begin{pmatrix} 1 & 0 \\ 0 & g_m \end{pmatrix} \mathbf{s}; \quad \mathbf{s} \in D_2,$$

where $g_m \equiv \text{sign}(\mathbf{b}_2 \circ \mathbf{b}_1(\tilde{\mathbf{s}}_m))$, equal to -1 if a reflection around the horizontal axis is needed, and equal to $+1$ otherwise. The fixed frame of reference is defined to be the domain $D_3 \equiv \{\mathbf{b}_3(\mathbf{s}) : \mathbf{s} \in D_2\}$.

Fixing the frame of reference can be useful when, for example, one is bootstrapping to do uncertainty quantification of the warped locations, because these warped locations are non-identifiable otherwise. Importantly, we have the following result when the covariance functions in the deformed space are solely functions of (scaled) distances.

Theorem 1. *Assume that the cross-covariance functions on the warped domain, $\tilde{C}_{ij,D}^o(\mathbf{h}; a_{ij})$, where $\{a_{ij}\}$ are scale parameters, are solely functions of $a_{ij}\|\mathbf{h}\|$, $\mathbf{h} \in \mathbb{R}^2$, $a_{ij} > 0$. Consider two cross-covariance matrix functions $\mathbf{C}_G^{(1)}(\cdot, \cdot)$ and $\mathbf{C}_G^{(2)}(\cdot, \cdot)$, where $C_{ij,G}^{(r)}(\mathbf{s}, \mathbf{u}) \equiv \tilde{C}_{ij,D}^o(\|\mathbf{f}^{(r)}(\mathbf{s}) - \mathbf{f}^{(r)}(\mathbf{u})\|; a_{ij}^{(r)})$, for $r = 1, 2$, and $\mathbf{s}, \mathbf{u} \in G$. If $\mathbf{C}_G^{(1)}(\cdot, \cdot) = \mathbf{C}_G^{(2)}(\cdot, \cdot)$, then $\mathbf{b}_0 \circ \mathbf{f}^{(1)}(\cdot) = \mathbf{b}_0 \circ \mathbf{f}^{(2)}(\cdot)$, where $\mathbf{b}_0(\cdot)$ is given by (3.8), and $a_{ij}^{(1)}\|\mathbf{f}^{(1)}(\mathbf{s}_l) - \mathbf{f}^{(1)}(\mathbf{s}_k)\| = a_{ij}^{(2)}\|\mathbf{f}^{(2)}(\mathbf{s}_l) - \mathbf{f}^{(2)}(\mathbf{s}_k)\|$. Conversely, if $\mathbf{b}_0 \circ \mathbf{f}^{(1)}(\cdot) = \mathbf{b}_0 \circ \mathbf{f}^{(2)}(\cdot)$ and $a_{ij}^{(1)}\|\mathbf{f}^{(1)}(\mathbf{s}_l) - \mathbf{f}^{(1)}(\mathbf{s}_k)\| = a_{ij}^{(2)}\|\mathbf{f}^{(2)}(\mathbf{s}_l) - \mathbf{f}^{(2)}(\mathbf{s}_k)\|$, for $i, j = 1, \dots, p$, then $\mathbf{C}_G^{(1)}(\cdot, \cdot) = \mathbf{C}_G^{(2)}(\cdot, \cdot)$.*

Proof: See Section S1.5 of the Supplementary Material.

Theorem 1 shows that, after homogenization using (3.8), locations warped using functions that yield the same cross-covariance matrix functions on the geographic domain must coincide. This result can be used to obtain a visual appreciation of the uncertainty in the estimated warping function when bootstrapping the warping parameters. Informally, after homogenization, two covariance functions that are similar should yield points that are in close proximity to one another, and vice versa. We use such a visual diagnostic in our simulation study in Section 4.1.

Theorem 1 also reveals that there is a one-to-one correspondence between the cross-covariance matrix function on the geographic domain and the scale parameters in the warped domain after homogenization. Specifically, $\tilde{a}_{ij} = a_{ij}^{(1)}\|\mathbf{f}^{(1)}(\mathbf{s}) - \mathbf{f}^{(1)}(\mathbf{u})\| = a_{ij}^{(2)}\|\mathbf{f}^{(2)}(\mathbf{s}) - \mathbf{f}^{(2)}(\mathbf{u})\|$, for $i, j = 1, \dots, p$, if and only if $\mathbf{C}_G^{(1)}(\mathbf{s}, \mathbf{u}) = \mathbf{C}_G^{(2)}(\mathbf{s}, \mathbf{u})$, for all $\mathbf{s}, \mathbf{u} \in G$. This leads to the following corollary, which shows that consistent inference of a transformation of the different processes' scale parameters in the warped domain can be made after homogenizing the warpings to a fixed frame of reference. This can be useful for validating our methods when the true warping function is known, as it is in the simulation study presented in Section 4.1.

Corollary 1. *Assume the conditions of Theorem 1 hold, and define $\tilde{a}_{ij} \equiv a_{ij}\|\mathbf{f}(\mathbf{s}_l) - \mathbf{f}(\mathbf{s}_k)\|$, for $i, j = 1, \dots, p$. Then, the set comprising the homogenized warping function and transformed scale parameters, $\{\mathbf{b}_0 \circ \mathbf{f}(\cdot), \{\tilde{a}_{ij}\}\}$, is identifiable. That is, two sets of parameters $\{\mathbf{b}_0 \circ \mathbf{f}^{(1)}(\cdot), \{\tilde{a}_{ij}^{(1)}\}\}$ and*

$\{\mathbf{b}_0 \circ \mathbf{f}^{(2)}(\cdot), \{\tilde{a}_{ij}^{(2)}\}\}$, where $\tilde{a}_{ij}^{(r)} = a_{ij}^{(r)} \|\mathbf{f}^{(r)}(\mathbf{s}_l) - \mathbf{f}^{(r)}(\mathbf{s}_k)\|$, for $r = 1, 2$, yield the same log-restricted likelihood function $\mathcal{L}^{(1)}(\cdot; \cdot)$ and $\mathcal{L}^{(2)}(\cdot; \cdot)$ if and only if they are identical. That is, $\{\mathbf{b}_0 \circ \mathbf{f}(\cdot), \{\tilde{a}_{ij}\}\}$ is identifiable in the sense of Kadane (1974).

Proof: See Section S1.5 of the Supplementary Material.

Fixing the frame of reference allows us to do uncertainty quantification on any warping function parameters and transformed scale parameters. While under certain conditions, REML estimators are asymptotically Gaussian (Cressie and Lahiri, 1996), we are not aware of an analytical form of the asymptotic distribution of the REML estimators for a nonstationary covariance model constructed through deformation. Hence, we use bootstrapping to make inference on these parameters. Bootstrapping with spatial data needs to be done with care, because the data are correlated; see Solow (1985) and Olea and Pardo-Iguzquiza (2011) for more discussion. A bootstrapping algorithm for quantifying the uncertainties of the parameters in model (3.7) is shown in Algorithm 1 in Section S3 of the Supplementary Material (Olea and Pardo-Iguzquiza, 2011). We use Algorithm 1 to visualize uncertainties on warped locations, and we use it for uncertainty quantification of parameter estimates in the simulation study of Section 4.1.

4 Data Illustrations

In this section, we use two illustrations to show the potential benefit of using multivariate DCSMs over conventional ones. In Section 4.1, we show the results from a study using data simulated from a symmetric nonstationary bivariate-covariance model. In Section 4.2, we show the results from a study using North Atlantic Ocean temperatures at two different depths. Section S4 of the Supplementary Material contains additional data illustrations using data simulated from an asymmetric nonstationary bivariate covariance model, from models with misspecified warping functions, and from a trivariate covariance model. It also contains an experiment using real maximum-and-minimum-temperature data in the United States. The code and data for reproducing the results from all our data illustrations are available from https://github.com/quantvu17/deepspat_multivar.

4.1 Simulated Symmetric Nonstationary Data

We first demonstrate the use of multivariate DCSMs on data simulated using a symmetric nonstationary bivariate covariance model. We simulated the bivariate data from a Gaussian multivariate DCSM, $\tilde{\mathbf{Y}}(\cdot)$, with constant mean (i.e., $q = 1$ and $\mathbf{x}(\cdot) = x_1(\cdot) = 1$ in (3.5), so that there are two intercepts, β_{11} and β_{21} , that need to be estimated). The data were simulated on an equally spaced 101×101 grid of the geographic domain $G \equiv [-0.5, 0.5] \times [-0.5, 0.5]$. The warping function we used was a composition of axial warping units, followed by a single-resolution radial basis function unit, followed by a Möbius transformation unit; see Zammit-Mangion et al. (2021) for a detailed description of these warping units. On the warped

Table 1: Average hold-out validation results, AIC, and the time required to fit for the simulation study in Section 4.1, where the measurement locations are randomly sampled 30 times from G .

	$\tilde{Y}_1(\cdot)$		$\tilde{Y}_2(\cdot)$		AIC	Time (s)
	RMSPE	CRPS	RMSPE	CRPS		
Model 4.1.1	0.404	0.221	0.089	0.048	887.0	823.5
Model 4.1.2	0.358	0.196	0.072	0.037	257.0	1545.4

domain, we modeled the covariances using a stationary isotropic multivariate parsimonious Matérn model. We randomly sampled 1000 locations from the grid, and used these as measurement locations.

We compared the predictions of the stationary parsimonious Matérn model (Model 4.1.1) with those of the multivariate DCSM (Model 4.1.2) in order to gauge the loss in prediction performance when the nonstationarity arising from the warping is ignored. After fitting Model 4.1.1 and Model 4.1.2 to the observations at the 1000 locations, we computed the predictions and prediction standard errors of the latent processes on the 101×101 grid. Figure 1 shows the true simulated fields, predictions, and prediction standard errors, from both models. As observed in the univariate case (Zammit-Mangion et al., 2021), we see that the DCSM can predict sharp features in the spatial fields, while the stationary parsimonious Matérn model smooths out such features. Furthermore, while the stationary parsimonious Matérn model produces prediction standard errors that are mostly unrelated to the process behavior (due to the stationarity assumption), the DCSM produces prediction standard errors that are highly reflective of the processes’ local anisotropies and scales. These visualizations illustrate the advantages of using a multivariate DCSM over a stationary multivariate model when the underlying processes are highly nonstationary.

To compare the predictive performance of the two models quantitatively, we calculated the predictive performance at the remaining 9201 locations using two commonly used scoring rules, namely the root mean squared prediction error (RMSPE) and the continuous-ranked probability score (CRPS) (Gneiting and Raftery, 2007). We repeated the procedure of random sampling 1000 locations and accessing the predictive performance 30 times. Table 1 summarizes the results, with averages from the 30 cross-validation studies, and also gives the average Akaike information criterion (AIC) from these studies. Figure S2 in the Supplementary Material shows box plots of the RMSPE and CRPS for both models across the 30 simulations. From the table and figure, it is clear that there is a large improvement in RMSPE and CRPS when using the DCSM (Model 4.1.2) over the stationary parsimonious Matérn model (Model 4.1.1). This improvement was expected because the data were generated from the highly nonstationary process. Using the more sophisticated nonstationary model comes at some computational cost. It took 1545.4 seconds, on average, to fit Model 4.1.2, almost twice as long as the 823.5 seconds it took to fit Model 4.1.1.

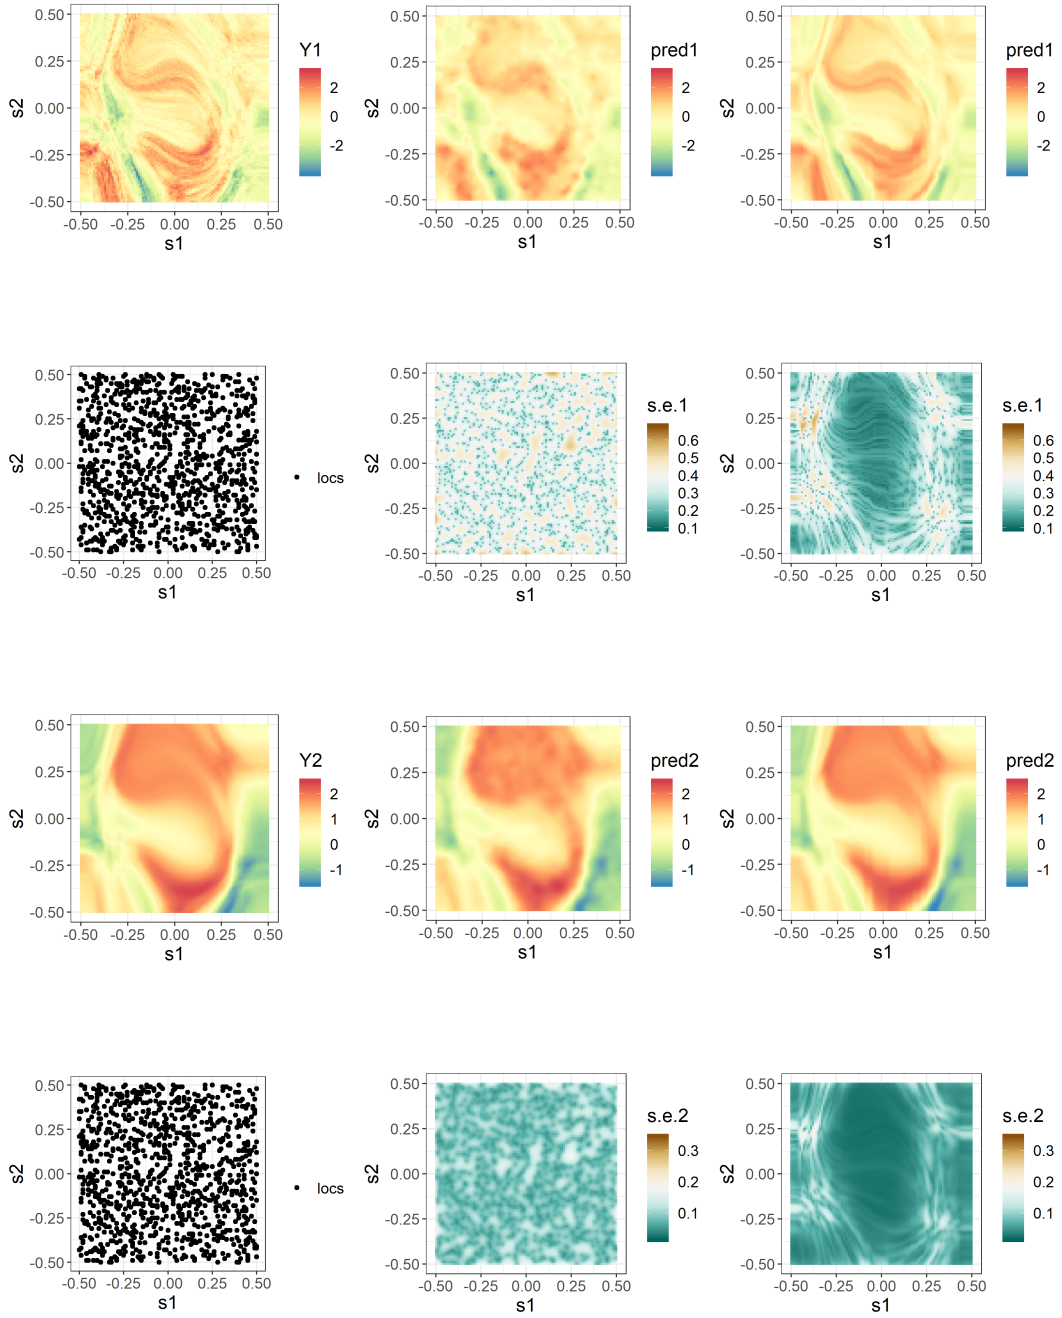


Figure 1: Comparison of predictions and prediction standard errors when using a bivariate stationary parsimonious Matérn model (Model 4.1.1) and a bivariate DCSM (Model 4.1.2) in the study of Section 4.1, where measurement locations were randomly sampled from G . First row: The process $\tilde{Y}_1(\cdot)$ (left panel), and the predictions obtained using Model 4.1.1 (center panel) and Model 4.1.2 (right panel). Second row: Locations of the measurement of $\tilde{Y}_1(\cdot)$ (left panel), and the prediction standard errors obtained when using Model 4.1.1 (center panel) and Model 4.1.2 (right panel). Third and fourth rows: Analogous to the first and second rows, respectively, for the process $\tilde{Y}_2(\cdot)$.




Figure 2: Measurement locations on the original domain G and the warped domain after homogenization, D_3 . Top row: Measurement locations on the original domain (left panel); true warped measurement locations (right panel). Bottom row: REML estimate of the warped measurement locations (left panel); bootstrap distribution of the warped measurement locations (right panel; for visualization purposes, only 100 bootstrap samples are shown).

We next used a bootstrap to examine the ability of the DCSM to recover the true parameters. We bootstrapped in a fixed frame of reference (via the homogenizing function) with 1000 bootstrap samples to quantify uncertainties on the model parameters using the method outlined in Section 3.3. Figure 2 shows the measurement locations in G , the measurement locations under the true warping function and homogenization, the measurement locations under the estimated warping function and homogenization, and the bootstrap samples of the warped locations. We see that the estimated warped locations and the bootstrap samples of the warped locations are similar to the warped locations under the true warping function. Specifically, important features, such as the contraction in the middle part of the domain, are recovered. Table S2 in the Supplementary Material lists the true cross-covariance matrix function parameters, along with their estimates and their 95% bootstrap confidence intervals. The REML estimates are relatively close to the true parameter values, and all of the 95% bootstrap confidence intervals of the model parameters contain the true values.

We also considered the case where the data are missing in a block, which is shown in Section S4.1 in the Supplementary Material.

4.2 Modeling Temperatures in the North Atlantic Ocean at Two Different Depths

We next consider sea temperatures in the North Atlantic Ocean at two very different depths: 0.5 m and 318.1 m. The data were obtained from the Copernicus Marine Environment Monitoring Service (CMEMS)¹¹. We analyzed temperatures on July 1, 2018, between 36.3°N–39.6°N and 60.0°W–63.3°W, with 1600 measurements on a 40 × 40 grid. Panels (1,1) and (3,1) in Figure 3 show the temperatures at the two depths, where we can see that there is a small amount of misalignment in the temperature processes, suggesting that the modeling of cross-covariance asymmetry may be important.

We considered the following models:

- Model 4.2.1: A bivariate stationary symmetric parsimonious Matérn model with only an intercept in the trend.
- Model 4.2.2: A bivariate symmetric DCSM, where the warping function $\mathbf{f}(\cdot)$ is a composition of axial warping units, a single-resolution radial basis function unit, and a Möbius transformation unit, and with Model 4.2.1 on the warped domain.
- Model 4.2.3: A bivariate asymmetric DCSM, where the aligning function $\mathbf{g}_2(\cdot)$ is an affine transformation (as described in Proposition 4) and the warping function is as in Model 4.2.2, and with Model 4.2.1 on the warped domain.

The predictive performance of these three models was first examined using a five-fold cross-validation study, where we randomly divided the 1600 measurement locations into five groups. The results are summarized in Table S3 in the Supplementary Material. We see that allowing for model nonstationarity and/or asymmetry does indeed result in improved predictions, but the observed improvement is not substantial.

We also considered the case where the data are missing in a block. Specifically, we assumed that we have all the measurements on the grid, except those between 37.5°N–38.2°N. Figure 3 shows the true fields and the predictions using Model 4.2.1, Model 4.2.2, and Model 4.2.3, while Table 2 shows the diagnostic results when predicting the temperature at the missing locations. The improvement of the bivariate DCSMs over the stationary symmetric Model 4.2.1 is evident in this case. Observe that the asymmetric version of the DCSM (Model 4.2.3) produces slightly better predictive diagnostics than those of the symmetric version (Model 4.2.2), illustrating the importance of being able to model asymmetry. Visualizations of the nonstationary and asymmetric structure of estimated cross-covariance matrix function are given in Figure S3 and Figure S4 in the Supplementary Material.

¹¹http://marine.copernicus.eu/services-portfolio/access-to-products/?option=com_csw&view=details&product_id=GLOBAL_ANALYSIS_FORECAST_PHY_001_024




Figure 3: Comparison of predictions when using a symmetric stationary parsimonious Matérn model (Model 4.2.1), a bivariate symmetric DCSM (Model 4.2.2), and a bivariate asymmetric DCSM (Model 4.2.3). First row: Original temperature observations at depth 0.5 m, Z_1 (first panel); predictions obtained using Model 4.2.1 (second panel), Model 4.2.2 (third panel), and Model 4.2.3 (fourth panel). Second row: Locations of the retained measurements (first panel); prediction standard errors obtained when using Model 4.2.1 (second panel), Model 4.2.2 (third panel), and Model 4.2.3 (fourth panel). Third and fourth rows: Analogous to the first and second rows, respectively, for Z_2 , the temperature observations at depth 318.1 m.

Table 2: Hold-out validation results, AIC, and the time required to fit for the ocean temperature data at depths 0.5 m and 318.1 m for the study in Section 4.2, where data were missing in the white block shown in the second and fourth rows of Figure 3.

	$T_{0.5}$		$T_{318.1}$		AIC	Time (s)
	RMSPE	CRPS	RMSPE	CRPS		
Model 4.2.1	0.450	0.226	0.463	0.236	-5274.7	1311.1
Model 4.2.2	0.234	0.136	0.301	0.176	-6891.7	2491.0
Model 4.2.3	0.228	0.126	0.301	0.166	-7892.1	4343.2

5 Conclusion

We have introduced a new class of cross-covariance matrix functions that are valid by construction and that capture both nonstationarity and asymmetry. Specifically, using p warping functions, each of which is constructed as a composition of elementary injective warping units, we model p -variate spatial processes that have nonstationary and asymmetric covariances on the geographic domain. These are modeled in terms of processes with stationary, symmetric, and possibly isotropic covariances on a warped domain. We also consider a special case where the same warping function is used for all p processes, resulting in a symmetric cross-covariance matrix function on the geographic domain. We show the benefit of using multivariate DCSMs over classical stationary models, such as the multivariate parsimonious Matérn model, through illustrations based on simulated data and real-world data.

There are a number of avenues that can be considered for future development of the proposed models. First, the models we propose do not consider nonstationarity in the variance parameters or in the cross-correlation parameters. Spatially varying variance parameters and cross-correlation parameters can be introduced, as in Kleiber and Nychka (2012) and Messick et al. (2017). Second, we use only the parsimonious multivariate Matérn model on the warped domain, but indeed any model could be used (e.g., one based on the cross-variogram). Third, we present data examples for two-dimensional space ($d = 2$), but our model can also be used in higher-dimensional space. The challenge is to find warping functions that are injective in a higher-dimensional domain. Furthermore, because the computational complexity of evaluating the likelihood function in the model is $O(N^3)$ for a data set of size N , the model needs to be modified when modeling large spatial data sets. Specifically, to deal with very large nonstationary spatial data sets, one would need to extend the model in such a way that it has a scalable structure on the deformed space. Several models that can deal with very large spatial data sets are summarized in Heaton et al. (2019). Finally, when using our multivariate spatial models, several decisions need to be made on the architecture (e.g., the number of layers of warping units, the order of the warping units, etc.), and further work needs to be carried out to determine how these decisions affect the predictive performance.

In conclusion, we show that multivariate DCSMs are easy to construct and fit from simple injective warping functions. We also show that they can provide superior predictive performance compared with that of conventional stationary models, particularly when data are missing over large regions.

Supplementary Material

Section S1 contains the proofs of Propositions 1–4, Theorem 1, and Corollary 1; Section S2 gives the log restricted likelihood and prediction formulas; Section S3 contains additional tables and figures; and Section S4 contains additional data illustrations.

Acknowledgements

Quan Vu was supported by a University Postgraduate Award from the University of Wollongong, Australia. Andrew Zammit-Mangion’s research was supported by an Australian Research Council Discovery Early Career Research Award (DECRA) DE180100203 and by Discovery Project DP190100180. Noel Cressie’s research was supported by Australian Research Council Discovery Projects DP150104576 and DP190100180, and by NSF grant SES-1132031 funded through the NSF-Census Research Network (NCRN) program. Cressie’s and Zammit-Mangion’s research was also supported by NASA ROSES grant 17-OCO2-17-0012. The authors are grateful to the Associate Editor and three referees whose suggestions led to a number of improvements in the manuscript. They would also like to thank Matt Moores for discussions related to this research.

Supplementary Material

S1 Proofs

S1.1 Proof of Proposition 1

Since $\mathbf{C}_D^o(\cdot)$ is symmetric, $C_{ij,D}^o(\cdot) = C_{ji,D}^o(\cdot)$, $i, j = 1, \dots, p$. Now consider any two locations $\mathbf{s}, \mathbf{u} \in G$. The covariance between $Y_i(\mathbf{s})$ and $Y_j(\mathbf{u})$ is given by

$$C_{ij,G}(\mathbf{s}, \mathbf{u}) = C_{ij,D}^o(\mathbf{f}(\mathbf{s}) - \mathbf{f}(\mathbf{u})) = C_{ji,D}^o(\mathbf{f}(\mathbf{s}) - \mathbf{f}(\mathbf{u})) = C_{ji,G}(\mathbf{s}, \mathbf{u}),$$

for $i, j = 1, \dots, p$, and therefore $\mathbf{C}_G(\cdot, \cdot)$ is symmetric.

S1.2 Proof of Proposition 2

Consider any two locations \mathbf{s}, \mathbf{u} in G . The covariance between $Y_i(\mathbf{s})$ and $Y_j(\mathbf{u})$, $i, j = 1, \dots, p$, is

$$\text{cov}(Y_i(\mathbf{s}), Y_j(\mathbf{u})) = C_{ij,G}(\mathbf{s}, \mathbf{u}) = C_{ij,D}^o(\mathbf{f}_i(\mathbf{s}) - \mathbf{f}_j(\mathbf{u})),$$

while the covariance between $Y_j(\mathbf{s})$ and $Y_i(\mathbf{u})$ is

$$\begin{aligned} \text{cov}(Y_j(\mathbf{s}), Y_i(\mathbf{u})) &= C_{ji,G}(\mathbf{s}, \mathbf{u}) = C_{ji,D}^o(\mathbf{f}_j(\mathbf{s}) - \mathbf{f}_i(\mathbf{u})) \\ &= C_{ij,D}^o(\mathbf{f}_j(\mathbf{s}) - \mathbf{f}_i(\mathbf{u})), \end{aligned}$$

since $\mathbf{C}_D^o(\cdot)$ is symmetric. Now, for $i \neq j$, there are many examples where $\mathbf{f}_i(\mathbf{s}) - \mathbf{f}_j(\mathbf{u}) \neq \mathbf{f}_j(\mathbf{s}) - \mathbf{f}_i(\mathbf{u})$, and hence $\text{cov}(Y_i(\mathbf{s}), Y_j(\mathbf{u})) \neq \text{cov}(Y_j(\mathbf{s}), Y_i(\mathbf{u}))$ for $i \neq j$. That is, the cross-covariance matrix function $\mathbf{C}_G(\cdot, \cdot)$ constructed through (3.2) is not necessarily symmetric.

S1.3 Proof of Proposition 3

By (3.2), we have that for $k = 1, \dots, n_i$, $l = 1, \dots, n_j$, and $i, j = 1, \dots, p$, $C_{ij,G}(\mathbf{s}_{ik}, \mathbf{s}_{jl}) = C_{ij,D}^o(\mathbf{f}_i(\mathbf{s}_{ik}) - \mathbf{f}_j(\mathbf{s}_{jl}))$. Therefore, for $i, j = 1, \dots, p$, we have that $\boldsymbol{\Sigma}_{ij,G} = (C_{ij,D}^o(\mathbf{f}_i(\mathbf{s}_{ik}) - \mathbf{f}_j(\mathbf{s}_{jl})) : k = 1, \dots, n_i, l = 1, \dots, n_j)$. Since $\mathbf{C}_D^o(\cdot)$ is valid, we have that for any $\mathbf{a} \in \mathbb{R}^N$, $\mathbf{a} \neq \mathbf{0}$, $\mathbf{a}'(C_{ij,D}^o(\mathbf{f}_i(\mathbf{s}_{ik}) - \mathbf{f}_j(\mathbf{s}_{jl})) : k = 1, \dots, n_i; l = 1, \dots, n_j; i, j = 1, \dots, p)\mathbf{a} \geq 0$, and hence $\mathbf{a}'\boldsymbol{\Sigma}_G\mathbf{a} \geq 0$. That is, $\boldsymbol{\Sigma}_G$ is nonnegative-definite.

S1.4 Proof of Proposition 4

Since $\mathbf{f}(\cdot)$ is the identity map,

$$\begin{aligned} C_{ij,G}(\mathbf{s}, \mathbf{u}) &= C_{ij,D}^o(\mathbf{g}_i(\mathbf{s}) - \mathbf{g}_j(\mathbf{u})); \quad i, j \neq k, \\ C_{kk,G}(\mathbf{s}, \mathbf{u}) &= C_{kk,D}^o(\mathbf{g}_k(\mathbf{s}) - \mathbf{g}_k(\mathbf{u})), \\ C_{ik,G}(\mathbf{s}, \mathbf{u}) &= C_{ik,D}^o(\mathbf{g}_i(\mathbf{s}) - \mathbf{g}_k(\mathbf{u})); \quad i \neq k, \end{aligned}$$

where $k \in \{2, \dots, p\}$. Write the affine transformation as $\mathbf{g}_k(\mathbf{s}) = \tilde{\mathbf{A}}\mathbf{s} + \tilde{\mathbf{d}}$, where $\tilde{\mathbf{A}}$ is a $d \times d$ matrix that is not generally equal to the identity matrix. Then,

$$\begin{aligned} C_{ij,G}(\mathbf{s}, \mathbf{u}) &= C_{ij,D}^o(\mathbf{s} - \mathbf{u}); \quad i, j \neq k, \\ C_{kk,G}(\mathbf{s}, \mathbf{u}) &= C_{kk,D}^o(\tilde{\mathbf{A}}\mathbf{s} - \tilde{\mathbf{A}}\mathbf{u}) = C_{kk,D}^o(\tilde{\mathbf{A}}(\mathbf{s} - \mathbf{u})), \\ C_{ik,G}(\mathbf{s}, \mathbf{u}) &= C_{ik,D}^o(\mathbf{s} - \tilde{\mathbf{A}}\mathbf{u} - \tilde{\mathbf{d}}); \quad i \neq k. \end{aligned}$$

As $\mathbf{C}_D^o(\cdot)$ is stationary, it follows that $C_{ij,G}(\mathbf{s}, \mathbf{u}), i, j \neq k$ and $C_{kk,G}(\mathbf{s}, \mathbf{u})$ are stationary, while $C_{ik,G}(\mathbf{s}, \mathbf{u}), i \neq k$, is nonstationary unless $\tilde{\mathbf{A}}$ is the identity matrix. That is, the cross-covariance function $C_{ik,G}(\cdot, \cdot)$ is not necessarily stationary.

S1.5 Proof of Theorem 1 and Corollary 1

Proof of Theorem 1:

By assumption, $\mathbf{C}_G^{(1)}(\mathbf{s}, \mathbf{u}) = \mathbf{C}_G^{(2)}(\mathbf{s}, \mathbf{u})$, for all $\mathbf{s}, \mathbf{u} \in G$, and hence by definition, $\tilde{C}_{ij,D}^o(\|\mathbf{f}^{(1)}(\mathbf{s}) - \mathbf{f}^{(1)}(\mathbf{u})\|; a_{ij}^{(1)}) = \tilde{C}_{ij,D}^o(\|\mathbf{f}^{(2)}(\mathbf{s}) - \mathbf{f}^{(2)}(\mathbf{u})\|; a_{ij}^{(2)})$, for all $\mathbf{s}, \mathbf{u} \in G$ and $i, j = 1, \dots, p$. That is, for each $\mathbf{s}, \mathbf{u} \in G$ and $i, j = 1, \dots, p$, we have that $a_{ij}^{(1)}\|\mathbf{f}^{(1)}(\mathbf{s}) - \mathbf{f}^{(1)}(\mathbf{u})\| = a_{ij}^{(2)}\|\mathbf{f}^{(2)}(\mathbf{s}) - \mathbf{f}^{(2)}(\mathbf{u})\|$, and hence $a_{ij}^{(1)}\|\mathbf{f}^{(1)}(\mathbf{s}_l) - \mathbf{f}^{(1)}(\mathbf{s}_k)\| = a_{ij}^{(2)}\|\mathbf{f}^{(2)}(\mathbf{s}_l) - \mathbf{f}^{(2)}(\mathbf{s}_k)\|$, for \mathbf{s}_l and \mathbf{s}_k two distinct points in G .

Hence, from the definition of $\mathbf{b}_1(\cdot)$ in Section 3.3,

$$\left\| \mathbf{b}_1 \circ \mathbf{f}^{(1)}(\mathbf{s}) - \mathbf{b}_1 \circ \mathbf{f}^{(1)}(\mathbf{u}) \right\| = \left\| \mathbf{b}_1 \circ \mathbf{f}^{(2)}(\mathbf{s}) - \mathbf{b}_1 \circ \mathbf{f}^{(2)}(\mathbf{u}) \right\|,$$

for all $\mathbf{s}, \mathbf{u} \in G$. Since $\mathbf{b}_2(\cdot)$ and $\mathbf{b}_3(\cdot)$ are distance-preserving transformations, we then have that

$$\left\| \mathbf{b}_0 \circ \mathbf{f}^{(1)}(\mathbf{s}) - \mathbf{b}_0 \circ \mathbf{f}^{(1)}(\mathbf{u}) \right\| = \left\| \mathbf{b}_0 \circ \mathbf{f}^{(2)}(\mathbf{s}) - \mathbf{b}_0 \circ \mathbf{f}^{(2)}(\mathbf{u}) \right\|.$$

Recall that three locations $\mathbf{s}_k, \mathbf{s}_l$, and \mathbf{s}_m are chosen in $G \subset \mathbb{R}^2$ such that $\mathbf{f}^{(r)}(\mathbf{s}_k), \mathbf{f}^{(r)}(\mathbf{s}_l)$, and $\mathbf{f}^{(r)}(\mathbf{s}_m)$ are not colinear; and $\mathbf{b}_0 \circ \mathbf{f}^{(r)}(\mathbf{s}_k) = (0, 0)'$, $\mathbf{b}_0 \circ \mathbf{f}^{(r)}(\mathbf{s}_l) = (1, 0)'$, and $b_{0,2} \circ \mathbf{f}^{(r)}(\mathbf{s}_m) > 0$, for $r = 1, 2$. Now, we have that for any two points $\mathbf{s}, \mathbf{u} \in G$, the distance $\left\| \mathbf{b}_0 \circ \mathbf{f}^{(r)}(\mathbf{s}) - \mathbf{b}_0 \circ \mathbf{f}^{(r)}(\mathbf{u}) \right\|$ does not depend on r . Because the two points $\mathbf{b}_0 \circ \mathbf{f}^{(r)}(\mathbf{s}_k)$ and $\mathbf{b}_0 \circ \mathbf{f}^{(r)}(\mathbf{s}_l)$ are fixed in D_3 , it follows that the distances $\left\| \mathbf{b}_0 \circ \mathbf{f}^{(r)}(\mathbf{s}_m) - \mathbf{b}_0 \circ \mathbf{f}^{(r)}(\mathbf{s}_k) \right\|$ and $\left\| \mathbf{b}_0 \circ \mathbf{f}^{(r)}(\mathbf{s}_m) - \mathbf{b}_0 \circ \mathbf{f}^{(r)}(\mathbf{s}_l) \right\|$ are also fixed. In \mathbb{R}^2 , when the distances from a point $\mathbf{b}_0 \circ \mathbf{f}^{(r)}(\mathbf{s}_m)$ to two fixed points $(0, 0)'$ and $(1, 0)'$ are constant, there exist two possible points \mathbf{s}_m , assuming $\mathbf{f}^{(r)}(\mathbf{s}_m)$ is not colinear with $\mathbf{f}^{(r)}(\mathbf{s}_k)$ and $\mathbf{f}^{(r)}(\mathbf{s}_l)$: one where $b_{0,2} \circ \mathbf{f}^{(r)}(\mathbf{s}_m) > 0$ and one where $b_{0,2} \circ \mathbf{f}^{(r)}(\mathbf{s}_m) < 0$. However, as we constrain $b_{0,2} \circ \mathbf{f}^{(r)}(\mathbf{s}_m) > 0$, $\mathbf{b}_0 \circ \mathbf{f}^{(r)}(\mathbf{s}_m)$ is unique. Since we have three fixed points $\mathbf{b}_0 \circ \mathbf{f}^{(r)}(\mathbf{s}_k), \mathbf{b}_0 \circ \mathbf{f}^{(r)}(\mathbf{s}_l)$, and $\mathbf{b}_0 \circ \mathbf{f}^{(r)}(\mathbf{s}_m)$, and a fixed set of distances $\left\| \mathbf{b}_0 \circ \mathbf{f}^{(r)}(\mathbf{s}) - \mathbf{b}_0 \circ \mathbf{f}^{(r)}(\mathbf{u}) \right\|$ for any two points \mathbf{s}, \mathbf{u} , then $\mathbf{b}_0 \circ \mathbf{f}^{(1)}(\mathbf{s}) = \mathbf{b}_0 \circ \mathbf{f}^{(2)}(\mathbf{s})$, for all $\mathbf{s} \in G$.

For the converse part of the proof, assume that $\mathbf{b}_0 \circ \mathbf{f}^{(1)}(\mathbf{s}) = \mathbf{b}_0 \circ \mathbf{f}^{(2)}(\mathbf{s})$, for all $\mathbf{s} \in G$. Then,

$$\left\| \mathbf{b}_1 \circ \mathbf{f}^{(1)}(\mathbf{s}) - \mathbf{b}_1 \circ \mathbf{f}^{(1)}(\mathbf{u}) \right\| = \left\| \mathbf{b}_1 \circ \mathbf{f}^{(2)}(\mathbf{s}) - \mathbf{b}_1 \circ \mathbf{f}^{(2)}(\mathbf{u}) \right\|,$$

for all $\mathbf{s}, \mathbf{u} \in G$, and therefore

$$\frac{\|\mathbf{f}^{(1)}(\mathbf{s}) - \mathbf{f}^{(1)}(\mathbf{u})\|}{\|\mathbf{f}^{(1)}(\mathbf{s}_l) - \mathbf{f}^{(1)}(\mathbf{s}_k)\|} = \frac{\|\mathbf{f}^{(2)}(\mathbf{s}) - \mathbf{f}^{(2)}(\mathbf{u})\|}{\|\mathbf{f}^{(2)}(\mathbf{s}_l) - \mathbf{f}^{(2)}(\mathbf{s}_k)\|}.$$

Because $a_{ij}^{(1)}\|\mathbf{f}^{(1)}(\mathbf{s}_l) - \mathbf{f}^{(1)}(\mathbf{s}_k)\| = a_{ij}^{(2)}\|\mathbf{f}^{(2)}(\mathbf{s}_l) - \mathbf{f}^{(2)}(\mathbf{s}_k)\|$, for $i, j = 1, \dots, p$, it follows that $a_{ij}^{(1)}\|\mathbf{f}^{(1)}(\mathbf{s}) - \mathbf{f}^{(1)}(\mathbf{u})\| = a_{ij}^{(2)}\|\mathbf{f}^{(2)}(\mathbf{s}) - \mathbf{f}^{(2)}(\mathbf{u})\|$, for all $\mathbf{s}, \mathbf{u} \in G$ and $i, j = 1, \dots, p$. Therefore, $\mathbf{C}_G^{(1)}(\mathbf{s}, \mathbf{u}) = \mathbf{C}_G^{(2)}(\mathbf{s}, \mathbf{u})$, for all $\mathbf{s}, \mathbf{u} \in G$.

Proof of Corollary 1:

Note that the Gaussian process model (3.5) is fully specified by its mean function and covariance function. Hence, its finite-dimensional distributions are solely a function of the mean and covariance-function parameters. Therefore, the log restricted likelihood function in (S2.1) where the mean-function parameters are profiled out, solely depends on covariance-function parameters and the data \mathbf{Z} . Now, suppose that two different sets each comprising a warping function and scale parameters, $\{\mathbf{f}^{(1)}(\cdot), \{a_{ij}^{(1)}\}\}$ and $\{\mathbf{f}^{(2)}(\cdot), \{a_{ij}^{(2)}\}\}$, yield the same log restricted likelihood function, for any set of measurement locations $\{\mathbf{s}_{ik} : k = 1, \dots, n_i; i = 1, \dots, p\} \subset G$. Then, this necessarily means that $\mathbf{C}_G^{(1)}(\mathbf{s}, \mathbf{u}) = \mathbf{C}_G^{(2)}(\mathbf{s}, \mathbf{u})$, for all $\mathbf{s}, \mathbf{u} \in G$ and, from the proof of Theorem 1 above, we see that this implies that $\mathbf{b}_0 \circ \mathbf{f}^{(1)}(\mathbf{s}) = \mathbf{b}_0 \circ \mathbf{f}^{(2)}(\mathbf{s})$, for all $\mathbf{s} \in G$, and that $a_{ij}^{(1)}\|\mathbf{f}^{(1)}(\mathbf{s}_l) - \mathbf{f}^{(1)}(\mathbf{s}_k)\| = a_{ij}^{(2)}\|\mathbf{f}^{(2)}(\mathbf{s}_l) - \mathbf{f}^{(2)}(\mathbf{s}_k)\|$, for $i, j = 1, \dots, p$. This in turn implies that the sets comprising the corresponding homogenized warping functions and transformed scale parameters, $\{\mathbf{b}_0 \circ \mathbf{f}^{(1)}(\cdot), \{\tilde{a}_{ij}^{(1)}\}\}$ and $\{\mathbf{b}_0 \circ \mathbf{f}^{(2)}(\cdot), \{\tilde{a}_{ij}^{(2)}\}\}$, where $\tilde{a}_{ij}^{(r)} = a_{ij}^{(r)}\|\mathbf{f}^{(r)}(\mathbf{s}_l) - \mathbf{f}^{(r)}(\mathbf{s}_k)\|$, $r = 1, 2$, are identical. Further, since Theorem 1 establishes an if-and-only-if relation, there is no other set, $\{\mathbf{b}_0 \circ \mathbf{f}^{(3)}(\cdot), \{\tilde{a}_{ij}^{(3)}\}\}$ say, that yields the same log restricted likelihood function, for any set of measurement locations $\{\mathbf{s}_{ik} : k = 1, \dots, n_i; i = 1, \dots, p\} \subset G$. This is because such a set would correspond to a different covariance function, $\mathbf{C}_G^{(3)}(\mathbf{s}, \mathbf{u})$ say, where $\mathbf{C}_G^{(3)}(\mathbf{s}, \mathbf{u}) \neq \mathbf{C}_G^{(1)}(\mathbf{s}, \mathbf{u})$ for some $\mathbf{s}, \mathbf{u} \in G$. Therefore, the set comprising a homogenized warping function and the transformed scale parameters, $\{\mathbf{b}_0 \circ \mathbf{f}(\cdot), \{\tilde{a}_{ij}\}\}$, where $\tilde{a}_{ij} \equiv a_{ij}\|\mathbf{f}(\mathbf{s}_l) - \mathbf{f}(\mathbf{s}_k)\|$ for $i, j = 1, \dots, p$, has a unique log restricted likelihood function associated with it; this set is thus identifiable (see Kadane, 1974, for more discussion on identifiability).

S2 Log Restricted Likelihood and Prediction Formulas

The log restricted likelihood for our model under the assumption of Gaussianity for \mathbf{Y} can be written as (Cressie and Lahiri, 1993),

$$\mathcal{L}(\boldsymbol{\theta}; \mathbf{Z}) = -\frac{N - pq}{2} \log(2\pi) + \frac{1}{2} \log |\mathbf{X}'\mathbf{X}| - \frac{1}{2} \log |\boldsymbol{\Sigma}_Z| - \frac{1}{2} \log |\mathbf{X}'\boldsymbol{\Sigma}_Z^{-1}\mathbf{X}| - \frac{1}{2} \mathbf{Z}'\boldsymbol{\Pi}\mathbf{Z}, \quad (\text{S2.1})$$

where $N = \sum_{i=1}^p n_i$, and $\boldsymbol{\Pi} = \boldsymbol{\Sigma}_Z^{-1} - \boldsymbol{\Sigma}_Z^{-1}\mathbf{X}(\mathbf{X}'\boldsymbol{\Sigma}_Z^{-1}\mathbf{X})^{-1}\mathbf{X}'\boldsymbol{\Sigma}_Z^{-1}$.

The estimate of $\boldsymbol{\beta}$, $\hat{\boldsymbol{\beta}}$, is given by

$$\hat{\boldsymbol{\beta}} = (\mathbf{X}'\hat{\boldsymbol{\Sigma}}_Z^{-1}\mathbf{X})^{-1}\mathbf{X}'\hat{\boldsymbol{\Sigma}}_Z^{-1}\mathbf{Z}, \quad (\text{S2.2})$$

where $\hat{\Sigma}_Z$ denotes Σ_Z evaluated at $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}$.

Treating the plug-in REML estimates as known parameters, the joint distribution of the data \mathbf{Z} and the process $\tilde{Y}_i(\cdot)$ evaluated at \mathbf{s}^* is

$$\begin{pmatrix} \mathbf{Z} \\ \tilde{Y}_i(\mathbf{s}^*) \end{pmatrix} \sim \text{Gau} \left(\begin{pmatrix} \mathbf{X} \\ \mathbf{x}_i^{*'} \end{pmatrix} \boldsymbol{\beta}, \begin{pmatrix} \Sigma_Z & \boldsymbol{\sigma}^*(\mathbf{s}^*) \\ \boldsymbol{\sigma}^{*'}(\mathbf{s}^*) & C_{ii,G}(\mathbf{s}^*, \mathbf{s}^*) \end{pmatrix} \right), \quad (\text{S2.3})$$

where $\mathbf{x}_i^* = (\mathbf{x}(\mathbf{s}^*)' I(j = i) : j = 1, \dots, p)'$, and $\boldsymbol{\sigma}^*(\mathbf{s}^*) = (C_{1i,G}(\mathbf{s}_{11}, \mathbf{s}^*), \dots, C_{1i,G}(\mathbf{s}_{1n_1}, \mathbf{s}^*), \dots, C_{pi,G}(\mathbf{s}_{p1}, \mathbf{s}^*), \dots, C_{pi,G}(\mathbf{s}_{pn_p}, \mathbf{s}^*))'$. From (S2.3), Gaussian conditioning yields

$$\begin{aligned} \text{E}(\tilde{Y}_i(\mathbf{s}^*) | \mathbf{Z}) &= \mathbf{x}_i^{*'} \boldsymbol{\beta} + \boldsymbol{\sigma}^{*'}(\mathbf{s}^*) \Sigma_Z^{-1} (\mathbf{Z} - \mathbf{X} \boldsymbol{\beta}), \\ \text{var}(\tilde{Y}_i(\mathbf{s}^*) | \mathbf{Z}) &= C_{ii,G}(\mathbf{s}^*, \mathbf{s}^*) - \boldsymbol{\sigma}^{*'}(\mathbf{s}^*) \Sigma_Z^{-1} \boldsymbol{\sigma}^*(\mathbf{s}^*). \end{aligned} \quad (\text{S2.4})$$

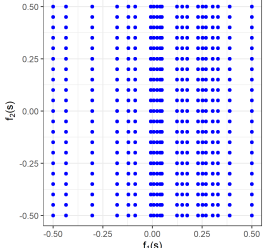
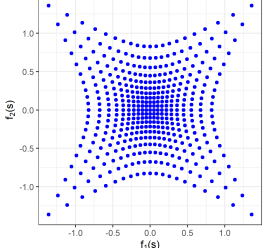
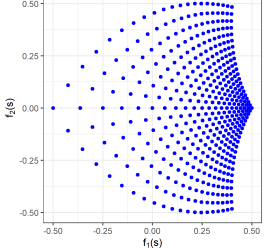
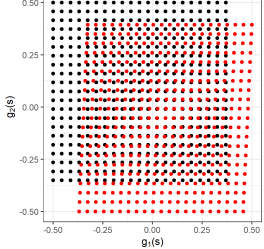
From (S2.4), it is also straightforward to make a probabilistic prediction of an observation at \mathbf{s}^* , say Z_i^* , since $\text{E}(Z_i^* | \mathbf{Z}) = \text{E}(\tilde{Y}_i(\mathbf{s}^*) | \mathbf{Z})$, and $\text{var}(Z_i^* | \mathbf{Z}) = \text{var}(\tilde{Y}_i(\mathbf{s}^*) | \mathbf{Z}) + \tau_i^2$.

S3 Additional Tables and Figures

Algorithm 1: Parameter bootstrapping for uncertainty quantification

- 1 Find the REML estimate $\hat{\boldsymbol{\theta}}$ by maximizing \mathcal{L} in (S2.1).
 - 2 Find the REML estimate $\hat{\boldsymbol{\beta}}$ using (S2.2).
 - 3 Find the Cholesky factor \mathbf{L} of $\hat{\Sigma}_Z$, that is, find \mathbf{L} such that $\mathbf{L}\mathbf{L}' = \hat{\Sigma}_Z$.
 - 4 Decorrelate the data: $\mathbf{Z}_0 = \mathbf{L}^{-1}(\mathbf{Z} - \mathbf{X}\hat{\boldsymbol{\beta}})$.
 - 5 Generate an uncorrelated bootstrap sample $\mathbf{Z}_{0,b}$ by sampling \mathbf{Z}_0 with replacement.
 - 6 Create the correlated bootstrap sample $\mathbf{Z}_b = \mathbf{L}\mathbf{Z}_{0,b} + \mathbf{X}\hat{\boldsymbol{\beta}}$.
 - 7 Find the bootstrap estimate $\hat{\boldsymbol{\theta}}_b$ and $\hat{\boldsymbol{\beta}}_b$ from \mathbf{Z}_b .
 - 8 Repeat from step 5 for B times to create a bootstrap sample of size B . (For most problems, $B \approx 1000$ should suffice.)
-

Table S1: Summary of the warping units. In each row, a warping unit is described. Parameters appearing in the functions composing the units are denoted using θ .

Type of deformation function	Functional form	Usage	Visualization
Axial warping	$\tilde{\mathbf{f}}(\mathbf{s}) = \begin{pmatrix} \tilde{f}(s_1) \\ s_2 \end{pmatrix} \text{ or } \begin{pmatrix} s_1 \\ \tilde{f}(s_2) \end{pmatrix},$ <p>where $\tilde{f}(s) = \sum_{i=1}^r w_i \phi_i(s)$, and where $\phi_1(s) = s$; $\phi_i(s) = \frac{1}{1 + \exp\{-\theta_1(s - \theta_2)\}}$, $i = 2, \dots, r$.</p>	warp space along one of the axes	
Radial basis function	$\tilde{\mathbf{f}}(\mathbf{s}) = \mathbf{s} + w \exp\{-\theta_2 \ \mathbf{s} - \theta_1\ ^2\}(\mathbf{s} - \theta_1)$	expand or contract space locally around the center of the basis function	
Möbius transformation	$\tilde{\mathbf{f}}(\mathbf{s}) = \begin{pmatrix} \text{Re}(\phi(\mathbf{s})) \\ \text{Im}(\phi(\mathbf{s})) \end{pmatrix}, \text{ where}$ $\phi(\mathbf{s}) = \frac{\theta_1 z(\mathbf{s}) + \theta_2}{\theta_3 z(\mathbf{s}) + \theta_4};$ $z(\mathbf{s}) = s_1 + i s_2; \theta_1, \theta_2, \theta_3, \theta_4 \in \mathbb{C}$	move points around fixed points (usually in circular paths)	
Affine transformation	$\tilde{\mathbf{g}}(\mathbf{s}) = \mathbf{A}\mathbf{s} + \mathbf{d}$	align processes with respect to the first process (using shifts and rotations)	




Figure S1: Illustration of the homogenizing function $\mathbf{b}_0(\cdot)$. Top left: A set of locations on the warped domain D , with the red, yellow, and green dots denoting $\mathbf{f}(\mathbf{s}_k)$, $\mathbf{f}(\mathbf{s}_l)$, and $\mathbf{f}(\mathbf{s}_m)$, respectively. Top right: The locations on the scaled and shifted domain D_1 . Bottom left: The locations on the scaled, shifted, and rotated domain D_2 . Bottom right: The locations on the scaled, shifted, rotated, and reflected, domain D_3 . This is the fixed frame of reference defined in Section 3.3.




Figure S2: Boxplots of RMSPE and CRPS for both models across 30 simulations in Section 4.1. The left two panels correspond to $\tilde{Y}_1(\cdot)$, and the right two panels to $\tilde{Y}_2(\cdot)$.

Table S2: True values, estimates, and 95% bootstrap confidence intervals (CI), of the model parameters for the simulation study in Section 4.1, where the measurement locations are randomly sampled from G .

Parameters	True values	REML estimates	95% bootstrap CI
ν_{11}	0.500	0.528	(0.235, 0.634)
ν_{22}	1.500	1.226	(0.929, 1.664)
σ_1	1.000	1.040	(0.775, 1.329)
σ_2	0.900	0.932	(0.763, 1.388)
ρ_{12}	0.450	0.392	(0.321, 0.537)
$1/\tilde{a}$	0.329	0.412	(0.274, 0.556)
τ_1	0.200	0.252	(0.178, 0.265)
τ_2	0.100	0.097	(0.092, 0.108)
β_{11}	0.000	-0.232	(-0.947, 0.415)
β_{21}	0.000	0.047	(-0.762, 0.741)

Table S3: Five-fold cross-validation results, AIC, and the time required to fit, for the ocean temperatures at depths 0.5 meters and 318.1 meters for the study in Section 4.2, where data were missing at random.

	$T_{0.5}$		$T_{318.1}$		AIC	Time (s)
	RMSPE	CRPS	RMSPE	CRPS		
Model 4.2.1	0.0661	0.0306	0.0265	0.0129	-4791.1	1338.8
Model 4.2.2	0.0584	0.0188	0.0280	0.0136	-6246.4	2545.8
Model 4.2.3	0.0666	0.0198	0.0241	0.0123	-6403.4	4455.2




Figure S3: Measurement locations under the estimated aligning function (left panel) and the homogenized warping function (right panel) for the ocean-temperatures data set in Section 4.2, where data are missing in a block.

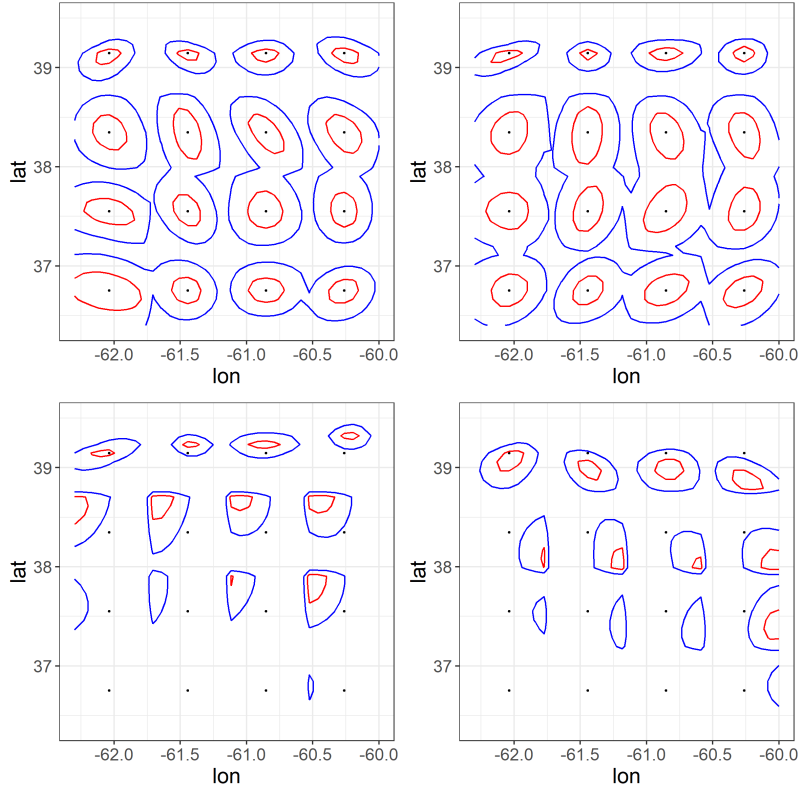


Figure S4: Contours illustrating the estimated covariance functions evaluated at a set of locations (black dots) for the data set in Section 4.2. Blue and red contours denote covariances equal to 0.4 and 0.8 times of $\hat{\sigma}_{ij}^2$, where $\hat{\sigma}_{ij}^2 = \hat{\sigma}_i^2$, for $i = j$, and $\hat{\sigma}_{ij}^2 = \hat{\rho}_{ij}\hat{\sigma}_i\hat{\sigma}_j$, for $i \neq j$, respectively. Top row: Marginal covariances of the first process, $C_{11,G}(\cdot, \cdot)$ (left) and second process, $C_{22,G}(\cdot, \cdot)$ (right). Bottom row: Cross-covariances of the first process with the second process, $C_{12,G}(\cdot, \cdot)$ (left) and of the second process with the first process, $C_{21,G}(\cdot, \cdot)$ (right). We see the estimated cross-covariances are asymmetric.

Table S4: Average hold-out validation results, AIC, and the time required to fit, for the simulation study in Section S4.1, where the measurement locations are randomly sampled 30 times from $G \setminus G_0$.

	$\tilde{Y}_1(\cdot)$		$\tilde{Y}_2(\cdot)$		AIC	Time (s)
	RMSPE	CRPS	RMSPE	CRPS		
Model 4.1.1	1.137	0.675	0.453	0.274	833.2	816.9
Model 4.1.2	0.713	0.424	0.261	0.167	310.4	1533.0

S4 Additional Data Illustrations

S4.1 Simulated Symmetric Nonstationary Data With A Missing Block

Following the simulation study in Section 4.1, we considered the case where the symmetric nonstationary data are missing in a block. This situation occurs often when observing environmental variables (for example, clouds could prevent a remote sensing instrument from collecting data over a large region). As in the study in Section 4.1, we sampled 1000 measurement locations at random 30 times, but this time on $G \setminus G_0$, where $G_0 \equiv [-0.28, -0.08] \times [-0.48, -0.28]$ (i.e., the block of data was omitted for both processes). Model 4.1.1 and Model 4.1.2 were then fitted to the data. Figure S5 shows the true simulated fields, the predictions, and the prediction standard errors from the two models. From Figure S5, we see that the predictions from the DCSM recover the salient features in the true fields despite the relatively large gap. The DCSM also produces relatively lower prediction standard errors in the unobserved region than the stationary parsimonious Matérn model. Table S4 shows the RMSPE and CRPS from the two models when predicting the latent process at the grid locations in G_0 and, again, it illustrates the improvement in RMSPE and CRPS that can be achieved when accounting for complex nonstationary properties of the process, even when the data have large gaps.

S4.2 Simulated Asymmetric Nonstationary Data

We now demonstrate inference for a bivariate DCSM from simulated asymmetric nonstationary data. We simulated bivariate data from a bivariate Gaussian DCSM $\tilde{Y}(\cdot)$ with constant mean (i.e., $q = 1$ and $\mathbf{x}(\cdot) = x(\cdot) = 1$) in (3.5), and cross-covariance matrix function as in (3.3). The data were simulated on the 101×101 grid defined on the geographic domain G . Now the *shared warping function* $\mathbf{f}(\cdot)$ is a composition of axial warping units, a single resolution radial basis function unit, and a Möbius transformation unit, while the *aligning function* $\mathbf{g}_2(\cdot)$ is an affine transformation. (Recall from Section 3.1 that we fix $\mathbf{g}_1(\cdot)$ to be the identity map.) On the warped domain, we use a stationary bivariate parsimonious Matérn model. As in Section 4.1, we randomly sampled without replacement 1000 locations from the grid and used them as measurement locations.

We compared the predictive performance of the following five models on the 101×101 grid on G .

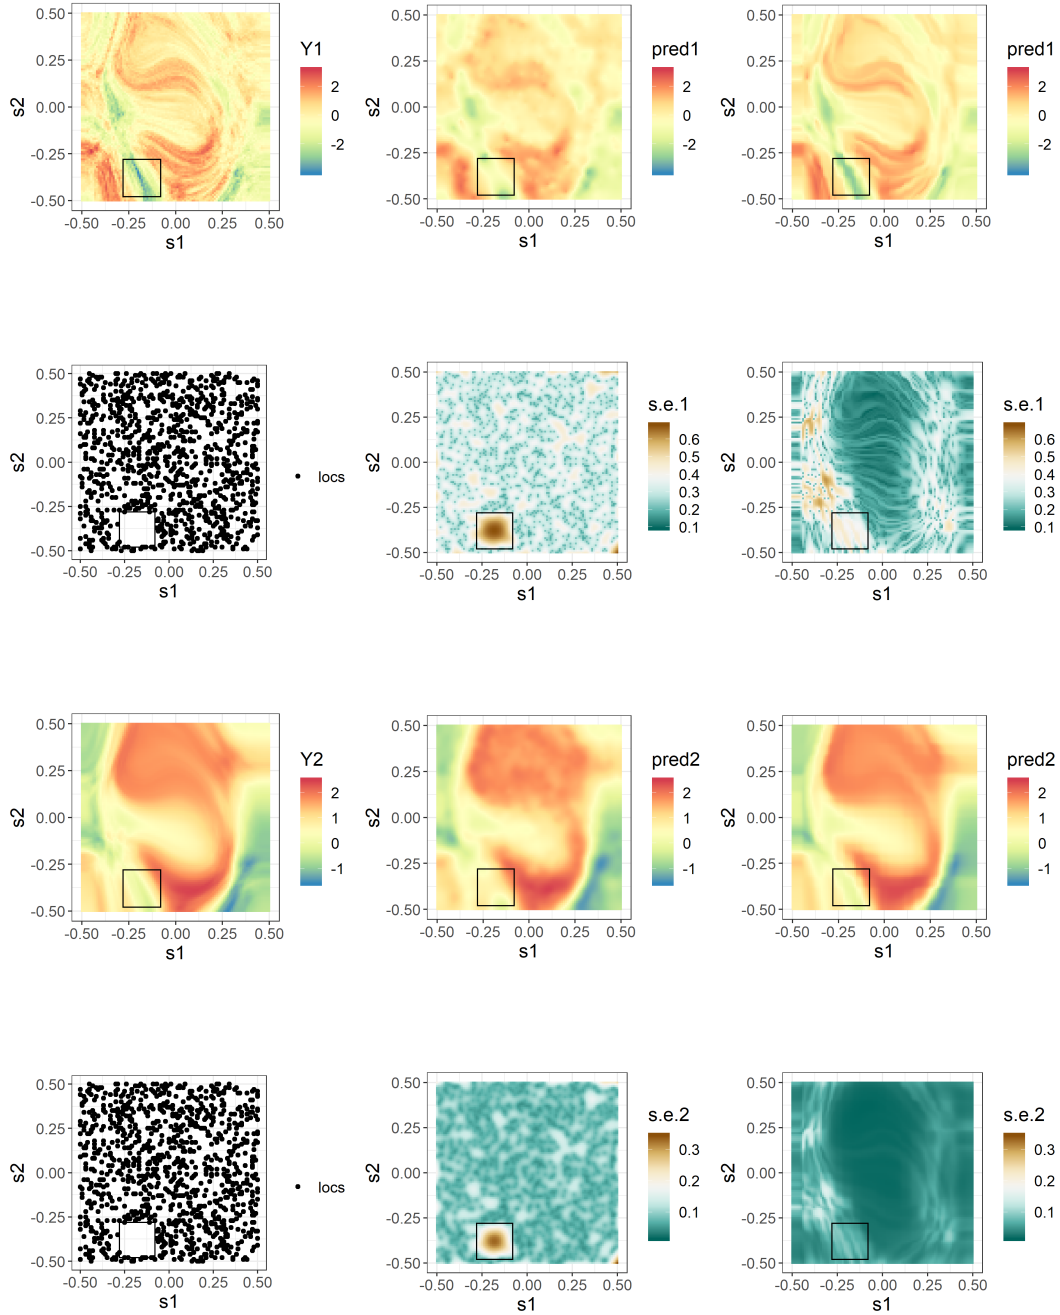


Figure S5: Same as Figure 1, but where the measurement locations are randomly sampled from $G \setminus G_0$, and where the hold-out region G_0 is enclosed by the black square.

- Model S4.2.1: A bivariate, stationary, symmetric, parsimonious Matérn model.
- Model S4.2.2: A bivariate, marginally stationary, asymmetric model with $\mathbf{f}(\cdot)$ the identity map and the aligning function $\mathbf{g}_2(\cdot)$ an affine transformation (as described in Proposition 4), with Model S4.2.1 on the warped domain.
- Model S4.2.3: A univariate DCSM for each of the processes, with the warping function $\mathbf{f}(\cdot)$ a composition of axial warping units, a single-resolution radial basis function, and a Möbius transformation unit, with the Matérn covariance model on the warped domain.
- Model S4.2.4: A bivariate symmetric DCSM, with the warping function as in Model S4.2.3, with Model S4.2.1 on the warped domain.
- Model S4.2.5: A bivariate asymmetric DCSM, with the aligning function $\mathbf{g}_2(\cdot)$ as in Model S4.2.2, the warping function as in Model S4.2.3, and Model S4.2.1 on the warped domain. This is the model from which the data were simulated.

Figure S6 shows the true simulated fields and the predicted fields from Model S4.2.2, Model S4.2.4, and Model S4.2.5. From Figure S6, we can see that Model S4.2.2 smooths out certain features (similar to the symmetric case), while Model S4.2.5 is able to reproduce sharper features than Model S4.2.3, illustrating that both nonstationarity and asymmetry could be important when modeling multivariate spatial processes. Figure S6 also shows the prediction standard errors for Model S4.2.2, Model S4.2.4, and Model S4.2.5. As in Section 4.1, we see that while there is no pattern in the prediction-standard-error map for Model S4.2.2, the DCSMs produce prediction standard errors that are reflective of the processes' local anisotropies and scales.

We can also compare, respectively, the estimated aligning function and the estimated warping function in Model S4.2.5 to the true aligning function and the true warping function in Model S4.2.5. Figure S7 shows the measurement locations under the true aligning and warping functions, and the measurement locations under the estimated aligning and warping functions. We see that the estimated aligning function generates a shift towards the east direction, which is similar to the true aligning function, which generates a shift towards the southeast direction. One can interpret that the second process needs to be shifted eastwards to align with the first process. The estimated warping function also retains important features of the true warping function, such as the contraction in the middle part of the domain. This can be interpreted that on the original domain G , the scale parameter in the middle region is smaller than the scale parameter in the boundary region.

In a similar manner to Section 4.1, we repeated the procedure of randomly sampling 1000 locations 30 times from G and computing predictive diagnostics. Table S5 displays the results from the cross-validation study, and Figure S8 shows the boxplots of the RMSPE and CRPS for the models across the 30 simulations. Model S4.2.5, which considers both nonstationarity and asymmetry, produces the best predictions in terms

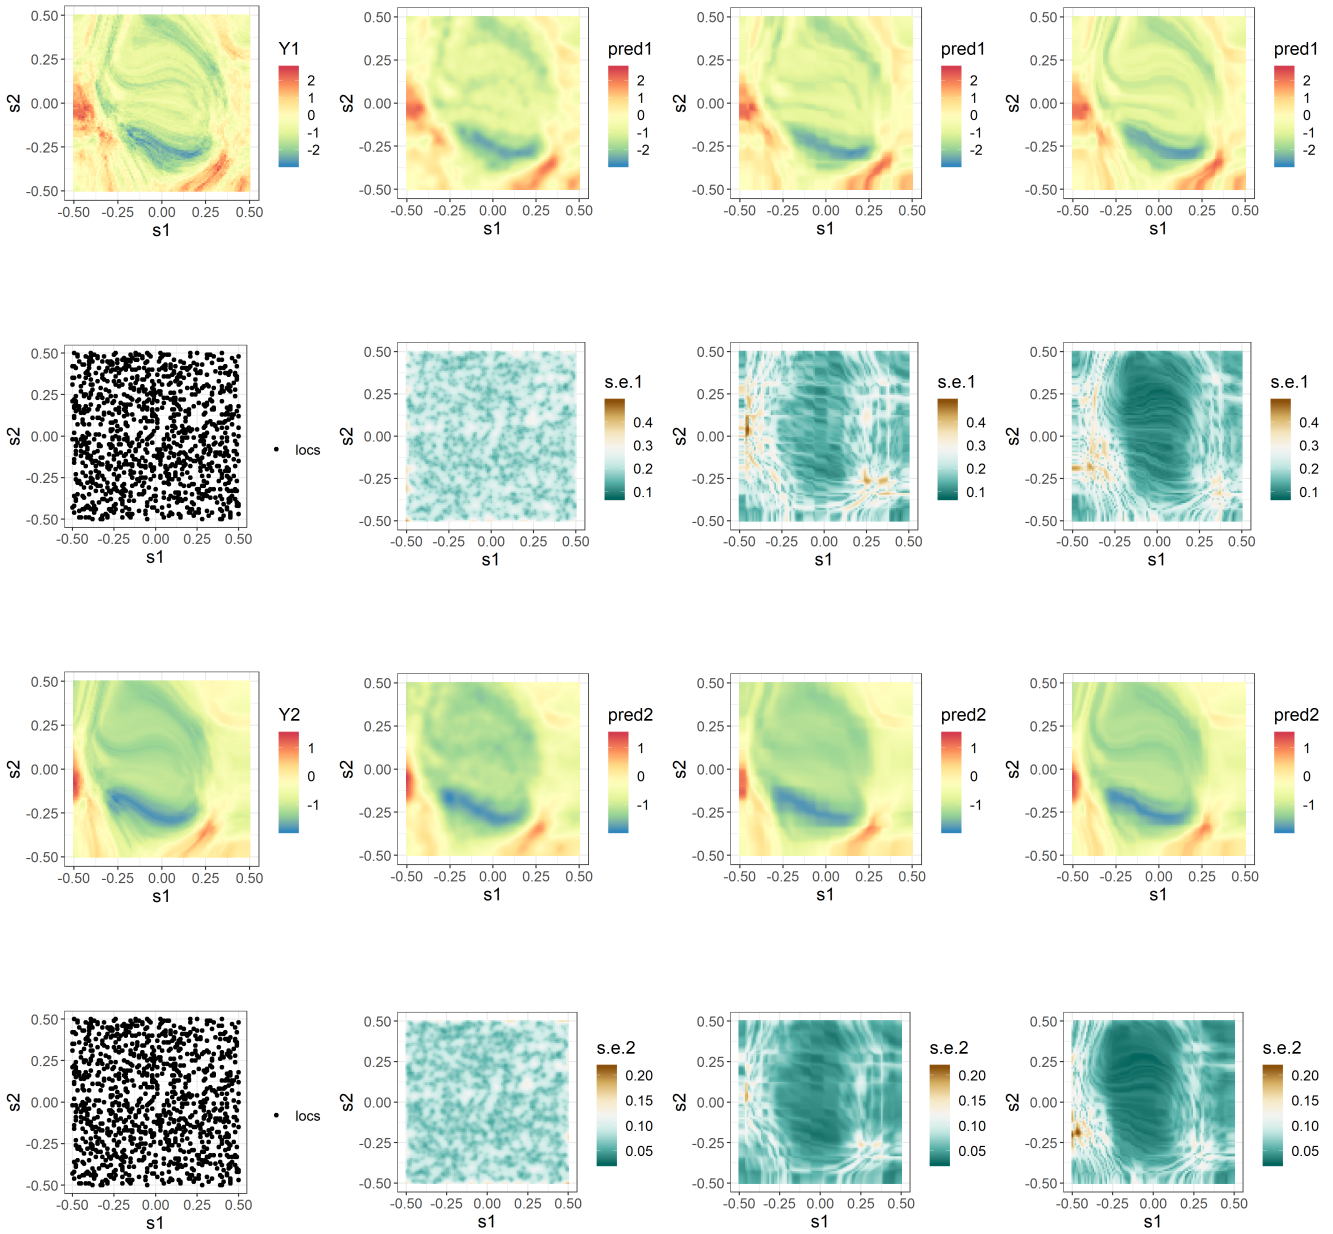


Figure S6: Comparison of predictions and prediction standard errors when using an asymmetric, stationary, parsimonious Matérn model (Model S4.2.2); a symmetric DCSM (Model S4.2.4); and an asymmetric DCSM (Model S4.2.5) in the study of Section S4.2, where measurement locations were randomly sampled without replacement from G . First row: The process $\tilde{Y}_1(\cdot)$ (first panel); predictions obtained using Model S4.2.2 (second panel), Model S4.2.4 (third panel) and Model S4.2.5 (fourth panel). Second row: Locations of the measurement of $\tilde{Y}_1(\cdot)$ (first panel); prediction standard errors obtained when using Model S4.2.2 (second panel), Model S4.2.4 (third panel) and Model S4.2.5 (fourth panel). Third and fourth rows: Analogous to the first and second rows, respectively, for the process $\tilde{Y}_2(\cdot)$.




Figure S7: Measurement locations under the aligning functions and the homogenized warping functions for the data set in Section S4.2. Top row: Measurement locations under the true aligning function (left panel) and the true warping function (right panel). Bottom row: Measurement locations under the estimated aligning function (left panel) and the estimated warping function (right panel) using Model S4.2.5.

of lowest RMSPE and CRPS, while the symmetric nonstationary model (Model S4.2.4) as well as the asymmetric stationary model (Model S4.2.2) are seen to yield slight improvements over the conventional symmetric, stationary model (Model S4.2.1). Surprisingly, even when accounting for nonstationarity in each process, the decoupled univariate model (Model S4.2.3) yields the worst predictions, showing the importance of a multivariate model.

S4.3 Simulated Data with Misspecified Warping Functions

In this section, we consider cases where the deformations are misspecified. Specifically, we present two cases: First, we consider data simulated from a bivariate, symmetric, stationary covariance function. Second, we consider data simulated from a bivariate, asymmetric, nonstationary covariance function with a misspecified warping function.

S4.3.1 Simulated Symmetric Stationary Data

In this section, we demonstrate the use of a bivariate DCSM with simulated symmetric stationary data. We simulated bivariate data from a bivariate, stationary, symmetric, parsimonious Matérn model. The

Table S5: Average hold-out validation results, AIC, and the time required to fit, for the simulation study in Section S4.2, where the measurement locations are randomly sampled 30 times from G .

	$\tilde{Y}_1(\cdot)$		$\tilde{Y}_2(\cdot)$		AIC	Time (s)
	RMSPE	CRPS	RMSPE	CRPS		
Model S4.2.1	0.304	0.169	0.091	0.048	386.3	815.3
Model S4.2.2	0.291	0.162	0.086	0.046	195.4	1188.6
Model S4.2.3	0.314	0.177	0.088	0.046	44.2	2025.7
Model S4.2.4	0.287	0.159	0.087	0.045	-3.8	1530.5
Model S4.2.5	0.269	0.149	0.080	0.041	-242.5	2654.1




Figure S8: Boxplots of RMSPE and CRPS for the five models across 30 simulations in Section S4.2. The top row corresponds to $\tilde{Y}_1(\cdot)$, and the bottom row to $\tilde{Y}_2(\cdot)$.

data were simulated on a 101×101 grid defined on the geographic domain G . We randomly sampled without replacement 1000 locations from the grid and used them as measurement locations.

We first examined the ability of the DCSM to retrieve the true warping function (in this case, the identity function). We used the same architecture in the DCSM as in Model S4.2.5 in Section S4.2. Figure S9 shows the measurement locations under the true aligning and warping function, and the measurement locations under the estimated aligning and warping function. We observe that both the aligning function and the warping function have been correctly estimated to be approximately the identity functions.




Figure S9: Measurement locations under the aligning functions and the homogenized warping functions for the data set in Section S4.3.1. Top row: Measurement locations under the true aligning function (left panel) and the true warping function (right panel). Bottom row: Measurement locations under the estimated aligning function (left panel) and the estimated warping function (right panel).

We also compared the predictive performance of the DCSM with the true model.

- Model S4.3.1.1: The bivariate, stationary, symmetric, parsimonious Matérn model. This is the model from which the data were simulated.
- Model S4.3.1.2: The same bivariate DCSM as Model S4.2.5 in Section S4.2.

Table S6 shows the cross-validation results of the simulation study. As expected, Model S4.3.1.2 produces worse RMSPE and CRPS than Model S4.3.1.1, but it does not perform much worse. Given that we use very complex deformations in Model S4.3.1.2, this shows that the DCSM is quite robust to overfitting.

S4.3.2 Simulated Asymmetric Nonstationary Data with a Misspecified Warping Function

In this section, we demonstrate the use of a DCSM with simulated asymmetric nonstationary data with a misspecified warping function. We simulated bivariate data from a bivariate Gaussian DCSM in a manner similar to Section S4.2, but now we use a different warping function, $\mathbf{f}(\mathbf{s}) = \mathbf{o} + (\mathbf{s} - \mathbf{o})\|\mathbf{s} - \mathbf{o}\|$ (Fouedjio et al., 2015), where $\mathbf{o} = (0, 0)'$. The data were simulated on a 101×101 grid defined on the

Table S6: Hold-out validation results, AIC, and the time required to fit, for the simulation study in Section S4.3.1, where the measurement locations are randomly sampled from G .

	$\tilde{Y}_1(\cdot)$		$\tilde{Y}_2(\cdot)$		AIC	Time (s)
	RMSPE	CRPS	RMSPE	CRPS		
Model S4.3.1.1	0.334	0.188	0.097	0.054	304.1	840.6
Model S4.3.1.2	0.346	0.195	0.108	0.060	437.1	2714.9

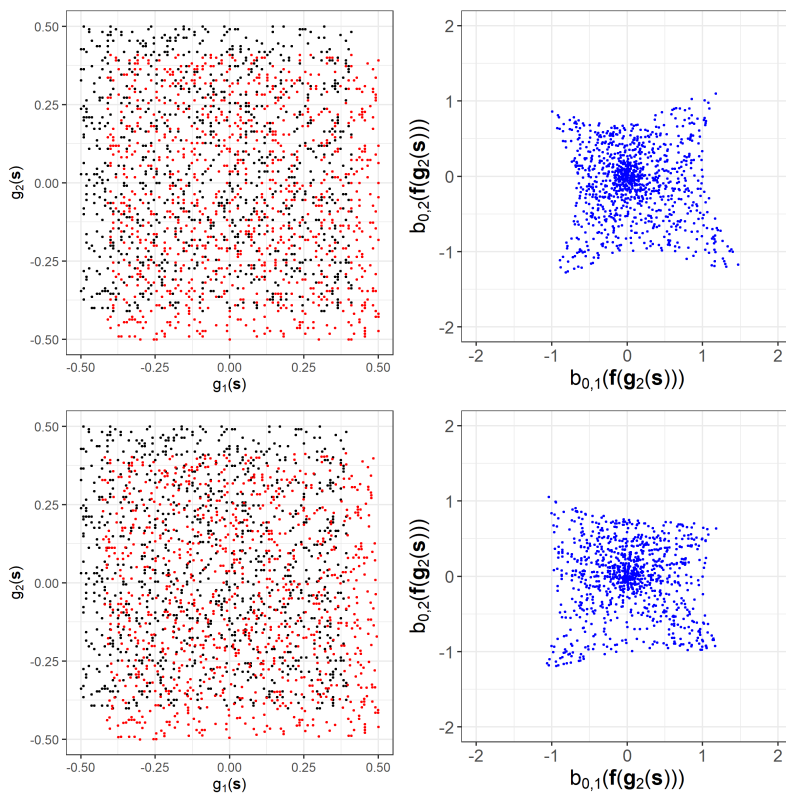


Figure S10: Measurement locations under the aligning functions and the homogenized warping functions for the data set in Section S4.3.2. Top row: Measurement locations under the true aligning function (left panel) and the true warping function (right panel). Bottom row: Measurement locations under the estimated aligning function (left panel) and the estimated warping function (right panel).

geographic domain G . We randomly sampled without replacement 1000 locations from G and used them as measurement locations.

As in Section S4.3.1, we examined the ability of the DCSM to retrieve the true warping function using the same model as Model S4.2.5 in Section S4.2. Figure S10 shows the measurement locations under the true aligning and warping function, and the measurement locations under the estimated aligning and warping function. We observe that the DCSM has correctly estimated the southeastern shift in the aligning function, and the contraction in the middle region of the warped domain.

Table S7: Hold-out validation results, AIC, and the time required to fit, for the simulation study in Section S4.4, where the measurement locations are randomly sampled from $G = [-0.5, 0.5] \times [-0.5, 0.5]$.

	$\tilde{Y}_1(\cdot)$		$\tilde{Y}_2(\cdot)$		$\tilde{Y}_3(\cdot)$		AIC	Time (s)
	RMSPE	CRPS	RMSPE	CRPS	RMSPE	CRPS		
Model S4.4.1	0.309	0.171	0.099	0.053	0.035	0.019	-1924.0	1385.4
Model S4.4.2	0.276	0.148	0.077	0.040	0.027	0.015	-2865.0	9047.0

S4.4 Simulated Trivariate Symmetric Nonstationary Data

In the previous data illustrations, we considered bivariate spatial data. In this section, we consider trivariate spatial data. We simulated trivariate data from a Gaussian DCSM, $\tilde{\mathbf{Y}}(\cdot)$, with constant mean. The data were simulated on an equally spaced 51×51 grid of the geographic domain, $G \equiv [-0.5, 0.5] \times [-0.5, 0.5]$ resulting in 2601 data. The warping function we used was a composition of axial warping units, followed by a single-resolution radial basis function unit, followed by a Möbius transformation unit. On the warped domain, we modeled the covariances using a trivariate stationary, isotropic parsimonious Matérn model. As in Section 4.1, we randomly sampled 1000 locations from the grid and used these as measurement locations.

We compared the predictive performance of the trivariate stationary parsimonious Matérn model (Model S4.4.1) to those of the trivariate DCSM (Model S4.4.2) by calculating the cross-validated predictive performance at the remaining 1601 locations using the RMSPE and the CRPS. Table S7 summarizes the results from this study. From this table, it is clear that, similar to the bivariate case in Section 4.1, there is a large improvement in RMSPE and CRPS when using the DCSM (Model S4.4.2) over the stationary parsimonious Matérn model (Model S4.4.1). The visualization of the estimated warping function is given in Figure S11. We can see that, similar to the bivariate case, the estimated warping function has retrieved important features such as the contraction in the middle part of the domain. However, we find that more iterations are needed, thus more computing time, to train the DCSM in the trivariate-data example.

S4.5 Modeling Maximum and Minimum Temperatures in Southwestern United States

In this section, we consider monthly maximum and minimum temperatures (which are calculated by averaging daily maximum and minimum temperatures over the month) in July 2018 at 909 weather stations over four southwestern states in the United States: Utah, Colorado, Arizona, and New Mexico. The data were extracted from the Global Summary of the Month data set of the National Oceanic and Atmospheric Administration²¹.

²¹<https://data.nodc.noaa.gov/cgi-bin/iso?id=gov.noaa.ncdc:C00946>




Figure S11: Measurement locations under the true (left panel) and estimated (right panel) homogenized warping functions for the data set in Section S4.4.

We compared the performance of the bivariate DCSMs to that of bivariate stationary Matérn models. Elevation is a potential covariate when modeling temperature. We considered two trend models, namely one with only an intercept (in which case $Y_1(\cdot)$ and $Y_2(\cdot)$ are highly nonstationary) and one with an intercept and elevation as a covariate (in which case $Y_1(\cdot)$ and $Y_2(\cdot)$ can be expected to be nonstationary but less so). Also, the maximum and minimum temperatures in a given month (here July 2018) can be treated as symmetric spatial processes since their interaction can be expected to be highly co-located. Thus, asymmetry of the cross-covariance matrices was not modeled in this data illustration. We considered the following four models:

- Model S4.5.1: A bivariate, stationary, parsimonious Matérn covariance model, and with only an intercept in the trend.
- Model S4.5.2: A bivariate, symmetric DCSM with the parsimonious bivariate Matérn covariance on the warped domain, and with only an intercept in the trend
- Model S4.5.3: A bivariate, stationary, parsimonious Matérn covariance model, and with an intercept and elevation as trend.
- Model S4.5.4: A bivariate, symmetric DCSM with the parsimonious bivariate Matérn covariance on the warped domain, and with an intercept and elevation as trend.

We used the same general architecture for the warping function of the bivariate DCSM as was used in the simulation study in Section 4.1.

We first examined the predictive performance of the four models when the data were missing at random by performing a five-fold cross-validation. We randomly chose 905 stations from the 909 stations for the study, and randomly divided the 905 stations into five groups of 181 stations each in order to carry out a five-fold cross-validation. Table S8 shows the results of this five-fold cross-validation study. We see a slight improvement in the predictive performance of the DCSMs over the corresponding stationary

Table S8: Five-fold cross-validation results, AIC, and the time required to fit, for the maximum and minimum temperature data in the southwestern USA for the study in Section S4.5, where data were missing at random.

	T_{max}		T_{min}		AIC	Time (s)
	RMSPE	CRPS	RMSPE	CRPS		
Model S4.5.1	3.17	1.74	2.44	1.33	7140.8	359.4
Model S4.5.2	3.09	1.69	2.31	1.26	7003.5	791.8
Model S4.5.3	1.07	0.58	2.02	1.14	5475.5	413.2
Model S4.5.4	1.04	0.57	1.98	1.11	5398.5	766.3

covariance models, on the order of 3–5% for RMSPE and CRPS. We also see a substantial improvement when elevation is included in the trend model, showing that addressing nonstationarity in the mean function is helpful in this application.

We next considered hold-out validation, where the data are missing in a block, and we held out 131 stations lying between 36°N–39°N and 104°W–108°W from all the 909 stations. The western part of the hold-out region is a mountainous area that extends into the non-hold-out region, while the eastern part of the hold-out region is an area with lower elevation. We then fitted the four models to the remaining 778 stations. Table S9 displays the hold-out validation results for predicting the maximum and minimum temperatures at the 131 hold-out stations, and Figure S12 shows the maps of predictions and prediction standard errors for Model S4.5.1 and Model S4.5.2.

From Table S9, the predictive performances of the bivariate DCSMs show a huge improvement in RMSPE and CRPS, on the order of 30% over those of the bivariate stationary covariance models when predicting maximum temperature. The visualization in Figure S12 further illustrates the utility of using DCSMs: With the intercept-only structure (i.e., constant mean) in the mean, the bivariate DCSM can predict lower temperatures on the western part of the hold-out region (that is, the mountainous areas), whereas the stationary model reverts to the constant mean. The DCSM also produces lower prediction standard errors over the hold-out region than the stationary model. These results corroborate those from the simulation study in Section 4.1, which showed that bivariate DCSMs can be useful even when data are missing over a large region. Figure S13 shows the estimated warping function in Model S4.5.2.




Figure S12: Comparison of predictions and prediction standard errors when using a bivariate symmetric, stationary, parsimonious Matérn model with constant mean (Model S4.5.1), and a bivariate symmetric DCSM with constant mean (Model S4.5.2). Data were left out of the region enclosed by the black rectangle. First row: Maximum temperature observations, Z_1 (left panel); predictions obtained using Model S4.5.1 (center panel), and Model S4.5.2 (right panel). Second row: Locations of the measurement of Z_1 (left panel); prediction standard errors obtained when using Model S4.5.1 (center panel), and Model S4.5.2 (right panel). Third and fourth rows: Analogous to the first and second rows, respectively, for the minimum temperature, Z_2 .

Table S9: Hold-out validation results, AIC, and the time required to fit, for the maximum and minimum temperature data in the southwestern USA for the study in Section S4.5, where data were missing in a block shown in Figure S12.

	T_{max}		T_{min}		AIC	Time (s)
	RMSPE	CRPS	RMSPE	CRPS		
Model S4.5.1	4.66	2.59	3.90	2.20	7610.1	418.7
Model S4.5.2	3.84	2.19	2.57	1.51	7474.6	915.0
Model S4.5.3	1.54	0.87	2.11	1.22	5858.4	475.0
Model S4.5.4	1.06	0.59	1.99	1.16	5787.2	898.5

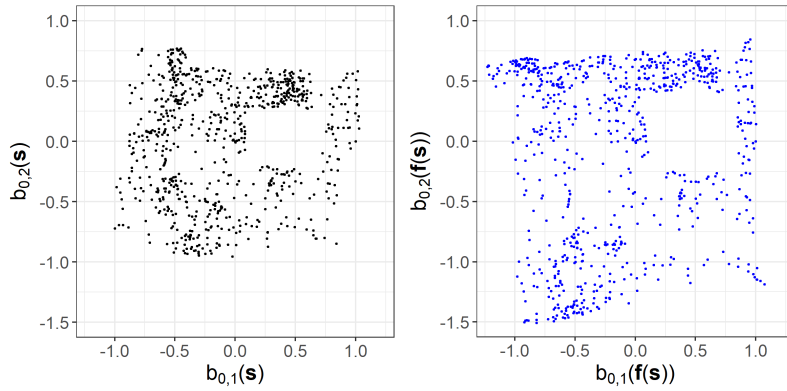


Figure S13: Measurement locations on the geographical domain and under the estimated homogenized warping function for the data set in Section S4.5.

References

- Allaire, J. J. and Tang, Y. (2019). *tensorflow: R Interface to 'TensorFlow'*. Online: Available from <https://github.com/rstudio/tensorflow>.
- Anderes, E. and Chatterjee, S. (2009). Consistent estimates of deformed isotropic Gaussian random fields on the plane. *The Annals of Statistics*, 37:2324–2350.
- Anderes, E. B. and Stein, M. L. (2008). Estimating deformations of isotropic Gaussian random fields on the plane. *The Annals of Statistics*, 36:719–741.
- Apanasovich, T. V. and Genton, M. G. (2010). Cross-covariance functions for multivariate random fields based on latent dimensions. *Biometrika*, 97:15–30.
- Apanasovich, T. V., Genton, M. G., and Sun, Y. (2012). A valid Matérn class of cross-covariance functions for multivariate random fields with any number of components. *Journal of the American Statistical Association*, 107:180–193.
- Calandra, R., Peters, J., Rasmussen, C. E., and Deisenroth, M. P. (2016). Manifold Gaussian processes for regression. In *Proceedings of the 2016 International Joint Conference on Neural Networks (IJCNN)*, pages 3338–3345. IEEE, Vancouver, BC, Canada.
- Cressie, N. and Johannesson, G. (2008). Fixed rank kriging for very large spatial data sets. *Journal of the Royal Statistical Society: Series B*, 70:209–226.
- Cressie, N. and Lahiri, S. N. (1993). The asymptotic distribution of REML estimators. *Journal of Multivariate Analysis*, 45:217–233.
- Cressie, N. and Lahiri, S. N. (1996). Asymptotics for REML estimation of spatial covariance parameters. *Journal of Statistical Planning and Inference*, 50:327–341.
- Cressie, N. and Zammit-Mangion, A. (2016). Multivariate spatial covariance models: A conditional approach. *Biometrika*, 103:915–935.
- Damian, D., Sampson, P. D., and Guttorp, P. (2001). Bayesian estimation of semi-parametric non-stationary spatial covariance structures. *Environmetrics*, 12:161–178.
- Fouedjio, F., Desassis, N., and Romary, T. (2015). Estimation of space deformation model for non-stationary random functions. *Spatial Statistics*, 13:45–61.
- Fuglstad, G.-A., Lindgren, F., Simpson, D., and Rue, H. (2015). Exploring a new class of non-stationary spatial Gaussian random fields with varying local anisotropy. *Statistica Sinica*, 25:115–133.

- Gelfand, A. E., Schmidt, A. M., Banerjee, S., and Sirmans, C. (2004). Nonstationary multivariate process modeling through spatially varying coregionalization. *Test*, 13:263–312.
- Gneiting, T., Kleiber, W., and Schlather, M. (2010). Matérn cross-covariance functions for multivariate random fields. *Journal of the American Statistical Association*, 105:1167–1177.
- Gneiting, T. and Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102:359–378.
- Goulard, M. and Voltz, M. (1992). Linear coregionalization model: Tools for estimation and choice of cross-variogram matrix. *Mathematical Geology*, 24:269–286.
- Gu, M. and Shen, W. (2020). Generalized probabilistic principal component analysis of correlated data. *Journal of Machine Learning Research*, 21:13.
- Gu, M., Wang, X., and Berger, J. O. (2018). Robust Gaussian stochastic process emulation. *Annals of Statistics*, 46:3038–3066.
- Heaton, M. J., Datta, A., Finley, A. O., Furrer, R., Guinness, J., Guhaniyogi, R., Gerber, F., Gramacy, R. B., Hammerling, D., Katzfuss, M., Lindgren, F., Nychka, D. W., Sun, F., and Zammit-Mangion, A. (2019). A case study competition among methods for analyzing large spatial data. *Journal of Agricultural, Biological and Environmental Statistics*, 24:398–425.
- Higdon, D., Swall, J., and Kern, J. (1999). Non-stationary spatial modeling. *Bayesian Statistics*, 6:761–768.
- Hildeman, A., Bolin, D., and Rychlik, I. (2019). Joint spatial modeling of significant wave height and wave period using the SPDE approach. *arXiv preprint, arXiv:1906.00286*.
- Hu, X. and Steinsland, I. (2016). Spatial modeling with system of stochastic partial differential equations. *Wiley Interdisciplinary Reviews: Computational Statistics*, 8:112–125.
- Kadane, J. B. (1974). The role of identification in Bayesian theory. In Fienberg, S. E. and Zellner, A., editors, *Studies in Bayesian Econometrics and Statistics*, pages 175–191. Amsterdam, The Netherlands.
- Kleiber, W. and Nychka, D. (2012). Nonstationary modeling for multivariate spatial processes. *Journal of Multivariate Analysis*, 112:76–91.
- Li, B. and Zhang, H. (2011). An approach to modeling asymmetric multivariate spatial covariance structures. *Journal of Multivariate Analysis*, 102:1445–1453.
- Lindgren, F., Rue, H., and Lindström, J. (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B*, 73:423–498.

- Majumdar, A. and Gelfand, A. E. (2007). Multivariate spatial modeling for geostatistical data using convolved covariance functions. *Mathematical Geology*, 39:225–245.
- Meiring, W., Monestiez, P., Sampson, P., and Guttorp, P. (1997). Developments in the modelling of non-stationary spatial covariance structure from space-time monitoring data. In Baafi, E. Y. and Schofield, N. A., editors, *Geostatistics Wollongong '96*, pages 162–173. Kluwer, Dordrecht, The Netherlands.
- Messick, R. M., Heaton, M. J., and Hansen, N. (2017). Multivariate spatial mapping of soil water holding capacity with spatially varying cross-correlations. *Annals of Applied Statistics*, 11:69–92.
- Nguyen, H., Cressie, N., and Braverman, A. (2017). Multivariate spatial data fusion for very large remote sensing datasets. *Remote Sensing*, 9:142.
- Olea, R. A. and Pardo-Iguzquiza, E. (2011). Generalized bootstrap method for assessment of uncertainty in semivariogram inference. *Mathematical Geosciences*, 43:203–228.
- Paciorek, C. J. and Schervish, M. J. (2006). Spatial modelling using a new class of nonstationary covariance functions. *Environmetrics*, 17:483–506.
- Perrin, O. and Monestiez, P. (1999). Modelling of non-stationary spatial structure using parametric radial basis deformations. In Gomez-Hernandez, J., Soares, A., and Froidevaux, R., editors, *geoENV II—Geostatistics for Environmental Applications*, pages 175–186. Springer, New York, NY.
- Porcu, E., Senoussi, R., Mendoza, E., and Bevilacqua, M. (2020). Reduction problems and deformation approaches to nonstationary covariance functions over spheres. *Electronic Journal of Statistics*, 14:890–916.
- Qadir, G. A. and Sun, Y. (2020). Semiparametric estimation of cross-covariance functions for multivariate random fields. *Biometrics*, in press, DOI:<https://doi.org/10.1111/biom.13323>.
- Sampson, P., Damian, D., and Guttorp, P. (2001). Advances in modeling and inference for environmental processes with nonstationary spatial covariance. In Monestiez, P., Allard, D., and Froidevaux, R., editors, *GeoENV III—Geostatistics for Environmental Applications*, pages 17–32. Springer, New York, NY.
- Sampson, P. D. and Guttorp, P. (1992). Nonparametric estimation of nonstationary spatial covariance structure. *Journal of the American Statistical Association*, 87:108–119.
- Schmidt, A. M. and O’Hagan, A. (2003). Bayesian inference for non-stationary spatial covariance structure via spatial deformations. *Journal of the Royal Statistical Society: Series B*, 65:743–758.
- Solow, A. R. (1985). Bootstrapping correlated data. *Mathematical Geology*, 17:769–775.

- Ver Hoef, J. M. and Barry, R. P. (1998). Constructing and fitting models for cokriging and multivariable spatial prediction. *Journal of Statistical Planning and Inference*, 69:275–294.
- Ver Hoef, J. M. and Cressie, N. (1993). Multivariable spatial prediction. *Mathematical Geology*, 25:219–240.
- Wackernagel, H. (2003). *Multivariate Geostatistics: An Introduction with Applications*. Springer, Berlin.
- Wiens, A., Kleiber, W., Barnhart, K. R., and Sain, D. (2020). Surface estimation for multiple misaligned point sets. *Mathematical Geosciences*, 52:527–542.
- Zammit-Mangion, A., Cressie, N., Ganesan, A. L., O’Doherty, S., and Manning, A. J. (2015a). Spatio-temporal bivariate statistical models for atmospheric trace-gas inversion. *Chemometrics and Intelligent Laboratory Systems*, 149:227–241.
- Zammit-Mangion, A., Ng, T. L. J., Vu, Q., and Filippone, M. (2021). Deep compositional spatial models. *Journal of the American Statistical Association*, in press, DOI:<https://doi.org/10.1080/01621459.2021.1887741>.
- Zammit-Mangion, A., Rougier, J., Schön, N., Lindgren, F., and Bamber, J. (2015b). Multivariate spatio-temporal modelling for assessing Antarctica’s present-day contribution to sea-level rise. *Environmetrics*, 26:159–177.
- Zhang, H. (2004). Inconsistent estimation and asymptotically equal interpolations in model-based geostatistics. *Journal of the American Statistical Association*, 99:250–261.