

NIASRA

NATIONAL INSTITUTE FOR APPLIED
STATISTICS RESEARCH AUSTRALIA



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

University of Wollongong, Australia

Working Paper

04-22

Basis-Function Models in Spatial Statistics

Noel Cressie, Matthew Sainsbury-Dale, and Andrew Zammit-Mangion

*Copyright © 2022 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 42215076, Fax: +61 2 42214998.

Email: karink@uow.edu.au

Basis-Function Models in Spatial Statistics

Noel Cressie, Matthew Sainsbury-Dale, and Andrew Zammit-Mangion

School of Mathematics and Applied Statistics, University of Wollongong, Australia, NSW
2522; email: ncressie@uow.edu.au

Abstract

Spatial statistics is concerned with the analysis of data that have spatial locations associated with them, and those locations are used to model statistical dependence between the data. The spatial data are treated as a single realisation from a probability model that encodes the dependence through both fixed effects and random effects, where randomness is manifest in the underlying spatial process and in the noisy, incomplete, measurement process. The focus of this review article is on the use of basis functions to provide an extremely flexible and computationally efficient way to model spatial processes that are possibly highly non-stationary. Several examples of basis-function models are provided to illustrate how they are used in Gaussian, non-Gaussian, multivariate, and spatio-temporal settings, with applications in geophysics. Our aim is to emphasise the versatility of these spatial statistical models and to demonstrate that they are now centre-stage in a number of application domains. The review concludes with a discussion and illustration of software currently available to fit spatial-basis-function models and implement spatial-statistical prediction.

Keywords— hierarchical statistical models, low-rank models, multivariate models, non-stationary covariance functions, spatial basis functions

1 Preface

This article divides naturally into four principal sections. Section 2 provides a succinct review of spatial statistics, its importance in scientific modelling and inference, particularly in the environmental sciences, and the role of basis-function models when the spatial data are ‘big’. Section 3 presents Gaussian and non-Gaussian spatial-basis-function models, with particular emphasis on their role in hierarchical statistical models and spatial prediction. Section 4 considers multivariate, warped, and spatio-temporal basis-function models. The review concludes with Section 5, which discusses software available, as of 2021, for spatial-basis-function modelling, and it uses a spatial data set of sea-surface temperatures for illustration.

2 Introduction to spatial statistics and basis-function models

There are two main subsections in Section 2, the first reviews spatial statistics and the second introduces basis-function models.

2.1 Spatial statistics, grand challenges, and spatial stochastic processes

In what follows in this subsection, we give a brief review of spatial statistics and its importance in scientific studies.

2.1.1 Spatial statistics

Spatial statistics is concerned with the statistical analysis (both exploratory and confirmatory) of data indexed by spatial locations or regions in a spatial domain of interest. A probabilistic framework is invoked that captures Tobler’s famous ‘first law of geography’ (Tobler, 1970): “Everything is related to everything else, but near things are more related than distant things.” Statistical independence within a spatial data set is almost never assumed, although one should not rule out an eventual statistical model where spatial data are independent.

In spatial statistics, models have traditionally been proposed directly on the spatial data, but this modelling approach does not recognise a latent scientific process behind the noisy and incomplete (over the spatial domain) data. We shall see in Section 2.1.2 that “what you see (data) is not what you want to get (process)” (Cressie & Wikle, 2011, p. xvi).

Grand challenges that need spatial statistics

Perhaps the greatest existential threat for humans is climate change; the planet is warming as its *Homo sapiens* are putting more and more greenhouse gases (in particular, carbon-based gases) into the atmosphere. The COP21 agreement signed in Paris in 2015 tasked each country to reduce their carbon emissions by an agreed-upon percentage by 2030, with the goal of capping global temperature increase at 2°C above the pre-industrial level. For a country to keep to its COP21 commitment, its government needs to know where the carbon sources are that can be decreased, and where the carbon sinks are that can be enhanced. These sources and sinks are more-or-less uncertain, and their estimation requires spatial-statistical methods (e.g., Michalak et al., 2004, Zammit-Mangion et al., 2021a).

Planet Earth also faces a resources ‘grand challenge’ of producing enough food, water, energy, and shelter for its inhabitants. Equally important for Earth’s future, we *Homo sapiens* must collectively recognise that the health of its diverse species is critical to our own health. For example, not only does logging of the Amazon region move carbon from the biosphere into the atmosphere where it is dangerous, it destroys habitat for plants and animals, severely impacting an ecosystem that is a source of food and medicine for our species. Spatial data sets collected by ecologists are at the centre of monitoring and assessment of critical ecosystem health (e.g., Hooten et al., 2017).

2.1.2 Spatial stochastic processes and their measurement

Models are a way to focus on the important aspects of a problem, and in this subsection we start with a univariate process of interest, defined on a domain D that is a subset of d -dimensional Euclidean space. A spatial stochastic process on D is written as

$$Y(\cdot) \equiv \{Y(\mathbf{s}) : \mathbf{s} \in D\}, \quad (1)$$

where $D \subset \mathbb{R}^d$ and, at $\mathbf{s} \equiv (s_1, \dots, s_d)^\top$, $Y(\mathbf{s})$ is a random quantity. On D , we define the Euclidean distance, $\|\mathbf{s} - \mathbf{u}\| \equiv \{\sum_{i=1}^d (s_i - u_i)^2\}^{1/2}$, between spatial locations \mathbf{s} and \mathbf{u} in D . The process $Y(\cdot)$ is well defined if all finite probability distributions of $(Y(\mathbf{s}_1), \dots, Y(\mathbf{s}_m))^\top$ exist for all $\{\mathbf{s}_1, \dots, \mathbf{s}_m\} \subset D$ and all $m \geq 1$, and they satisfy Kolmogorov’s consistency conditions (e.g., Tao, 2011, Sec. 2.4). Cressie (1993, Ch. 1) and Banerjee et al. (2004, Ch. 1) give three main cases for D , resulting in geostatistical processes (D fixed, uncountable with Lebesgue measure $|D| > 0$), lattice spatial processes (D fixed, countable), and marked point processes (D a countable subset of a fixed domain). See also Cressie (1993, Sec. 1.1) for an integrative spatial model covering all these cases, where *both* $Y(\cdot)$ and D in Equation 1 are assumed random.

Spatial stochastic process: A countable or uncountable set of random variables indexed by locations in a spatial domain.

In the rest of this review, we concentrate on basis-function models for the geostatistical case. However, basis-function models certainly have a place in spatial statistics for lattice data (e.g., Bradley et al., 2016b, 2020) and point patterns (e.g., Cseke et al. 2016, Simpson et al. 2016, Hooten et al. 2017, Ch. 4).

To introduce the spatial data $\mathbf{Z} \equiv (Z(\mathbf{s}_1), \dots, Z(\mathbf{s}_n))^T$ into the probability model, we specify the *data model*, $[\mathbf{Z} | Y(\cdot)]$, which is the probability distribution of \mathbf{Z} given $Y(\cdot)$. Then the joint distribution of the latent spatial process $Y(\cdot)$ and the spatial data \mathbf{Z} is

$$[Y(\cdot), \mathbf{Z}] = [\mathbf{Z} | Y(\cdot)][Y(\cdot)], \quad (2)$$

where henceforth the probability model $[Y(\cdot)]$ of the latent process $Y(\cdot)$ given by Equation 1, is referred to as the *process model*. In Equation 2 and elsewhere, the bracket notation is used: For generic random quantities A and B , $[A, B]$ denotes the joint probability distribution of A and B ; $[A | B]$ denotes the conditional probability distribution of A given B ; and $[A]$ denotes the marginal probability distribution of A .

As explained in Section 2.1.1, the data \mathbf{Z} are incomplete (i.e., do not cover all of the spatial domain D) and noisy (i.e., contain measurement error), which is captured in the data model, $[\mathbf{Z} | Y(\cdot)]$. A common assumption made in the data model is that of conditional independence, since usually the act of measurement made at one location is statistically independent of the act of measurement at another. Henceforth in this review, the data model is assumed to exhibit this conditional independence, namely

$$[\mathbf{Z} | Y(\cdot)] = \prod_{i=1}^n [Z(\mathbf{s}_i) | Y(\mathbf{s}_i)]. \quad (3)$$

The spatial dependence one expects from ‘Tobler’s first law’ is captured probabilistically in the process model $[Y(\cdot)]$. Then the spatial-statistical dependence in the spatial data \mathbf{Z} that Tobler articulated is inherited from $[Y(\cdot)]$ through the expression,

$$[\mathbf{Z}] = \int [\mathbf{Z} | Y(\cdot)][Y(\cdot)]dY(\cdot),$$

where the integral is used formally to represent marginalisation of the joint probability measure of $Y(\cdot)$ and \mathbf{Z} . The combination of the probability model $[Y(\cdot)]$ and Equation 3 defines a *hierarchical spatial-statistical model* (e.g., Wikle & Berliner, 2007).

In this review, our principal interest is in spatial prediction of $Y(\mathbf{s}_0)$ at a given spatial location $\mathbf{s}_0 \in D$. To do this, we use the predictive distribution, $[Y(\mathbf{s}_0) | \mathbf{Z}]$, which is given by Bayes’ Rule,

$$[Y(\mathbf{s}_0) | \mathbf{Z}] = [\mathbf{Z} | Y(\mathbf{s}_0)][Y(\mathbf{s}_0)]/[\mathbf{Z}]. \quad (4)$$

When both the data \mathbf{Z} and the process $Y(\cdot)$ are Gaussian, this predictive distribution can be computed using standard Gaussian identities (see Section 3.1). When the data and/or the process model are non-Gaussian, approximations or Monte Carlo sampling-based techniques are typically needed for inferences such as spatial prediction (see Section 3.2).

2.2 Basis-function representations of spatial processes and covariance functions

In the subsections that follow, we use basis functions to represent spatial processes and their corresponding spatial covariance functions.

2.2.1 Basis-function representations of spatial processes

From a mathematical point of view, a basis of a function space is a collection of elements for which any function in the space can be represented as a linear combination

Bayes’ Rule: For generic random quantities A and B , Bayes’ Rule relies on the simple relationship, $[A | B] = [B | A][A]/[B]$.

of these basis elements. Relevant to this review, a countable basis $\{\phi_j(\cdot) : j = 1, 2, \dots\}$ of the space of square-integrable functions, \mathcal{F} , has the property that any function $f(\cdot) \in \mathcal{F}$ can be represented as

$$f(\cdot) = \sum_{j \geq 1} a_j \phi_j(\cdot),$$

Inner product: The inner-product operation for two elements $f(\cdot)$ and $g(\cdot)$ of a function space whose functions are square integrable on the domain D , is $\int_D f(\mathbf{s})g(\mathbf{s})d\mathbf{s}$.

Orthogonality: $f(\cdot)$ and $g(\cdot)$ are *orthogonal* if their inner product is zero.

Orthonormality: $f(\cdot)$ and $g(\cdot)$ are *orthonormal* if they are orthogonal and the inner products of $f(\cdot)$ and $g(\cdot)$ with themselves are both equal to 1.

in an L_2 sense. Furthermore, if the elements of $\{\phi_j(\cdot) : j = 1, 2, \dots\}$ are mutually orthonormal in terms of the inner product, then, straightforwardly, $a_j = \int_D f(\mathbf{s})\phi_j(\mathbf{s})d\mathbf{s}$ for $j \geq 1$. Orthogonality of a basis is attractive but not necessary for a good representation of a function's behaviour. For example, spline bases are not usually orthogonal and, while many wavelet bases are, there may be computational reasons for using non-orthogonal wavelets.

Now we switch our point of view from mathematical to statistical: Consider a probability measure on the space of spatial functions on D defined by a finite random linear combination of basis functions plus an independent error term (e.g., Solo, 2002). The probability measure comes from the representation

$$Y(\cdot) = \sum_{j=1}^r \alpha_j \phi_j(\cdot) + \delta(\cdot), \quad (5)$$

where $\sum_{j=1}^r \alpha_j \phi_j(\cdot)$ is the basis-function term consisting of a linear combination of the basis functions $\{\phi_j(\cdot) : j = 1, \dots, r\}$; the coefficients $\{\alpha_j : j = 1, \dots, r\}$ are in general dependent random variables; and the error term, $\delta(\cdot) \equiv \{\delta(\mathbf{s}) : \mathbf{s} \in D\}$, is a stochastic process independent of the basis-function coefficients. The inclusion of $\delta(\cdot)$ in Equation 5 accounts statistically for the error incurred by using only a finite number of basis functions. As we discuss below, it preserves variability of $Y(\cdot)$ and, in that sense, Equation 5 is a statistical model, not an approximation, of spatial variability.

The basis functions in Equation 5 may be defined by truncations of a countable basis, or they may simply be functions thought to be important for representing the spatial variability. For example, basis functions can come in the form of splines (Wahba, 1990), wavelets (Vidakovic & Müller, 1999), bisquare functions (Cressie & Johannesson, 2008), Wendland functions (Nychka et al., 2015), and finite elements (Lindgren et al., 2011). The basis functions of Banerjee et al. (2008) depend on covariance-function parameters (unlike the examples given above), which results in slower computations when fitting and predicting (Bradley et al., 2016a). The resulting process model $[Y(\cdot)]$, where $Y(\cdot)$ is given by Equation 5, has been called the spatial random effects (SRE) model (Cressie & Johannesson, 2008), which is generalised below to a spatial mixed effects model and used extensively in spatial prediction from very large data sets (Zammit-Mangion & Cressie, 2021). A review of such low-rank representations can be found in Wikle (2010).

In the hierarchical model (HM), $Y(\cdot)$ represents the scientific process, and hence $\{\text{var}(Y(\mathbf{s})) : \mathbf{s} \in D\}$ should be conserved, no matter what basis-function representation is chosen. Equation 5 conserves variability in that $\text{var}(Y(\mathbf{s})) = \text{var}(\sum_{j=1}^r \alpha_j \phi_j(\mathbf{s})) + \text{var}(\delta(\mathbf{s}))$ (Cressie & Johannesson, 2008). That is, Equation 5 is a *model*, not an approximation for $Y(\cdot)$. Clearly, $\delta(\cdot)$ is a critical component, since $\text{var}(Y(\mathbf{s})) \geq \text{var}(\sum_{j=1}^r \alpha_j \phi_j(\mathbf{s}))$ for all $\mathbf{s} \in D$. Of course, the model for $\delta(\cdot)$ changes if the basis-function component changes, since the variances and covariances of the total, $Y(\cdot)$, are fixed by the underlying science.

Throughout this article, the basis functions are known, and the randomness is in their coefficients $\{\alpha_j : j = 1, \dots, r\}$. There is another type of representation we shall mention briefly, known as spatial factor analysis (Christensen & Amemiya, 2002, 2003). The models are defined for multivariate spatial processes, where the factors are spatial functions that are random, and the coefficients are deterministic but have to

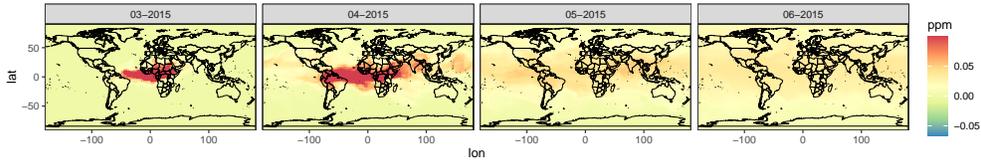


Figure 1: A spatio-temporal basis function, in units of parts per million (ppm) of CO_2 , generated using the transport model MOZART in response to a pulse of CO_2 flux in Northern Africa throughout the month of March 2015. The plots show the CO_2 concentration in ppm on March 15, April 15, May 15, and June 15, 2015, respectively.

satisfy critical identifiability conditions; further discussion is given in Sections 4.1 and 4.3.2.

Now basis functions are often multi-resolutional (e.g., wavelets, bisquares) but some of them could be physical. A physical basis function could be some readily available geophysical quantity, such as the elevation of the terrain as a function of $\mathbf{s} \in D$ when modelling (say) maximum temperature in a geographic region D . In more complicated cases, a basis function could be the output of a numerical model, such as an atmospheric-transport model. For example, **Figure 1** summarises the three-dimensional spatio-temporal output of the Model for OZone And Related chemical Tracers (MOZART), a numerical atmospheric-transport model. The output is simulated CO_2 changes in the atmosphere in response to a pulse of emissions in central Africa throughout the month of March 2015. Several of these basis functions can be constructed by ‘driving’ the atmospheric model with pulses at different times and in different regions, which can then be used to model a temporal sequence of spatial fields of CO_2 (in parts per million). This approach to basis-function modelling of CO_2 is ubiquitous in the atmospheric sciences (e.g., Enting, 2002), and it was developed into a fully Bayesian statistical framework by Zammit-Mangion et al. (2021a). Another example of physically motivated basis functions was provided by Wikle et al. (2001), who used the equatorial normal mode orthogonal basis functions to model tropical ocean surface winds.

The general basis-function model with non-zero mean $\mu(\cdot)$ can be written as

$$Y(\cdot) = \mu(\cdot) + \sum_{j=1}^r \alpha_j \phi_j(\cdot) + \delta(\cdot) = \mathbf{x}(\cdot)^\top \boldsymbol{\beta} + \boldsymbol{\phi}(\cdot)^\top \boldsymbol{\alpha} + \delta(\cdot), \quad (6)$$

where the second equality is imposed when the mean is expressed as a linear regression on a p -dimensional vector of covariates $\mathbf{x}(\cdot)$; that is, $\mu(\cdot) = \mathbf{x}(\cdot)^\top \boldsymbol{\beta}$. Equation 6 is called a *spatial mixed effects model*, since now the fixed effects $\mathbf{x}(\cdot)^\top \boldsymbol{\beta}$ are added to the spatial random effects expressed in Equation 5. In Equation 6, we have written the linear combination of basis functions as $\sum_{j=1}^r \alpha_j \phi_j(\cdot) = \boldsymbol{\phi}(\cdot)^\top \boldsymbol{\alpha}$, where $\boldsymbol{\phi}(\cdot) \equiv (\phi_1(\cdot), \dots, \phi_r(\cdot))^\top$ and $\boldsymbol{\alpha} \equiv (\alpha_1, \dots, \alpha_r)^\top$. Functions such as a constant or trends in the components of \mathbf{s} , are clearly covariates, but physical functions such as elevation could be a covariate (fixed effect) or a basis function (random effect), depending on the application. Basis functions whose coefficients are deterministic are part of the fixed effects and are considered as covariates; overfitting of the fixed effects can be handled by regularization, such as with the least absolute shrinkage and selection operator (lasso; Tibshirani, 1996), and this has a Bayesian interpretation.

In what follows, Equation 6 will be the process model in the HM. Data $\mathbf{Z} = (Z(\mathbf{s}_1), \dots, Z(\mathbf{s}_n))^\top$ are an imperfect ‘look’ at $Y(\cdot)$, and the principal inferential goal is to predict $Y(\cdot)$ over D with known or model-derived (in practice, estimated) uncertainty.

2.2.2 Spatial covariance functions for basis-function models

Positive-semidefinite:

A covariance function $C(\cdot, \cdot)$ is positive-semidefinite if $\sum_{i=1}^m \sum_{j=1}^m C(\mathbf{s}_i, \mathbf{s}_j) a_i \bar{a}_j \geq 0$ for any $\mathbf{s}_i \in \mathbb{R}^d$, $i = 1, \dots, m$, any complex $\{a_i : i = 1, \dots, m\}$, and any integer $m > 0$. (Here, \bar{a}_j is the complex conjugate of a_j .)

The covariance function of $Y(\cdot)$ is defined as

$$C_Y(\mathbf{s}, \mathbf{u}) \equiv \text{cov}(Y(\mathbf{s}), Y(\mathbf{u})), \quad \mathbf{s}, \mathbf{u} \in D. \quad (7)$$

The covariance function is a key functional parameter in spatial statistics, as it serves as the primary mechanism for capturing the dependence one expects from ‘Tobler’s first law’. For a spatial-statistical model to be valid, the covariance function is required to be positive-semidefinite. Recalling that the term $\delta(\cdot)$ is a stochastic process independent of the basis-function coefficients $\boldsymbol{\alpha}$, the covariance function implied by Equation 6 is

$$\begin{aligned} C_Y(\mathbf{s}, \mathbf{u}) &= \text{cov}(\mathbf{x}(\mathbf{s})^\top \boldsymbol{\beta} + \boldsymbol{\phi}(\mathbf{s})^\top \boldsymbol{\alpha} + \delta(\mathbf{s}), \mathbf{x}(\mathbf{u})^\top \boldsymbol{\beta} + \boldsymbol{\phi}(\mathbf{u})^\top \boldsymbol{\alpha} + \delta(\mathbf{u})) \\ &= \boldsymbol{\phi}(\mathbf{s})^\top \mathbf{K} \boldsymbol{\phi}(\mathbf{u}) + C_\delta(\mathbf{s}, \mathbf{u}), \end{aligned} \quad (8)$$

where $\mathbf{K} \equiv \text{cov}(\boldsymbol{\alpha}, \boldsymbol{\alpha})$, and $C_\delta(\mathbf{s}, \mathbf{u})$ is (usually) a simple covariance function that describes fine-scale dependence not captured by the basis-function term. The basis-function model is valid: Because \mathbf{K} is a covariance matrix, which is a positive-semidefinite matrix, the first term in Equation 8 is positive-semidefinite, and the second term is positive-semidefinite by definition; hence, their sum is positive-semidefinite.

A classical spatial-modelling assumption is covariance stationarity and, possibly, covariance isotropy. However, these assumptions are typically only realistic over small spatial domains. In addition to guaranteeing validity, basis-function models can yield non-stationary covariance functions through the covariance matrix \mathbf{K} or through the basis functions themselves. Often, \mathbf{K} is chosen to be of a relatively simple parametric form with minimal parameters to be estimated, in which case the burden of capturing non-stationarity falls to the basis functions.

To illustrate this point, **Figure 2** shows three sets of basis functions on the one-dimensional spatial domain, $D = [0, 1]$, and the associated covariance functions that arise when using a simple covariance matrix \mathbf{K} . The first two sets of basis functions consist of regularly-located bisquare basis functions that have compact support; the first set has (almost) non-overlapping basis functions, while the second set has basis functions with significant overlap. For these two sets of basis functions, \mathbf{K} was set to be the covariance matrix of a simple first-order autoregressive process. The non-overlapping bisquare basis functions produce undesirable artefacts in the covariance function; the overlapping bisquare basis functions yield a ‘pseudo-stationary’ covariance function. The third set of basis functions are highly irregular, and they are analogous to physically motivated basis functions, such as elevation, or the output of an atmospheric transport model. Here, \mathbf{K} was simply set to the identity matrix. Despite \mathbf{K} having a simple structure in this third set, the resulting covariance function is clearly non-stationary. This example shows two important results: First, including a degree of overlap when using basis functions with local scope can reduce undesirable artefacts in the covariance function; and second, the covariance function can be non-stationary even if \mathbf{K} has a simple parametric form, provided the basis functions are irregular in some way. In Section 4.2, we review a recent approach that warps the spatial domain D and allows highly non-stationary covariance functions to be constructed using simple structures for \mathbf{K} and simple basis functions.

For stationary covariance functions $C^*(\mathbf{h})$, continuity and differentiability properties as $\mathbf{h} \rightarrow \mathbf{0}$ are directly related to the smoothness properties of the sample paths of the underlying spatial process (e.g., Cressie 1993, p. 60; Stein 1999). These are important for determining which parameters of $C^*(\cdot)$ can be estimated consistently (Stein, 1999). Since basis-function models are non-stationary, these theoretical results are not applicable here. Under quite restrictive assumptions applied only in

Covariance stationarity:

a covariance function is stationary if it satisfies $C(\mathbf{s}, \mathbf{u}) = C^*(\mathbf{h})$, where the vector $\mathbf{h} \equiv \mathbf{u} - \mathbf{s}$ is the spatial lag between two locations \mathbf{u} and \mathbf{s} in \mathbb{R}^d .

Covariance isotropy:

a covariance function is isotropic if it depends only on the distance $\|\mathbf{h}\|$.

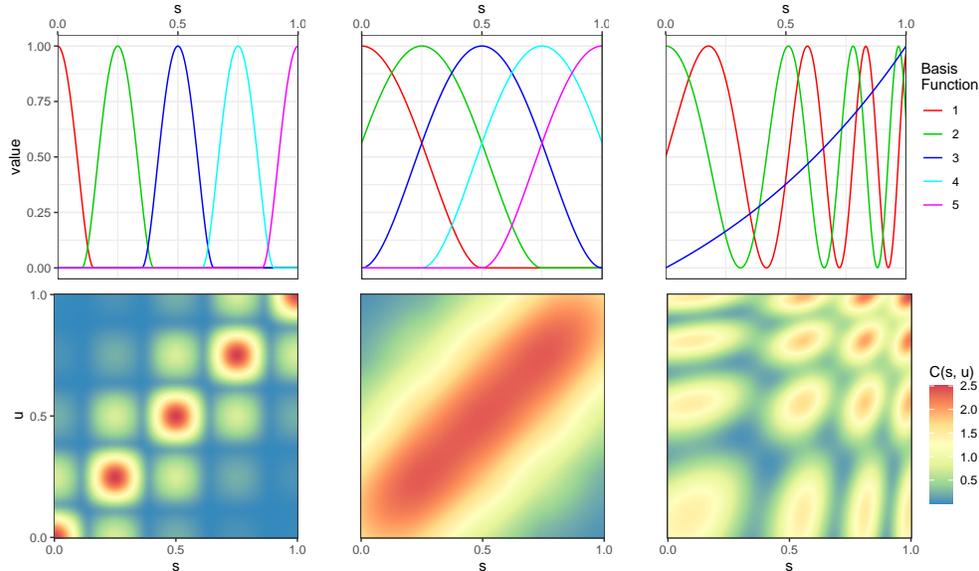


Figure 2: Three sets of bisquare basis functions (top row) on the one-dimensional spatial domain, $D = [0, 1]$, and the associated covariance functions that arise (bottom row). The first two sets of basis-function coefficients follow a first-order autoregressive model, while the third set of basis-function coefficients are modelled as independent.

one-dimensional space, Stein (2014) critiqued spatial-basis-function models. He assumed that the true spatial model is a periodic stationary process, which is then approximated by a basis-function model that results from the truncation of the spectral representation of $\text{cov}(\mathbf{Z}, \mathbf{Z})$. That is, he assumed that the basis functions are the orthonormal eigenvectors of $\text{cov}(\mathbf{Z}, \mathbf{Z})$, and he added a white-noise error process to preserve the variability of the true stationary model. While interesting in its own right, this set-up differs substantially from that used in this review, where the true model is a non-stationary basis-function model that uses any basis functions (orthonormal or not) defined over all $\mathbf{s} \in D \subset \mathbb{R}^d$ for $d \geq 1$.

There is a need for more general asymptotic results that would identify which parameters in \mathbf{K} can be estimated consistently. Here, conditions on the smoothness (or roughness) of the basis functions will be critical, since they determine the smoothness of the underlying non-stationary process $Y(\cdot)$ in Equation 5 and, as discussed above, smoothness is key (when $Y(\cdot)$ is stationary).

In addition to facilitating non-stationary covariance functions, it will be shown in the following section that the quadratic form in Equation 8 results in substantial computational savings when conducting spatial-statistical inference, particularly optimal spatial prediction. These computational savings make basis-function models well placed to handle the ‘big data’ often encountered in modern spatial statistics.

3 Gaussian and non-Gaussian basis-function models

Section 3 is divided into Sections 3.1, 3.2, and 3.3, which describe well established methodologies using basis-function representations. In Sections 3.1 and 3.2, we review basis-function representations in univariate Gaussian and non-Gaussian processes, respectively, and we consider spatial prediction of the underlying process $Y(\cdot)$. In Section 3.3, we discuss the empirical and Bayesian approaches to inference on the parameters of the basis-function model presented in Section 2.1.2, along with their implications

for spatial prediction of $Y(\cdot)$.

3.1 Basis-function representations in Gaussian spatial processes

In this subsection, we assume that both the process model and the data model are Gaussian. This so-called ‘Gau-Gau’ case is the most straightforward to work with in spatial statistics, since the likelihood, predictive means, and predictive variances are all derived from Gaussian distributions whose mean vector and covariance matrix are known. However, with increasing sample sizes in many applications, direct computation of the means and covariances has become increasingly problematic. The difficulty arises from the inversion of the $n \times n$ matrix $\mathbf{C}_Z \equiv \text{cov}(\mathbf{Z}, \mathbf{Z})$, the covariance matrix of the spatial data $\mathbf{Z} \equiv (\mathbf{Z}(\mathbf{s}_1), \dots, \mathbf{Z}(\mathbf{s}_n))^\top$ which, in general, is $O(n^3)$ in computational complexity. However, when the process model has a representation in terms of r basis functions, it will be seen below that inverting \mathbf{C}_Z is generally $O(nr^2)$ in computational complexity; when r is fixed, this results in $O(n)$ complexity.

Gaussian process: $\{Y(\mathbf{s}) : \mathbf{s} \in D\}$ is a Gaussian process if $\sum_{i=1}^m Y(\mathbf{s}_i)a_i$ is a Gaussian random variable for any $\{\mathbf{s}_i : i = 1, \dots, m\} \in D$, any real $\{a_i : i = 1, \dots, m\}$, and any integer $m > 0$.

3.1.1 Gaussian data and Gaussian process (Gau-Gau) model

It will be convenient to write the model in terms of vectors and matrices. Define \mathbf{Y} , $\boldsymbol{\delta}$, and $\boldsymbol{\epsilon}$ as we defined \mathbf{Z} in the introduction to this subsection, that is, as vectors of $Y(\cdot)$, $\delta(\cdot)$, and $\epsilon(\cdot)$ respectively, evaluated at each observation location $\{\mathbf{s}_i : i = 1, \dots, n\}$. Define the $n \times p$ matrix $\mathbf{X} \equiv (\mathbf{x}(\mathbf{s}_1), \dots, \mathbf{x}(\mathbf{s}_n))^\top$ and the $n \times r$ matrix $\boldsymbol{\Phi} \equiv (\boldsymbol{\phi}(\mathbf{s}_1), \dots, \boldsymbol{\phi}(\mathbf{s}_n))^\top$ whose i th rows respectively consist of the regression covariates and the spatial basis functions evaluated at $\{\mathbf{s}_i : i = 1, \dots, n\}$. Then the Gau-Gau model is

$$\mathbf{Z} = \mathbf{Y} + \boldsymbol{\epsilon}, \tag{9}$$

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\Phi}\boldsymbol{\alpha} + \boldsymbol{\delta}, \tag{10}$$

where $\boldsymbol{\beta}$ is a p -dimensional vector of fixed but generally unknown regression coefficients, $\boldsymbol{\alpha}$ is an r -dimensional Gaussian random vector of basis-function coefficients, $\boldsymbol{\delta}$ is an n -dimensional Gaussian random vector of process-model errors, and $\boldsymbol{\epsilon} \equiv (\epsilon_1, \dots, \epsilon_n)^\top$ is an n -dimensional random vector of data-model errors (i.e., measurement errors) independent of \mathbf{Y} . Parameters of this Gau-Gau model include $\boldsymbol{\beta}$, $\mathbf{K} \equiv \text{cov}(\boldsymbol{\alpha}, \boldsymbol{\alpha})$, $\mathbf{C}_\delta \equiv \text{cov}(\boldsymbol{\delta}, \boldsymbol{\delta})$, and $\mathbf{C}_\epsilon \equiv \text{cov}(\boldsymbol{\epsilon}, \boldsymbol{\epsilon})$. We discuss inference on the parameters in Section 3.3.

The components of Equations 9 and 10 are shown for $D = [0, 1] \subset \mathbb{R}^1$ in Figure 3. The top of the figure shows: bisquare spatial basis functions, $\phi_1(\cdot), \dots, \phi_{13}(\cdot)$, for two resolutions at equally spaced locations within a resolution. The bottom of the figure shows: the process $Y(\mathbf{s})$ for $\mathbf{s} \in D$ (thin black line), generated by Equation 6 and used to construct \mathbf{Y} in Equation 10, with $\boldsymbol{\beta} = \mathbf{0}$, $\boldsymbol{\alpha} \sim \text{Gau}(\mathbf{0}, \mathbf{K})$ for \mathbf{K} defined by an exponential covariance function, and $\boldsymbol{\delta} = \mathbf{0}$; the data \mathbf{Z} (black dots) generated by Equation 9 with $\boldsymbol{\epsilon} \sim \text{Gau}(\mathbf{0}, \sigma_\epsilon^2 \mathbf{I})$; and the optimal predictor $Y^*(\mathbf{s}_0) \equiv \mathbb{E}(Y(\mathbf{s}_0) | \mathbf{Z})$, for $\mathbf{s}_0 \in D$ (red dashed line) and the 90% pointwise prediction intervals based on the predictive standard deviation $\{\text{var}(Y(\mathbf{s}_0) | \mathbf{Z})\}^{1/2}$, for $\mathbf{s}_0 \in D$ (grey shaded region), obtained from Equation 12 below.

Gau-Gau spatial-basis-function models are ubiquitous in the literature; see, for example, Cressie & Johannesson (2008), Finley et al. (2009), Wikle (2010), Lindgren et al. (2011), Sang & Huang (2012), Katzfuss (2013), Nychka et al. (2015), Tzeng & Huang (2018), and Zammit-Mangion & Rougier (2020).

3.1.2 Spatial prediction

For spatial prediction, the quantity of interest is the predictive distribution of the underlying Gaussian spatial process given the data. Hence, we consider $[Y(\mathbf{s}_0) | \mathbf{Z}]$

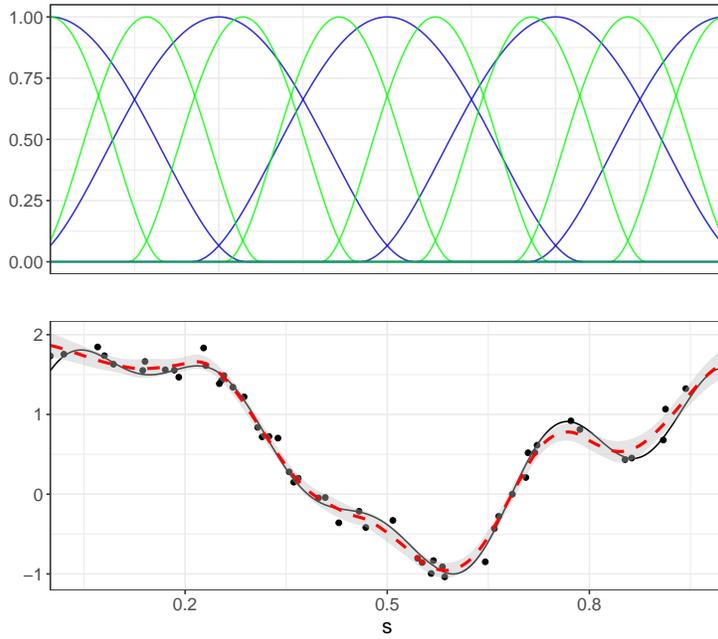


Figure 3: A simple illustration of the Gau-Gau basis-function model in Equations 9 and 10. (Top) $r = 13$ bisquare basis functions, five of which are at a coarse resolution (blue lines) while eight are at a finer resolution (green lines). (Bottom) Hidden process $Y(s)$ for $s \in [0, 1]$ (thin black line), data Z (black dots), optimal predictor $Y^*(s_0) \equiv \mathbb{E}(Y(s_0) | Z)$, for $s_0 \in [0, 1]$ (red dashed line) and the 90% pointwise prediction intervals based on the predictive standard deviation $\{\text{var}(Y(s_0) | Z)\}^{1/2}$, for $s_0 \in [0, 1]$ (grey shaded region).

for each prediction location $\mathbf{s}_0 \in D$. Note that prediction of $Y(\cdot)$ over subsets of D is not considered here. Recall from Equation 4 that the predictive distribution is given by $[Y(\mathbf{s}_0) | \mathbf{Z}] \propto [\mathbf{Z} | Y(\mathbf{s}_0)][Y(\mathbf{s}_0)]$, which is the product of two Gaussian distributions. By ‘completing the square,’ it is straightforward to see that $[Y(\mathbf{s}_0) | \mathbf{Z}]$ is also Gaussian. Computing the predictive mean and the predictive variance of $[Y(\mathbf{s}_0) | \mathbf{Z}]$ is perhaps most easily performed by considering the joint distribution,

$$\begin{bmatrix} Y(\mathbf{s}_0) \\ \mathbf{Z} \end{bmatrix} \sim \text{Gau} \left(\begin{bmatrix} \mathbf{x}(\mathbf{s}_0)^\top \boldsymbol{\beta} \\ \mathbf{X}\boldsymbol{\beta} \end{bmatrix}, \begin{bmatrix} c(\mathbf{s}_0) & \mathbf{c}_0^\top \\ \mathbf{c}_0 & \mathbf{C}_Z \end{bmatrix} \right), \quad (11)$$

where $c(\mathbf{s}_0) \equiv \text{var}(Y(\mathbf{s}_0))$, $\mathbf{c}_0^\top \equiv \text{cov}(Y(\mathbf{s}_0), \mathbf{Z}) = \text{cov}(Y(\mathbf{s}_0), \mathbf{Y})$, and $\mathbf{C}_Z \equiv \text{cov}(\mathbf{Z}, \mathbf{Z})$. The predictive mean vector and the predictive covariance matrix are obtainable in closed form; as a result,

$$Y(\mathbf{s}_0) | \mathbf{Z} \sim \text{Gau}(\mathbf{x}(\mathbf{s}_0)^\top \boldsymbol{\beta} + \mathbf{c}_0^\top \mathbf{C}_Z^{-1} (\mathbf{Z} - \mathbf{X}\boldsymbol{\beta}), c(\mathbf{s}_0) - \mathbf{c}_0^\top \mathbf{C}_Z^{-1} \mathbf{c}_0). \quad (12)$$

The optimal predictor of the hidden value $Y(\mathbf{s}_0)$ is $Y^*(\mathbf{s}_0) \equiv \mathbb{E}(Y(\mathbf{s}_0) | \mathbf{Z})$ given by Equation 12. That is,

$$Y^*(\mathbf{s}_0) = \mathbf{x}(\mathbf{s}_0)^\top \boldsymbol{\beta} + \mathbf{c}_0^\top \mathbf{C}_Z^{-1} (\mathbf{Z} - \mathbf{X}\boldsymbol{\beta}). \quad (13)$$

The $n \times n$ matrix, \mathbf{C}_Z^{-1} , appears in the expression for both the mean and the variance of the predictive distribution in Equation 12, and it represents a potential computational bottleneck when the sample size is large. Indeed, even in the classical stationary model, which assumes a stationary covariance function for $Y(\cdot)$, \mathbf{C}_Z generally does not have a special structure that can be exploited for computational savings, and the full cost of an $O(n^3)$ matrix inversion must be met. This renders traditional spatial-prediction techniques computationally infeasible for large data sets. However, there are considerable computational efficiencies afforded by spatial-basis-function models.

The basis-function model given by Equation 10 results in the following expressions for the quantities in Equations 11 and 12: $c(\mathbf{s}_0) = \boldsymbol{\phi}(\mathbf{s}_0)^\top \mathbf{K} \boldsymbol{\phi}(\mathbf{s}_0) + C_\delta(\mathbf{s}_0, \mathbf{s}_0)$, $\mathbf{c}_0^\top = \boldsymbol{\phi}(\mathbf{s}_0)^\top \mathbf{K} \boldsymbol{\Phi}^\top + \text{cov}(\delta(\mathbf{s}_0), \boldsymbol{\delta})$ and, after marginalisation,

$$\mathbf{C}_Z = \boldsymbol{\Phi} \mathbf{K} \boldsymbol{\Phi}^\top + \mathbf{C}_\delta + \mathbf{C}_\epsilon. \quad (14)$$

The optimal predictor $Y^*(\mathbf{s}_0)$ for the basis-function model is obtained by substituting these quantities into Equation 13.

Typically, \mathbf{C}_ϵ is a diagonal matrix since it corresponds to independent measurement error, while \mathbf{C}_δ is often assumed to be diagonal, corresponding to fine-scale variation that is almost spatially uncorrelated (see Ma & Kang, 2020, for an exception); here, for simplicity, we assume that \mathbf{C}_δ is diagonal. Now write $\boldsymbol{\xi} \equiv \boldsymbol{\delta} + \boldsymbol{\epsilon}$ and, because of the independence implied by the HM structure, $\mathbf{C}_\xi = \mathbf{C}_\delta + \mathbf{C}_\epsilon$, which is also a diagonal covariance matrix. Then Equation 14 may be written as

$$\mathbf{C}_Z = \boldsymbol{\Phi} \mathbf{K} \boldsymbol{\Phi}^\top + \mathbf{C}_\xi, \quad (15)$$

where \mathbf{C}_ξ is diagonal with positive entries. Applying the Sherman–Morrison–Woodbury matrix identity (Searle & Henderson, 1981) to Equation 15, we have

$$\mathbf{C}_Z^{-1} = \mathbf{C}_\xi^{-1} - \mathbf{C}_\xi^{-1} \boldsymbol{\Phi} (\mathbf{I} + \mathbf{K} \boldsymbol{\Phi}^\top \mathbf{C}_\xi^{-1} \boldsymbol{\Phi})^{-1} \mathbf{K} \boldsymbol{\Phi}^\top \mathbf{C}_\xi^{-1}. \quad (16)$$

Equation 16 involves inverting two matrices, an $r \times r$ matrix and the possibly-very-large but, importantly, easy-to-invert diagonal matrix \mathbf{C}_ξ . We have already remarked that fixing r to be a manageable number circumvents the so-called ‘big n ’ problem since the computational complexity of inverting \mathbf{C}_Z drops from $O(n^3)$ to $O(nr^2)$. Hence, for low-rank models where $r \ll n$, basis-function models afford scalable, $O(n)$, spatial predictions.

Marginalisation: The operation of ‘integrating out’ random variables, in this case the random effects, in a probability model.

Low-rank models: models where the number of basis functions is (usually substantially) less than the number of data points.

3.2 Basis-function representations in non-Gaussian spatial processes

Many spatial problems of interest deal with distinctly non-Gaussian data (e.g., counts, binary responses, extreme values), which give rise to non-Gaussian process models. In this subsection, we review basis-function models that account for non-Gaussianity.

3.2.1 Trans-Gaussian spatial processes

The most straightforward way to construct a non-Gaussian spatial process is via a transformation $g(\cdot)$ of a spatial Gaussian process. We refer to these as *trans-Gaussian spatial processes*. Consider a spatial process, $Y(\cdot) = \{Y(\mathbf{s}) : \mathbf{s} \in D\}$ and a monotonic non-linear transformation $g(\cdot)$; if the process $W(\cdot) \equiv g(Y(\cdot))$ is Gaussian, then $Y(\cdot) = g^{-1}(W(\cdot))$ is a (non-Gaussian) trans-Gaussian process.

Among the most popular transformations is the *Box-Cox transformation*, $g_{\text{BC}}(\cdot)$, originally applied to spatial models by De Oliveira et al. (1997). The transformation is given by

$$g_{\text{BC}}(Y(\cdot)) \equiv \begin{cases} \frac{Y(\cdot)^\lambda - 1}{\lambda}, & \lambda \neq 0, \\ \log Y(\cdot), & \lambda = 0, \end{cases}$$

where $Y(\cdot)$ has been standardised to have no units and, for a given λ , $Y(\cdot)^\lambda$ is well defined. Setting $\lambda = 0$ in the Box-Cox transformation yields a *lognormal spatial process* provided $Y(\cdot) > 0$, while setting $\lambda = 1$ yields a conventional Gaussian spatial process.

Several other transformations have been considered to date in the context of spatial-process modelling. These include the square-root transform (Johns et al., 2003); the power transform (Allcroft & Glasbey, 2003); and the Tukey g-h transform (Xu & Genton, 2017). There are other innovative ways to construct non-Gaussian processes: For example, the process $(W(\cdot) + \lambda|U|)/\sqrt{V}$, where $\lambda > 0$, $U \sim \text{Gau}(0, 1)$, $V \sim \text{Gamma}(\nu/2, \nu/2)$, and ν is a positive integer, is a skew- t spatial process, which has been used to model skewed processes with heavy tails (e.g., Tagle et al., 2020).

The spatial models presented in the literature given just above did not employ a basis-function representation, although there is no reason why they could not have done so. Indeed, a basis-function representation for $W(\cdot)$ is particularly attractive in this setting, since its dimension-reduction facilitates estimation and prediction for non-Gaussian processes, which is typically very computationally intensive. Zammit-Mangion et al. (2016) used a low-rank representation of $g_{\text{BC}}(Y(\cdot))$ to model and predict methane sources $Y(\cdot)$, which are non-negative and skewed. There, each basis function was a two-dimensional rectangular function, equal to 1 over a small area, and equal to 0 otherwise. That is, 0-1 basis functions were used to represent $Y(\cdot)$.

3.2.2 Spatial generalised linear models

The generalised linear model (GLM, Nelder & Wedderburn, 1972, McCullagh & Nelder, 1989) was originally used to model non-spatial, non-Gaussian data. The data model assumes that, given their respective mean responses and in some cases given a dispersion parameter, the observations are conditionally independent and come from probability distributions in the *exponential family*. The process model uses a monotonic, differentiable *link function* to model the transformed mean response in terms of a linear combination of covariates. The generalised linear mixed model (GLMM) extends the GLM by introducing random effects into the process model.

The *spatial GLMM* (Diggle et al., 1998) is defined as follows. Independently for

Exponential family: A set of probability distributions that include the Gaussian, Poisson, binomial, and gamma distributions. Members of this family cover a range of supports and incorporate a wide variety of mean-variance relationships.

data locations $\{\mathbf{s}_i : i = 1, \dots, n\}$,

$$Z(\mathbf{s}_i) | Y(\mathbf{s}_i), \gamma \sim \text{EF}(Y(\mathbf{s}_i), \gamma), \quad (17)$$

$$g(Y(\mathbf{s}_i)) = \mathbf{x}(\mathbf{s}_i)^\top \boldsymbol{\beta} + \eta(\mathbf{s}_i), \quad (18)$$

where $Y(\cdot)$ is the ‘mean’ parameter of the exponential family, EF corresponds to a probability distribution in the exponential family with dispersion parameter γ , $g(\cdot)$ is a link function, the regression coefficients $\boldsymbol{\beta}$ are fixed but unknown parameters associated with covariates $\mathbf{x}(\cdot)$, and $\eta(\cdot)$ is a mean-zero spatially dependent *random* process, typically a Gaussian process. Diggle et al. (1998) assumed a stationary Gaussian process for $\eta(\cdot)$ and used Markov chain Monte Carlo (MCMC) methods for inference and prediction.

MCMC: Markov chain Monte Carlo is a class of algorithms often used to sample from an intractable (posterior) probability distribution.

Since this seminal work, the spatial GLMM framework has been used extensively, and some authors have chosen to model the spatial process $\eta(\cdot)$ in Equation 18 with the basis-function statistical model

$$\eta(\cdot) = \sum_{j=1}^r \alpha_j \phi_j(\cdot) + \delta(\cdot).$$

In a remote-sensing, ‘big data’ application, Sengupta & Cressie (2013) showed the computational advantages of using the spatial GLMM framework; Equation 17 was modelled with a gamma distribution, $g(\cdot) = \log(\cdot)$, and $g(Y(\cdot))$ in Equation 18 was given by the basis-function model, Equation 6. The basis-function coefficients were modelled as a Gaussian random vector with mean zero and an unconstrained covariance matrix. Then Laplace approximations in an expectation-maximisation algorithm (see Section 3.3) resulted in maximum likelihood estimates of the unknown parameters. Finally, the estimates replaced the unknown parameters, and the empirical predictive distribution was generated using an MCMC algorithm.

Empirical predictive distribution: The predictive distribution of process values given the data, with parameter estimates replacing the unknown parameters.

Lee & Park (2020) took a scalable basis-function approach for modelling non-stationary, non-Gaussian spatial data using the spatial GLMM framework. First, a clustering algorithm partitioned the spatial domain D into disjoint subregions and then, for each subregion, a thin-plate-spline basis-function model was used independently of the other subregions. Finally, the global process was represented as a weighted sum of the local processes. However, they did not include fine-scale variation in their local basis-function models (i.e., they put $\delta(\cdot) \equiv 0$), and they modelled the basis-function coefficients as independent.

3.2.3 Spatial prediction

The quantity of interest is the predictive distribution of the hidden spatial process given the data, namely $[Y(\mathbf{s}_0) | \mathbf{Z}]$, where $\mathbf{s}_0 \in D$ is a given prediction location. Bayes’ Rule can be used to obtain this distribution (see Equation 4), which in this non-Gaussian setting is typically not available analytically. Samples from the predictive distribution may be obtained via computational algorithms such as MCMC or via analytic approximations such as the Laplace approximation (e.g., Sengupta et al., 2016). Similar computational efficiencies in sampling from the predictive distribution in the Gaussian case, presented in Section 3.1.2, are afforded by basis-function representations of non-Gaussian spatial processes. Software developments over the past decade, particularly in the open-source language R (R Core Team, 2021), have led to many options for spatial prediction of non-Gaussian processes using basis-function representations; see Section 5.1 for an overview of currently available software.

Laplace approximation: A method where a second-order Taylor-series approximation replaces the unobserved components of the complete-data likelihood function, with a multivariate Gaussian distribution, facilitating marginalisation.

3.3 Inference in a hierarchical spatial statistical model (HM)

In the subsections that follow, we describe two types of hierarchical model: One type (EHM) considers the parameters fixed, unknown, and to be estimated, and the other

type (BHM) puts a Bayesian prior distribution on the parameters and computes a posterior distribution of the unknown spatial process and the parameters. In both cases, inference on the unknown spatial process is obtained from the predictive distribution.

3.3.1 The EHM and the BHM

Thus far, the presence of parameters θ in the HM have been de-emphasised, but they are there! If the parameters are treated as fixed but unknown, Equation 2 is re-written as:

$$[Y(\cdot), \mathbf{Z} | \theta] = [\mathbf{Z} | Y(\cdot), \theta][Y(\cdot) | \theta]. \quad (19)$$

Further, define the likelihood as

$$L(\theta; \mathbf{Z}) \equiv [\mathbf{Z} | \theta] = \int [\mathbf{Z} | Y(\cdot), \theta][Y(\cdot) | \theta] dY(\cdot). \quad (20)$$

Let Θ denote the parameter space; then the *maximum likelihood estimator (MLE)* of θ is defined as $\hat{\theta} \equiv \arg \sup_{\theta \in \Theta} L(\theta; \mathbf{Z})$. Inference on θ proceeds by approximating $\text{var}(\hat{\theta})$ with the inverse of Fisher's information matrix. This approach to handling the unknown parameters by obtaining an estimate $\hat{\theta}$, yields what Cressie & Wikle (2011) called an *empirical hierarchical model (EHM)* given by:

$$\begin{aligned} \text{Data model:} & \quad [\mathbf{Z} | Y(\cdot), \hat{\theta}], \\ \text{Process model:} & \quad [Y(\cdot) | \hat{\theta}]. \end{aligned}$$

If, on the other hand, a probability distribution $[\theta]$ is used to encode uncertainty in θ , Equation 2 is rewritten as:

$$[Y(\cdot), \mathbf{Z}, \theta] = [\mathbf{Z} | Y(\cdot), \theta][Y(\cdot) | \theta][\theta]. \quad (21)$$

The distribution $[\theta]$ is sometimes referred to as the *parameter model* (Berliner, 1996) but, most commonly, it is simply referred to as the *parameter prior*. In this case, inferences on unknowns $Y(\cdot)$ and θ are carried out using the posterior distribution obtained from Bayes' Rule:

$$[Y(\cdot), \theta | \mathbf{Z}] = [Y(\cdot), \theta, \mathbf{Z}] / [\mathbf{Z}], \quad (22)$$

where the joint distribution on the right-hand side of Equation 22 is given by Equation 21, and $[\mathbf{Z}]$ is the 'normalising constant'. This approach to handling the unknown parameters θ yields the *Bayesian hierarchical model (BHM)* given by

$$\begin{aligned} \text{Data model:} & \quad [\mathbf{Z} | Y(\cdot), \theta], \\ \text{Process model:} & \quad [Y(\cdot) | \theta], \\ \text{Parameter model:} & \quad [\theta]. \end{aligned}$$

3.3.2 Inference in EHM and BHM

To help with the exposition in this subsection, we replace generic θ with $\theta \equiv \{\theta_D, \theta_P\}$, where the data model and the process model are written as $[\mathbf{Z} | Y(\cdot), \theta_D]$ and $[Y(\cdot) | \theta_P]$ to emphasise that different parameters control different levels of the HM.

In the EHM, inference on the unknown $Y(\cdot)$ is treated separately and differently from inference on the unknowns θ_D and θ_P . Recall that the likelihood is $L(\theta_D, \theta_P; \mathbf{Z}) \equiv [\mathbf{Z} | \theta_D, \theta_P]$ which, assuming θ is fixed but unknown, is also the difficult-to-evaluate normalising constant, but now viewed as a function of the parameters. The MLE could be obtained directly if there were ways to evaluate and maximise $[\mathbf{Z} | \theta_D, \theta_P]$; however, this is rarely possible when the HM is not Gau-Gau

Fisher's information matrix: The *score* vector is the derivative of the log-likelihood function with respect to the parameters. Fisher's information matrix is the variance of the score vector.

as in Section 3.1. One approach that circumvents this problem is the expectation-maximisation (E-M) algorithm (Dempster et al., 1977, McLachlan & Krishnan, 2007), which consists of an E-step followed by an M-step, and the steps are repeated in an iterative manner. Specifically, the l th iteration is:

$$\begin{aligned} \text{E-step:} & \text{ Calculate } \mathbb{E}(\log([\mathbf{Z} | Y(\cdot), \boldsymbol{\theta}_D][Y(\cdot) | \boldsymbol{\theta}_P]) | \mathbf{Z}, \hat{\boldsymbol{\theta}}^{(l-1)}) \equiv q(\boldsymbol{\theta} | \hat{\boldsymbol{\theta}}^{(l-1)}), \\ \text{M-step:} & \text{ Find the } \boldsymbol{\theta} \text{ that maximises } q(\boldsymbol{\theta} | \hat{\boldsymbol{\theta}}^{(l-1)}); \text{ call this } \hat{\boldsymbol{\theta}}^{(l)}. \end{aligned} \quad (23)$$

With a suitable starting value $\hat{\boldsymbol{\theta}}^{(0)}$ and a convergence criterion (usually depending on the closeness of successive parameter estimates and/or successive values of q in the iteration), the E-M algorithm yields an MLE, $\hat{\boldsymbol{\theta}} \equiv \{\hat{\boldsymbol{\theta}}_D, \hat{\boldsymbol{\theta}}_P\}$. It has been used to fit basis-function models in a spatial setting by Sengupta et al. (2016) and in a spatio-temporal setting by Dewar et al. (2009) and Katzfuss & Cressie (2011).

Now, from Bayes' Rule (see Equation 22),

$$[Y(\cdot) | \mathbf{Z}, \hat{\boldsymbol{\theta}}] \propto [\mathbf{Z} | Y(\cdot), \hat{\boldsymbol{\theta}}_D][Y(\cdot) | \hat{\boldsymbol{\theta}}_P]. \quad (24)$$

Using an optimality criterion of minimising the mean-squared prediction error, the best predictor of $Y(\mathbf{s}_0)$ is $\mathbb{E}(Y(\mathbf{s}_0) | \mathbf{Z}, \hat{\boldsymbol{\theta}})$ with uncertainty quantified by, for example, $\text{var}(Y(\mathbf{s}_0) | \mathbf{Z}, \hat{\boldsymbol{\theta}})$. This is implemented for spatial-basis-function models in, for example, Cressie & Johannesson (2008), Cressie & Kang (2010), Sengupta & Cressie (2013), and Zammit-Mangion & Cressie (2021). In the EHM, the variability due to $\hat{\boldsymbol{\theta}}$ is ignored; then closed-form formulas, approximations, or Monte Carlo sampling (e.g., MCMC) can be applied directly to find the predictive distribution with $\hat{\boldsymbol{\theta}}$ substituted for $\boldsymbol{\theta}$, as in Equation 24.

Stein (2014) gave results on estimation for rather specialised basis-function models. He found that maximum-likelihood estimators sometimes performed poorly, however his set-up was very different than that assumed in this review; see Section 2.2.2 for further details. He also discussed spatial prediction using basis-function models, although he did not use a HM and hence predicted $Z(\mathbf{s}_0)$ instead of $Y(\mathbf{s}_0)$. His conclusions about spatial prediction were more circumspect; he suggested a simulation study be done to compare kriging (i.e., spatial) predictors from basis-function models, with other spatial statistical predictors. Bradley et al. (2016a) used a validation approach to make such a comparison from small, medium, and large data sets, and the basis-function models performed very well, as they did under the common-task-framework design in Heaton et al. (2019).

All inference for the BHM is in terms of the posterior distribution given by Equation 22. However, the model, $Y(\cdot) = \mathbf{x}(\cdot)^\top \boldsymbol{\beta} + \boldsymbol{\phi}(\cdot)^\top \boldsymbol{\alpha} + \delta(\cdot)$, given by Equation 6 may have an identifiability problem when a prior distribution is put on $\boldsymbol{\beta}$, since $\boldsymbol{\alpha}$ is by default random. The noisy and incomplete data \mathbf{Z} could have trouble distinguishing fixed effects from random effects. In the context of lattice data observed on a lattice process, Griffith (2000) proposed using a class of basis functions orthogonal to the column space of \mathbf{X} (given in Equation 10). These are known as *Moran's I basis functions*, in reference to the functional form of Moran's I statistic (Cliff & Ord, 1981). This approach has been taken up by Hodges & Reich (2010), Paciorek (2010), Hughes & Haran (2013), and Bradley et al. (2015), and the use of these basis functions in Equations 9 and 10 has become known as restricted spatial regression (RSR). In the geostatistical context, which is central to this review, Hanks et al. (2015) show that RSR does not necessarily alleviate the confounding of the fixed and random effects in the BHM. Since $\boldsymbol{\beta}$ is fixed in the EHM, it is less susceptible to this problem.

As in the EHM, inference in the BHM is hindered by not knowing the normalising constant, here $[\mathbf{Z}]$. One popular approach that does not require its knowledge is MCMC; see, for example, Nychka et al. (2002), Paciorek (2007), Wikle (2010), Kang & Cressie (2011), Eidsvik et al. (2012), and Zammit-Mangion & Rougier (2020) for

the use of MCMC in a spatial-basis-function-model setting. Approximate Bayesian techniques are also used in basis-function settings, two of the most popular being integrated nested Laplace approximations (INLA; e.g., Rue et al., 2009) and variational Bayes (VB; e.g., Cseke et al., 2016). In contrast to the EHM, variability due to uncertainty in $\boldsymbol{\theta}$ is explicitly accounted for in a BHM.

4 Multivariate, warped, and spatio-temporal basis-function models

In Section 3, we focused on univariate Gaussian and non-Gaussian spatial processes with basis-function components. In this section, we present a generalisation of basis-function models to multivariate spatial processes, non-stationary univariate processes where the spatial domain is warped, and spatio-temporal processes.

4.1 Basis-function representations in multivariate spatial models

Multivariate spatial models give more realistic representations of a complex world, where multiple processes interact and do not behave independently from one another. Basis-function models in a multivariate context have been used for modelling interacting geophysical processes in Antarctica (Zammit-Mangion et al., 2015), and for modelling dependence along the vertical direction in CO₂ fields (Nguyen et al., 2017).

For illustration, we consider the mean-zero bivariate Gaussian process, $\mathbf{Y}(\cdot) \equiv \{Y(\mathbf{s}) : \mathbf{s} \in D\}$, where $\mathbf{Y}(\mathbf{s}) \equiv (Y_1(\mathbf{s}), Y_2(\mathbf{s}))^\top$ for $\mathbf{s} \in D$; then $Y_1(\cdot) \equiv \{Y_1(\mathbf{s}) : \mathbf{s} \in D\}$ is a mean-zero univariate spatial process, and likewise for $Y_2(\cdot)$. The covariance function for $Y_1(\cdot)$, $C_{11}(\cdot, \cdot)$, captures the within-process spatial dependence of $Y_1(\cdot)$, and likewise for $C_{22}(\cdot, \cdot)$. The concurrent between-process spatial dependence is given by $C_{12}(\mathbf{s}, \mathbf{s})$ for $\mathbf{s} \in D$, and the ‘‘spatially lagged’’ between-process spatial dependence is given by

$$C_{12}(\mathbf{s}, \mathbf{u}) \equiv \text{cov}(Y_1(\mathbf{s}), Y_2(\mathbf{u})), \quad \mathbf{s}, \mathbf{u} \in D. \quad (25)$$

From Equation 25, $C_{12}(\mathbf{s}, \mathbf{u}) = C_{21}(\mathbf{u}, \mathbf{s})$ for $\mathbf{s}, \mathbf{u} \in D$, although the special symmetry relation, $C_{12}(\mathbf{s}, \mathbf{u}) = C_{21}(\mathbf{s}, \mathbf{u})$ for $\mathbf{s}, \mathbf{u} \in D$, is not true in general.

One cannot use any set of functions $C_{11}(\cdot, \cdot)$, $C_{22}(\cdot, \cdot)$, $C_{12}(\cdot, \cdot)$, and $C_{21}(\cdot, \cdot)$ to model the within and between spatial-statistical dependence. Clearly, both $C_{11}(\cdot, \cdot)$ and $C_{22}(\cdot, \cdot)$ have to be positive-semidefinite, as explained in Section 2.2. However, the role of $C_{12}(\cdot, \cdot)$ and $C_{21}(\cdot, \cdot)$ is delicate, since they must bind the processes together in such a way that the matrix of covariance functions,

$$\begin{bmatrix} C_{11}(\cdot, \cdot) & C_{12}(\cdot, \cdot) \\ C_{21}(\cdot, \cdot) & C_{22}(\cdot, \cdot) \end{bmatrix}, \quad (26)$$

is positive-semidefinite (Genton & Kleiber, 2015, Cressie & Zammit-Mangion, 2016).

Now, modelling the multivariate spatial process $\mathbf{Y}(\cdot)$ using basis functions, we have

$$Y_i(\cdot) = \boldsymbol{\phi}_i(\cdot)^\top \boldsymbol{\alpha}_i + \delta_i(\cdot), \quad i = 1, 2, \quad (27)$$

where $\{\delta_1(\cdot), \delta_2(\cdot)\}$ are independent mean-zero Gaussian processes with covariance functions $\{C_{\delta_1}(\cdot, \cdot), C_{\delta_2}(\cdot, \cdot)\}$; $\boldsymbol{\alpha}_i \equiv (\alpha_{i1}, \dots, \alpha_{ir_i})^\top$ and $\boldsymbol{\phi}_i(\cdot) \equiv (\phi_{i1}(\cdot), \dots, \phi_{ir_i}(\cdot))^\top$ are r_i -dimensional vectors for $i = 1, 2$; and

$$\begin{bmatrix} \boldsymbol{\alpha}_1 \\ \boldsymbol{\alpha}_2 \end{bmatrix} \sim \text{Gau}(\mathbf{0}, \mathbf{K}), \text{ for } \mathbf{K} \equiv \begin{bmatrix} \mathbf{K}_{11} & \mathbf{K}_{12} \\ \mathbf{K}_{21} & \mathbf{K}_{22} \end{bmatrix}.$$

Multivariate spatial Gaussian process in D : A multivariate spatial process where all linear combinations of any of the processes at any of the spatial locations in D , are Gaussian.

Then the elements of the 2×2 matrix in Equation 26 are given by

$$C_{ij}(\mathbf{s}, \mathbf{u}) = \begin{cases} \phi_i(\mathbf{s})^\top \mathbf{K}_{ij} \phi_j(\mathbf{u}) + C_{\delta_i}(\mathbf{s}, \mathbf{u}), & i = j, \\ \phi_i(\mathbf{s})^\top \mathbf{K}_{ij} \phi_j(\mathbf{u}), & i \neq j, \end{cases} \quad (28)$$

for $\mathbf{s}, \mathbf{u} \in D$ and $i, j = 1, 2$. Notice that the matrix of covariance functions with its elements given by Equation 28, is positive-semidefinite when \mathbf{K} is positive-semidefinite. However, the cross-covariance functions $C_{12}(\cdot, \cdot)$ and $C_{21}(\cdot, \cdot)$ may not satisfy the special symmetry relation, $C_{12}(\mathbf{s}, \mathbf{u}) = C_{21}(\mathbf{s}, \mathbf{u})$, which is a strength of multivariate spatial-basis-function models.

Basis-function models for a univariate spatial process are always valid; see Section 2.2.2 where their positive-semidefiniteness is established. The same is true for multivariate basis-function models defined by Equation 27, as long as the matrix \mathbf{K} is positive-semidefinite. In contrast, the multivariate Matérn cross-covariance function requires special conditions on parameters in order for Equation 26 to be positive-semidefinite (Gneiting et al., 2010). Another strong feature of covariance functions derived from univariate basis-function models is that they are generally non-stationary. From Equation 28, this important feature carries over to multivariate basis-function models.

Recall that we have previously mentioned spatial factor analytic models (Christensen & Amemiya, 2002, 2003). These multivariate models look like basis-function models, but the factors are random spatial functions, and their coefficients are deterministic and must satisfy critical identifiability conditions; also see Section 4.3.2.

There is a *conditional approach to multivariate spatial modelling* based on the decomposition

$$[Y_1(\cdot), Y_2(\cdot)] = [Y_2(\cdot) \mid Y_1(\cdot)][Y_1(\cdot)], \quad (29)$$

where $[Y_2(\cdot) \mid Y_1(\cdot)]$ is shorthand for $[Y_2(\cdot) \mid \{Y_1(\mathbf{s}) : \mathbf{s} \in D\}]$; see Cressie & Zammit-Mangion (2016). The advantage of this conditional approach in Equation 29, in contrast to the joint approach in Equation 26, is the ease of building positive-semidefiniteness conditions into the bivariate spatial model, as it only requires specification of two valid covariance functions, namely $C_{11}(\mathbf{s}, \mathbf{u})$ and $C_{2|1}(\mathbf{s}, \mathbf{u}) \equiv \text{cov}(Y_2(\mathbf{s}), Y_2(\mathbf{u}) \mid Y_1(\cdot))$. For basis-function models given by Equation 27, the random coefficients α_1 and α_2 of the basis functions $\phi_1(\cdot)$ and $\phi_2(\cdot)$, respectively, have means equal to zero, and hence

$$\begin{aligned} \alpha_2 \mid \alpha_1 &\sim \text{Gau}(\mathbf{A}\alpha_1, \mathbf{K}_{2|1}), \\ \alpha_1 &\sim \text{Gau}(\mathbf{0}, \mathbf{K}_{11}), \end{aligned} \quad (30)$$

where \mathbf{A} is an $r_2 \times r_1$ matrix that describes the dependence of α_2 on α_1 , and $\mathbf{K}_{2|1} \equiv \text{cov}(\alpha_2, \alpha_2 \mid \alpha_1)$. Then the iterated-covariance relationship gives

$$\begin{aligned} \text{var}(\alpha_2) &\equiv \mathbf{K}_{22} = \mathbf{K}_{2|1} + \mathbf{A}\mathbf{K}_{11}\mathbf{A}^\top, \\ \text{cov}(\alpha_1, \alpha_2) &\equiv \mathbf{K}_{12} = \mathbf{K}_{11}\mathbf{A}^\top, \end{aligned}$$

and the joint $(r_1 + r_2) \times (r_1 + r_2)$ covariance matrix is

$$\mathbf{K} \equiv \begin{bmatrix} \mathbf{K}_{11} & \mathbf{K}_{12} \\ \mathbf{K}_{21} & \mathbf{K}_{22} \end{bmatrix} = \begin{bmatrix} \mathbf{K}_{11} & \mathbf{K}_{11}\mathbf{A}^\top \\ \mathbf{A}\mathbf{K}_{11} & \mathbf{K}_{2|1} + \mathbf{A}\mathbf{K}_{11}\mathbf{A}^\top \end{bmatrix}. \quad (31)$$

Hence, the conditional approach to building bivariate basis-function models in spatial statistics is to model $\{\mathbf{K}_{ij}\}$ in Equation 28 indirectly through three matrix parameters, \mathbf{K}_{11} ($r_1 \times r_1$, positive-semidefinite), $\mathbf{K}_{2|1}$ ($r_2 \times r_2$, positive-semidefinite), and \mathbf{A} ($r_2 \times r_1$, real matrix). Contrast this with the joint approach that requires three matrix parameters \mathbf{K}_{11} , \mathbf{K}_{22} , and \mathbf{K}_{12} , but they are constrained so that \mathbf{K} , the full $(r_1 + r_2) \times (r_1 + r_2)$ covariance matrix, is positive-semidefinite.

The Gaussian data model is

$$Z_i(\mathbf{s}_{ij}) = Y_i(\mathbf{s}_{ij}) + \epsilon_{ij}, \quad j = 1, \dots, n_i, \quad i = 1, 2, \quad (32)$$

where the measurement errors $\{\epsilon_{ij}\}$ are independent Gaussian random variables with mean zero. Optimal prediction of $Y_2(\cdot)$, say, given bivariate data $\{\mathbf{Z}_1, \mathbf{Z}_2\}$, will be more precise than the univariate prediction of $Y_2(\cdot)$ given just \mathbf{Z}_2 , since the bivariate basis-function model exploits the dependence between coefficients $\boldsymbol{\alpha}_2$ and $\boldsymbol{\alpha}_1$. If, in fact, $\mathbb{E}(\boldsymbol{\alpha}_2 | \boldsymbol{\alpha}_1) = \mathbf{0} = \mathbb{E}(\boldsymbol{\alpha}_2)$ (i.e., $\mathbf{A} = \mathbf{0}$), then the basis-function models for $Y_1(\cdot)$ and $Y_2(\cdot)$ are independent, and hence the extra data \mathbf{Z}_1 beyond the data \mathbf{Z}_2 does not improve prediction of $Y_2(\cdot)$.

4.2 Warping the spatial locations

We return to univariate basis-function models, but now the main aim is to model highly non-stationary fields. Using simple basis functions with local support induces sparsity in the basis-function matrix $\boldsymbol{\Phi}$, and using a simple parametric form for the covariance matrix \mathbf{K} (or the precision matrix \mathbf{K}^{-1}) can dramatically reduce the number of parameters. Exploiting this is central to what are commonly known as Fixed Rank Kriging (FRK) models (Zammit-Mangion & Cressie, 2021), LatticeKrig models (Nychka et al., 2015), and stochastic partial differential equation (SPDE) models (Lindgren et al., 2011); further discussion of these models is given in Section 5.1. However, as seen in **Figure 2**, this approach may not allow substantial spatial non-stationarity to be captured.

A straightforward way to model highly non-stationary fields with basis-function models is first to *warp* the domain and then to construct a spatial-basis-function model with \mathbf{K} having a simple structure on the warped domain. In the seminal work of Sampson & Guttorp (1992), warping was done using multi-dimensional scaling and thin-plate splines. The central idea behind warping is as follows: Suppose that a spatial process has a highly non-stationary covariance function $C_{D_G}(\mathbf{s}, \mathbf{u})$ on the original, *geographic*, domain D_G . Then assume the existence of a warping function $\mathbf{f}(\cdot) : D_G \rightarrow D_W$ that maps from the geographic domain D_G to the *warped* (or deformed) domain D_W , such that $C_{D_G}(\mathbf{s}, \mathbf{u}) = C_{D_W}(\mathbf{f}(\mathbf{s}), \mathbf{f}(\mathbf{u}))$, where typically $C_{D_W}(\cdot, \cdot)$ is a stationary, isotropic, covariance function of known parametric form on the warped domain D_W .

There are several ways in which one can warp space, but we focus here on basis-function models. Note that spatial warping functions are necessarily multivariate since they have a multi-dimensional output (e.g., two spatial coordinates in \mathbb{R}^2). To the best of our knowledge, the earliest basis-function warping models were due to Smith (1996), who modelled $\mathbf{f}(\cdot)$ using radial basis functions, and Perrin & Monestiez (1999), who considered multiple radial basis functions linked through composition. More recently, a class of machine-learning models has emerged, known as manifold Gaussian processes, where the warping is carried out using a deep neural network. While neural networks can be treated as basis-function models, they tend to exhibit pathologies in applications in low dimensions, such as in spatial settings (e.g., Dunlop et al., 2018). The *deep compositional spatial model* of Zammit-Mangion et al. (2021b) circumvents this by forcing the mapping in each layer of the deep neural network to be injective, which precludes the possibility of ‘space-folding’, where space is mapped onto itself and can create undesirable spatial-prediction artefacts (Schmidt & O’Hagan, 2003). These deep models bear strong connections to deep Gaussian processes (e.g., Damianou & Lawrence, 2013, Hensman & Lawrence, 2014), where the warping layers are modelled as (typically) low-rank Gaussian processes.

To show the potential benefit of this type of warping using injective mappings with a basis-function component, consider elevation data from the CryoSat-2 instrument over the Northern Antarctic Peninsula. Elevation data is important when assessing the

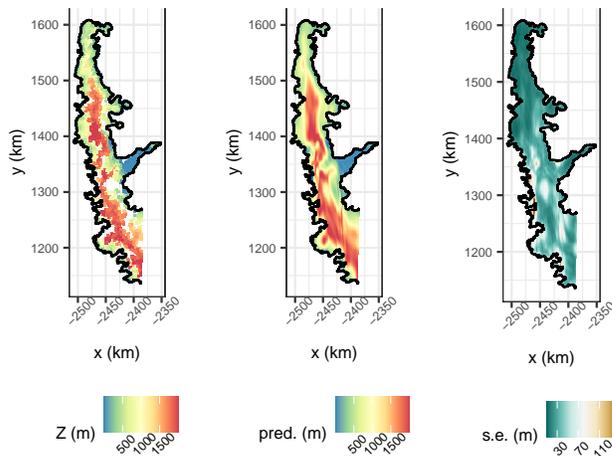


Figure 4: Spatial prediction using basis-function models of elevation (in metres) in the Northern Antarctic Peninsula. (First panel from the left) CryoSat-2 data on elevation between 2010 and 2017. (Second panel) Prediction using `deepspat`. (Third panel) Prediction standard errors obtained using `deepspat`. Only data and predictions for the Antarctic mainland (black line) are shown.

behaviour of an ice sheet, as subsequent losses or gains in observed elevation generally indicate ice-sheet-mass loss or gain, which leads to sea-level rise or fall, respectively. In **Figure 4**, first panel, we show incomplete and noisy CryoSat-2 elevation data on this part of the Antarctic mainland between 2010 and 2017. The second panel in **Figure 4** shows predicted elevation over the whole Northern Antarctic Peninsula using the R package `deepspat`¹ (Zammit-Mangion et al., 2021b), which implements the deep compositional spatial model with a basis-function component. The third panel shows the prediction standard errors. Note how the predictions and prediction standard errors display a certain ‘distortion’ that is not typically obtainable from spatial statistical models whose covariance functions on the geographic domain approximate a stationary covariance function. Further examples are given in Zammit-Mangion et al. (2021b).

4.3 Basis-function models in spatio-temporal statistics

In this subsection, we briefly review basis-function models in spatio-temporal statistics. Recall **Figure 1**, which shows the output of the transport model MOZART in response to a pulse of CO₂ flux in Northern Africa throughout the month of March 2015; in fact, this figure shows a spatio-temporal basis function. We invoke a spatio-temporal HM with data and process indexed by \mathbf{s} (space) and t (time). Without loss of generality, we assume that the data are collected at times $t \in \{1, \dots, T\}$ and, at each time point, we have n_t measurements collected across space; define

$$\mathbf{Z}_t \equiv (Z(\mathbf{s}_{t1}), \dots, Z(\mathbf{s}_{tn_t}))^\top, \quad t = 1, \dots, T.$$

The full data vector is then $\mathbf{Z} \equiv (\mathbf{Z}_1^\top, \dots, \mathbf{Z}_T^\top)^\top$. Again we focus on the Gau-Gau case, whereby both the data model and the underlying process are Gaussian, and we simply note that the modelling of non-Gaussian spatio-temporal processes follows in an analogous fashion to that given in Section 3.2 (e.g., Wikle et al., 2019, Sec. 4.5).

In the following subsections, we present a *descriptive* approach (Section 4.3.1) and a *dynamic* approach (Section 4.3.2) to spatio-temporal statistics, through the lens of

¹<https://github.com/andrewzm/deepspat>

spatio-temporal-basis-function models.

4.3.1 Descriptive (joint) approach

The *descriptive approach* focuses on describing spatio-temporal dependence in a process $\{Y(\mathbf{s}, t) : \mathbf{s} \in D, t = 1, 2, \dots\}$ through its mean function and its covariance function. The mean function is $\mu(\mathbf{s}, t) \equiv \mathbb{E}(Y(\mathbf{s}, t))$, which is often modelled as $\mathbf{x}(\mathbf{s}, t)^\top \boldsymbol{\beta}$ for spatio-temporal covariates $\mathbf{x}(\mathbf{s}, t)$. The covariance function is $C(\mathbf{s}, \mathbf{s}^*; t, t^*) \equiv \text{cov}(Y(\mathbf{s}, t), Y(\mathbf{s}^*, t^*))$. The challenges associated with constructing valid, realistic spatial covariance functions are exacerbated in the spatio-temporal setting, since time is not just another dimension and needs to be treated differently to the spatial dimensions. See Wikle et al. (2019, Ch. 4) for a lengthier review than the one given below.

Classical spatio-temporal models assume *second-order* (or *weak*) stationarity in space and time. A further classical assumption that is often made is *separability* between space and time. In the case of regular spatio-temporal data, where the data are observed at the same n spatial locations at each time point, separability leads to significant computational savings, since the nT -dimensional covariance matrix, $\mathbf{C}_Z \equiv \text{cov}(\mathbf{Z}, \mathbf{Z})$, can be written as a Kronecker product $\mathbf{C}_Z^{(s)} \otimes \mathbf{C}_Z^{(t)}$ of an n -dimensional covariance matrix $\mathbf{C}_Z^{(s)}$ and a T -dimensional covariance matrix $\mathbf{C}_Z^{(t)}$, facilitating computation of \mathbf{C}_Z^{-1} .

Basis-function models serve as an important tool for constructing valid, non-stationary, non-separable, non-fully symmetric, spatio-temporal covariance functions. In the spatial case, one uses known spatial basis functions, with associated random coefficients. In a spatio-temporal setting, one may use spatial basis functions with temporally indexed random coefficients (e.g., Wikle & Cressie, 1999), or temporal basis functions with spatially indexed random coefficients (e.g., Sansó et al., 2008), or simply spatio-temporal basis functions with random coefficients that are not indexed spatially or temporally (e.g., Zammit-Mangion & Cressie, 2021). In this review, we present this latter case but note that similar comments can be made for the other two. In the descriptive approach, the extension of Equation 6 to a spatio-temporal setting is, for $\mathbf{s} \in D$ and $t = 1, 2, \dots$,

$$Y(\mathbf{s}, t) = \mathbf{x}(\mathbf{s}, t)^\top \boldsymbol{\beta} + \boldsymbol{\phi}(\mathbf{s}, t)^\top \boldsymbol{\alpha} + \delta(\mathbf{s}, t), \quad (33)$$

where each term is defined analogously to its spatial counterpart; in particular, $\boldsymbol{\phi}(\cdot, \cdot) \equiv (\phi_1(\cdot, \cdot), \dots, \phi_r(\cdot, \cdot))^\top$ are spatio-temporal basis functions with associated random coefficients $\boldsymbol{\alpha}$, and $\delta(\cdot, \cdot)$ represents a spatio-temporal error process. From Equation 33, the covariance function of $Y(\cdot, \cdot)$ is, for $\mathbf{s}, \mathbf{s}^* \in D$ and $t, t^* = 1, 2, \dots$,

$$C(\mathbf{s}, \mathbf{s}^*; t, t^*) = \boldsymbol{\phi}(\mathbf{s}, t)^\top \mathbf{K} \boldsymbol{\phi}(\mathbf{s}^*, t^*) + \text{cov}(\delta(\mathbf{s}, t), \delta(\mathbf{s}^*, t^*))$$

which, in general, is non-stationary and non-separable.

Now consider Equation 33 evaluated over the location-time pairs where observations occur. In matrix-vector form, we have, $\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\Phi}\boldsymbol{\alpha} + \boldsymbol{\delta}$, where each term is defined analogously to its spatial counterpart; see Equation 10. In a Gau-Gau model, the data model is, $\mathbf{Z} = \mathbf{Y} + \boldsymbol{\epsilon}$, where $\boldsymbol{\epsilon}$ denotes independent Gaussian error; see Equation 9. The covariance matrix of the data is, $\mathbf{C}_Z = \boldsymbol{\Phi}\mathbf{K}\boldsymbol{\Phi}^\top + \mathbf{C}_\epsilon$, where recall that $\mathbf{C}_\epsilon \equiv \text{cov}(\boldsymbol{\delta}, \boldsymbol{\delta}) + \text{cov}(\boldsymbol{\epsilon}, \boldsymbol{\epsilon})$ (typically assumed to be a diagonal matrix). In the low-rank case, the Sherman–Morrison–Woodbury matrix identity facilitates computation of the critically important matrix inverse \mathbf{C}_Z^{-1} present in all inferences, without the need for separability and a Kronecker product.

Choosing spatio-temporal basis functions is a task that is typically more difficult, and more consequential, than choosing spatial basis functions. One straightforward approach is through the tensor product of a set of spatial basis functions and a set

Stationarity: A spatio-temporal covariance function is stationary if it can be written in terms of spatial and temporal lags, that is, $C(\mathbf{s}, \mathbf{s}^*; t, t^*) = C^*(\mathbf{h}; \tau)$, where $\mathbf{h} \equiv \mathbf{s}^* - \mathbf{s}$ and $\tau \equiv t^* - t$.

Separability: A stationary spatio-temporal covariance function is space-time separable if it can be written as $C^*(\mathbf{h}, \tau) = C^{(s)}(\mathbf{h})C^{(t)}(\tau)$. It is guaranteed to be valid (i.e., positive-semidefinite) provided both the spatial covariance function, $C^{(s)}(\cdot)$, and the temporal covariance function, $C^{(t)}(\cdot)$, are valid.

Kronecker product: The Kronecker product, $\mathbf{U} \otimes \mathbf{V}$, has (i, j) th block, $u_{ij}\mathbf{V}$. Then $(\mathbf{U} \otimes \mathbf{V})^{-1} = \mathbf{U}^{-1} \otimes \mathbf{V}^{-1}$.

Full symmetry: A spatio-temporal covariance function that satisfies $C(\mathbf{s}, \mathbf{s}^*; t, t^*) = C(\mathbf{s}, \mathbf{s}^*; t^*, t)$ is said to be fully symmetric. All separable covariance functions are fully symmetric; the converse is not true.

Tensor product:
The tensor product of (a_1, \dots, a_k) and (b_1, \dots, b_l) is $(a_1 b_1, \dots, a_1 b_l, \dots, a_k b_1, \dots, a_k b_l)$.

of temporal basis functions (e.g., Zammit-Mangion & Cressie, 2021); importantly, the process model does not exhibit the restrictive separability property between space and time when the tensor product is used.

4.3.2 Dynamic (conditional) approach

In the *dynamic approach*, the extension of Equation 6 to a (discrete-time) spatio-temporal setting is

$$Y_t(\mathbf{s}) = \mathbf{x}_t(\mathbf{s})^\top \boldsymbol{\beta} + \boldsymbol{\phi}(\mathbf{s})^\top \boldsymbol{\alpha}_t + \delta_t(\mathbf{s}), \quad \mathbf{s} \in D, \quad t = 1, 2, \dots, \quad (34)$$

where the different notation, $Y_t(\mathbf{s}) \equiv Y(\mathbf{s}, t)$, is used to reflect a spatial process evolving in time $t \in \{1, \dots, T\}$; $\boldsymbol{\phi}(\cdot)$ are known spatial basis functions associated with *dynamically evolving coefficients* $\boldsymbol{\alpha}_t$; and $\delta_t(\cdot)$ is an error process that is independent in time. Typically, a first-order Markov assumption on $\{\boldsymbol{\alpha}_t : t = 1, \dots, T\}$ is made, in which case the basis-function coefficients evolve according to

$$\boldsymbol{\alpha}_t = \mathbf{M}\boldsymbol{\alpha}_{t-1} + \boldsymbol{\omega}_t, \quad \text{for } t = 2, \dots, T, \quad (35)$$

where \mathbf{M} is referred to as the transition or propagator matrix; $\boldsymbol{\omega}_t \sim \text{Gau}(\mathbf{0}, \mathbf{C}_\omega)$ is independent of $\boldsymbol{\alpha}_{t-1}$ and independent in time; and \mathbf{C}_ω is referred to as the innovation matrix. Typically, \mathbf{M} is chosen to be non-normal (i.e., $\mathbf{M}^\top \mathbf{M} \neq \mathbf{M}\mathbf{M}^\top$), so that the model may capture the phenomenon known as *transient growth* (e.g., Wikle et al., 2019, p. 223).

Transient growth: A process that is globally stable, but may have local periods of explosive growth, is said to exhibit transient growth.

Now consider the process model evaluated over the location-time pairs where observations occur. In matrix-vector form, we have the process model,

$$\mathbf{Y}_t = \mathbf{X}_t \boldsymbol{\beta} + \boldsymbol{\Phi} \boldsymbol{\alpha}_t + \boldsymbol{\delta}_t, \quad t = 1, \dots, T, \quad (36)$$

where each term is defined analogously to its spatial counterpart, and the coefficients $\{\boldsymbol{\alpha}_t\}$ evolve dynamically according to Equation 35. In a Gau-Gau model, the data model is,

$$\mathbf{Z}_t = \mathbf{Y}_t + \boldsymbol{\epsilon}_t, \quad t = 1, \dots, T, \quad (37)$$

where $\boldsymbol{\epsilon}_t$ denotes Gaussian measurement error. Various generalisations of Equation 37 are possible; for instance, one may relate \mathbf{Z}_t and \mathbf{Y}_t via a linear or non-linear mapping, such as discussed in Cressie & Wikle (2011, Sec. 7.1). Basis functions have often been used in dynamic spatio-temporal models, such as in Wikle & Cressie (1999), Wikle et al. (2001), Stroud et al. (2001), Sahu & Mardia (2005), Sansó et al. (2008), Dewar et al. (2009), Cressie et al. (2010), Freestone et al. (2011), Cressie & Wikle (2011), Katzfuss & Cressie (2011), Zammit-Mangion et al. (2012), Bradley et al. (2018), and Bradley et al. (2019). See Wikle et al. (2019, Ch. 5) for a longer review than the one given above.

In this review of spatio-temporal-basis-function models, we have focused on spatial models with known basis functions and random coefficients evolving in time. However, as noted in Section 2.2.1, there is another type of spatial-basis-function model that considers the basis functions to be random and their coefficients fixed. In the spatio-temporal setting, this type of model has been used for analysing Gaussian data (Lopes et al., 2008) and, using the GLMM framework, non-Gaussian data (Lopes et al., 2011).

5 Software with an application

This section contains two subsections. In Section 5.1, we discuss some of the spatial-statistical software currently available for fitting and predicting with basis-function models. In Section 5.2, we give a geophysical application where six software packages are used for analysing the same sea-surface temperature data set.

5.1 Software for spatial-basis-function modelling

In this discussion of software that is available as of 2021 for spatial statistics with basis-function models, we focus primarily on implementations in the R language (R Core Team, 2021).

FRK (Zammit-Mangion & Cressie, 2021) is a user-friendly R package for prediction of large spatial and spatio-temporal data sets. In the spatial setting, the process model it uses is given by Equation 6, with $\delta(\cdot)$ included. A recent version of **FRK** (Sainsbury-Dale et al., 2021) caters for both Gaussian and non-Gaussian data models by employing a basis-function representation of the spatial GLMM; users can specify their own basis functions, and the default is the multi-resolution bisquare basis functions. **FRK** uses spatial/spatio-temporal basic areal units (BAUs) to handle differing data supports (both point-referenced and areal) and any user-specified prediction regions.

INLA (Lindgren & Rue, 2015) is a general-purpose R package for approximate inference with latent Gaussian models. From a spatial-statistics point of view, **INLA** is geared to efficiently model and fit random-fields that are solutions to SPDEs (Lindgren et al., 2011). Specifically, the package uses finite elements (which are basis functions) to approximate the solutions to these SPDEs. **INLA** is able to cater for a number of non-Gaussian distributions, and it can accommodate different forms of geographically referenced data, as well as spatio-temporal data. **INLA** is a general-purpose package, not specifically designed with spatial prediction in mind; a recently developed package aimed at facilitating spatial modelling and fitting with **INLA** is **inlabru** (Bachl et al., 2019).

LatticeKrig (Nychka et al., 2015) is a user-friendly R package for spatial prediction with large spatial data sets. It uses Wendland basis functions (that have compact support) and a Markov assumption to construct a sparse precision matrix that describes the dependence between the coefficients of the basis functions. This results in efficient computations and potentially the use of a large number of basis functions. **LatticeKrig** is limited to spatial, point-referenced, Gaussian data. Its model puts $\delta(\cdot) = 0$ in Equation 6, which means that it omits fine-scale spatial variation. Consequently, the finest scale of the process is limited to the finest resolution of the basis functions used.

mgcv (Wood, 2017) can be used to fit and predict with spatial or spatio-temporal data using generalised additive models (GAMs), which rely on constructing smooth functions of the covariates. In a spatial setting, the covariates include spatial location. The software typically yields excellent computational times, and it can cater for a wide range of types of non-Gaussian data. The package is designed to handle point-referenced data.

The multi-resolitional approximation (MRA) of Katzfuss (2017) is a scalable approach based on a tree-structured partitioning of the spatial domain D into blocks upon which predictive-process basis functions are defined. Huang et al. (2002) have a similar basis-function model but with 0-1 basis functions. Both exhibit spatial ‘blockiness’; see Tzeng et al. (2005) for a way to remove this. An advantage of these approaches is that they allow for easy parallelisation, which is not necessarily true for other dimension-reducing approaches. Spatial prediction using MRA has been implemented efficiently in C++ by Huang et al. (2019).

5.2 Application to mapping sea-surface temperatures

We now illustrate use of these software packages by way of a comparative study. As a benchmark, we used traditional kriging, implemented with the R package **gstat** (Pebesma, 2004). Our comparison uses the data analysed by Zammit-Mangion & Rougier (2020), which consists of global sea-surface temperature (SST) data obtained from the Visible Infrared Imaging Radiometer Suite (VIIRS) on board the Suomi National Polar-orbiting Partnership (Suomi NPP) weather satellite (Cao et al., 2013).

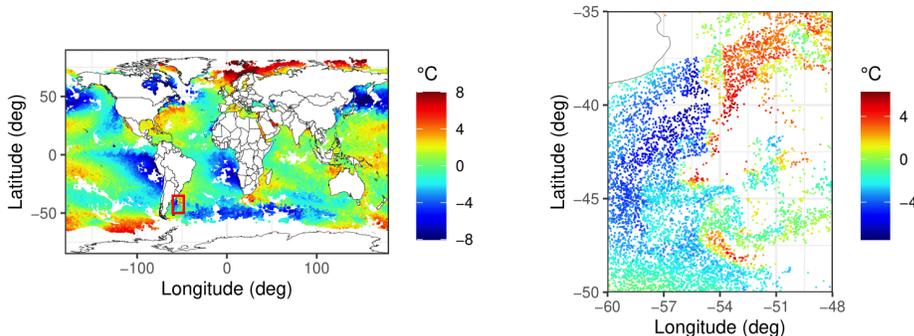


Figure 5: SST residuals used for training the models. Left panel: Global view of residuals, with the $12^\circ \times 15^\circ$ box corresponding to the Brazil-Malvinas confluence zone marked in red. Right panel: Residuals in the Brazil-Malvinas confluence zone (note the different colour scale). Note that the testing data set is not shown, but is similar to the training set in terms of regions of data paucity and trend.

As in Zammit-Mangion & Rougier (2020), spatial modelling was carried out on the residuals from a linear model with covariates given by an intercept, the latitude coordinate, and the square of the latitude coordinate. This detrended data set is shown in **Figure 5**, left panel. In this study, we focus on a region of the ocean known as the Brazil-Malvinas confluence zone; see **Figure 5**, right panel. We divide the data within this region into a training and testing data set, each consisting of approximately 8000 observations.

Brazil-Malvinas confluence zone:

An energetic region of the ocean just off the coast of Argentina and Uruguay, named after the two currents that meet at this location, namely the warm Brazil current and the cold Malvinas current.

The packages used in this study required several modelling decisions, which had to be made in a way that balanced predictive performance and run time. We took a systematic approach to the model-selection phase by splitting the training data set in two, using half for model fitting and the other half for model validation. In this way, we were able to experiment with the large number of arguments required by each package and choose the combination that was the best in terms of predictive performance and run time. Typically, the most important choice for basis-function representations is the number of basis functions to use: for **FRK**, we used a total of 9210 bisquare basis functions at 4 resolutions; for **INLA**, we used 11,857 basis functions; for **LatticeKrig**, we used a total of 30,482 Wendland basis functions at 4 resolutions; for **mgcv**, we used 2250 knot locations. (In **mgcv**, we used the `bam()` function, which is similar to the generalised additive model function `gam()` but optimised for large data sets.) For **MRA**, we used 2 subregions in each partitioning, 49 knots in each subregion before the finest resolution, and we let the software determine an appropriate number of resolutions based on the data. For **gstat**, we used the popular stationary Matérn covariance function, with parameters estimated by fitting the Matérn variogram to the empirical variogram. The Matérn covariance function is $C_\nu(\mathbf{h}) = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\sqrt{2\nu} \frac{\|\mathbf{h}\|}{\rho} \right) \times K_\nu \left(\sqrt{2\nu} \frac{\|\mathbf{h}\|}{\rho} \right)$, where $\Gamma(\cdot)$ is the gamma function, $K_\nu(\cdot)$ is the Bessel function of the second kind of order ν , and ρ and ν are the so-called range and smoothness parameters, respectively.

Table 1 summarises the results on the testing data set. The results are similar across all methods; this is reinforced by **Figure 6**, which shows predictions and prediction standard errors over the entire domain of interest. In unobserved regions, the prediction standard errors from all packages are similar, apart from those from **mgcv**, which are smaller.

Interestingly, all methods perform similarly to **gstat** on this example, which is using simple kriging for prediction of the residual process. Indeed, on data sets of sizes 10,000 or less, there is often little benefit in using basis functions for spatial

Table 1: Comparison of spatial predictions of SST residuals in the Brazil-Malvinas confluence zone using the five R packages `FRK`, `INLA`, `LatticeKrig`, `mgcv`, and `gstat`, and using a C++ implementation of MRA provided by Huang et al. (2019). The diagnostics are the root mean squared prediction error (RMSPE), the empirical coverage (COV90) and the interval score (IS90) from a prediction interval with a nominal coverage of 0.9, the continuous ranked probability score (CRPS), and the total run time needed to conduct inference. Note that we omit the run time for the MRA software as it was run on a different computer than the other methods.

Method	RMSPE	COV90	IS90	CRPS	Run Time (Minutes)
FRK	0.46	0.88	2.21	0.23	8.10
INLA	0.46	0.93	2.18	0.24	0.86
LatticeKrig	0.45	0.91	2.09	0.23	9.63
mgcv	0.45	0.90	2.12	0.23	5.73
MRA	0.44	0.92	2.06	0.23	*
gstat	0.44	0.91	2.07	0.22	7.56

prediction over kriging, unless the basis functions are specifically chosen to reproduce a non-stationary feature in the data. However, there are still benefits to the basis-function approach; in particular, if the sample size had been an order of magnitude larger, `gstat` would not have been computationally feasible (unless local kriging were used).

Implementing from scratch the methods used in this study would have been an arduous undertaking. Fortunately, thanks to the open-source software R and the many contributions from its community, our implementation of these complicated methods involved only a few lines of code. For the code used in this study, please see https://github.com/msainsburydale/ARSIA_BasisFunctionModels_code.

6 Epilogue

This review of basis-function models in spatial statistics has highlighted the reach and potential these models have for analysing spatial and spatio-temporal data. With technological advances in data collection, storage, and management, and with a realisation that the ‘Where’ (and ‘When’) question is highly relevant to answering important ‘Why’ questions, these non-stationary and computationally scalable spatial statistical models offer a path forward. We have emphasised here the importance of not only producing a parameter estimate or a spatial prediction, but also of quantifying uncertainties as part of these inferences. Sections 2–5 establish basis-function modelling in specific areas of spatial statistics but, in principle, any area of spatial statistics that involves spatial-covariance modelling could straightforwardly incorporate basis-function models.

ACKNOWLEDGEMENTS

Noel Cressie’s and Andrew Zammit-Mangion’s research was supported by an Australian Research Council (ARC) Discovery Project, DP190100180. Andrew Zammit-Mangion’s research was also supported by an ARC Discovery Early Career Research Award, DE180100203. Matthew Sainsbury-Dale’s research was supported by an Australian Government Research Training Program Scholarship. The authors would like to thank Ann Stavert for generating the MOZART output shown in **Figure 1**; Stephen Chuter and Geoffrey Dawson for providing the processed CryoSat-2 data, originally

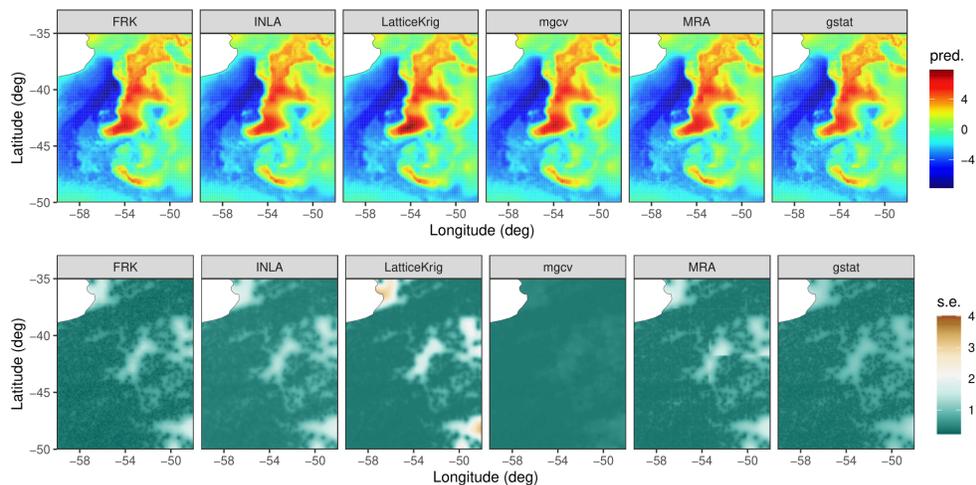


Figure 6: Spatial prediction using kriging and basis-function models of SST surrounding the Brazil–Malvinas confluence zone. The first row contains the spatial predictions, while the second contains the prediction standard errors. The first column corresponds to the results obtained using FRK; the second to INLA; the third to `LatticeKrig`; the fourth to `mgcv`; the fifth to MRA; and the sixth to `gstat`. The predictions and standard errors are similar for all methods. Note that blocks are visible if one “zooms in” on the MRA predictions or standard errors when viewing the pdf form of this article; this “blockiness” is a well known feature of multi-resolutional approximations.

made available by the European Space Agency, which was used in the example of Section 4.2; Yuliya Marchetti for providing the sea-surface temperature data set in the case study of Section 5.2; and Yi Cao for running the MRA software for the example of Section 5.2. We are also grateful to an anonymous reviewer whose suggestions enhanced the presentation of our article.

References

- Allcroft DJ, Glasbey CA. 2003. A latent Gaussian Markov random-field model for spatiotemporal rainfall disaggregation. *Journal of the Royal Statistical Society: Series C* 52:487–498
- Bachl FE, Lindgren F, Borchers DL, Illian JB. 2019. `inlabru`: an R package for Bayesian spatial modelling from ecological survey data. *Methods in Ecology and Evolution* 10:760–766
- Banerjee S, Carlin BP, Gelfand AE. 2004. Hierarchical Modeling and Analysis for Spatial Data, 2nd ed. London, UK: Chapman and Hall
- Banerjee S, Gelfand AE, Finley AO, Sang H. 2008. Gaussian predictive process models for large spatial data sets. *Journal of the Royal Statistical Society: Series B* 70:825–848
- Berliner LM. 1996. Hierarchical Bayesian time series models. In *Proceedings of the XVth Workshop on Maximum Entropy and Bayesian Methods*. Dordrecht, NL: Kluwer Academic Publishers, 15–22
- Bradley J, Holan S, Wikle C. 2015. Multivariate spatio-temporal models for high-dimensional areal data with application to longitudinal employer-household dynamics. *Annals of Applied Statistics* 9:1761–1791

- Bradley JR, Cressie N, Shi T. 2016a. A comparison of spatial predictors when datasets could be very large. *Statistics Surveys* 10:100–131
- Bradley JR, Holan SH, Wikle CK. 2018. Computationally efficient multivariate spatio-temporal models for high-dimensional count-valued data (with discussion). *Bayesian Analysis* 13:253–310
- Bradley JR, Holan SH, Wikle CK. 2020. Bayesian hierarchical models with conjugate full-conditional distributions for dependent data from the natural exponential family. *Journal of the American Statistical Association* 115:2037–2052
- Bradley JR, Wikle CK, Holan SH. 2016b. Bayesian spatial change of support for count-valued survey data with application to the American Community Survey. *Journal of the American Statistical Association* 111:472–487
- Bradley JR, Wikle CK, Holan SH. 2019. Spatio-temporal models for big multinomial data using the conditional multivariate logit beta distribution. *Journal of Time Series Analysis* 50:363–382
- Cao C, Xiong J, Blonski S, Liu Q, Uprety S, et al. 2013. Suomi NPP VIIRS sensor data record verification, validation, and long-term performance monitoring. *Journal of Geophysical Research: Atmospheres* 118:11–664
- Christensen WF, Amemiya Y. 2002. Latent variable analysis of multivariate spatial data. *Journal of the American Statistical Association* 97:302–317
- Christensen WF, Amemiya Y. 2003. Modeling and prediction for multivariate spatial factor analysis. *Journal of Statistical Planning and Inference* 115:543–564
- Cliff A, Ord J. 1981. *Spatial Processes – Models and Applications*. London, UK: Pion
- Cressie N. 1993. *Statistics for Spatial Data*, rev. ed. Hoboken, NJ: Wiley
- Cressie N, 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, Kang EL. 2010. High-resolution digital soil mapping: kriging for very large datasets. In *Proximal Soil Sensing*, eds. RA Viscarra Rossel, AB McBratney, B Minasny. New York, NY: Springer-Verlag, 49–63
- Cressie N, Shi T, Kang EL. 2010. Fixed rank filtering for spatio-temporal data. *Journal of Computational and Graphical Statistics* 19:724–745
- Cressie N, Wikle CK. 2011. *Statistics for Spatio-Temporal Data*. Hoboken, NJ: Wiley
- Cressie N, Zammit-Mangion A. 2016. Multivariate spatial covariance models: a conditional approach. *Biometrika* 103:915–935
- Cseke B, Zammit-Mangion A, Heskens T, Sanguinetti G. 2016. Sparse approximate inference for spatio-temporal point process models. *Journal of the American Statistical Association* 111:1746–1763
- Damianou A, Lawrence N. 2013. Deep Gaussian processes. In *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics*, eds. CM Carvalho, P Ravikumar, vol. 31 of *Proceedings of Machine Learning Research*. Scottsdale, AZ, 207–215
- De Oliveira V, Kedem B, Short DA. 1997. Bayesian prediction of transformed Gaussian random fields. *Journal of the American Statistical Association* 92:1422–1433

- Dempster AP, Laird NM, Rubin DB. 1977. Maximum likelihood from incomplete data via the EM algorithm (with discussion). *Journal of the Royal Statistical Society: Series B* 39:1–38
- Dewar M, Scerri K, Kadiramanathan V. 2009. Data-driven spatio-temporal modeling using the integro-difference equation. *IEEE Transactions on Signal Processing* 57:83–91
- Diggle PJ, Tawn JA, Moyeed RA. 1998. Model-based geostatistics (with discussion). *Journal of the Royal Statistical Society: Series C* 47:299–350
- Dunlop MM, Girolami M, Stuart AM, Teckentrup AL. 2018. How deep are deep Gaussian processes? *Journal of Machine Learning Research* 19:1–46
- Eidsvik J, Finley AO, Banerjee S, Rue H. 2012. Approximate Bayesian inference for large spatial datasets using predictive process models. *Computational Statistics and Data Analysis* 56:1362–1380
- Enting IG. 2002. Inverse Problems in Atmospheric Constituent Transport. Cambridge, UK: Cambridge University Press
- Finley AO, Sang H, Banerjee S, Gelfand AE. 2009. Improving the performance of predictive process modeling for large datasets. *Computational Statistics & Data Analysis* 53:2873–2884
- Freestone DR, Aram P, Dewar M, Scerri K, Grayden DB, Kadiramanathan V. 2011. A data-driven framework for neural field modeling. *NeuroImage* 56:1043–1058
- Genton MG, Kleiber W. 2015. Cross-covariance functions for multivariate geostatistics. *Statistical Science* 30:147–163
- Gneiting T, Kleiber W, Schlather M. 2010. Matérn cross-covariance functions for multivariate random fields. *Journal of the American Statistical Association* 105:1167–1177
- Griffith D. 2000. A linear regression solution to the spatial autocorrelation problem. *Journal of Geographical Systems* 2:141–156
- Hanks E, Schliep E, Hooten M, Hoeting J. 2015. Restricted spatial regression in practice: geostatistical models, confounding, and robustness under model misspecification. *Environmetrics* 26:243–254
- Heaton MJ, Datta A, Finley AO, Furrer R, Guinness J, et al. 2019. A case study competition among methods for analyzing large spatial data. *Journal of Agricultural, Biological and Environmental Statistics* 24:398–425
- Hensman J, Lawrence ND. 2014. Nested variational compression in deep Gaussian processes. *arXiv:1412.1370*
- Hodges J, Reich B. 2010. Adding spatially-correlated errors can mess up the fixed effect you love. *Annals of Applied Statistics* 64:325–334
- Hooten MB, Johnson DS, McClintock BT, Morales JM. 2017. Animal Movement: Statistical Models for Telemetry Data. Boca Raton, FL: CRC Press
- Huang H, Blake LR, Hammerling DM. 2019. Pushing the limit: a hybrid parallel implementation of the multi-resolution approximation for massive data. *arXiv:1905.00141*
- Huang HC, Cressie N, Gabrosek J. 2002. Fast, resolution-consistent spatial prediction of global processes from satellite data. *Journal of Computational and Graphical Statistics* 11:63–88

- Hughes J, Haran M. 2013. Dimension reduction and alleviation of confounding for spatial generalized linear mixed models. *Journal of the Royal Statistical Society: Series B* 75:139–159
- Johns CJ, Nychka D, Kittel TGF, Daly C. 2003. Infilling sparse records of spatial fields. *Journal of the American Statistical Association* 98:796–806
- Kang EL, Cressie N. 2011. Bayesian inference for the spatial random effects model. *Journal of the American Statistical Association* 106:972–983
- Katzfuss M. 2013. Bayesian nonstationary spatial modeling for very large datasets. *Environmetrics* 24:189–200
- Katzfuss M. 2017. A multi-resolution approximation for massive spatial datasets. *Journal of the American Statistical Association* 112:201–214
- Katzfuss M, Cressie N. 2011. Spatio-temporal smoothing and EM estimation for massive remote-sensing data sets. *Journal of Time Series Analysis* 32:430–446
- Lee BS, Park J. 2020. A scalable partitioned approach to model massive nonstationary non-Gaussian spatial datasets. *arXiv:2001.09111*
- Lindgren F, Rue H. 2015. Bayesian spatial modelling with R-INLA. *Journal of Statistical Software* 63(19):1–25
- Lindgren F, Rue H, Lindström J. 2011. An explicit link between Gaussian fields and Gaussian Markov random fields: The stochastic partial differential equation approach (with discussion). *Journal of the Royal Statistical Society: Series B* 73:423–498
- Lopes HF, Gamerman D, Salazar E. 2011. Generalized spatial dynamic factor models. *Computational Statistics and Data Analysis* 55:1319–1330
- Lopes HF, Salazar E, Gamerman D. 2008. Spatial dynamic factor analysis. *Bayesian Analysis* 3:759–792
- Ma P, Kang EL. 2020. A fused Gaussian process model for very large spatial data. *Journal of Computational and Graphical Statistics* 29:479–489
- McCullagh P, Nelder JA. 1989. *Generalized Linear Models*, 2nd ed. London, UK: Chapman & Hall
- McLachlan GJ, Krishnan T. 2007. *The EM Algorithm and Extensions*, 2nd ed. Hoboken, NJ: Wiley
- Michalak AM, Bruhwiler L, Tans PP. 2004. A geostatistical approach to surface flux estimation of atmospheric trace gases. *Journal of Geophysical Research: Atmospheres* 109:D14109
- Nelder JA, Wedderburn RWM. 1972. Generalized linear models. *Journal of the Royal Statistical Society: Series A* 135:370–384
- Nguyen H, Cressie N, Braverman A. 2017. Multivariate spatial data fusion for very large remote sensing datasets. *Remote Sensing* 9:142–161
- Nychka D, Bandyopadhyay S, Hammerling D, Lindgren F, Sain S. 2015. A multiresolution Gaussian process model for the analysis of large spatial datasets. *Journal of Computational and Graphical Statistics* 24:579–599
- Nychka D, Wikle C, Royle JA. 2002. Multiresolution models for nonstationary spatial covariance functions. *Statistical Modelling* 2:315–331

- Paciorek C. 2010. The importance of scale for spatial-confounding bias and precision of spatial regression estimators. *Statistical Science* 25:107–125
- Paciorek CJ. 2007. Bayesian smoothing with Gaussian processes using Fourier basis functions in the spectralGP package. *Journal of Statistical Software* 19(2):1–38
- Pebesma EJ. 2004. Multivariable geostatistics in S: the gstat package. *Computers and Geosciences* 30:683–691
- Perrin O, Monestiez P. 1999. Modelling of non-stationary spatial structure using parametric radial basis deformations. In *GeoENV II—Geostatistics for Environmental Applications*, eds. J Gómez-Hernández, A Soares, R Froidevaux. New York, NY: Springer, 175–186
- R Core Team. 2021. R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria
- Rue H, Martino S, Chopin N. 2009. Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations (with discussion). *Journal of the Royal Statistical Society: Series B* 71:319–392
- Sahu SK, Mardia KV. 2005. A Bayesian kriged Kalman model for short-term forecasting of air pollution levels. *Journal of the Royal Statistical Society: Series C* 54:223–244
- Sainsbury-Dale M, Zammit-Mangion A, Cressie N. 2021. Modelling, fitting, and prediction with non-gaussian spatial and spatio-temporal data using frk. *arXiv:2110.02507*
- Sampson PD, Guttorp P. 1992. Nonparametric estimation of nonstationary spatial covariance structure. *Journal of the American Statistical Association* 87:108–119
- Sang H, Huang JZ. 2012. A full scale approximation of covariance functions for large spatial data sets. *Journal of the Royal Statistical Society: Series B* 74:111–132
- Sansó B, Schmidt AM, Nobre AA. 2008. Bayesian spatio-temporal models based on discrete convolutions. *Canadian Journal of Statistics* 36:239–258
- Schmidt AM, 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
- Searle SR, Henderson HV. 1981. On deriving the inverse of a sum of matrices. *Society for Industrial and Applied Mathematics Review* 23:53–60
- Sengupta A, Cressie N. 2013. Hierarchical statistical modelling of big spatial datasets using the exponential family of distributions. *Spatial Statistics* 4:14–44
- Sengupta A, Cressie N, Kahn BH, Frey R. 2016. Predictive inference for big, spatial, non-Gaussian data: MODIS cloud data and its change-of-support. *Australian and New Zealand Journal of Statistics* 58:15–45
- Simpson D, Illian JB, Lindgren F, Sørbye SH, Rue H. 2016. Going off grid: computationally efficient inference for log-Gaussian Cox processes. *Biometrika* 103:49–70
- Smith RL. 1996. Estimating nonstationary spatial correlations. Tech. rep., Cambridge University, Cambridge, UK. Online: Available from <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.40.5988&rep=rep1&type=pdf>

- Solo V. 2002. Identification of a noisy stochastic heat equation with the EM algorithm. In *Proceedings of the 41st IEEE Conference on Decision and Control*. Las Vegas, NV: IEEE, 4505–4508
- Stein ML. 1999. *Interpolation of Spatial Data: Some Theory for Kriging*. New York, NY: Springer-Verlag
- Stein ML. 2014. Limitations on low rank approximations for covariance matrices of spatial data. *Spatial Statistics* 8:1–19
- Stroud JR, Müller P, Sansó B. 2001. Dynamic models for spatiotemporal data. *Journal of the Royal Statistical Society: Series B* 63:673–689
- Tagle F, Genton MG, Yip A, Mostamandi S, Stenchikov G, Castruccio S. 2020. A high-resolution bilevel skew-t stochastic generator for assessing Saudi Arabia’s wind energy resources. *Environmetrics* 31:1–17
- Tao T. 2011. *An Introduction to Measure Theory*. Providence, RI: American Mathematical Society
- Tibshirani R. 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B* 58:267–288
- Tobler WR. 1970. A computer movie simulating urban growth in the Detroit region. *Economic Geography* 46:234–240
- Tzeng S, Huang HC. 2018. Resolution adaptive fixed rank kriging. *Technometrics* 60:198–208
- Tzeng S, Huang HC, Cressie N. 2005. A fast, optimal spatial-prediction method for massive datasets. *Journal of the American Statistical Association* 100:1343–1357
- Vidakovic B, Müller P. 1999. An introduction to wavelets. In *Bayesian Inference in Wavelet-Based Models*, eds. B Vidakovic, P Müller. New York, NY: Springer, 1–18
- Wahba G. 1990. *Spline Models for Observational Data*. Philadelphia, PA: Society for Industrial and Applied Mathematics
- Wikle CK. 2010. Low-rank representations for spatial processes. In *Handbook of Spatial Statistics*, eds. AE Gelfand, PJ Diggle, M Fuentes, P Guttorp. Boca Raton, FL: Chapman & Hall/CRC, 107–118
- Wikle CK, Berliner LM. 2007. A Bayesian tutorial for data assimilation. *Physica D: Nonlinear Phenomena* 230:1–16
- Wikle CK, Cressie N. 1999. A dimension-reduced approach to space-time Kalman filtering. *Biometrika* 86:815–829
- Wikle CK, Milliff RF, Nychka D, Berliner LM. 2001. Spatiotemporal hierarchical Bayesian modeling tropical ocean surface winds. *Journal of the American Statistical Association* 96:382–397
- Wikle CK, Zammit-Mangion A, Cressie N. 2019. *Spatio-Temporal Statistics with R*. Boca Raton, FL: Chapman & Hall/CRC
- Wood SN. 2017. *Generalized Additive Models: An Introduction with R*. Boca Raton, FL: Chapman & Hall/CRC, 2nd ed.
- Xu G, Genton MG. 2017. Tukey g-and-h random fields. *Journal of the American Statistical Association* 112:1236–1249

- Zammit-Mangion A, Bertolacci M, Fisher J, Stavert A, Rigby ML, et al. 2021a. WOM-BAT: A fully Bayesian global flux-inversion framework. *Geosci. Model Dev. Disc.* <https://doi.org/10.5194/gmd-2021-181>
- Zammit-Mangion A, Cressie N. 2021. FRK: An R package for spatial and spatio-temporal prediction with large datasets. *Journal of Statistical Software* 98(4):1–48
- Zammit-Mangion A, Cressie N, Ganesan AL. 2016. Non-Gaussian bivariate modelling with application to atmospheric trace-gas inversion. *Spatial Statistics* 18:194–220
- Zammit-Mangion A, Ng TLG, Vu Q, Filippone M. 2021b. Deep compositional spatial models. *Journal of the American Statistical Association* doi: 10.1080/01621459.2021.1887741
- Zammit-Mangion A, Rougier J. 2020. Multi-scale process modelling and distributed computation for spatial data. *Statistics and Computing* 30:1609–1627
- Zammit-Mangion A, Rougier J, Schön N, Lindgren F, Bamber J. 2015. Multivariate spatio-temporal modelling for assessing Antarctica’s present-day contribution to sea-level rise. *Environmetrics* 26:159–177
- Zammit-Mangion A, Sanguinetti G, Kadiramanathan V. 2012. Variational estimation in spatiotemporal systems from continuous and point-process observations. *IEEE Transactions on Signal Processing* 60:3449–3459