

Spatial deformation for non-stationary extremal dependence

J. Richards^{1*} and J. L. Wadsworth²

¹STOR-i Centre for Doctoral Training, Department of Mathematics and Statistics, Lancaster University LA1 4YR, United Kingdom

²Department of Mathematics and Statistics, Lancaster University LA1 4YW, United Kingdom

*Correspondence to: j.flett@lancaster.ac.uk

January 19, 2021

Abstract

Modelling the extremal dependence structure of spatial data is considerably easier if that structure is stationary. However, for data observed over large or complicated domains, non-stationarity will often prevail. Current methods for modelling non-stationarity in extremal dependence rely on models that are either computationally difficult to fit or require prior knowledge of covariates. Sampson and Guttorp (1992) proposed a simple technique for handling non-stationarity in spatial dependence by smoothly mapping the sampling locations of the process from the original geographical space to a latent space where stationarity can be reasonably assumed. We present an extension of this method to a spatial extremes framework by considering least squares minimisation of pairwise theoretical and empirical extremal dependence measures. Along with some practical advice on applying these deformations, we provide a detailed simulation study in which we

propose three spatial processes with varying degrees of non-stationarity in their extremal and central dependence structures. The methodology is applied to Australian summer temperature extremes and UK precipitation to illustrate its efficacy compared to a naive modelling approach.

Keywords— non-stationary spatial dependence; extremal dependence; spatial deformation; max-stable processes

1 Introduction

Statistical methodology for spatial extremes can increasingly handle data sampled at more observation locations. If these observations are taken over large domains with complex features, then there is a strong chance that the data will exhibit spatial non-stationarity in both the marginal distributions and dependence structure. Marginal non-stationarity can often be dealt with by site-wise modelling and transformation. However, there are currently few methods to deal with non-stationarity in extremal dependence structures, and a typical approach is to falsely assume stationarity when fitting spatial extremes models. This may be appropriate when modelling data sampled over small and/or homogeneous regions in space, but as we will illustrate through the examples in Section 3, this assumption is not realistic for many datasets with larger spatial domains.

Beyond site-wise transformation of margins, marginal non-stationarity can be handled by jointly modelling marginal parameters as functions of covariates. This can either be achieved parametrically (Mannshardt-Shamseldin et al., 2010, Davison et al., 2012, Ribatet, 2013, Eastoe, 2019) or semiparametrically (Jonathan et al., 2014, Ross et al., 2017, Youngman, 2019, 2020, Zanini et al., 2020) through the use of splines. Another widely applied approach is the use of Bayesian hierarchical models, in which the marginal parameters are assumed to come from some non-stationary latent process (Casson and Coles, 1999, Cooley et al., 2007, Sang and Gelfand, 2010, Opitz et al., 2018).

Non-stationarity in the spatial dependence structure has been studied by Huser and Genton (2016) in the context of max-stable models, through incorporation of a non-stationary variogram. However, this approach requires knowledge of relevant covariates, and asymptotically dependent max-stable models for spatial extremes have been shown to be too inflexible for many spatial datasets

(Wadsworth and Tawn, 2012, Davison et al., 2013, Huser et al., 2017, Huser and Wadsworth, 2019). Another approach is to assume local stationarity for model fitting, see Blanchet and Creutin (2017), Castro-Camilo and Huser (2019). This framework is well-suited to modelling processes with short-range dependence but is unlikely to fully capture dependence at large distances. Cooley et al. (2007) and Blanchet and Davison (2011) account for non-stationarity by transforming their spatial domain of interest to some new ‘climate space’ in which observation locations with similar characteristics are grouped closer together. Again, this approach requires access to relevant covariates and a deeper understanding of the processes which are being modelled.

In this work we develop a computationally quick and simple method, which does not require prior knowledge of covariates and which can be applied before fitting any model suited to spatial extremes. Our method uses spatial deformation and is based on the work of Sampson and Guttorp (1992) and Smith (1996), which has not been fully adapted for use in a spatial extremes framework. The deformation methodology may reveal physical features and/or covariates that can be incorporated into a spatial extremes model, removing the need for models with complex dependence structures.

Wadsworth and Tawn (2019) applied the deformation method of Smith (1996) before fitting a conditional spatial extremes model to the same Australian summer temperatures data that we explore in Section 4.1. However, because this method is not tailored to extremal dependence, it was necessary to assume that patterns in non-stationarity were similar for both the extremal and non-extremal dependence structures. Youngman (2020) and Chevalier et al. (2017) provide extensions of the Sampson and Guttorp (1992) methodology and fit models for spatial extremes using deformations: a Gaussian process using a censored pairwise likelihood and a max-stable model, respectively. Although these models may be reasonable for some processes, use of either puts restrictions on the types of dependence that the process can exhibit. We look to develop a method that makes no strong assumptions on the extremal dependence structure.

The remainder of this section provides an overview of existing methodology for spatial deformation and modelling of spatial extremes. Our developments of the spatial deformation methodology are detailed in Section 2. We present a simulation study in Section 3, which is usually absent from the literature on spatial deformations. This study is used to convey that our adaptations to the deformation methodology are necessary when considering extremal dependence and that our method can be used

for different processes with a wide range of extremal dependence structures. Finally, we apply our method to temperature and precipitation datasets in Section 4, and conclude with a discussion in Section 5.

1.1 Non-stationary spatial processes

The spatial deformation approach for handling non-stationarity in spatial processes was first proposed by Sampson and Guttorp (1992) and Guttorp and Sampson (1994), with further developments by Meiring et al. (1997), see Sampson (2010, Ch. 9.5). The underlying principle of their approach is that a smooth non-linear transformation can be used to map the sampling locations of a process from a geographical plane, or G-plane, to some latent space, which they name a D-plane, or dispersion plane. Within the D-plane, the dependence structure of the process is assumed to be both stationary and isotropic, and the usual statistical inferences can be made using stationary geostatistical models. To obtain the D-plane, optimisation techniques are used to minimise some objective function which is associated with a stationary geostatistical model. Here Sampson and Guttorp (1992) use multi-dimensional scaling and a stationary spatial dispersion function, whereas further work proposed by Smith (1996) uses the likelihood for a stationary Gaussian process. Our approach is to change this objective function for one which is associated with a stationary spatial extremes model, such as the max-stable, or inverted max-stable, processes.

We begin by assuming we have realisations $\mathbf{Z} = \{Z_1, \dots, Z_N\}$ from a spatial field observed at sampling locations s_1, \dots, s_d , and so we have $Z_k = \{Z_k(s_1), \dots, Z_k(s_d)\}$ for all $k = 1, \dots, N$. We require some smooth mapping function from the G-plane to the D-plane, given by $f(s_i) = s_i^*$ for $i = 1, \dots, d$, where $s_i = (x_i, y_i)$ and $s_i^* = (x_i^*, y_i^*)$ are the corresponding locations in the D-plane. Both Sampson and Guttorp (1992) and Smith (1996) propose the use of thin-plate splines to achieve this mapping. However, we note that under certain conditions on the correlation structure, analytical forms for $f(\cdot)$ do exist. Perrin and Meiring (1999) prove that this mapping is identifiable assuming differentiability of the stationary and isotropic correlation function used for fitting and Perrin and Senoussi (2000) derive analytical forms for $f(\cdot)$ under the same assumption, with extensions to anisotropic correlation structures. As these results are available only for correlation functions, and not for extremal dependence functions, we instead use the more flexible thin-plate spline approach.

A thin-plate spline is a mapping function $f(\cdot)$, passing through a finite number of data points $f_i^* = f^*(x_i, y_i)$, ($i = 1, \dots, n$), minimising the bending energy

$$J(f) = \iint_{\mathbb{R}^2} \left\{ \left(\frac{\partial^2 f}{\partial x^2} \right)^2 + 2 \left(\frac{\partial^2 f}{\partial x \partial y} \right)^2 + \left(\frac{\partial^2 f}{\partial y^2} \right)^2 \right\} dx dy.$$

Here we have denoted f^* the ‘true’ function that we wish to estimate with the thin-plate spline, f , and f_i^* are observations. Green and Silverman (1994) give a solution to this problem in the form

$$f(x, y) = a + bx + cy + \sum_{i=1}^n \delta_i g_i(x, y), \quad (1)$$

where

$$\sum_{i=1}^n \delta_i = \sum_{i=1}^n \delta_i x_i = \sum_{i=1}^n \delta_i y_i = 0, \quad (2)$$

and $g_i(x, y) = h_i^2 \log h_i$, with h_i the Euclidean distance between (x, y) and (x_i, y_i) . This represents f as the sum of linear terms and n radial basis functions with centres at the observed data locations (x_i, y_i) and the constraints are in place to ensure that the system of equations does not become overdetermined. An interpolating spline satisfies $f_i^* = f(x_i, y_i)$ for all $i = 1, \dots, n$, whereas we desire a smoothing spline; this can be created by minimising

$$S(f) = \sum_{i=1}^n \{f_i^* - f(x_i, y_i)\}^2 + \alpha J(f),$$

for some smoothing parameter $\alpha > 0$. Sampson and Guttorp (1992) give a method for estimating α in the context of multidimensional scaling, but here we take the approach of Smith (1996), who uses a restricted representation of (1) instead. A subset of m radial basis functions is used and so we let $\delta_i = 0$ for all $i \notin \{i_1, \dots, i_m\}$. The choice of this subset is discussed in Section 3.

The function in (1) maps \mathbb{R}^2 to \mathbb{R} , so the spline is applied twice with different parameter estimates to produce both components. Smith (1996) gives a parametrisation as

$$f^{(1)}(x, y) = b_1^2 x + \rho b_1 b_2 y + \sum_{i=1}^m \delta_i^{(1)} g_i(x, y) \quad (3)$$

$$f^{(2)}(x, y) = b_2^2 y + \rho b_1 b_2 x + \sum_{i=1}^m \delta_i^{(2)} g_i(x, y), \quad (4)$$

where $b_1 > 0$, $b_2 > 0$, $\rho \in \mathbb{R}$ and each of the sequences $\delta^{(1)}$, $\delta^{(2)}$ satisfy the constraint in (2). The introduction of the parameters b_1, b_2 and ρ is to ensure that the model is invariant under orthogonal rotations when $m = 0$. Overall, this yields a spline with $2m - 3$ free parameters whenever $m \geq 3$.

The resulting spline is then used to map the sampling locations s_i to locations s_i^* in a latent space. Parameters are estimated by minimising some objective function provided by a stationary model. As previously mentioned, Sampson and Guttorp (1992) use a stationary spatial dispersion model and multidimensional scaling, the details of which are not given here. Instead, we focus on the approach by Smith (1996), who uses a stationary Gaussian likelihood. It is assumed that $(Z(s_1^*), \dots, Z(s_d^*)) \sim N_d(\mu, \Omega)$, where μ and Ω are the mean vector and a stationary covariance matrix, respectively. As we are only interested in measuring the dependence structure, it is assumed that the means and variances at each location are known. Analysis is then simplified to only considering the minimisation of the negative log likelihood given by

$$-\log L(\Omega) = \frac{N}{2} \log |\Omega| + \frac{N-1}{2} \text{tr} \left(\Omega^{-1} \hat{\Omega} \right), \quad (5)$$

where Ω and $\hat{\Omega}$ are the theoretical, and sample, correlation matrices and $\text{tr}(\cdot)$ and $|\cdot|$ are the trace and determinant operators, respectively. The entries of the theoretical correlation matrix are produced by using a stationary covariance function. Smith (1996) uses the Matérn covariance function, and so

$$\Omega_{ij} = \frac{1}{2^{\theta_2-1} \Gamma(\theta_2)} \left(\frac{2h_{ij}^* \sqrt{\theta_2}}{\theta_1} \right)^{\theta_2} K_{\theta_2} \left(\frac{2h_{ij}^* \sqrt{\theta_2}}{\theta_1} \right), \quad (6)$$

where $\theta_1 > 0$, $\theta_2 > 0$ and $K_{\theta_2}(\cdot)$ is the modified Bessel function of the second kind of order θ_2 and $h_{ij}^* = \|s_i^* - s_j^*\|$ is the Euclidean distance between locations s_i^* and s_j^* in the D-plane. It is noted that θ_1 can be set to 1 as the spatial scaling of the locations is controlled by the spline.

1.2 Spatial extremes

Before describing an extension of the spatial deformation methodology tailored to spatial extremes, we first provide a brief review of methods for modelling spatial extremes.

1.2.1 Max-stable and inverted max-stable processes

Max-stable processes were introduced by de Haan (1984) and developed further by Smith (1990) and Schlather (2002), who suggested models that were first fitted by pairwise composite likelihood in Padoan et al. (2010). They are usually described by a spectral construction. Suppose $\{r_i; i \geq 1\}$ are points of a Poisson process on $(0, \infty)$ with unit intensity. Let $S \subseteq \mathbb{R}^2$ be a spatial index set, and $\{W_i(s); s \in S, i \geq 1\}$ be independent and identically distributed copies of a non-negative stochastic process satisfying $\mathbb{E}[W(s)] = 1$. Then

$$Z(s) = \max_{i \geq 1} \{W_i(s)/r_i\} \quad (7)$$

is a max-stable process with unit Fréchet margins. The d -dimensional joint distribution function for Z is

$$\Pr\{Z(s_1) \leq z_1, \dots, Z(s_d) \leq z_d\} = \exp\{-V(z_1, \dots, z_d)\}, \quad (8)$$

where the exponent is

$$V(z_1, \dots, z_d) = \mathbb{E} \left[\max \left\{ \frac{W(s_1)}{z_1}, \dots, \frac{W(s_d)}{z_d} \right\} \right]. \quad (9)$$

Careful specification of the stochastic process $W(s)$ leads to a limited selection of parametric models for the max-stable process. A particularly flexible model is the Brown-Resnick model (Brown and Resnick, 1977, Kabluchko et al., 2009). This involves specifying $W(s) = \exp\{U(s) - \gamma^*(s, 0)\}$ for $U(s)$ a centred Gaussian process with semivariogram $\gamma^*(\cdot, \cdot)$ and where $U(0) = 0$ almost surely. This leads to a 2-dimensional joint distribution with exponent function

$$V(z_i, z_j) = \frac{1}{z_i} \Phi \left\{ \frac{a}{2} - \frac{1}{a} \log \left(\frac{z_i}{z_j} \right) \right\} + \frac{1}{z_j} \Phi \left\{ \frac{a}{2} - \frac{1}{a} \log \left(\frac{z_j}{z_i} \right) \right\}, \quad (10)$$

where $a = [2\gamma^*(s_i, s_j)]^{1/2}$ and $\Phi(\cdot)$ denotes the standard normal distribution function. Note that for a stationary and isotropic Brown-Resnick process, $\gamma^*(s_i, s_j)$ is dependent on $h_{ij} = \|s_i - s_j\|$ only. For clarity, we write $\gamma(h_{ij})$ when Z is stationary and isotropic, and $\gamma^*(s_i, s_j)$, otherwise. Representations for (10) in higher dimensions exist (see Huser and Davison (2013) or Wadsworth and Tawn (2014)), but due to their computational complexity, inference for max-stable processes is typically done pairwise, providing a reasonable balance between computation time and efficiency.

Max-stable processes are inherently asymptotically dependent, or perfectly independent. That is, $Z(s_i)$ and $Z(s_j)$ are asymptotically dependent, or perfectly independent, for all $s_i, s_j \in S$. Here we characterise asymptotic dependence using the upper tail index χ (Joe, 1997). Assuming $Z(s_i) \sim F_i, Z(s_j) \sim F_j$, we have

$$\chi(s_i, s_j) = \lim_{q \rightarrow 1} \Pr\{F_i\{Z(s_i)\} > q | F_j\{Z(s_j)\} > q\}, \quad (11)$$

where the process is asymptotically independent at locations s_i and s_j if $\chi(s_i, s_j) = 0$, and asymptotically dependent otherwise. Here we write $\chi(s_i, s_j)$ as Z is not necessarily stationary; henceforth, we write $\chi(h_{ij})$ for $h_{ij} = \|s_i - s_j\|$ when it is assumed that χ is a function of distance only. As this measure is theoretically non-zero at all spatial lags for any max-stable process exhibiting positive spatial association ie., $\chi(s_i, s_j) > 0$ for all $s_i, s_j \in S$, we require other modelling approaches to deal with processes that may exhibit asymptotic independence.

Wadsworth and Tawn (2012) introduced the inverted max-stable process as that obtained by applying a monotonically decreasing marginal transformation to a max-stable process. For example, with Z as defined in (7), taking $Y(s) = 1/Z(s)$ gives an inverted max-stable process with exponential margins and joint survival function

$$\Pr\{Y(s_1) \geq y_1, \dots, Y(s_d) \geq y_d\} = \exp\{-V(1/y_1, \dots, 1/y_d)\}, \quad (12)$$

where V is as given in (9). Such a process is asymptotically independent with $\chi(s_i, s_j) = 0$ for all $s_i \neq s_j$, but can accommodate a variety of flexible extremal dependence structures exhibiting positive association. The dependence in asymptotically independent processes may be characterised by a pre-limiting version of (11). Specifically, under an assumption of hidden regular variation (Ledford and Tawn, 1996, Resnick, 2002),

$$\chi_q(s_i, s_j) = \Pr\{F_i\{Z(s_i)\} > q | F_j\{Z(s_j)\} > q\} = L(1 - q)(1 - q)^{1/\eta(s_i, s_j) - 1}, \quad (13)$$

with $L(\cdot)$ slowly varying at 0 and $\eta(s_i, s_j) \in (0, 1]$ the coefficient of tail dependence. For an inverted max-stable process, $\chi_q(s_i, s_j) = (1 - q)^{V(1,1) - 1}$.

We fit both max-stable and inverted max-stable models after applying our deformation method for non-stationary spatial extremes. Note that although max-stable processes are typically taken to represent the limiting behaviour of maxima, in practice they, along with inverted max-stable processes, can be used for all extreme values through specification of a censored likelihood; see Section 2.5. Inference on these models can then be used to determine the efficacy of our deformation method.

1.2.2 Conditional extremes

An alternative approach to modelling spatial extremes is to condition on the behaviour of the process when it is extreme at a single site. Here we give a brief overview of modelling the extremal behaviour of the process at two sites using this approach. For a full characterisation, see Wadsworth and Tawn (2019) or Shooter et al. (2019). We suppress some of the notation used by Wadsworth and Tawn (2019) and Shooter et al. (2019) as we are only considering a discrete pairwise fit, that we will employ in Section 4 as a diagnostic measure. For further details of the discrete approach, see Heffernan and Tawn (2004). Winter et al. (2016) apply this same methodology to a dataset of Australian temperatures, which we revisit in Section 4.1.

We begin by assuming that $\{X(s) : s \in S \subset \mathbb{R}^2\}$ is a stationary and isotropic process with exponential-tailed marginals and denote $X(s_i) = X_i$. Conditioning on $X_i = x_i > u$ being large and considering $X_j, i \neq j$, Heffernan and Tawn (2004) assume that there exist normalising functions $a(x_i) : \mathbb{R} \rightarrow \mathbb{R}, b(x_i) : \mathbb{R} \rightarrow \mathbb{R}_+$, for which

$$\lim_{x_i \rightarrow \infty} [\Pr(X_j \leq a(x_i) + b(x_i)z | X_i = x_i)] = G(z),$$

where G is non-degenerate. Re-writing $Z = \{X_j - a(x_i)\}/b(x_i)$ as the standardised residual, and making the assumption that the limit holds above some high threshold u , we have

$$\Pr(Z \leq z | X_i = x_i) = G(z), \quad x_i > u,$$

where $X_i | X_i > u \sim \text{Exp}(1)$ is independent of Z . Inference on G is often simplified by making the working assumption that $Z \sim N(\mu, \sigma^2)$ and using a specified parametric form for the normalising functions $a(\cdot), b(\cdot)$. For positively dependent data, we simplify the normalising functions to $a(x_i) = \alpha x_i$

for $\alpha \in [0, 1]$ and $b(x_i) = x_i^\beta$ for $\beta \in [0, 1)$. The bivariate form of the conditional model can thus be expressed

$$X_j | (X_i = x_i) = \alpha x_i + x_i^\beta Z, \quad x_i > u.$$

The conditional model holds some useful advantages over joint modelling using max-stable, or inverted max-stable, processes. For one, it is able to handle both asymptotically dependent, or asymptotically independent, data. Parameter estimates for α and β can indicate the nature of the dependence between X_j and X_i . For example, asymptotic dependence between X_j and X_i is implied by estimates $\alpha = 1, \beta = 0$. Within the class of asymptotically independent variables, $\alpha < 1, \beta > 0$, with $\alpha = \beta = 0$ giving near extremal independence.

The spatial extensions of this model (Wadsworth and Tawn, 2019, Shooter et al., 2019) specify α and β as functions of distance between sites, when the underlying process is stationary and isotropic. As such, we can use these parameter estimates as diagnostics, to determine whether our deformation method has created a process that has a more stationary extremal dependence structure. We are motivated to use these estimates as our deformation method does not use a conditional extremes approach for fitting.

2 Spatial deformation for extremes

In this section, we discuss our adaptations of the deformation methodology for application in a spatial extremes framework. We begin in Section 2.1 by proposing a new objective function to that of (5). Instead, we consider minimising the difference between theoretical and empirical χ measures, where the former are produced through specification of a stationary max-stable dependence structure for the process in the D-plane. This does not in fact mean that this method will not work for asymptotically independent data; on the contrary, in Sections 2.2 and 2.3 we show that the model choice for $\chi(\cdot)$ is somewhat arbitrary and a single, simple parametric form works well for both classes of extremal dependence. Section 2.4 follows with some practical advice for choosing the anchor points used in estimating the thin-plane spline and we conclude with details of model fitting and selection using censored pairwise likelihoods in Section 2.5. To assess the efficacy of the deformations we produce, we fit full max-stable, and inverted max-stable, dependence models.

2.1 Objective function

To adapt the methodology of Sampson and Guttorp (1992) and Smith (1996) to better suit a spatial extremes framework, we change the objective function given in (5) to the Frobenius norm of the difference between theoretical and empirical pairwise dependence matrices $\mathbf{X} := [\chi(h_{ij}^*)]$ and $\hat{\mathbf{X}} := [\hat{\chi}(h_{ij}^*)]$. That is, we estimate the parameters of the thin plate spline through computing

$$\min \|\mathbf{X} - \hat{\mathbf{X}}\|_F = \min \sqrt{\sum_{i=1}^d \sum_{j=1}^d \left\{ \chi(h_{ij}^*) - \hat{\chi}(h_{ij}^*) \right\}^2}, \quad (14)$$

where $\chi(h_{ij}^*)$, defined in (11), is the upper tail index calculated between the process at locations s_i^* and s_j^* in the D-plane and $\hat{\chi}(h_{ij}^*)$ is its empirical estimate. Recall that we assume stationarity in the D-plane, and so write $\chi(h_{ij}^*)$, rather than $\chi(s_i^*, s_j^*)$. In practice, this measure cannot be estimated in the limit as $q \rightarrow 1$. As such, we estimate $\hat{\chi}(h_{ij}^*)$ by fixing some high threshold $q < 1$ and calculating

$$\hat{\chi}_q(h_{ij}^*) = \Pr\{\hat{F}_i\{Z(s_i^*)\} > q | \hat{F}_j\{Z(s_j^*)\} > q\} = \Pr\{\hat{F}_i\{Z(s_i)\} > q | \hat{F}_j\{Z(s_j)\} > q\}, \quad (15)$$

where $\hat{F}_k(\cdot)$ is the empirical distribution of observations $Z(s_k^*) = Z(s_k)$. Under asymptotic dependence, we assume that $\chi_q(h_{ij}^*) \equiv \chi(h_{ij}^*)$ for large enough q . Under asymptotic independence, although $\chi_q(h^*) \rightarrow 0$ as $q \rightarrow 1$, we typically have $\chi_q(h^*) > 0$ for $q < 1$ and spatial structure in this measure that makes it informative about non-stationarity.

We now focus on a choice of function $\chi(h^*)$, which we only require to be monotonically decreasing from 1 to 0. This leaves several options, including specific parametric forms for $\chi(h^*)$ and $\chi_q(h^*)$ from max-stable, and inverted max-stable, processes. We remark that while we have used χ to measure extremal dependence, other extremal dependence measures exist, and can also be used in this framework. For example, the coefficient of tail dependence, $\eta(h_{ij}^*)$, from (13) can also be used to characterise the strength of asymptotic independence in extremes. This can be estimated separately from $\chi(h_{ij}^*)$, however, we found that due to the high variance of the estimator for $\eta(h_{ij}^*)$, it was often outperformed by using $\chi(h_{ij}^*)$.

2.2 Asymptotic dependence versus asymptotic independence

As a parametric model for $\chi(h^*)$ we take the form implied by the stationary Brown-Resnick process,

$$\chi(h_{ij}^*) = 2 - \theta(h_{ij}^*) = 2 - 2\Phi \left\{ \frac{[2\gamma(h_{ij}^*)]^{1/2}}{2} \right\}, \quad (16)$$

where $\theta(\cdot)$ is the extremal coefficient function (Schlather and Tawn, 2003) and $\theta(h_{ij}^*) = V(1, 1)$, with $V(\cdot, \cdot)$ defined in (10). The semivariogram $\gamma(h_{ij}^*)$ controls the dependence of the max-stable field and a typical choice for the semivariogram would be

$$\gamma(h_{ij}^*) = (h_{ij}^*/\lambda)^\kappa, \quad (17)$$

where $\lambda > 0$ is a scaling parameter and $\kappa \in (0, 2]$ is a smoothing parameter. Note that setting $\kappa = 2$ yields the Smith process (Smith, 1990), a special case of the Brown-Resnick process. As previously mentioned when discussing the Smith (1996) methodology for spatial deformation, we can set the scaling parameter λ to 1, as the spatial scaling of locations is controlled by the deformation itself. Note that the motivation for using the Brown-Resnick process as a parametric model is that $\chi(h^*) \rightarrow 0$ as $h^* \rightarrow \infty$, unlike other popular parametric models. For a stationary inverted Brown-Resnick process, we have

$$\chi_q(h_{ij}^*) = (1 - q)^{\theta(h_{ij}^*) - 1}. \quad (18)$$

We denote the dependence measures in (16) and (18) as χ^{BR} and χ_q^{IBR} , respectively. Note that although these two measures have different parametric forms, and are applicable to different dependence structures, they often approximate each other very closely when used within a deformation framework; this is illustrated in Figure 1. Here we create deformations for a simulated dataset as described in Section 3.1 using both χ^{BR} and χ_q^{IBR} . The plots show that both methods give very similar deformations when considering the non-stationarity in the $\chi(h_{ij}^*)$ estimates. This seems to be the case for both asymptotically dependent and asymptotically independent data. Hence, for the sake of simplicity we only use χ^{BR} to create deformations in the case studies in Section 4, as it appears to be flexible enough to capture non-stationarity in both classes of extremal dependence.




Figure 1: Comparison of deformations created using both parametric forms χ^{BR} and χ_q^{IBR} for $\chi(\cdot)$ for both max-stable data (left) and inverted max-stable data (right). Plots show empirical $\chi(h_{ij}^*)$ estimates against distance, where the black triangles correspond to those created using $\chi(\cdot)$ given by (16) and green triangles for those created using (18). The blue and red lines give the fitted function from (16) and (18), respectively. Distances are normalised so that the maximum distance is consistent between deformations.

2.3 Choice of parametric model for $\chi(h^*)$

We have also found that the function $\chi(h^*)$ from a Brown-Resnick process is sufficiently flexible to create suitable deformations for a variety of different extremal dependence structures. This is for similar reasons to above; different functions $\chi(h^*)$ which decrease to zero as $h^* \rightarrow \infty$ can approximate each other well. To illustrate this, we also considered the Gaussian-Gaussian process (Wadsworth and Tawn, 2012), which encompasses different dependence structures to the Brown-Resnick process, but for which $\chi(h^*) \rightarrow 0$ as $h^* \rightarrow \infty$. Its theoretical form is

$$\chi^{GG}(h^*) = 1 - \frac{1}{2} \int_{\mathbb{R}^2} \{\phi(u)^2 - 2\rho(h^*)\phi(u)\phi(u - \tilde{h}) + \phi(u - \tilde{h})^2\}^{1/2} du,$$

where $\rho(h^*)$ is a stationary correlation function and $\tilde{h} = (h^*, h^*)^T$ and $\phi(\cdot)$ is the bivariate Gaussian density function with mean 0 and covariance matrix $\Sigma = \text{diag}(\sigma^2, \sigma^2)$. Note that using a Matérn correlation function given in (6) with parameters $\theta_1 > 0$ and $\theta_2 > 0$, this function has one extra parameter than $\chi^{BR}(h^*)$, namely $\sigma > 0$.

We chose not to use this parametric form for $\chi(h^*)$, due to the high computational cost required

to compute the double integral for each pair of locations. However, we have found that the deformation method described in Section 2 appears fairly robust to the choice of $\chi(h^*)$. As Supplementary Figure S1 in the Supplementary Material shows, the much simpler $\chi^{BR}(h^*)$ can approximate the more complex $\chi^{GG}(h^*)$ very closely for much of $h^* \in \mathbb{R}_+$.

2.4 Practical aspects for creating deformations

We now comment on practical aspects of creating the deformations, including choosing a subset of radial basis functions for the thin-plate spline and reducing the chances of producing a non-bijective transformation.

We found that there is no simple robust method for picking the number m , or configuration, of the anchor points used in the deformation splines given in (3). As detailed in Sampson and Guttorp (1992), there is a trade-off in picking m . Larger values provide “better” deformations, in the sense that the objective function to be calculated is lower and the deformations seem to capture more of the non-stationarity in the process. However, this comes at the price of computational cost, the risk of over-fitting and the phenomenon in which the D-plane folds on to itself. This provides a non-bijective transformation, which is physically unrealistic. Iovleff and Perrin (2004) detail an approach to ensure that the deformation is always bijective through use of a simulated annealing algorithm, with later extensions by Youngman (2020). These approaches add further constraints into the modelling procedure, which we have chosen to avoid. Instead we use a more heuristic approach for avoiding non-bijectivity.

We begin by randomly sampling m_0 initial anchor points with index set given by $I_0 = \{i_1, \dots, i_{m_0}\}$. There is no single best way to choose I_0 ; however, we found that ensuring that the anchor points are spread out over the spatial domain helped to create better deformations. Performing a deformation with I_0 yields parameter estimates $\hat{\psi}_0 = (\hat{b}_1, \hat{b}_2, \hat{\rho}, \hat{\kappa}, \hat{\delta}_4^{(1)}, \hat{\delta}_4^{(2)}, \dots, \hat{\delta}_{m_0}^{(1)}, \hat{\delta}_{m_0}^{(2)})$. Recall that we have parameters $\delta_i^{(1)}, \delta_i^{(2)}$ indexed by $i \geq 4$ as those indexed by $i = 1, 2, 3$ are uniquely determined by the constraints given in (2). If the deformation for I_0 is bijective, we create a new set of indices $I_1 = \{I_0, i_{m_0+1}\}$, where i_{m_0+1} is sampled from the remaining indices. A deformation is then created using I_1 , but with initial parameters in the optimisation program given by $\hat{\psi}_1 = \{\hat{\psi}_0, \delta_{m_0+1}^{(1)} = 0, \delta_{m_0+1}^{(2)} = 0\}$. This ensures that the initial input into the optimisation program creates a deformation that is already bijective. We

then continue in this fashion until we have created a deformation using m^* anchor points. Bijectivity is checked by eye.

Using this approach reduces the chances of the D-plane folding as m increases and provides a deformation with m^* anchor points. Here we set m^* as approximately a quarter of the sampling locations as we have not found a clear way to optimize this aspect. Typically this approach can be used for a number of initial index sets. However, in the interest of reducing computational cost, the simulation studies in Section 3 are conducted using the same initial index set for each deformation method. We also ensure that the new index sampled at each iteration is consistent across different samples, processes and deformation methods.

2.5 Model fitting and selection

To determine whether the deformation has created a process that is more stationary in the extremal dependence structure, and to compare between deformation methods, we look to fitting max-stable and inverted max-stable models to the data using the sampling locations in both the G-plane and the D-plane. In Section 1, the computational complexities of the max-stable and inverted max-stable models were discussed. To accommodate for this, we take a pairwise composite likelihood approach and assume independence between pairs (Padoan et al., 2010). The joint distribution for a Brown-Resnick process is given in (8) and the joint survival function for an inverted Brown-Resnick process is given in (12). Note that the former is on standard Fréchet margins, whereas the latter is on standard exponential. To compare between the asymptotically dependent and asymptotically independent structures provided by the two models, we calculate all likelihoods on exponential margins, by first using a site-wise empirical transformation.

Given realisations $\{z_1, \dots, z_N\}$ from a spatial field, observed at sampling locations s_1, \dots, s_d , the censored composite likelihood is

$$L_{CL}(\lambda, \kappa) = \prod_{i=1}^N L_{CL}(\lambda, \kappa; z_i) = \prod_{i=1}^N \prod_{k=2}^d \prod_{l < k} g_u(z_i(s_k), z_i(s_l); \lambda, \kappa), \quad (19)$$

where

$$g_u(z_i(s_k), z_i(s_l); \lambda, \kappa) = \begin{cases} f(z_i(s_k), z_i(s_l); \lambda, \kappa) & \text{if } \min(z_i(s_k), z_i(s_l)) > u, \\ \frac{\partial}{\partial z_i(s_k)} F(z_i(s_k), u; \lambda, \kappa) & \text{if } z_i(s_k) > u, z_i(s_l) \leq u, \\ \frac{\partial}{\partial z_i(s_l)} F(u, z_i(s_l); \lambda, \kappa) & \text{if } z_i(s_k) \leq u, z_i(s_l) > u, \\ F(u, u; \lambda, \kappa) & \text{if } \max(z_i(s_k), z_i(s_l)) \leq u, \end{cases} \quad (20)$$

with u a high threshold and $F(\cdot)$ and $f(\cdot)$ the bivariate joint distribution and density functions for the model. Note that although we set $\lambda = 1$ when producing the deformation, here we treat it as a free parameter. Although the likelihoods give a good indication of the performance of the deformation methods, we use the Composite Likelihood version of the Akaike Information Criterion (CLAIC) for model selection. As given in Varin et al. (2011), the CLAIC is

$$-2\{\log L(\hat{\lambda}, \hat{\kappa}) - \text{tr}(J(\hat{\lambda}, \hat{\kappa})H^{-1}(\hat{\lambda}, \hat{\kappa}))\}, \quad (21)$$

where $(\hat{\lambda}, \hat{\kappa})$ are the maximum likelihood estimates from (19), $H(\cdot)$ is the Hessian matrix and $J(\cdot)$ is the variance of the score function, i.e.

$$J(\hat{\lambda}, \hat{\kappa}) = \text{var} \left(\nabla \log L_{CL}(\hat{\lambda}, \hat{\kappa}) \right) = \text{var} \left(\sum_{i=1}^N \nabla \log L_{CL}(\hat{\lambda}, \hat{\kappa}; z_i) \right).$$

In practice, we estimate $J(\cdot)$ by using numerical methods to find $\Delta_i = \nabla \log L_{CL}(\hat{\lambda}, \hat{\kappa}; z_i)$, and then estimating the variance of the score function by setting a block of length $b < N$ and computing

$$\hat{J}(\hat{\lambda}, \hat{\kappa}) = \frac{N}{b} \times \text{var} \left(\sum_{i=1}^b \Delta_i, \dots, \sum_{i=N-b+1}^N \Delta_i \right). \quad (22)$$

The block sizes are chosen such that each block of data is more reasonably assumed approximately independent. This is usually specific to the data and will be given alongside any results.

3 Simulation study

We conduct three simulation studies to illustrate the efficacy of the deformation framework for modelling extremal dependence of non-stationary spatial processes. These studies are designed to highlight

the following:

- When fitting a stationary model to the extremal dependence of non-stationary spatial data, using a deformation method will improve the fit when compared to using the original sampling locations in the G-plane;
- The deformation methodology described in Section 2.1 is more effective than the original Smith (1996) method when modelling non-stationary extremal dependence, as the latter is tailored towards modelling dependence in the body of the data rather than the extremes;
- It is often necessary to use a deformation method that is tailored explicitly to extremal dependence, rather than dependence throughout the body; especially for processes that exhibit different degrees of non-stationarity throughout their extremal and central dependence structures.

In order to illustrate these points, we consider five different processes. These processes are chosen as they each exhibit different behaviour in their respective extremal dependence structures. In Section 3.1, we consider two processes: a non-stationary Brown-Resnick process and a non-stationary inverted Brown-Resnick process. In Section 3.2, we consider two more processes which are both mixtures of stationary and non-stationary processes. We term these max-mixture process and one exhibits asymptotic dependence whilst the other exhibits asymptotic independence. A final process is considered in Section 3.3, which is an asymptotically independent Gaussian mixture process.

For each setting, we begin with a sample of 1000 realisations of a spatial process. For this sample, we create four separate deformations using the procedure set out in Section 2.4. The first two deformations are created using the approach detailed in Section 2.1; with χ^{BR} from (16) and χ_q^{IBR} from (18) as the dependence measures used in the objective function in (14). The latter two are correlation-based deformation methods: one of these is the original Smith (1996) methodology, while the other method replaces $\chi(h_{ij}^*)$ in (14) with pairwise correlation $\rho(h_{ij}^*)$ as the dependence measure, and replaces the theoretical $\chi(h^*)$ function with the stationary Matérn correlation function detailed in (6). Note that in both of the latter two methods, correlation is estimated on a Gaussian marginal scale, and for the former two methods, we set $q = 0.9$ in (15) and (18).

As detailed in Section 2.5, we evaluate the efficacy of each of the four deformations by fitting a model to the extremal dependence of the sample. We fit the same dependence model five times: once using the sampling locations in the original G-plane and then once for each of the respective D-plane sampling locations given from the four deformations. For each fitted model, we calculate the CLAIC given in (21). Ordering of the CLAIC allows us to determine which deformation method (if any) was the most effective in accounting for the non-stationarity in that sample. As the underlying process from which the sample is drawn is known, we fit a stationary extremal dependence model of an appropriate class. That is, for processes that are asymptotically dependent, we fit a stationary Brown-Resnick model, and for processes that are asymptotically independent, we fit a stationary inverted Brown-Resnick model.

This procedure is repeated for 50 different samples of a single process. In this simulation study, each deformation for each sample is created using the same anchor points. For each sample, we determine which deformation method was the most effective and the proportion of times this occurred over all samples is reported, with the results in Tables 1, 2 and 3. These results show that stationary dependence models for non-stationary spatial processes routinely provide a better fit if the deformation methodology is used as a preprocessing step. We also show that the original Smith (1996) deformation is outperformed by our extensions.

3.1 Non-stationary Brown-Resnick and inverted Brown-Resnick process

The first setting we consider consists of replications of a non-stationary Brown-Resnick, and inverted Brown-Resnick, process sampled at 64 equally spaced locations on $[-1, 1] \times [-1, 1]$. We use a non-stationary variogram in the exponent function in (16) to ensure that $\chi(h_{ij})$ is not simply a function of distance. In the context of non-stationary Gaussian processes, Fouedjio et al. (2015) propose a semivariogram of the form $\gamma^*(s_i, s_j)$ where

$$\gamma^*(s_i, s_j) = \gamma(\|\psi(s_i) - \psi(s_j)\|), \quad (23)$$

and

$$\psi(s) = o + (s - o)\|s - o\|$$

is a radial basis function with some centre point o and $\gamma(\cdot)$ is the stationary and isotropic semivariogram given in (17). The use of the radial basis function $\psi(s)$ within this semivariogram causes pairs that are closer to o to be more strongly dependent than those pairs that are further away. From (16) and (17), the Brown-Resnick process with this semivariogram has theoretical $\chi(s_i, s_j)$ given by

$$\chi(s_i, s_j) = 2 - 2\Phi \left\{ \frac{\|\psi(s_i) - \psi(s_j)\|^{\kappa/2}}{\lambda^{\kappa/2}\sqrt{2}} \right\}, \quad (24)$$

for locations s_i, s_j and $\kappa \in (0, 2], \lambda > 0$. For this study, we take the centre o to be the origin and use scale and shape parameters $\lambda = 2$ and $\kappa = 0.8$ in (24). To illustrate the process a high resolution realisation is given in Supplementary Figure S2. Simulations are produced using the method of Dieker and Mikosch (2015).

Table 1 contains some interesting results. Most notably, in all cases a deformation has aided in model fitting when compared to using the original simulation grid. For both the max-stable, and inverted max-stable, cases, improvements on the efficacy of the original Smith (1996) method are made by utilising the Frobenius norm in the objective function. However, it is not entirely clear whether use of an extremal dependence measure for creating deformations is necessary in this case. We often found that deforming the space using measures for dependence throughout the distribution created better deformations than those using extremal dependence measures. We believe that this is because the variance of the estimator for $\rho(h_{ij}^*)$ is much lower than that of $\chi(h_{ij}^*)$, as we use all of the data to estimate correlation, and that there are strong similarities in patterns of spatial non-stationarity for the central- and extremal-dependence structures of this process. We next consider other processes with more complicated dependence structures.

3.2 Max-mixture process

We now consider the hybrid dependence model, detailed in full by Wadsworth and Tawn (2012). Let $X(s)$ be a max-stable process and $Y(s)$ an asymptotically independent spatial process, both with standard Fréchet margins. For $\omega \in [0, 1]$, $H(s) = \max\{\omega X(s), (1 - \omega)Y(s)\}$ is an asymptotically

dependent spatial process with standard Fréchet margins. In particular, we take $X(s)$ to be the non-stationary Brown-Resnick process detailed in Section 3.1 and $Y(s)$ to be a marginally transformed stationary Gaussian process with the Matérn correlation structure detailed in (6).

It can be shown that the theoretical $\chi(h_{ij})$ values for $H(s)$ are the same as for $X(s)$, but multiplied by ω . There is no closed form for the correlation for $H(s)$ on the Gaussian scale. Computationally, it can be shown that it is a mixture of the correlation from both $X(s)$ and $Y(s)$. As such, we would expect the extremal dependence and central dependence of $H(s)$ to be mixtures of those coming from $X(s)$ and $Y(s)$, with different amounts of mixing occurring for both. We set ω to be 0.3 and take $(\theta_1, \theta_2) = (1, 1.2)$ in (6).

By construction of $H(s)$, taking its reciprocal creates an asymptotically independent process on standard exponential margins, as with the inverted max-stable process. As in Section 3.1, the simulation study is repeated separately for the asymptotically dependent and asymptotically independent mixtures. The results are given in Table 2.

In contrast to the results given in Table 1, Table 2 shows a clearer need for an extremal dependence-based approach when creating deformations for a process that exhibits more complicated dependence structures. Here this max-mixture process is designed to represent a process with a mixture of stationarity in both the extremal dependence and dependence throughout the distribution. We now consider a process that has non-stationary extremal dependence, but is nearly stationary in the body.

3.3 Gaussian mixture process

With previous simulations, we found it is sometimes sufficient to simply use measures of central dependence when deforming the spatial domain to create a process with a more stationary extremal dependence structure. This is because the central- and extremal-dependence structures of these processes are closely related and using either approach typically creates similar deformations. In applications, we may find that these structures are not so closely related. As such, we are motivated to consider a process that is designed to have completely different dependence in the body to the tails.

Let $Y_S(s), Y_{NS}(s)$ be stationary and non-stationary Gaussian processes, respectively, each with

standard Gaussian margins. We then consider the process

$$Y^*(s) = \begin{cases} Y_S(s), & \text{if } \Phi(Y(s_0)) \leq p \\ Y_{NS}(s), & \text{if } \Phi(Y(s_0)) > p \end{cases}, \quad (25)$$

where $s_0 \in S$ is a fixed location, $\Phi(\cdot)$ is the standard Gaussian cdf, and $p \in [0, 1]$ is a probability. By specifying $Y^*(s)$ in this manner, we create a process with an extremal dependence structure determined mostly by the correlation structure of Y_{NS} and with dependence through the body determined mostly by Y_S . Simulation of this process is simple; we draw $Y(s_0) \sim N(0, 1)$ and then simulate the rest of the field conditioning on that value and whether $\Phi(Y(s_0)) \leq p$ or $\Phi(Y(s_0)) > p$.

For this particular study, we use replications of this Gaussian mixture sampled at 81 equally spaced locations in $[-1, 1] \times [-1, 1]$. We take s_0 to be the origin and $p = 0.9$. Both Y_S and Y_{NS} are specified to have the Matérn correlation structure given in (6), with respective parameter sets $\boldsymbol{\theta}^{(S)} = (\theta_1^{(S)}, \theta_2^{(S)})$ and $\boldsymbol{\theta}^{(NS)} = (\theta_1^{(NS)}, \theta_2^{(NS)}, o)$. Note that $\boldsymbol{\theta}^{(NS)}$ contains an extra parameter as we use the difference of the radial basis functions given in (23) and detailed by (Fouedjio et al., 2015) as a measure of pairwise distance, rather than Euclidean distance. The parameters for this study are set to $\boldsymbol{\theta}^{(S)} = (2, 1)$ and $\boldsymbol{\theta}^{(NS)} = (2, 0.8, (0, 0))$. Results are given in Table 3.

Table 3 highlights a clear need for extremal dependence-based methods when creating deformations for processes that have different patterns of non-stationarity in their central- and extremal dependence structures. In contrast to the results given in the previous studies, here using $\chi(h_{ij}^*)$ or $\chi_q(h_{ij}^*)$ is always favoured.

4 Case studies

We present two case studies using our deformation methodology. In both cases, we follow the procedure set out in Section 2.4. However, as we consider relatively large spatial domains we use Great Earth distance in place of Euclidean distance for h and h^* . We consider 30 different initial index sets, taking the best deformation over all sets. Here we define the best deformation to be that which provides the lowest objective value in (14) whilst remaining a bijective mapping. When using extremal dependence

measures, we focus on deformations based on χ^{BR} only, following the justification in Section 2.2. We then fit max-stable and inverted max-stable models to the data in the G-plane and D-plane, comparing the model fits using CLAIC estimates. For both studies, all pairs of sampling locations are used in model fitting and the block size in (22) corresponds to a season. We propose two diagnostics for scrutinising the model fits and deformations.

4.1 Australian summer temperatures

Data consist of daily summer (DJF) maximum near-surface air temperatures taken from the HadGHCND global gridded dataset (Caesar et al., 2006) and interpolated to 72 grid point locations covering Australia, for the period 1957-2014. Previous analysis of this data has been conducted using the multivariate conditional extremes model, detailed in Section 1.2.2 (Winter et al., 2016) and its spatial extension (Wadsworth and Tawn, 2019). Figure 2 shows the original sampling locations and estimated pairwise $\chi(h_{ij})$ against distances. We estimate $\chi(h_{ij})$ by setting $q = 0.98$ in (15). The deformation was produced using $m^* = 18$, i.e. a quarter of the original sampling locations. These are presented as the blue points on Figures 2 and 3, where the latter figure depicts the sampling locations in the D-plane. Figure 3 also presents $\hat{\chi}(h_{ij}^*)$ against distance in the deformed space. We observe that the deformation has created a process that appears to be much more stationary with regards to the $\chi(h_{ij}^*)$ estimates in the new coordinates.

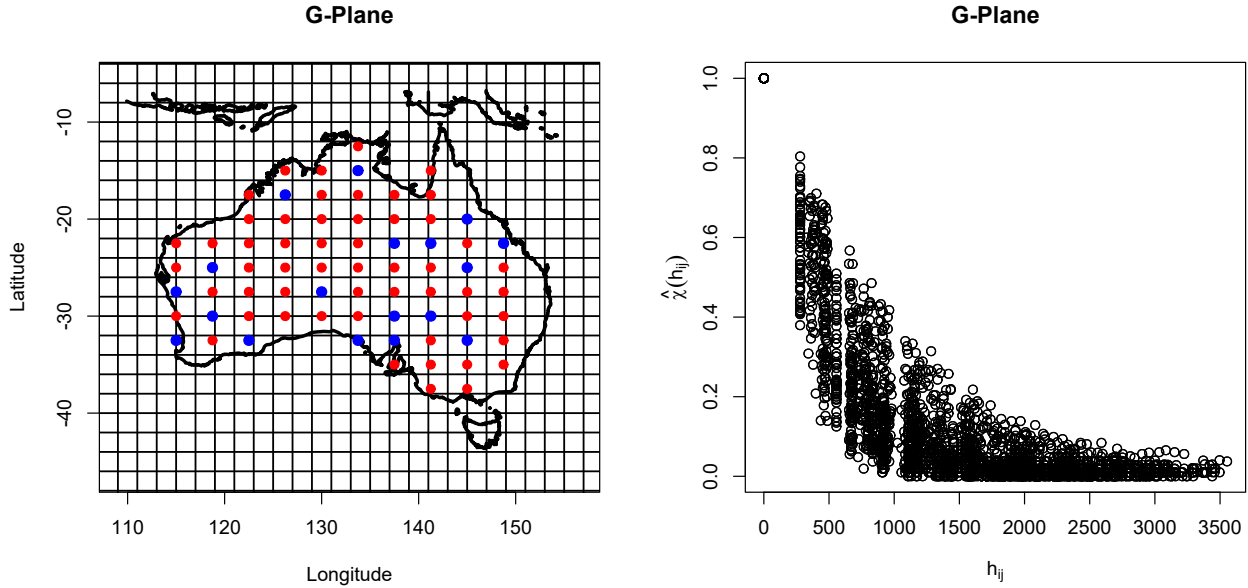


Figure 2: Australia summer temperatures. Left: the original 72 sampling locations. The blue points are the anchor points used for the thin-plate splines. Right: empirical $\chi(h_{ij})$ measures against distance (km). Estimates $\hat{\chi}(h_{ij})$ are calculated above a threshold given by the 98% empirical quantile.

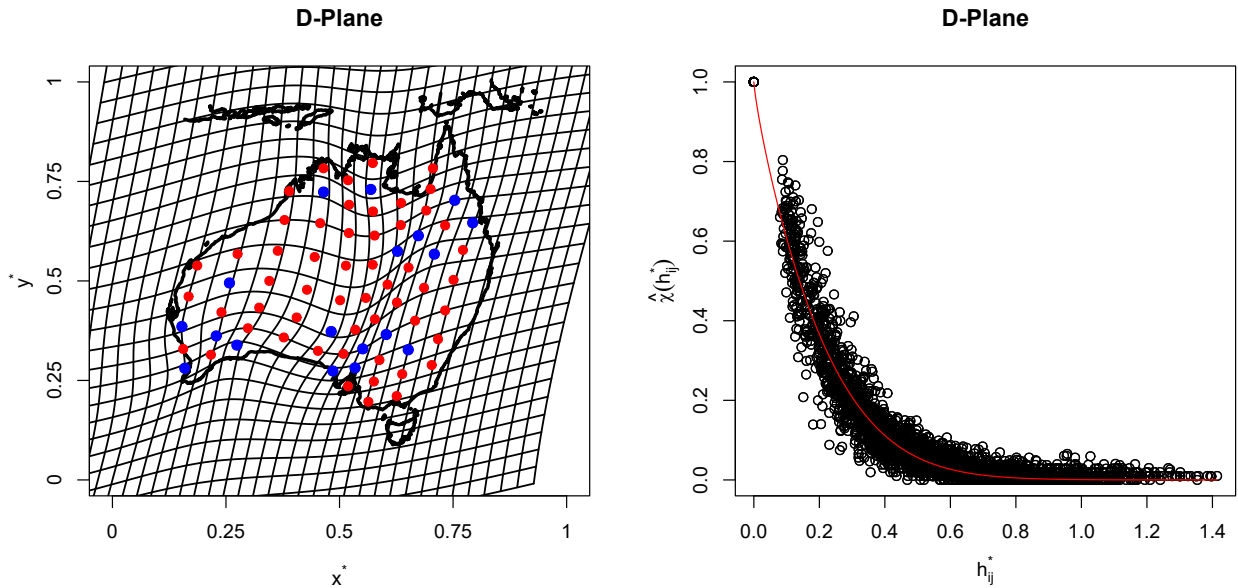


Figure 3: Australia summer temperatures. Left: the 72 sampling locations in the D-plane. The blue points are the anchor points used for the thin-plate splines. The coordinates have been scaled to $[0, 1] \times [0, 1]$, which equals the aspect ratio of the left plot in Figure 2. Right: empirical $\chi(h_{ij}^*)$ measures against distance in the D-plane. The red line gives the fitted function $\chi(h^*)$ used in the deformation.

The fits of the max-stable and inverted max-stable models are summarised in Table 4. The CLAIC estimates suggest that a max-stable model is more appropriate for the data. This becomes even more apparent when we consider that fitting an inverted Brown-Resnick model yields an inverted Smith model as the best fit. These processes are typically quite smooth and often provide unrealistic representations of actual data. However, we note that when naively fitting models on the G-plane, the inverted Smith model provided the lowest CLAIC estimate. This is further evidence that non-stationarity in this data should be incorporated into the modelling procedure.

We use two diagnostics to scrutinise the deformation and the model fit. As our deformation method is tailored to $\chi(h_{ij}^*)$, we seek to use other extremal dependence measures to verify that the resulting deformation is not subject to overfitting. To do this, the conditional extremes model described in Section 1.2.2 is fitted pairwise and the parameter estimates are used to calculate the conditional expectation of one variable when the other variable is at the modelling threshold u , taken as the 98% quantile of the marginal distribution. For each pair, $(X(s_i), X(s_j)), i \neq j$, we have

$$\mathbb{E}[X(s_j)|X(s_i) = u] = \hat{\alpha}u + u^{\hat{\beta}}\hat{\mu},$$

where $(\hat{\alpha}, \hat{\beta}, \hat{\mu})$ are the maximum likelihood estimates for the model. For a stationary and isotropic process, we would expect this measure to be a smooth function of Euclidean distance. The conditional expectation is plotted against distance for both the process on the G-plane and the D-plane.

A second diagnostic is used to evaluate the best model fit in the D-plane. As we have used $\chi(h_{ij}^*)$ to create the deformations, we compare the theoretical triple-wise χ , which we denote $\chi(s_i^*, s_j^*, s_k^*) = \chi(s_i, s_j, s_k)$, from the model fits against empirical estimates. The triple-wise χ is defined as

$$\begin{aligned} \chi(s_i^*, s_j^*, s_k^*) &= \lim_{q \rightarrow 1} \Pr[F_i\{Z(s_i^*)\} > q, F_j\{Z(s_j^*)\} > q | F_k\{Z(s_k^*)\} > q] \\ &= \lim_{q \rightarrow 1} \Pr[F_i\{Z(s_i)\} > q, F_j\{Z(s_j)\} > q | F_k\{Z(s_k)\} > q] = \chi(s_i, s_j, s_k) \end{aligned}$$

for $i \neq j \neq k$. For a Brown-Resnick process, the theoretical value for this measure is

$$\chi(s_i^*, s_j^*, s_k^*) = 3 - V_2(1, 1; i, j) - V_2(1, 1; i, k) - V_2(1, 1; j, k) + V_3(1, 1, 1),$$

where $V_2(\cdot, \cdot; l, m)$ is the pairwise exponent given in (10) and $V_3(\cdot, \cdot, \cdot)$ is the triple-wise exponent measure, for which the parametric form is given in Huser and Davison (2013); recall that if the process is stationary, both of these are functions of Euclidean distance. A similar parametrisation can be given for $\chi_q(s_i^*, s_j^*, s_k^*)$ for an inverted Brown-Resnick process, which is $\chi_q(s_i^*, s_j^*, s_k^*) = (1 - q)^{V_3(1,1,1)-1}$.

Standard errors for empirical estimates of $\chi_q(s_i^*, s_j^*, s_k^*)$ are estimated using a stationary bootstrap (Politis and Romano, 1994). We begin by drawing a random block size B from a geometric distribution with mean K . The bootstrap sample for locations s_i, s_j, s_k , $i \neq j \neq k$ is built by drawing a random starting time τ and creating a block of observations

$$\{\mathbf{z}_\tau^*, \dots, \mathbf{z}_{\tau+B-1}^*\}, \quad \text{where } \mathbf{z}_t^* = \{z_t(s_i), z_t(s_j), z_t(s_k)\},$$

which we add to the bootstrap sample. This procedure is repeated and the bootstrap is built up iteratively until it has length n . We then estimate $\chi(s_i^*, s_j^*, s_k^*)$ for that sample and repeat for a number of samples. When choosing locations to compare empirical and theoretical values of $\chi(s_i^*, s_j^*, s_k^*)$, we take advantage of the gridded structure of the coordinates in the G-plane, and ensure that each set of points share roughly the same configuration and pairwise distances. This is used to evaluate the stationarity of the dependence structure on the original G-plane, as we would expect the empirical values of $\chi(s_i^*, s_j^*, s_k^*)$ to be consistently similar across sets of locations with the same configuration.

Diagnostics for the deformations and best model fit are given in Figure 4. For the estimation of $\chi(s_i^*, s_j^*, s_k^*)$, 30 sets of three adjacent locations along the north/south transect in the G-plane are randomly selected and a stationary bootstrap with mean block size $K = 14$ and 1000 samples is used to create 95% confidence intervals for the empirical estimates of $\chi(s_i^*, s_j^*, s_k^*)$. Empirical estimates of $\chi(s_i^*, s_j^*, s_k^*)$ are calculated above the 98% quantile. The right panel of Figure 4 displays estimates for the conditional expectation from the conditional extremes model, where distances are normalised so that the average distance is equal for both the values in the G-plane and the D-plane.

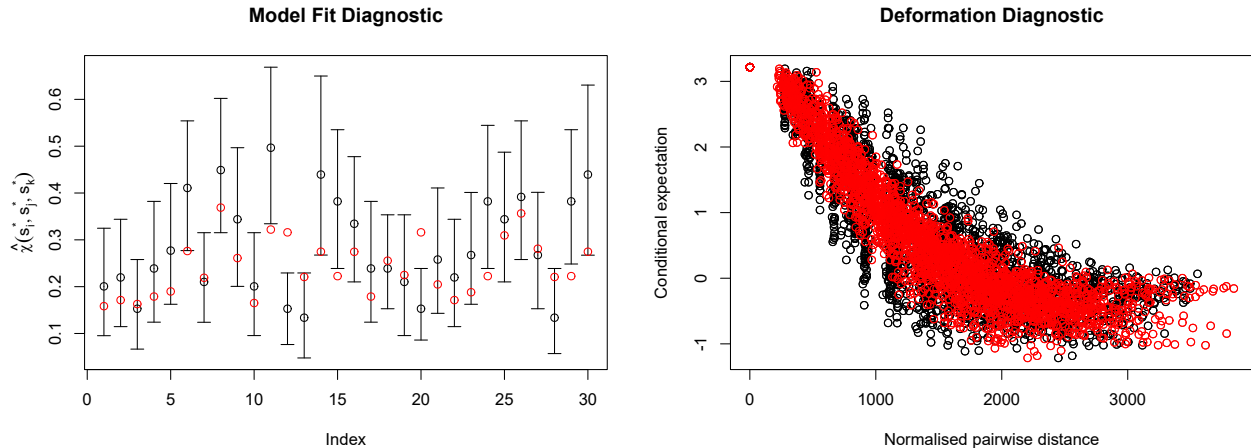


Figure 4: Australian summer temperatures diagnostics. Left: estimates of $\chi(s_i^*, s_j^*, s_k^*)$ (black dots) and 95% confidence intervals using the stationary bootstrap. Red dots are the respective theoretical values suggested by the model fit. Right: conditional expectation from conditional extremes model. Red points denote estimates for the process on the D-plane; black points are those on the G-plane.

The diagnostic based on $\chi(s_i^*, s_j^*, s_k^*)$ from Figure 4 suggests that a max-stable model is a reasonable fit for the data in the deformed space, as the patterns of the theoretical $\chi(s_i^*, s_j^*, s_k^*)$ values follow the empirical estimates. The large variability in the bootstrap estimates across sets of locations with similar configurations suggests that the process on the original plane is highly non-stationary. Estimates from the conditional extremes model provide further evidence that the deformation has produced something more stationary with regards to the dependence structure, especially at smaller distances. The use of a measure for extremal dependence that is not used for fitting lends credibility to the $\chi(h_{ij}^*)$ plot in Figure 3 and suggests that the deformation has worked well.

4.2 UK precipitation rate

Data consist of hourly precipitation rate (mm/day) observed at locations on two 10×10 grids; the first is centred in Snowdonia, Wales and the second is centred in the Scottish Highlands. Observations are taken from the UK climate projections 2018 (UKCP18) (Lowe et al., 2018) which contain values produced at hourly intervals on $2.2 \times 2.2 \text{ km}^2$ grid boxes between the years 1980 and 2000. We have treated the centre of each grid box as a sampling location and we take every fifth grid box to create the 10×10 grid of sampling locations. Observations are aggregated to 12-hr intervals, beginning at

12pm, and to remove the seasonal effect often observed in precipitation data, we have taken only winter observations (DJF). This leaves 3600 observations at each sampling location.

Figure 5 shows both sets of original sampling locations and their respective estimates of $\chi(h_{ij})$ against distances. In both cases, we estimate $\chi(h_{ij})$ by setting $q = 0.95$ in (15). Both deformations are produced using $m^* = 25$ and these are presented as the blue points in Figure 5. Figure 6 presents both deformations and estimates of $\hat{\chi}(h_{ij}^*)$ against distance in the respective deformed spaces. We observe that both deformations have created a process that appears to be much more stationary with regards to their respective $\chi(h_{ij}^*)$ estimates in the new coordinates. In both cases, deformations are more prominent around areas of higher elevation.

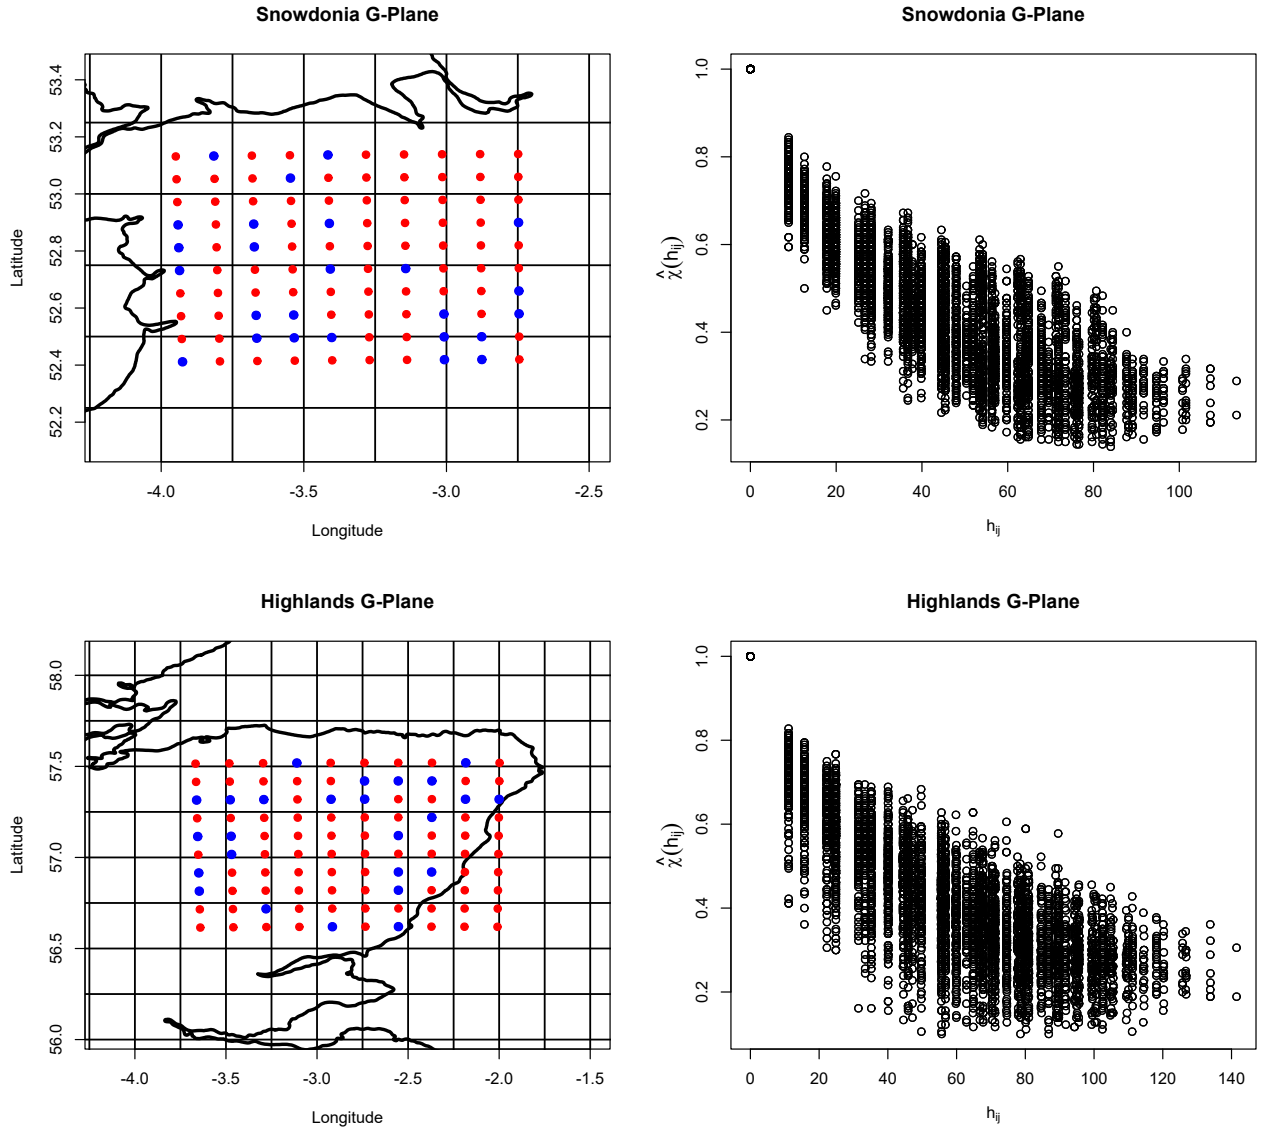


Figure 5: Top row: Snowdonia. Bottom row: Scottish Highlands. Left: the original 100 sampling locations. The blue points are the anchor points used for the thin-plate splines. Right: empirical $\chi(h_{ij})$ measures against distance (km) in the respective G-planes. Estimates $\hat{\chi}(h_{ij})$ are calculated above a threshold given by the 95% empirical quantile.

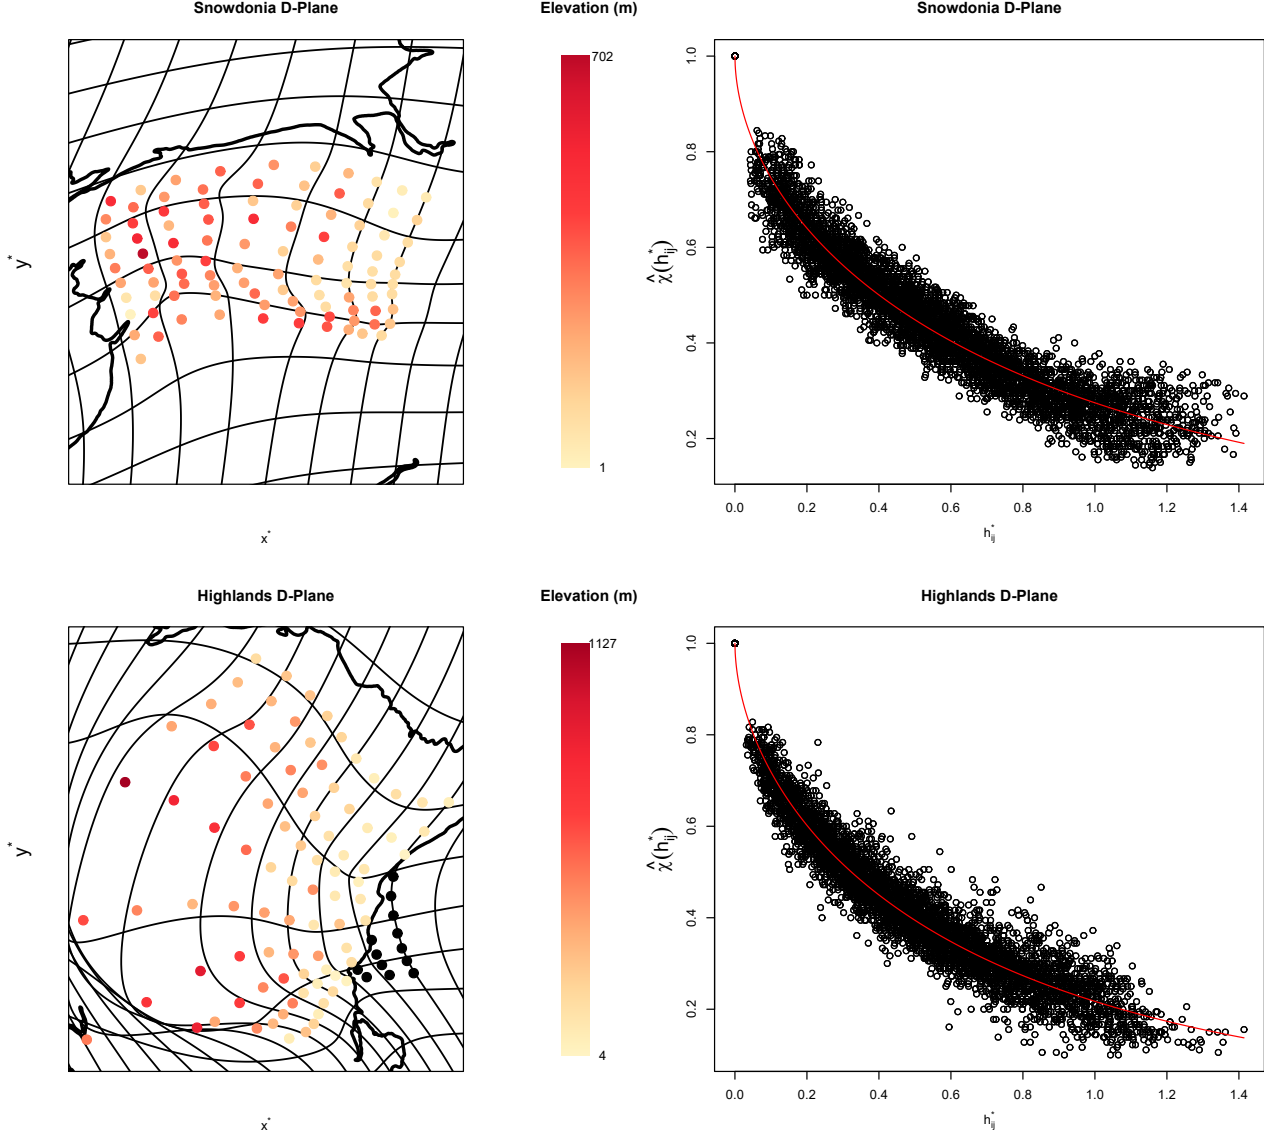


Figure 6: Top row: Snowdonia. Bottom row: Scottish Highlands. Left: the 100 sampling locations in their respective D-planes. The points are coloured such that darker points correspond to sampling locations with higher elevation and black points correspond to locations over sea. The coordinates have been scaled to $[0, 1] \times [0, 1]$, which equals the aspect ratio of the left plots in Figure 5. Right: empirical $\chi(h_{ij}^*)$ measures against distance in the D-plane. The red line gives the fitted function $\chi(h^*)$ used in the deformation.

Table 5 summarises the fits for the Brown-Resnick and inverted Brown-Resnick models for both sets of sampling locations. The CLAIC estimates in Table 5 suggest that an inverted max-stable model is the most appropriate for both the Snowdonia and Highlands data. Both see improved fits using the sampling locations mapped to the respective D-planes. In Figures 7 and 8, we present diagnostics for

the deformations and best model fits using the same measures described in Section 4.1. As the best fitting model for both datasets is the inverted Brown-Resnick process, Figure 7 compares empirical estimates and model-based values of $\chi_q(s_i^*, s_j^*, s_k^*)$ with $q = 0.95$. Confidence intervals for the empirical estimates of $\chi_q(s_i^*, s_j^*, s_k^*)$ are calculated by randomly selecting 30 sets of three adjacent locations along the east/west transect and using stationary bootstrap with mean block size $K = 14$ and 1000 samples. For the diagnostic given in Figure 8, the 95% quantile is used for fitting the conditional extremes model and we plot the pairwise conditional expectation estimates against distance. Distances are normalised so that the average distance is equal for both the values in the G-plane and the D-plane.




Figure 7: UK precipitation model fit diagnostics. Estimates of $\chi_q(s_i^*, s_j^*, s_k^*)$ (black dots) with $q = 0.95$ and 95% confidence intervals using the stationary bootstrap. Red dots are the respective theoretical values suggested by the model fits. Left: Snowdonia. Right: Highlands.




Figure 8: UK precipitation deformation diagnostics. Conditional expectation from conditional extremes model. Red points denote estimates for the process on the D-plane; black points are those on the G-plane. Left: Snowdonia. Right: Highlands.

Figure 7 shows that the inverted max-stable model gives a relatively good fit to the extremal dependence of both datasets with sampling locations mapped to their respective D-planes, but the fit appears better for the Scottish Highlands. The low variability in the $\chi_q(s_i^*, s_j^*, s_k^*)$ estimates suggests that the original process may not be highly non-stationary. The pairwise conditional expectation estimates in Figure 8 suggest that both deformations have produced a more stationary process, albeit more so in the case of the Snowdonia D-plane. The small change in the Highlands estimates may suggest that overfitting to the $\chi(h_{ij})$ values has occurred, especially when compared to the Snowdonia estimates. This may also explain the stronger agreement of the $\chi_q(s_i^*, s_j^*, s_k^*)$ measures in Figure 8. To investigate the possibility of overfitting, we recreated the diagnostic using deformations created with fewer anchor points, but this did not show any improvements.

5 Discussion

In this paper, we presented a simple yet effective approach to modelling non-stationary extremal dependence. This approach extends that of Sampson and Guttorp (1992) and Smith (1996) to be applicable for modelling extremal dependence, rather than dependence throughout the body. We do this by replacing the objective function in these methods with the Frobenius norm of the difference between empirical, and theoretical, pairwise dependency matrices, with the theoretical measures coming from a

stationary dependence model. Although most of our focus is on $\chi(h_{ij}^*)$ as the dependence measure, we have also shown that this is easily replaced by other measures, such as $\chi_q(h_{ij}^*)$ and correlation. Model selection is carried out using pairwise composite likelihoods and CLAIC estimation and we propose diagnostics for evaluating these model fits.

We presented two case studies; in each scenario, we showed that when modelling the extremal dependence of the data using stationary models, better fits are provided using our methodology. Here we have fit very simple models to the data. However, in practice these deformations may be used as a pre-processing step to reveal covariates or orography that can be incorporated into the modelling procedure. Two diagnostics were introduced and used to provide evidence that our approach has produced a process which is more stationary with regards to the extremal dependence.

As with many areas of extreme value analysis, there is a bias-variance trade-off present when estimating χ_q . Using values of q closer to 1 puts greater focus on extremal dependence at the expense of increased variance of the estimator. In Sections 3 and 4, we choose q close to 1 whilst preserving some initial spatial structure observed in the χ estimates. However, if q is too high then it is possible that any structure is masked by the high variability of the estimators and the deformation methodology is likely to fail in such circumstances. We have not considered the effect of estimator variability on the deformation, but note this could form a future research direction.

A further issue that could be considered is the possible non-bijectivity of the mapping used in the deformation. We detail an approach to reduce this in Section 3, however, this method is not particularly robust. Bijectivity of deformations must be checked by eye which can become cumbersome when a large number are produced. To avoid this necessary supervision, the G-plane can be represented as a Delaunay triangulation, see Iovleff and Perrin (2004) and Youngman (2020). Incorporating this extra computational aspect into the model adds to the complexity, and so as to preserve the simplicity of our approach, we leave this as a future consideration.

Acknowledgements

We thank the associate editor and two referees for comments that improved the manuscript.

We gratefully acknowledge funding through the STOR-i Doctoral Training Centre and Engineering and Physical Sciences Research Council (grant EP/L015692/1 and fellowship

EP/P002838/1). The authors are grateful to Simon Brown and Robert Shooter of the Met Office Hadley Centre, UK, for access to data and preliminary code.

Data Accessibility

The data that supports the findings in this study are available in the accompanying R package, `sdfEXTREME`, which can be found at <https://github.com/Jbrich95/sdfExtreme>.

References

- Blanchet, J. and Creutin, J.-D. (2017). Co-occurrence of extreme daily rainfall in the French Mediterranean region. *Water Resources Research*, 53(11):9330–9349.
- Blanchet, J. and Davison, A. C. (2011). Spatial modeling of extreme snow depth. *The Annals of Applied Statistics*, 5(3):1699–1725.
- Brown, B. M. and Resnick, S. I. (1977). Extreme Values of Independent Stochastic Processes. *Journal of Applied Probability*, 14(4):732–739.
- Caesar, J., Alexander, L., and Vose, R. (2006). Large-scale changes in observed daily maximum and minimum temperatures: Creation and analysis of a new gridded data set. *Journal of Geophysical Research: Atmospheres*, 111(D5).
- Casson, E. and Coles, S. (1999). Spatial regression models for extremes. *Extremes*, 1(4):449–468.
- Castro-Camilo, D. and Huser, R. (2019). Local likelihood estimation of complex tail dependence structures, applied to U.S. precipitation extremes. *Journal of the American Statistical Association*, 0(0):1–29.
- Chevalier, C., Ginsbourger, D., and Martius, O. (2017). Modeling non-stationary extreme dependence with stationary max-stable processes and multidimensional scaling. *arXiv e-prints*, page arXiv:1711.01878.
- Cooley, D., Nychka, D., and Naveau, P. (2007). Bayesian spatial modeling of extreme precipitation return levels. *Journal of the American Statistical Association*, 102(479):824–840.

- Davison, A. C., Huser, R., and Thibaud, E. (2013). Geostatistics of dependent and asymptotically independent extremes. *Mathematical Geosciences*, 45(5):19. 511–529.
- Davison, A. C., Padoan, S. A., and Ribatet, M. (2012). Statistical Modeling of Spatial Extremes. *Statistical Science*, 27(2):161–186.
- de Haan, L. (1984). A spectral representation for max-stable processes. *Ann. Probab.*, 12(4):1194–1204.
- Dieker, A. B. and Mikosch, T. (2015). Exact simulation of Brown-Resnick random fields at a finite number of locations. *Extremes*, 18(2):301–314.
- Eastoe, E. F. (2019). Nonstationarity in peaks-over-threshold river flows: A regional random effects model. *Environmetrics*, 30(5):e2560. e2560 env.2560.
- Fouedjio, F., Desassis, N., and Romary, T. (2015). Estimation of space deformation model for non-stationary random functions. *Spatial Statistics*, 13:45–61.
- Green, P. and Silverman, B. (1994). *Nonparametric regression and generalized linear models: a roughness penalty approach*. Chapman and Hall, United Kingdom.
- Guttorp, P. and Sampson, P. D. (1994). Methods for estimating heterogeneous spatial covariance functions with environmental applications. *Handbook of statistics*, 12:661–689.
- Heffernan, J. E. and Tawn, J. A. (2004). A conditional approach for multivariate extreme values (with discussion). *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 66(3):497–546.
- Huser, R. and Davison, A. C. (2013). Composite likelihood estimation for the Brown–Resnick process. *Biometrika*, 100(2):511–518.
- Huser, R. and Genton, M. G. (2016). Non-Stationary Dependence Structures for Spatial Extremes. *Journal of Agricultural, Biological, and Environmental Statistics*, 21(3):470–491.
- Huser, R., Opitz, T., and Thibaud, E. (2017). Bridging asymptotic independence and dependence in spatial extremes using gaussian scale mixtures. *Spatial Statistics*, 21:166 – 186.
- Huser, R. and Wadsworth, J. L. (2019). Modeling spatial processes with unknown extremal dependence class. *Journal of the American Statistical Association*, 114(525):434–444.

- Iovleff, S. and Perrin, O. (2004). Estimating a nonstationary spatial structure using simulated annealing. *Journal of Computational and Graphical Statistics*, 13(1):90–105.
- Joe, H. (1997). *Multivariate models and multivariate dependence concepts*. CRC Press.
- Jonathan, P., Randell, D., Wu, Y., and Ewans, K. (2014). Return level estimation from non-stationary spatial data exhibiting multidimensional covariate effects. *Ocean Engineering*, 88:520 – 532.
- Kabluchko, Z., Schlather, M., and de Haan, L. (2009). Stationary max-stable fields associated to negative definite functions. *Annals of Probability*, 37(5):2042–2065.
- Ledford, A. W. and Tawn, J. A. (1996). Statistics for near independence in multivariate extreme values. *Biometrika*, 83(1):169–187.
- Lowe, J. A., Bernie, D., Bett, P., Bricheno, L., Brown, S., Calvert, D., Clark, R., Eagle, K., Edwards, T., Fosser, G., et al. (2018). UKCP18 science overview report. *Met Office Hadley Centre: Exeter, UK*.
- Mannshardt-Shamseldin, E. C., Smith, R. L., Sain, S. R., Mearns, L. O., and Cooley, D. (2010). Downscaling extremes: A comparison of extreme value distributions in point-source and gridded precipitation data. *The Annals of Applied Statistics*, 4(1):484–502.
- Meiring, W., Monestiez, P., Sampson, P. D., and Guttorp, P. (1997). Developments in the modelling of non stationary spatial covariance structure from space-time monitoring data. Unpublished manuscript. <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.56.8670>.
- Opitz, T., Huser, R., Bakka, H. C., and Rue, H. (2018). INLA goes extreme: Bayesian tail regression for the estimation of high spatio-temporal quantiles. *Extremes*, 21:441–462.
- Padoan, S. A., M., and Sisson, S. A. (2010). Likelihood-based inference for max-stable processes. *Journal of the American Statistical Association*, 105(489):263–277.
- Perrin, O. and Meiring, W. (1999). Identifiability for non-stationary spatial structure. *Journal of Applied Probability*, 36(4):1244–1250.
- Perrin, O. and Senoussi, R. (2000). Reducing non-stationary random fields to stationarity and isotropy using a space deformation. *Statistics & Probability Letters*, 48(1):23 – 32.

- Politis, D. N. and Romano, J. P. (1994). The stationary bootstrap. *Journal of the American Statistical Association*, 89(428):1303–1313.
- Resnick, S. (2002). Hidden regular variation, second order regular variation and asymptotic independence. *Extremes*, 5(4):303–336.
- Ribatet, M. (2013). Spatial extremes: Max-stable processes at work. *Journal de la Société Française de Statistique*, 154(2):156–177.
- Ross, E., Randell, D., Ewans, K., Feld, G., and Jonathan, P. (2017). Efficient estimation of return value distributions from non-stationary marginal extreme value models using bayesian inference. *Ocean Engineering*, 142:315 – 328.
- Sampson, P. D. (2010). Spatial deformation models. In Gelfand, A. E., Diggle, P., Guttorp, P., and Fuentes, M., editors, *Handbook of Spatial Statistics*, chapter Constructions for Nonstationary Spatial Processes, pages 124–126. CRC press.
- Sampson, P. D. and Guttorp, P. (1992). Nonparametric Estimation of Nonstationary Spatial Covariance Structure. *Journal of the American Statistical Association*, 87(417):108–119.
- Sang, H. and Gelfand, A. E. (2010). Continuous spatial process models for spatial extreme values. *Journal of Agricultural, Biological, and Environmental Statistics*, 15(1):49–65.
- Schlather, M. (2002). Models for stationary max-stable random fields. *Extremes*, 5(1):33–44.
- Schlather, M. and Tawn, J. A. (2003). A dependence measure for multivariate and spatial extreme values: Properties and inference. *Biometrika*, 90(1):139–156.
- Shooter, R., Ross, E., Tawn, J. A., and Jonathan, P. (2019). On spatial conditional extremes for ocean storm severity. *Environmetrics*, 30(6).
- Smith, R. L. (1990). Max-stable processes and spatial extremes. Unpublished manuscript. <https://rls.sites.oasis.unc.edu/postscript/rs/spatex.pdf>.
- Smith, R. L. (1996). Estimating nonstationary spatial correlations. Unpublished manuscript. <https://rls.sites.oasis.unc.edu/postscript/rs/nonstationary.pdf>.

- Varin, C., Reid, N., and Firth, D. (2011). An overview of composite likelihood methods. *Statistica Sinica*, 21(1):5–42.
- Wadsworth, J. L. and Tawn, J. (2019). Higher-dimensional spatial extremes via single-site conditioning. *arXiv e-prints*, arXiv:1912.06560.
- Wadsworth, J. L. and Tawn, J. A. (2012). Dependence modelling for spatial extremes. *Biometrika*, 99(2):253–272.
- Wadsworth, J. L. and Tawn, J. A. (2014). Efficient inference for spatial extreme value processes associated to log-Gaussian random functions. *Biometrika*, 101(1):1–15.
- Winter, H., Tawn, J. A., and Brown, S. (2016). Modelling the effect of the El Nino-Southern Oscillation on extreme spatial temperature events over Australia. *Annals of Applied Statistics*, 10(4):2075–2101.
- Youngman, B. D. (2019). Generalized additive models for exceedances of high thresholds with an application to return level estimation for U.S. wind gusts. *Journal of the American Statistical Association*, 114(528):1865–1879.
- Youngman, B. D. (2020). Flexible models for nonstationary dependence: Methodology and examples. *arXiv e-prints*, arXiv:2001.06642.
- Zanini, E., Eastoe, E., Jones, M. J., Randell, D., and Jonathan, P. (2020). Flexible covariate representations for extremes. *Environmetrics*, 31(5):e2624.

Process (G-plane)	Fitted Model (D-plane)	Deformation Method	Proportion of lowest CLAIC	
Non-stationary Brown-Resnick	Stationary Brown-Resnick	None	0	0.34
		χ^{BR}	0.22	
		χ_q^{IBR}	0.12	
		ρ	0.44	
		Smith (1996)	0.22	
Non-stationary Inverted Brown-Resnick	Stationary Inverted Brown-Resnick	None	0	0.56
		χ^{BR}	0.24	
		χ_q^{IBR}	0.32	
		ρ	0.28	
		Smith (1996)	0.16	

Table 1: Proportion of lowest CLAIC estimates provided by fitting models to deformations for 50 realisations of non-stationary Brown-Resnick and inverted Brown-Resnick processes. The CLAIC has been estimated with a block size of $b = 1$, corresponding to temporal independence. Composite likelihoods are estimated with the threshold in (20) as the 90% empirical quantile, which is also used for estimating $\chi(h_{ij}^*)$ in (15).

Process (G-plane)	Fitted Model (D-plane)	Deformation Method	Proportion of lowest CLAIC	
Asymptotically-dependent Max-mixture	Stationary Brown-Resnick	None	0.06	0.78
		χ^{BR}	0.14	
		χ_q^{IBR}	0.64	
		ρ	0.16	
		Smith (1996)	0	
Asymptotically-independent Max-mixture	Stationary Inverted Brown-Resnick	None	0	0.90
		χ^{BR}	0.42	
		χ_q^{IBR}	0.48	
		ρ	0.06	
		Smith (1996)	0.04	

Table 2: Proportion of lowest CLAIC estimates provided by fitting models to deformations of 50 realisations of asymptotically dependent and asymptotically independent max-mixture processes. The CLAIC has been estimated with a block size of $b = 1$, corresponding to temporal independence. Composite likelihoods are estimated with the threshold in (20) as the 90% empirical quantile, which is also used for estimating $\chi(h_{ij}^*)$ in (15).

Process (G-plane)	Fitted Model (D-plane)	Deformation Method	Proportion of lowest CLAIC
Gaussian Mixture	Stationary Inverted Brown-Resnick	None	0
		χ^{BR}	0.08
		χ_q^{IBR}	0.92
		ρ	0
		Smith (1996)	0

Table 3: Proportion of lowest CLAIC estimates provided by fitting models to deformations of 50 realisations of the Gaussian mixture process, see (25). The CLAIC has been estimated with a block size of $b = 1$, corresponding to temporal independence. Composite likelihoods are estimated with the threshold in (20) as the 90% empirical quantile, which is also used for estimating $\chi(h_{ij}^*)$ in (15).

	Model	Negative Composite Log-Likelihood ($\times 10^6$)	$(\hat{\kappa}, \hat{\lambda})$ (2 d.p.)	CLAIC ($\times 10^7$)
G-Plane	IMSP*	3.078	(2.00, 1048.20)	6.157
	MSP	3.078	(1.59, 358.30)	6.157
D-Plane	IMSP*	3.074	(2.00, 2.61)	6.148
	MSP	3.073	(1.71, 0.95)	6.146

Table 4: Model parameters and diagnostics for the Australian summer temperatures data. Composite likelihoods are estimated with the threshold in (20) as the 98% empirical quantile. (* estimated using Smith process likelihood). CLAIC and negative composite log-likelihood estimates are given to four significant figures.

		Model	Negative Composite Log-Likelihood ($\times 10^6$)	$(\hat{\kappa}, \hat{\lambda})$ (2 d.p.)	CLAIC ($\times 10^7$)
Snowdonia	G-Plane	IMSP	8.023	(1.40, 111.84)	1.605
		MSP	8.050	(1.00, 25.96)	1.610
	D-Plane	IMSP	8.011	(1.29, 3.33)	1.602
		MSP	8.037	(0.93, 0.69)	1.607
Highlands	G-Plane	IMSP	8.099	(1.25, 143.77)	1.620
		MSP	8.124	(0.87, 27.37)	1.625
	D-Plane	IMSP	8.076	(1.30, 3.34)	1.615
		MSP	8.099	(0.93, 0.69)	1.620

Table 5: Model parameters and diagnostics for the UK precipitation data. Composite likelihoods are estimated with the threshold in (20) as the 95% empirical quantile. CLAIC and negative composite log-likelihood estimates are given to four significant figures.