

An M-estimator of spatial tail dependence

John Einmahl¹, Anna Kiriliouk², Andrea Krajina³ and Johan Segers²

¹*Department of Econometrics & OR and CentER, Tilburg University, P.O. Box 90153, 5000 LE Tilburg, The Netherlands. E-mail: j.h.j.einmahl@uvt.nl*

²*Institut de Statistique, Biostatistique et Sciences Actuarielles, Université catholique de Louvain, Voie du Roman Pays 20, B-1348 Louvain-la-Neuve, Belgium. E-mail:*

anna.kiriliouk@uclouvain.be, johan.segers@uclouvain.be

³*Institute of Mathematical Stochastics, University of Göttingen, Göttingen, Germany. E-mail: andrea.krajina@mathematik.uni-goettingen.de*

Abstract

Tail dependence models for distributions attracted to a max-stable law are fitted using observations above a high threshold. To cope with spatial, high-dimensional data, a rank-based M-estimator is proposed relying on bivariate margins only. A data-driven weight matrix is used to minimize the asymptotic variance. Empirical process arguments show that the estimator is consistent and asymptotically normal. Its finite-sample performance is assessed in simulation experiments involving popular max-stable processes perturbed with additive noise. An analysis of wind speed data from the Netherlands illustrates the method.

Keywords: Brown–Resnick process; exceedances; multivariate extremes; ranks; spatial statistics; stable tail dependence function

JEL codes: C13, C14

1 Introduction

Max-stable random processes have become the standard for modelling extremes of environmental quantities, such as wind speed, precipitation, or snow depth. In such a context, data are modelled as realizations of spatial processes, observed at a finite number of locations. The statistical problem then consists of modelling the joint tail of a multivariate distribution. This problem can be divided into two separate issues: modelling the marginal distributions and modelling the dependence structure. A popular practice is to transform the marginals to an appropriate form and to fit a max-stable model to componentwise monthly or annual maxima using composite likelihood methods. This is done either in a frequentist setting (Padoan et al., 2010; Davison et al., 2012) or in a Bayesian one (Reich and Shaby, 2012; Cooley et al., 2012). Alternatively, Yuen and Stoev (2013) propose an M-estimator based on finite-dimensional cumulative distribution functions. Popular parametric models for max-stable processes include the ones proposed by Smith (1990), Schlather (2002) and Kabluchko et al. (2009), going back to Brown and Resnick (1977). Recent review articles on spatial extremes include Cooley et al. (2012), Davison et al. (2012) and Ribatet (2013).

The above approaches consider block maxima, whereas more information can be extracted from the data by using all data vectors of which at least one component is large. Although

threshold-based methods are common in multivariate extreme value theory, in spatial extremes they are only starting to be developed. A first example is de Haan and Pereira (2006), where several one- and two-dimensional models for spatial extremes are proposed. Another parametric model for spatial tail dependence is introduced in Buishand et al. (2008). The parameter estimator is shown to be asymptotically normal and the method is applied to daily rainfall data. In Huser and Davison (2014), a pairwise censored likelihood is used to analyze space-time extremes. The method is applied to an extension of Schlather’s model. Another study of space-time extremes can be found in Davis et al. (2013), where asymptotic normality of the pairwise likelihood estimators of the parameters of a Brown–Resnick process is proven for a jointly increasing number of spatial locations and time points. In Jeon and Smith (2012), bivariate threshold exceedances are modelled using a composite likelihood procedure. Asymptotic normality of the estimator is obtained by assuming second-order regular variation for the distribution function that is in the max-domain of attraction of an extreme value distribution. A numerical study comparing two distinct approaches for composite likelihoods can be found in Bacro and Gaetan (2013). In Wadsworth and Tawn (2013), a censored Poisson process likelihood is considered in order to simplify the likelihood expressions in the Brown–Resnick process and in Engelke et al. (2014), the distribution of extremal increments of processes that are in the max-domain of attraction of the Brown–Resnick process is investigated. Finally, in Bienvenüe and Robert (2014), a censored likelihood procedure is used to fit high-dimensional extreme value models for which the tail dependence function has a particular representation.

The above methods all require estimation of the tails of the marginal distributions. This is not necessarily an easy task if the number of variables is large. Moreover, they are likelihood-based and therefore cannot be used to fit, e.g., spectrally discrete max-stable models (Wang and Stoev, 2011).

The aim of this paper is to propose a new method for fitting multivariate tail dependence models to high-dimensional data arising for instance in spatial statistics. No likelihoods come into play as our approach relies on the stable tail dependence function, which is related to the upper tail of the underlying cumulative distribution function. The method is threshold-based in the sense that a data point is considered to be extreme if the rank of at least one component is sufficiently high. The only assumption is that, upon standardization of the margins, the underlying distribution is attracted to a parametrically specified multivariate extreme value distribution, see (2.2) below.

By reducing the data to their ranks, the tails of the univariate marginal distributions need not be estimated. Indeed, the marginal distributions are not even required to be attracted to an extreme value distribution. Another advantage of the rank-based approach is that the estimator is invariant under monotone transformations of the margins, notably for Box–Cox type of transformations.

Our starting point is Einmahl et al. (2012), where an M-estimator for a parametrically modelled tail dependence function in dimension d is derived. However, that method crucially relies on d -dimensional integration, which becomes intractable in high dimensions. This is why we consider tail dependence functions of pairs of variables only. Our estimator is constructed as the minimizer of the distance between a vector of integrals of parametric pairwise tail dependence functions and the vector of their empirical counterparts. The asymptotic variance of the estimator can be minimized by replacing the Euclidean distance by a quadratic form based on a weight matrix estimated from the data. In the simulation studies we will compute estimates in dimensions up to 100.

We show that our estimator is consistent under minimal assumptions and asymptotically

normal under an additional condition controlling the growth of the threshold. In our analysis, we take into account the variability stemming from the rank transformation, the randomness of the threshold, the random weight matrix and, in particular, the fact that the max-stable model is only an approximation in the tail.

A final point worth noticing is the generality of our methodology. Where many studies focus on a concrete parametric (tail) model, ours is generic and makes weak assumptions only. Also, the field of application of extreme value analysis in high dimensions is not restricted to environmental studies: see for example Dematteo et al. (2013), where a spectral clustering approach is introduced and applied to gas pressure data in the shipping industry.

The paper is organized as follows. Section 2 presents the necessary background on multivariate extreme value theory and extremes of stochastic processes. Section 3 contains the definition of the pairwise M-estimator and the main theoretical results on consistency and asymptotic normality, as well as the practical aspects of the choice of the weight matrix. In Section 4 the tail dependence functions of the anisotropic Brown–Resnick process and the Smith model are presented, as well as several simulation studies: two for a large number of locations, illustrating the computational feasibility of the estimator in high dimensions, and one for a smaller number of locations, presenting the benefits of the weight matrix. In addition, we compare the performance of our estimator to the one proposed in Engelke et al. (2014). Finally, we present an application to wind speed data from the Netherlands. Proofs and computations are deferred to the appendices.

2 Background

2.1 Multivariate extreme value theory

Let $\mathbf{X}_i = (X_{i1}, \dots, X_{id})$, $i \in \{1, \dots, n\}$, be independent random vectors in \mathbb{R}^d with common continuous distribution function F and marginal distribution functions F_1, \dots, F_d . Write $M_{nj} = \max_{i=1, \dots, n} X_{ij}$ for $j = 1, \dots, d$. We say that F is in the max-domain of attraction of an extreme-value distribution G if there exist sequences $a_{nj} > 0$ and $b_{nj} \in \mathbb{R}$ for $j = 1, \dots, d$ such that

$$\lim_{n \rightarrow \infty} \mathbb{P} \left[\frac{M_{n1} - b_{n1}}{a_{n1}} \leq x_1, \dots, \frac{M_{nd} - b_{nd}}{a_{nd}} \leq x_d \right] = G(\mathbf{x}), \quad \mathbf{x} \in \mathbb{R}^d. \quad (2.1)$$

The margins, G_1, \dots, G_d , of G are univariate extreme value distributions and the function G is determined by

$$G(\mathbf{x}) = \exp \{ -\ell(-\log G_1(x_1), \dots, -\log G_d(x_d)) \},$$

where $\ell : [0, \infty)^d \rightarrow [0, \infty)$ is called the stable tail dependence function. The distribution function of $(1/\{1 - F_j(X_{1j})\})_{j=1, \dots, d}$ is in the max-domain of attraction of the extreme-value distribution $G_0(\mathbf{z}) = \exp \{ -\ell(1/z_1, \dots, 1/z_d) \}$, $\mathbf{z} \in (0, \infty)^d$, and we can retrieve the function ℓ via

$$\ell(\mathbf{x}) = \lim_{t \downarrow 0} t^{-1} \mathbb{P}[1 - F_1(X_{11}) \leq tx_1 \text{ or } \dots \text{ or } 1 - F_d(X_{1d}) \leq tx_d], \quad \mathbf{x} \in [0, \infty)^d. \quad (2.2)$$

Note that G_0 has unit Fréchet margins, $G_{0,j}(z_j) = \exp(-1/z_j)$ for $z_j > 0$ and $j = 1, \dots, d$. If F is already an extreme-value distribution, then it is attracted by itself. Otherwise, relation (2.1) is equivalent to relation (2.2) and convergence of the d marginal distributions in (2.1).

From now on we will only assume the weaker relation (2.2) instead of (2.1), making no assumptions on the marginal distributions F_1, \dots, F_d except for continuity. The function ℓ is convex, homogeneous of order one and satisfies $\ell(0, \dots, 0, x_j, 0, \dots, 0) = x_j$ for $j = 1, \dots, d$. We assume that ℓ belongs to some parametric family $\{\ell(\cdot; \theta) : \theta \in \Theta\}$, with $\Theta \subset \mathbb{R}^p$. There are numerous such parametric models and new families of models continue to be invented. We will see some examples of parametric stable tail dependence functions in Section 4. For more examples and background on multivariate extreme value theory, see Coles (2001), Beirlant et al. (2004), or de Haan and Ferreira (2006).

2.2 Extremes of stochastic processes

Max-stable processes arise in the study of component-wise maxima of random processes rather than of random vectors. Let \mathcal{S} be a compact subset of \mathbb{R}^2 and let $\mathbb{C}(\mathcal{S})$ denote the space of continuous, real-valued functions on \mathcal{S} , equipped with the supremum norm $\|f\|_\infty = \sup_{\mathbf{s} \in \mathcal{S}} |f(\mathbf{s})|$ for $f \in \mathbb{C}(\mathcal{S})$. The restriction to \mathbb{R}^2 is for convenience only. In the applications to spatial data that we have in mind, \mathcal{S} will represent the region of interest.

Consider independent copies $\{X_i(\mathbf{s})\}_{\mathbf{s} \in \mathcal{S}}$ for $i \in \{1, \dots, n\}$ of a process $\{X(\mathbf{s})\}_{\mathbf{s} \in \mathcal{S}}$ in $\mathbb{C}(\mathcal{S})$. Then X is in the max-domain of attraction of the max-stable process Z if there exist sequences of continuous functions $a_n(\mathbf{s}) > 0$ and $b_n(\mathbf{s})$ such that

$$\left\{ \frac{\max_{i=1, \dots, n} X_i(\mathbf{s}) - b_n(\mathbf{s})}{a_n(\mathbf{s})} \right\}_{\mathbf{s} \in \mathcal{S}} \xrightarrow{w} \{Z(\mathbf{s})\}_{\mathbf{s} \in \mathcal{S}}, \quad \text{as } n \rightarrow \infty,$$

where \xrightarrow{w} denotes weak convergence in $\mathbb{C}(\mathcal{S})$; see de Haan and Lin (2001) for a full characterization of max-domain of attraction conditions for the case $\mathcal{S} = [0, 1]$. A max-stable process Z is called simple if its marginal distribution functions are all unit Fréchet.

Although our interest lies in the underlying stochastic processes X_i , data are always obtained on a finite subset of \mathcal{S} only, i.e., at fixed locations $\mathbf{s}_1, \dots, \mathbf{s}_d$. As a consequence, statistical inference is based on a sample of d -dimensional random vectors. The finite-dimensional distributions of Z are multivariate extreme value distributions. This brings us back to the ordinary, multivariate setting.

3 M-estimator

3.1 Estimation

As in Section 2.1, let $\mathbf{X}_1, \dots, \mathbf{X}_n$ be an independent random sample from a d -variate distribution F with continuous margins and with stable tail dependence function ℓ , see equation (2.2). Assuming that ℓ belongs to a parametric family $\{\ell(\cdot; \theta) : \theta \in \Theta\}$ with $\Theta \subset \mathbb{R}^p$, the goal is to estimate the parameter vector θ .

To this end, we first define a nonparametric estimator of ℓ . Let R_{ij}^n denote the rank of X_{ij} among X_{1j}, \dots, X_{nj} for $j = 1, \dots, d$. Replacing F and F_1, \dots, F_d in (2.2) by their empirical counterparts and replacing t by k/n yields

$$\widehat{\ell}_{n,k}(\mathbf{x}) := \frac{1}{k} \sum_{i=1}^n \mathbb{1} \left\{ R_{i1}^n > n + \frac{1}{2} - kx_1 \text{ or } \dots \text{ or } R_{id}^n > n + \frac{1}{2} - kx_d \right\}. \quad (3.1)$$

For the estimator to be consistent, we need $k = k_n \in \{1, \dots, n\}$ to depend on n in such a way that $k \rightarrow \infty$ and $k/n \rightarrow 0$ as $n \rightarrow \infty$. The estimator was originally defined in the bivariate case in Huang (1992) and Drees and Huang (1998).

Let θ_0 denote the true parameter value, $\ell = \ell(\cdot; \theta_0)$, and let $g = (g_1, \dots, g_q)^T : [0, 1]^d \rightarrow \mathbb{R}^q$ with $q \geq p$ denote a column vector of integrable functions. In Einmahl et al. (2012), an M-estimator of θ_0 is defined by

$$\hat{\theta}'_n := \arg \min_{\theta \in \Theta} \sum_{m=1}^q \left(\int_{[0,1]^d} g_m(\mathbf{x}) \left\{ \hat{\ell}_{n,k}(\mathbf{x}) - \ell(\mathbf{x}; \theta) \right\} d\mathbf{x} \right)^2. \quad (3.2)$$

Under suitable conditions, the estimator $\hat{\theta}'_n$ is consistent and asymptotically normal. The use of ranks via the nonparametric estimator in (3.1) permits to avoid having to fit a model to the (tails of the) marginal distributions. In fact, the only assumption on F , the existence of the stable tail dependence function ℓ in (2.2), is even weaker than the assumption that F belongs to the maximal domain of attraction of a max-stable distribution.

However, the approach is ill-adapted to the spatial setting, where data are gathered from dozens of locations. In high dimensions, the computation of $\hat{\theta}'_n$ becomes infeasible due to the presence of d -dimensional integrals in the objective function in (3.2).

Akin to composite likelihood methods, we opt for a pairwise approach, minimizing over quadratic forms of vectors of two-dimensional integrals. Let q represent the number of pairs of locations that we wish to take into account, so that $p \leq q \leq d(d-1)/2$. Let π be the function from $\{1, \dots, q\}$ to $\{1, \dots, d\}^2$ that describes these pairs, that is, for $m \in \{1, \dots, q\}$, we have $\pi(m) = (\pi_1(m), \pi_2(m)) = (u, v)$ with $1 \leq u < v \leq d$. In the spatial setting (cf. Section 2.2), the indices u and v correspond to locations \mathbf{s}_u and \mathbf{s}_v respectively.

The bivariate margins of the stable tail dependence function $\ell(\cdot; \theta)$ and the nonparametric estimator in (3.1) are given by

$$\begin{aligned} \ell_{\pi(m)}(x_{\pi_1(m)}, x_{\pi_2(m)}; \theta) &= \ell_{uv}(x_u, x_v; \theta) := \ell(0, \dots, 0, x_u, 0, \dots, 0, x_v, 0, \dots, 0; \theta), \\ \hat{\ell}_{n,k,\pi(m)}(x_{\pi_1(m)}, x_{\pi_2(m)}) &= \hat{\ell}_{n,k,uv}(x_u, x_v) := \hat{\ell}_{n,k}(0, \dots, 0, x_u, 0, \dots, 0, x_v, 0, \dots, 0), \end{aligned}$$

respectively. Consider the random $q \times 1$ column vector

$$L_{n,k}(\theta) := \left(\int_{[0,1]^2} \left\{ \hat{\ell}_{n,k,\pi(m)}(x_{\pi_1(m)}, x_{\pi_2(m)}) - \ell_{\pi(m)}(x_{\pi_1(m)}, x_{\pi_2(m)}; \theta) \right\} dx_{\pi_1(m)} dx_{\pi_2(m)} \right)_{m=1}^q.$$

Let $\hat{\Omega}_n \in \mathbb{R}^{q \times q}$ be a symmetric, positive definite, possibly random matrix. Define

$$f_{n,k,\hat{\Omega}_n}(\theta) := L_{n,k}(\theta)^T \hat{\Omega}_n L_{n,k}(\theta), \quad \theta \in \Theta.$$

The pairwise M-estimator of θ is defined as

$$\hat{\theta}_n := \arg \min_{\theta \in \Theta} f_{n,k,\hat{\Omega}_n}(\theta) = \arg \min_{\theta \in \Theta} \left\{ L_{n,k}(\theta)^T \hat{\Omega}_n L_{n,k}(\theta) \right\}. \quad (3.3)$$

The simplest choice for $\hat{\Omega}_n$ is just the $q \times q$ identity matrix I_q , yielding

$$f_{n,k,I_q}(\theta) = \sum_{(u,v)} \left(\int_{[0,1]^2} \left\{ \hat{\ell}_{n,k,uv}(x_u, x_v) - \ell_{uv}(x_u, x_v; \theta) \right\} dx_u dx_v \right)^2. \quad (3.4)$$

Note the similarity of this objective function with the one for the original M-estimator in equation (3.2). The role of the matrix $\widehat{\Omega}_n$ is to be able to assign data-driven weights to quantify the size of the vector of discrepancies $L_{n,k}(\theta)$ via a generalized Euclidian norm. As we will see in Section 3.2, a judicious choice of this matrix will allow to minimize the asymptotic variance.

Comparison with Einmahl et al. (2012).

A natural question that arises is whether the quality of estimation decreases when making the step from the d -dimensional estimator $\widehat{\theta}'_n$ in (3.2) to the pairwise estimator $\widehat{\theta}_n$ in (3.3). We will demonstrate for the multivariate logistic model that this is not the case, necessarily in a dimension where $\widehat{\theta}'_n$ can be computed. The d -dimensional logistic model has stable tail dependence function

$$\ell(\mathbf{x}; \theta) = \left(x_1^{1/\theta} + \dots + x_d^{1/\theta} \right)^\theta,$$

with $\theta \in [0, 1]$. We simulate 200 samples of size $n = 1500$ from the logistic model in dimension $d = 5$ with parameter value $\theta = 0.5$ and we assess the quality of our estimates via the (empirical) bias and root mean squared error (RMSE) for $k \in \{40, 80, \dots, 320\}$. The dashed lines in Figure 1 show the bias and RMSE for the M-estimator of Einmahl et al. (2012) with weight function $g \equiv 1$. The results are the same as in Einmahl et al. (2012, Figure 1). The solid lines show the bias and RMSE for the pairwise M-estimator with $q = 9$ and $\widehat{\Omega}_n$ as in Section 3.2. We see that the pairwise estimator performs somewhat better in terms of bias and also has the lower minimal RMSE, for $k = 160$.




Figure 1: Bias and RMSE for estimators of $\theta = 0.5$ for the logistic model, for the 5-dimensional M-estimator (dashed lines) and the pairwise M-estimator (solid lines); 200 replications of $n = 1500$.

3.2 Asymptotic results and choice of the weight matrix

We show consistency and asymptotic normality of the rank-based pairwise M-estimator. Moreover, we provide a data-driven choice for $\widehat{\Omega}_n$ which minimizes the asymptotic covariance matrix of the limiting normal distribution. Results for the construction of confidence regions and hypothesis tests are presented as well.

A quantity related to the stable tail dependence function ℓ is the exponent measure Λ , which is a measure on $[0, \infty]^d \setminus \{(\infty, \dots, \infty)\}$ determined by

$$\Lambda(\{\mathbf{w} \in [0, \infty]^d : w_1 \leq x_1 \text{ or } \dots \text{ or } w_d \leq x_d\}) = \ell(\mathbf{x}), \quad \mathbf{x} \in [0, \infty)^d.$$

Let W_Λ be a mean-zero Gaussian process, indexed by the Borel sets of $[0, \infty]^d \setminus \{(\infty, \dots, \infty)\}$ and with covariance function

$$\mathbb{E}[W_\Lambda(A_1) W_\Lambda(A_2)] = \Lambda(A_1 \cap A_2),$$

where A_1, A_2 are Borel sets in $[0, \infty]^d \setminus \{(\infty, \dots, \infty)\}$. For $\mathbf{x} \in [0, \infty)^d$, define

$$\begin{aligned} W_\ell(\mathbf{x}) &= W_\Lambda(\{\mathbf{w} \in [0, \infty]^d \setminus \{(\infty, \dots, \infty)\} : w_1 \leq x_1 \text{ or } \dots \text{ or } w_d \leq x_d\}), \\ W_{\ell,j}(x_j) &= W_\ell(0, \dots, 0, x_j, 0, \dots, 0), \quad j = 1, \dots, d. \end{aligned}$$

Let $\dot{\ell}_j$ be the partial derivative of ℓ with respect to x_j , and define

$$B(\mathbf{x}) := W_\ell(\mathbf{x}) - \sum_{j=1}^d \dot{\ell}_j(\mathbf{x}) W_{\ell,j}(x_j), \quad \mathbf{x} \in [0, \infty)^d.$$

For $m \in \{1, \dots, q\}$ with $\pi(m) = (\pi_1(m), \pi_2(m)) = (u, v)$, put

$$B_{\pi(m)}(x_{\pi_1(m)}, x_{\pi_2(m)}) = B_{uv}(x_u, x_v) := B(0, \dots, 0, x_u, 0, \dots, 0, x_v, 0, \dots, 0).$$

Also define the mean-zero random column vector

$$\tilde{B} := \left(\int_{[0,1]^2} B_{\pi(m)}(x_{\pi_1(m)}, x_{\pi_2(m)}) dx_{\pi_1(m)} dx_{\pi_2(m)} \right)_{m=1}^q.$$

The law of \tilde{B} is zero-mean Gaussian and its covariance matrix $\Gamma(\theta) \in \mathbb{R}^{q \times q}$ depends on θ via the model assumption $\ell = \ell(\cdot; \theta)$. For pairs $\pi(m) = (u, v)$ and $\pi(m') = (u', v')$, we can obtain the (m, m') -th entry of $\Gamma(\theta)$ by

$$\Gamma_{(m,m')}(\theta) = \mathbb{E}[\tilde{B}_m \tilde{B}_{m'}] = \int_{[0,1]^4} \mathbb{E}[B_{uv}(x_u, x_v) B_{u'v'}(x_{u'}, x_{v'})] dx_u dx_v dx_{u'} dx_{v'}. \quad (3.5)$$

Define $\psi : \Theta \rightarrow \mathbb{R}^q$ by

$$\psi(\theta) := \left(\int_{[0,1]^2} \ell_{\pi(m)}(x_{\pi_1(m)}, x_{\pi_2(m)}; \theta) dx_{\pi_1(m)} dx_{\pi_2(m)} \right)_{m=1}^q. \quad (3.6)$$

Assuming θ is an interior point of Θ and ψ is differentiable in θ , let $\dot{\psi}(\theta) \in \mathbb{R}^{q \times p}$ denote the total derivative of ψ at θ .

Theorem 3.1 (Existence, uniqueness and consistency). *Let $\{\ell(\cdot; \theta) : \theta \in \Theta\}$, $\Theta \in \mathbb{R}^p$, be a parametric family of d -variate stable tail dependence functions and let $(\pi(m))_{m=1}^q$, with $p \leq q \leq d(d-1)/2$, be q distinct pairs in $\{1, \dots, d\}$ such that the map ψ in (3.6) is a homeomorphism from Θ to $\psi(\Theta)$. Let the d -variate distribution function F have continuous margins and stable tail dependence function $\ell(\cdot; \theta_0)$ for some interior point $\theta_0 \in \Theta$. Let $\mathbf{X}_1, \dots, \mathbf{X}_n$ be an iid sample from F . Assume also that*

(C1) *ψ is twice continuously differentiable on a neighbourhood of θ_0 and $\dot{\psi}(\theta_0)$ is of full rank;*

(C2) *there exists a symmetric, positive definite matrix Ω such that $\widehat{\Omega}_n \xrightarrow{\mathbb{P}} \Omega$.*

Then with probability tending to one, the minimizer $\widehat{\theta}_n$ of $f_{n,k,\widehat{\Omega}_n}$ exists and is unique. Moreover,

$$\widehat{\theta}_n \xrightarrow{\mathbb{P}} \theta_0, \quad \text{as } n \rightarrow \infty.$$

Let $\Delta_{d-1} = \{\mathbf{w} \in [0, 1]^d : w_1 + \dots + w_d = 1\}$ denote the unit simplex in \mathbb{R}^d .

Theorem 3.2 (Asymptotic normality). *If in addition to the assumptions of Theorem 3.1*

(C3) *$t^{-1}\mathbb{P}[1 - F_1(X_{11}) \leq tx_1 \text{ or } \dots \text{ or } 1 - F_d(X_{1d}) \leq tx_d] - \ell(\mathbf{x}; \theta_0) = O(t^\alpha)$ uniformly in $\mathbf{x} \in \Delta_{d-1}$ as $t \downarrow 0$ for some $\alpha > 0$;*

(C4) *$k = o(n^{2\alpha/(1+2\alpha)})$ and $k \rightarrow \infty$ as $n \rightarrow \infty$,*

then

$$\sqrt{k}(\widehat{\theta}_n - \theta_0) \xrightarrow{d} \mathcal{N}_p(0, M(\theta_0))$$

where, for $\theta \in \Theta$ such that $\dot{\psi}(\theta)$ is of full rank,

$$M(\theta) := (\dot{\psi}(\theta)^T \Omega \dot{\psi}(\theta))^{-1} \dot{\psi}(\theta)^T \Omega \Gamma(\theta) \Omega \dot{\psi}(\theta) (\dot{\psi}(\theta)^T \Omega \dot{\psi}(\theta))^{-1}. \quad (3.7)$$

The proofs of Theorems 3.1 and 3.2 are deferred to Appendix A.

An asymptotically optimal choice for the random weight matrix $\widehat{\Omega}_n$ would be one for which the limit Ω minimizes the asymptotic covariance matrix in (3.7) with respect to the positive semi-definite partial ordering on the set of symmetric matrices. This minimization problem shows up in other contexts as well, and its solution is well-known: provided $\Gamma(\theta)$ is invertible, the minimum is attained at $\Omega = \Gamma(\theta)^{-1}$, the matrix $M(\theta)$ simplifying to

$$M_{\text{opt}}(\theta) = (\dot{\psi}(\theta)^T \Gamma(\theta)^{-1} \dot{\psi}(\theta))^{-1}, \quad (3.8)$$

see for instance Abadir and Magnus (2005, page 339). However, this choice of the weight matrix requires the knowledge of θ , which is unknown. The solution consists of computing the optimal weight matrix evaluated at a preliminary estimator of θ .

For $\theta \in \Theta$, let H_θ be the spectral measure related to $\ell(\cdot; \theta)$ (de Haan and Resnick, 1977; Resnick, 1987): it is a finite measure defined on the unit simplex Δ_{d-1} and it satisfies

$$\ell(\mathbf{x}; \theta) = \int_{\Delta_{d-1}} \max_{j=1, \dots, d} \{w_j x_j\} H_\theta(d\mathbf{w}), \quad \mathbf{x} \in [0, \infty)^d.$$

Corollary 3.3 (Optimal weight matrix). *In addition to the assumptions of Theorem 3.2, assume the following:*

(C5) *for all θ in the interior of Θ , the matrix $\Gamma(\theta)$ in (3.5) has full rank;*

(C6) *the mapping $\theta \mapsto H_\theta$ is weakly continuous at θ_0 .*

Assume $\hat{\theta}_n^{(0)}$ converges in probability to θ_0 and let $\hat{\theta}_n$ be the pairwise M-estimator with weight matrix $\hat{\Omega}_n = \Gamma(\hat{\theta}_n^{(0)})^{-1}$. Then, with M_{opt} as in (3.8), we have

$$\sqrt{k}(\hat{\theta}_n - \theta_0) \xrightarrow{d} \mathcal{N}_p(0, M_{\text{opt}}(\theta_0)), \quad n \rightarrow \infty.$$

For any choice of the positive definite matrix Ω in (3.7), the difference $M(\theta) - M_{\text{opt}}(\theta)$ is positive semi-definite.

In view of Corollary 3.3, we propose the following two-step procedure:

1. Compute the pairwise M-estimator $\hat{\theta}_n^{(0)}$ with the weight matrix equal to the identity matrix, i.e., by minimizing f_{n,k,I_q} in (3.4).
2. Calculate the pairwise M-estimator $\hat{\theta}_n$ by minimizing $f_{n,k,\hat{\Omega}_n}$ with $\hat{\Omega}_n = \Gamma^{-1}(\hat{\theta}_n^{(0)})$.

We will see in Section 4.2 that this choice of $\hat{\Omega}_n$ indeed reduces the estimation error.

Calculating $M(\theta)$ can be a challenging task. The matrix $\Gamma(\theta)$ can become quite large since for a d -dimensional model, the maximal number of pairs is $d(d-1)/2$. In practice we will choose a smaller number of pairs: we will see in Section 4.2 that this even may have positive influence on the quality of our estimator. The entries of $\Gamma(\theta)$ are four-dimensional integrals of $\mathbb{E}[B_{uv}(x_u, x_v)B_{u'v'}(x_{u'}, x_{v'})]$ for $\pi(m) = (u, v)$ and $\pi(m') = (u', v')$, $m = 1, \dots, q$. Appendix B shows how to express the entries of $\Gamma(\theta)$ in the function ℓ and gives details on the implementation.

Finally, we present results that can be used for the construction of confidence regions and hypothesis tests.

Corollary 3.4. *If the assumptions from Corollary 3.3 are satisfied, then*

$$k(\hat{\theta}_n - \theta_0)^T M(\hat{\theta}_n)^{-1} (\hat{\theta}_n - \theta_0) \xrightarrow{d} \chi_p^2, \quad \text{as } n \rightarrow \infty.$$

Let $r < p$ and $\theta = (\theta_1, \theta_2) \in \Theta$ with $\theta_1 \in \mathbb{R}^{p-r}$ and $\theta_2 \in \mathbb{R}^r$. Suppose we want to test $\theta_2 = \theta_2^*$ against $\theta_2 \neq \theta_2^*$. Write $\hat{\theta}_n = (\hat{\theta}_{1n}, \hat{\theta}_{2n})$ and let $M_2(\theta)$ be the $r \times r$ matrix corresponding to the lower right corner of $M(\theta)$.

Corollary 3.5. *If the assumptions from Corollary 3.3 are satisfied and if $\theta_0 = (\theta_1, \theta_2^*) \in \Theta$ for some θ_1 , then*

$$k(\hat{\theta}_{2n} - \theta_2^*)^T M_2(\hat{\theta}_{1n}, \theta_2^*)^{-1} (\hat{\theta}_{2n} - \theta_2^*) \xrightarrow{d} \chi_r^2.$$

We will not prove these corollaries here, since their proofs are straightforward extensions of those in Einmahl et al. (2012, Corollary 4.3; Corollary 4.4).

4 Spatial models

4.1 Theory and definitions

The isotropic Brown–Resnick process on $\mathcal{S} \subset \mathbb{R}^2$ is given by

$$Z(\mathbf{s}) = \max_{i \in \mathbb{N}} \xi_i \exp \{ \epsilon_i(\mathbf{s}) - \gamma(\mathbf{s}) \}, \quad \mathbf{s} \in \mathcal{S},$$

where $\{\xi_i\}_{i \geq 1}$ is a Poisson process on $(0, \infty]$ with intensity measure $\xi^{-2} d\xi$ and $\{\epsilon_i(\cdot)\}_{i \geq 1}$ are independent copies of a Gaussian process with stationary increments, $\epsilon(0) = 0$, variance $2\gamma(\cdot)$, and semi-variogram $\gamma(\cdot)$. The process with $\gamma(\mathbf{s}) = (\|\mathbf{s}\|/\rho)^\alpha$ appears as the only limit of (rescaled) maxima of stationary and isotropic Gaussian random fields (Kablichko et al., 2009); here $\rho > 0$ and $0 < \alpha \leq 2$. Since isotropy is not a reasonable assumption for most spatial applications, we follow Blanchet and Davison (2011) and Engelke et al. (2014) and introduce a transformation matrix V defined by

$$V := V(\beta, c) := \begin{bmatrix} \cos \beta & -\sin \beta \\ c \sin \beta & c \cos \beta \end{bmatrix}, \quad \beta \in [0, \pi/2), c > 0,$$

and a transformed space $\mathcal{S}' = \{V^{-1}\mathbf{s} : \mathbf{s} \in \mathcal{S}\}$, so that an isotropic process on \mathcal{S} is transformed to an anisotropic process on \mathcal{S}' . For $\mathbf{s}' \in \mathcal{S}'$ we focus on the anisotropic Brown–Resnick process

$$Z_V(\mathbf{s}') := Z(V\mathbf{s}') = \max_{i \in \mathbb{N}} \xi_i \exp \{ \epsilon_i(V\mathbf{s}') - \gamma(V\mathbf{s}') \}, \quad (4.1)$$

whose semi-variogram is defined by

$$\gamma_V(\mathbf{s}') := \gamma(V\mathbf{s}') = \left[\mathbf{s}'^T \frac{V^T V}{\rho^2} \mathbf{s}' \right]^{\alpha/2}.$$

The pairwise stable tail dependence function for a pair (u, v) , corresponding to locations $(\mathbf{s}'_u, \mathbf{s}'_v)$, is given by

$$\ell_{uv}(x_u, x_v) = x_u \Phi \left(\frac{a_{uv}}{2} + \frac{1}{a_{uv}} \log \frac{x_u}{x_v} \right) + x_v \Phi \left(\frac{a_{uv}}{2} + \frac{1}{a_{uv}} \log \frac{x_v}{x_u} \right),$$

where $a_{uv} := \sqrt{2\gamma_V(\mathbf{s}'_u - \mathbf{s}'_v)}$ and Φ is the standard normal distribution function. Observe that the choice $\alpha = 2$ leads to

$$a_{uv}^2 = 2\gamma(V(\mathbf{s}'_u - \mathbf{s}'_v)) = (\mathbf{s}'_u - \mathbf{s}'_v)^T \Sigma^{-1} (\mathbf{s}'_u - \mathbf{s}'_v), \quad \text{for some } \Sigma = \begin{bmatrix} \sigma_{11} & \sigma_{12} \\ \sigma_{12} & \sigma_{22} \end{bmatrix},$$

where Σ represents any valid 2×2 covariance matrix. This submodel is known as the Gaussian extreme value process or simply the Smith model (Smith, 1990). We will present simulation studies for processes in the domain of attraction, in the sense of (2.2), of both the Smith model and the anisotropic Brown–Resnick process.

To calculate the weight matrix $\Gamma(\theta)^{-1}$, we will need to compute integrals over the four-dimensional margins of the stable tail dependence function, see (3.5) and Appendix B. In Huser and Davison (2013) the following representation is given for $\ell(x_1, \dots, x_d; \theta)$ for general d . If Z_V is defined as in (4.1) then for $\mathbf{s}_{u_1}, \dots, \mathbf{s}_{u_d} \in \mathbb{R}^2$

$$\ell_{u_1, \dots, u_d}(x_1, \dots, x_d) = \sum_{i=1}^d x_i \Phi_{d-1}(\eta^{(i)}(1/x); R^{(i)}),$$

where

$$\eta^{(i)}(x) = (\eta_1^{(i)}(x_1, x_i), \dots, \eta_{i-1}^{(i)}(x_{i-1}, x_i), \eta_{i+1}^{(i)}(x_{i+1}, x_i), \dots, \eta_d^{(i)}(x_d, x_i)) \in \mathbb{R}^{d-1},$$

$$\eta_j^{(i)}(x_j, x_i) = \sqrt{\frac{\gamma_V(s_{u_i} - s_{u_j})}{2}} + \frac{\log(x_j/x_i)}{\sqrt{2\gamma_V(s_{u_i} - s_{u_j})}} \in \mathbb{R},$$

and $R^{(i)} \in \mathbb{R}^{(d-1) \times (d-1)}$ is the correlation matrix with entries

$$R_{jk}^{(i)} = \frac{\gamma_V(s_{u_i} - s_{u_j}) + \gamma_V(s_{u_i} - s_{u_k}) - \gamma_V(s_{u_j} - s_{u_k})}{2\sqrt{\gamma_V(s_{u_i} - s_{u_j})\gamma_V(s_{u_i} - s_{u_k})}}, \quad j, k = 1, \dots, d; j, k \neq i.$$

4.2 Simulation studies

In order to study the performance of the pairwise M-estimator when the underlying distribution function F satisfies (2.2) for a function ℓ corresponding to the max-stable models described before, we generate random samples from Brown–Resnick processes and Smith models perturbed with additive noise. If $\mathbf{Z} = (Z_1, \dots, Z_d)$ is a max-stable process observed at d locations, then we consider

$$X_j = Z_j + \epsilon_j, \quad j = 1, \dots, d,$$

where ϵ_j are independent half normally distributed random variables, corresponding to the absolute value of a normally distributed random variable with standard deviation $1/2$. All simulations are done in R (R Core Team, 2013). Realizations of Z are simulated using the SpatialExtremes package (Ribatet et al., 2013).

Perturbed max-stable processes on a large grid.

Assume that we have $d = 100$ locations on a 10×10 unit distance grid. We simulate 500 samples of size $n = 500$ from the perturbed Smith model with parameters

$$\Sigma = \begin{bmatrix} 1.0 & 0.5 \\ 0.5 & 1.5 \end{bmatrix},$$

and from a perturbed anisotropic Brown–Resnick process with parameters $\alpha = 1$, $\rho = 3$, $\beta = 0.5$ and $c = 0.5$. Instead of estimating ρ , β , and c directly, we estimate the three parameters of the matrix

$$\mathcal{T} = \begin{bmatrix} \tau_{11} & \tau_{12} \\ \tau_{12} & \tau_{22} \end{bmatrix} = \rho^{-2} V(\beta, c)^T V(\beta, c).$$

In practice, this parametrization, which is in line with the one of the Smith model, often yields better results. We study the bias and root mean squared error (RMSE) for $k \in \{25, 50, 75, 100\}$. We compare the estimators for two sets of pairs: one containing all pairs ($q = 4950$) and one containing only pairs of neighbouring locations ($q = 342$). Although the first option may sound like a time-consuming procedure, estimation of the parameters for one sample takes about 20 seconds for the Smith model and less than two minutes for the anisotropic Brown–Resnick process. We let the weight matrix Ω be the $q \times q$ identity matrix, since for so many pairs a data-driven computation of the optimal weight matrix is too

time-consuming. Figure 2 shows the bias and RMSE of $(\sigma_{11}, \sigma_{22}, \sigma_{12})$ for the Smith model. We see that great improvements are achieved by using only pairs of neighbouring locations and that the thus obtained estimator performs well. Using all pairs causes the parameters to have a large positive bias, which translates into a high RMSE. In general, distant pairs often lead to less dependence and hence less information about ℓ and its parameters. Observe that small values of k are preferable, i.e. $k = 25$ or $k = 50$.

Figure 3 shows the bias and RMSE of the pairwise M-estimators of (α, ρ, β, c) for the anisotropic Brown–Resnick process. We see again that using only pairs of neighbouring locations improves the quality of estimation. The corresponding estimators perform well for the estimation of α , β , and c . The lesser performance when estimating ρ seems to be inherent to the Brown–Resnick process and appears regardless of the estimation procedure: see for example Engelke et al. (2014) or Wadsworth and Tawn (2013), who both report a positive bias of ρ for small sample sizes. Compared to those for the Smith model, the values of k for which the estimation error is smallest are higher, i.e., $k = 50$ or $k = 75$.

A perturbed Brown–Resnick process on a small grid with optimal weight matrix.

We consider $d = 12$ locations on an equally spaced unit distance 4×3 grid. We simulate 500 samples of size $n = 1000$ from an anisotropic Brown–Resnick process with parameters $\alpha = 1.5$, $\rho = 1$, $\beta = 0.25$ and $c = 1.5$. We study the bias, standard deviation, and RMSE for $k \in \{25, 75, 125\}$. Three estimation methods are compared: one containing all pairs ($q = 66$), one containing only pairs of neighbouring locations ($q = 29$), and one using optimal weight matrices chosen according to the two-step procedure described after Corollary 3.3, based on the 29 pairs of neighbouring locations. In line with the theory the weighted estimators have lower (or equal) standard deviation (and RMSE) than the unweighted estimators. The difference is clearest for low k for α and ρ .

Comparison with Engelke et al. (2014).

To compare the pairwise M-estimator with the one from Engelke et al. (2014), we consider the setting used in the simulation study of the latter paper: we simulate 500 samples of size 8000 of the *univariate* Brown–Resnick process on an equidistant grid on the interval $[0, 3]$ with step size 0.1. The parameters of the model are $(\alpha, \rho) = (1, 1)$. We estimate the unknown parameters for $k = 500$ and 140 pairs, so that the locations of the selected pairs are at most a distance 0.5 apart. We use the identity weight matrix, since in this particular setting the weight matrix is very large and, as far as we could tell from some preliminary experiments, it leads to only a small reduction in estimation error. Asymptotically we see a reduction of the standard deviations of about 13% for α and 3% for ρ . In Figure 5 below, the results are presented in the form of boxplots, to facilitate comparison with Figure 2 in Engelke et al. (2014). Our procedure turns out to perform equally well for the estimation of α and only slightly worse when estimating ρ . It is to be kept in mind that, whereas the method in Engelke et al. (2014) is tailor-made for the Brown–Resnick process, our method is designed to work for general parametric models.

4.3 Application: speeds of wind gusts

Using extreme value theory to estimate the frequency and magnitude of extreme wind events or to estimate the return levels for (extremely) long return periods is not a novelty in the fields of meteorology and climatology. Numerous research papers published in the last 20–25




Figure 2: Bias and RMSE for estimators of $\sigma_{11} = 1$ (top), $\sigma_{22} = 1.5$ (middle), and $\sigma_{12} = 1.5$ (bottom) for the perturbed 100-dimensional Smith model with identity weight matrix; 500 replications of $n = 500$.

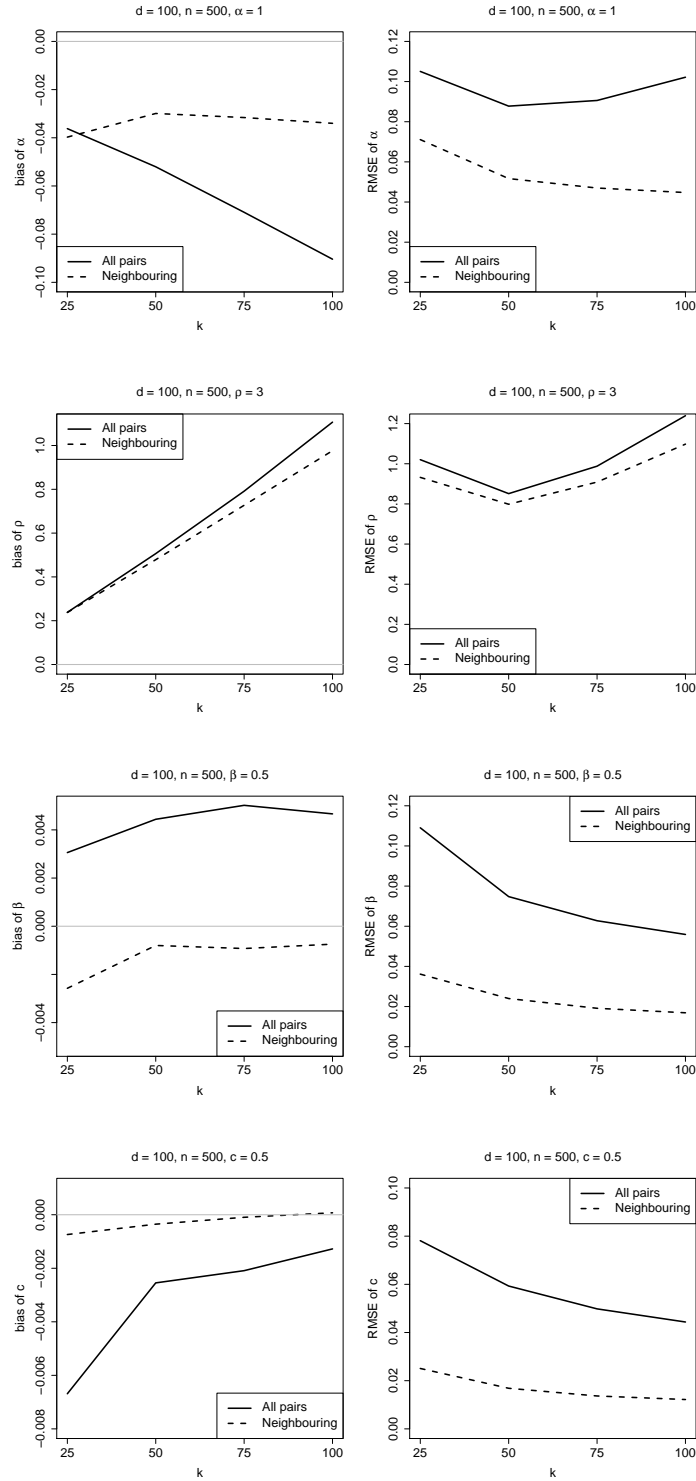


Figure 3: Bias and RMSE for estimators of $\alpha = 1$ (top), $\rho = 3$ (top middle), $\beta = 0.5$ (bottom middle) and $c = 0.5$ (bottom) for the perturbed 100-dimensional Brown–Resnick process with identity weight matrix; 500 replications of $n = 500$.




Figure 4: Bias, standard deviation and RMSE for estimators of $\alpha = 1.5$ (top), $\rho = 1$ (top middle), $\beta = 0.25$ (bottom middle) and $c = 1.5$ (bottom) for the perturbed 12-dimensional Brown–Resnick process; 500 replications of $n = 1000$.

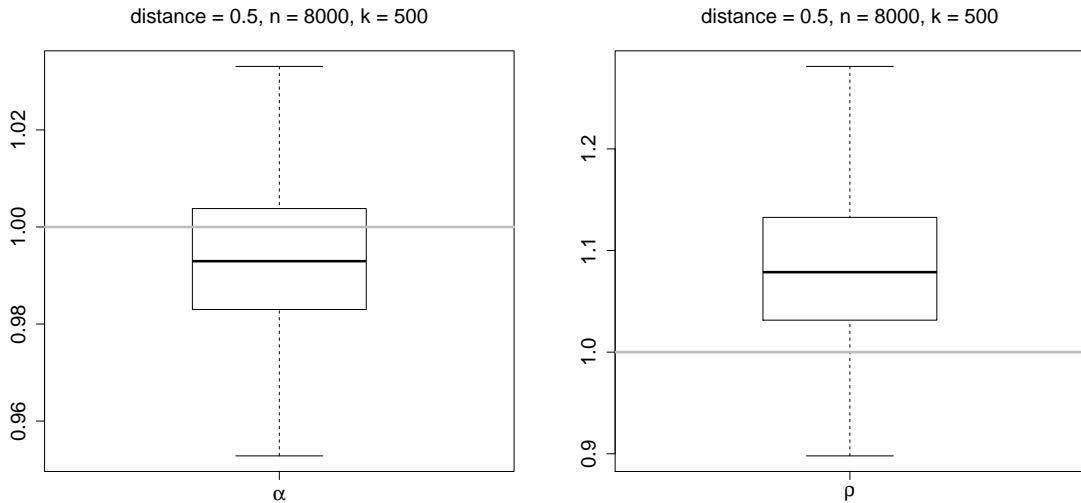


Figure 5: Boxplots of estimators of $\alpha = 1$ (left) and $\rho = 1$ (right) for a univariate Brown–Resnick process on the interval $[0, 3]$ with $d = 31$; 500 replications of $n = 8000$, $k = 500$.

years are applying methods from extreme value theory to treat those estimation problems, see, for example, Karpa and Naess (2013); Ceppi et al. (2008); Palutikof et al. (1999) and the references therein. However, until very recently, all statistical approaches were univariate. The scientific and computational advancements nowadays facilitate the usage of high-dimensional or spatial models. In Engelke et al. (2014) and Oesting et al. (2013), for instance, Brown–Resnick processes are used to model wind speed data.

We consider a data set from the Royal Netherlands Meteorological Institute (KNMI), consisting of the daily maximal speeds of wind gusts, which are measured in 0.1 m/s. The data are observed at 35 weather stations in the Netherlands, over the time period from January 1, 1990 to May 16, 2012. The data set is freely available from http://www.knmi.nl/climatology/daily_data/selection.cgi. Due to the strong influence of the sea on the wind speeds in the coastal area, we only consider the inland stations, of which we removed three stations with more than 1000 missing observations. The thus obtained 22 stations and the remaining amount of missing data per station are shown in the left panel of Figure 6. We aggregate the daily maxima to three-day maxima in order to minimize temporal dependence and we also restrict our observation period to the summer season (June, July and August) to obtain more or less equally distributed data. Finally, we use the “complete deletion approach” for the remaining missing data and obtain a data set with $n = 563$ observations.

We consider the function ℓ corresponding to the Brown–Resnick process (see Section 4.1) and we first test, based on the $q = 29$ pairs of stations that are at most 50 kilometers apart, if the isotropic process suffices for the above data. In the reparametrization introduced in Section 4.2, the case $\tau_{11} = \tau_{22}$ and $\tau_{12} = 0$ corresponds to isotropy. Using Corollary 3.5 and the test statistic

$$k(\hat{\tau}_{11} - \hat{\tau}_{22}, \hat{\tau}_{12}) M_2(\hat{\alpha}, \hat{\tau}_{11} + \hat{\tau}_{22}, 0, 0)^{-1} (\hat{\tau}_{11} - \hat{\tau}_{22}, \hat{\tau}_{12})^T,$$

with $k = 50$, we obtain a value of 0.227, corresponding to a p -value of 0.89, so we can not

reject the null hypothesis and we proceed to estimate the parameters corresponding to the isotropic Brown–Resnick process, using the optimal weight matrix chosen according to the two-step procedure described after Corollary 3.3. The estimates, with standard errors in parentheses, are $\alpha = 0.408$ (0.171) and $\rho = 0.634$ (0.344) for $k = 50$. We see that the Smith model would not fit these data well since α is much smaller than 2.

To visually assess the goodness-of-fit, we plot, in the right panel of Figure 6, the 29 nonparametric estimates of the extremal coefficient function $\ell(1, 1)$ and the values computed from the model (the curve) against the estimated distances $a_{uv} = \sqrt{2\gamma_V(\mathbf{s}_u - \mathbf{s}_v)}$ for the 29 pairs of stations. It is more in line with our M-estimator, which uses integration over $[0, 1]^2$, to focus on the center $(1/2, 1/2)$ instead of the vertex $(1, 1)$ of the unit square. Hence, we use the homogeneity of ℓ to replace $\ell(1, 1)$ with $2\ell(1/2, 1/2)$ and then estimate the latter with $2\hat{\ell}_{n,k}(1/2, 1/2)$, see (3.1). The nonparametric estimates of $\ell(1, 1)$ in the figure are obtained in this way. Observe that these estimates of $\ell(1, 1)$ have a reasonably high variability, even when considering only a small interval for the distances a_{uv} . This explains why the curve cannot fit the points very well.

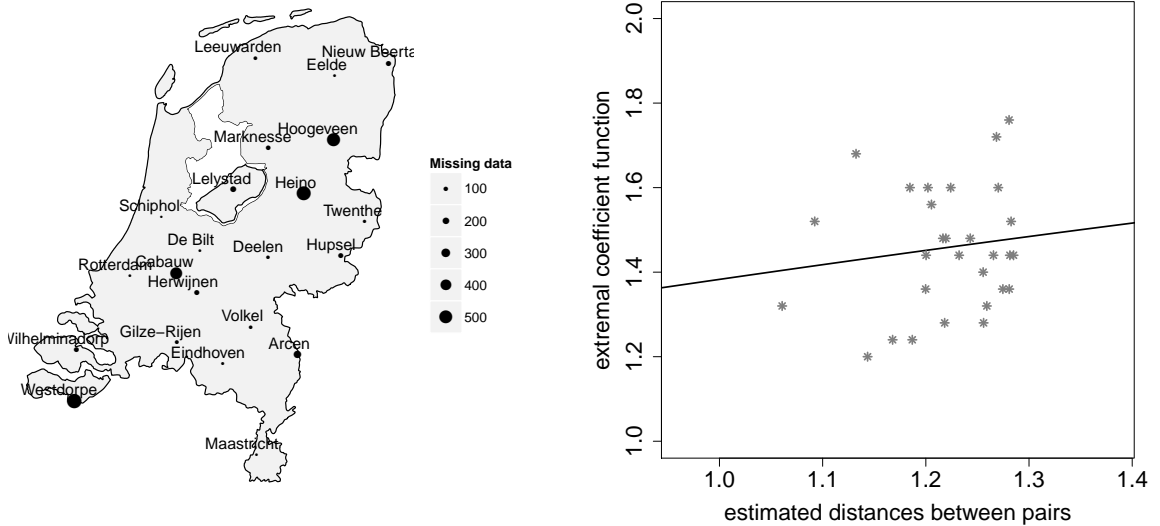


Figure 6: KNMI weather stations (left). Estimates of the extremal coefficient function (right).

A Proofs

The notations are as in Section 3. Let $\hat{\Theta}_n$ denote the (possibly empty) set of minimizers of the function

$$f_{n,k,\hat{\Omega}_n}(\theta) = L_{n,k}(\theta)^T \hat{\Omega}_n L_{n,k}(\theta) =: \|L_{n,k}(\theta)\|_{\hat{\Omega}_n}^2.$$

Write δ_0 for the Dirac measure concentrated at zero. Recall that to each $m \in \{1, \dots, q\}$ there corresponds a pair of indices $\pi(m) = (u, v)$ with $1 \leq u < v \leq d$. Let $\mu = (\mu_1, \dots, \mu_q)^T$ denote a column vector of measures on \mathbb{R}^d whose m -th element is defined as

$$\mu_m(d\mathbf{x}) = \mu_m(dx_1 \times \dots \times dx_d) = \mu_{m1}(dx_1) \times \dots \times \mu_{md}(dx_d) := dx_u dx_v \prod_{j \neq u,v} \delta_0(dx_j),$$

so that μ_{mj} is the Lebesgue measure if $j = u$ or $j = v$, and μ_{mj} is the Dirac measure at zero for $j \neq u, v$. Using the measures μ_m allows us to write

$$L_{n,k}(\theta) = \left(\int_{[0,1]^d} \left\{ \widehat{\ell}_{n,k}(\mathbf{x}) - \ell(\mathbf{x}; \theta) \right\} \mu_m(d\mathbf{x}) \right)_{m=1}^q = \int \widehat{\ell}_{n,k} \mu - \psi(\theta).$$

Lemma A.1. *If $0 < \lambda_{n,1} \leq \dots \leq \lambda_{n,q}$ and $0 < \lambda_1 \leq \dots \leq \lambda_q$ denote the ordered eigenvalues of the symmetric matrices $\widehat{\Omega}_n$ and $\Omega \in \mathbb{R}^{q \times q}$, respectively, then, as $n \rightarrow \infty$,*

$$\widehat{\Omega}_n \xrightarrow{\mathbb{P}} \Omega \quad \text{implies} \quad (\lambda_{n,1}, \dots, \lambda_{n,q}) \xrightarrow{\mathbb{P}} (\lambda_1, \dots, \lambda_q).$$

Proof of Lemma A.1. The convergence $\widehat{\Omega}_n \xrightarrow{\mathbb{P}} \Omega$ elementwise implies $\|\widehat{\Omega}_n - \Omega\| \xrightarrow{\mathbb{P}} 0$ for any matrix norm $\|\cdot\|$ on $\mathbb{R}^{q \times q}$. If we take the spectral norm $\|\Omega\|$ (i.e., $\|\Omega\|^2$ is the largest eigenvalue of $\Omega^T \Omega$), then Weyl's perturbation theorem (Jiang, 2010, page 145) states that

$$\max_{i=1, \dots, q} |\lambda_{n,i} - \lambda_i| \leq \|\widehat{\Omega}_n - \Omega\|,$$

so that the desired result follows immediately. \square

By the diagonalization of $\widehat{\Omega}_n$ in terms of its eigenvectors and eigenvalues, the norm $\|\cdot\|_{\widehat{\Omega}_n}$ is equivalent to the Euclidian norm $\|\cdot\|$ in the sense that

$$\lambda_{n,1} \|L_{n,k}(\theta)\|^2 \leq \|L_{n,k}(\theta)\|_{\widehat{\Omega}_n}^2 \leq \lambda_{n,q} \|L_{n,k}(\theta)\|^2.$$

Proof of Theorem 3.1. Let $\varepsilon_0 > 0$ be such that the closed ball $B_{\varepsilon_0}(\theta_0) = \{\theta : \|\theta - \theta_0\| \leq \varepsilon_0\}$ is a subset of Θ ; such an ε_0 exists since θ_0 is an interior point of Θ . Fix $\varepsilon > 0$ such that $0 < \varepsilon \leq \varepsilon_0$. We show first that

$$\mathbb{P}[\widehat{\Theta}_n \neq \emptyset \text{ and } \widehat{\Theta}_n \subset B_\varepsilon(\theta_0)] \rightarrow 1, \quad n \rightarrow \infty. \quad (\text{A.1})$$

Because ψ is a homeomorphism, there exists $\delta > 0$ such that for $\theta \in \Theta$, $\|\psi(\theta) - \psi(\theta_0)\| \leq \delta$ implies $\|\theta - \theta_0\| \leq \varepsilon$. Equivalently, for every $\theta \in \Theta$ such that $\|\theta - \theta_0\| > \varepsilon$ we have $\|\psi(\theta) - \psi(\theta_0)\| > \delta$. Define the event

$$A_n = \left\{ \left\| \psi(\theta_0) - \int \widehat{\ell}_{n,k} \mu \right\| \leq \frac{\delta \sqrt{\lambda_{n,1}}}{2 + \sqrt{\lambda_{n,q}}} \right\}.$$

If $\theta \in \Theta$ is such that $\|\theta - \theta_0\| > \varepsilon$, then on the event A_n , we have

$$\begin{aligned} \|L_{n,k}(\theta)\|_{\widehat{\Omega}_n} &= \left\| \psi(\theta_0) - \psi(\theta) - \left(\psi(\theta_0) - \int \widehat{\ell}_{n,k} \mu \right) \right\|_{\widehat{\Omega}_n} \\ &\geq \|\psi(\theta_0) - \psi(\theta)\|_{\widehat{\Omega}_n} - \left\| \psi(\theta_0) - \int \widehat{\ell}_{n,k} \mu \right\|_{\widehat{\Omega}_n} \\ &\geq \sqrt{\lambda_{n,1}} \|\psi(\theta_0) - \psi(\theta)\| - \sqrt{\lambda_{n,q}} \left\| \psi(\theta_0) - \int \widehat{\ell}_{n,k} \mu \right\| \\ &> \delta \sqrt{\lambda_{n,1}} - \delta \frac{\sqrt{\lambda_{n,1} \lambda_{n,q}}}{2 + \sqrt{\lambda_{n,q}}} = \frac{2\delta \sqrt{\lambda_{n,1}}}{2 + \sqrt{\lambda_{n,q}}}. \end{aligned}$$

It follows that on A_n ,

$$\inf_{\theta: \|\theta - \theta_0\| > \varepsilon} \|L_{n,k}(\theta)\|_{\widehat{\Omega}_n} \geq \frac{2\delta\sqrt{\lambda_{n,1}}}{2 + \sqrt{\lambda_{n,q}}} > \left\| \psi(\theta_0) - \int \widehat{\ell}_{n,k} \mu \right\| \geq \inf_{\theta: \|\theta - \theta_0\| \leq \varepsilon} \left\| \psi(\theta) - \int \widehat{\ell}_{n,k} \mu \right\|.$$

The infimum on the right-hand side is actually a minimum since ψ is continuous and $B_\varepsilon(\theta_0)$ is compact. Hence on A_n the set $\widehat{\Theta}_n$ is non-empty and $\widehat{\Theta}_n \subset B_\varepsilon(\theta_0)$.

To show (A.1), it remains to be shown that $\mathbb{P}[A_n] \rightarrow 1$ as $n \rightarrow \infty$. Uniform consistency of $\widehat{\ell}_{n,k}$ for $d = 2$ was shown in Huang (1992); see also de Haan and Ferreira (2006, page 237). The proof for $d > 2$ is a straightforward extension. By the continuous mapping theorem, it follows that $\int \widehat{\ell}_{n,k} \mu$ is consistent for $\int \ell \mu = \psi(\theta_0)$. By Lemma A.1, $\lambda_{n,m}$ is consistent for λ_m for all $m \in \{1, \dots, q\}$. This yields $\mathbb{P}[A_n] \rightarrow 1$ and thus (A.1).

Next we wish to prove that, with probability tending to one, $\widehat{\Theta}_n$ has exactly one element, i.e., the function $f_{n,k,\widehat{\Omega}_n}$ has a unique minimizer. To do so, we will show that there exists $\varepsilon_1 \in (0, \varepsilon_0]$ such that, with probability tending to one, the Hessian of $f_{n,k,\widehat{\Omega}_n}$ is positive definite on $B_{\varepsilon_1}(\theta_0)$ and thus $f_{n,k,\widehat{\Omega}_n}$ is strictly convex on $B_{\varepsilon_1}(\theta_0)$. In combination with (A.1) for $\varepsilon \in (0, \varepsilon_1]$, this will yield the desired conclusion.

For $\theta \in \Theta$, define the symmetric $p \times p$ matrix $\mathcal{H}(\theta; \theta_0)$ by

$$(\mathcal{H}(\theta; \theta_0))_{i,j} := 2 \left(\frac{\partial \psi(\theta)}{\partial \theta_j} \right)^T \Omega \left(\frac{\partial \psi(\theta)}{\partial \theta_i} \right) - 2 \left(\frac{\partial^2 \psi(\theta)}{\partial \theta_j \partial \theta_i} \right) \Omega (\psi(\theta_0) - \psi(\theta))$$

for $i, j \in \{1, \dots, p\}$. The map $\theta \mapsto \mathcal{H}(\theta; \theta_0)$ is continuous and

$$\mathcal{H}(\theta_0) := \mathcal{H}(\theta_0; \theta_0) = 2 \dot{\psi}(\theta_0)^T \Omega \dot{\psi}(\theta_0),$$

is a positive definite matrix. Let $\|\cdot\|$ denote a matrix norm. By an argument similar to that in the proof of Lemma A.1, there exists $\eta > 0$ such that every symmetric matrix $A \in \mathbb{R}^{p \times p}$ with $\|A - \mathcal{H}(\theta_0)\| \leq \eta$ has positive eigenvalues and is therefore positive definite. Let $\varepsilon_1 \in (0, \varepsilon_0]$ be sufficiently small such that the second-order partial derivatives of ψ are continuous on $B_{\varepsilon_1}(\theta_0)$ and such that $\|\mathcal{H}(\theta; \theta_0) - \mathcal{H}(\theta_0)\| \leq \eta/2$ for all $\theta \in B_{\varepsilon_1}(\theta_0)$.

Let $\mathcal{H}_{n,k,\widehat{\Omega}_n}(\theta) \in \mathbb{R}^{p \times p}$ denote the Hessian matrix of $f_{n,k,\widehat{\Omega}_n}$. Its (i, j) -th element is

$$\begin{aligned} (\mathcal{H}_{n,k,\widehat{\Omega}_n}(\theta))_{ij} &= \frac{\partial^2}{\partial \theta_j \partial \theta_i} \left[L_{n,k}(\theta)^T \widehat{\Omega}_n L_{n,k}(\theta) \right] = \frac{\partial}{\partial \theta_j} \left[-2L_{n,k}(\theta)^T \widehat{\Omega}_n \frac{\partial \psi(\theta)}{\partial \theta_i} \right] \\ &= 2 \left(\frac{\partial \psi(\theta)}{\partial \theta_j} \right)^T \widehat{\Omega}_n \left(\frac{\partial \psi(\theta)}{\partial \theta_i} \right) - 2 \left(\frac{\partial^2 \psi(\theta)}{\partial \theta_j \partial \theta_i} \right) \widehat{\Omega}_n L_{n,k}(\theta). \end{aligned}$$

Since $L_{n,k}(\theta) = \int \widehat{\ell}_{n,k} \mu - \psi(\theta)$ and since $\int \widehat{\ell}_{n,k} \mu$ converges in probability to $\psi(\theta_0)$, we obtain

$$\sup_{\theta \in B_{\varepsilon_1}(\theta_0)} \|\mathcal{H}_{n,k,\widehat{\Omega}_n}(\theta) - \mathcal{H}(\theta; \theta_0)\| \xrightarrow{\mathbb{P}} 0, \quad n \rightarrow \infty. \quad (\text{A.2})$$

By the triangle inequality, it follows that

$$\mathbb{P} \left[\sup_{\theta \in B_{\varepsilon_1}(\theta_0)} \|\mathcal{H}_{n,k,\widehat{\Omega}_n}(\theta) - \mathcal{H}(\theta_0)\| \leq \eta \right] \rightarrow 1, \quad n \rightarrow \infty.$$

In view of our choice for η , this implies that, with probability tending to one, $\mathcal{H}_{n,k}(\theta)$ is positive definite for all $\theta \in B_{\varepsilon_1}(\theta_0)$, as required. \square

Proof of Theorem 3.2. First note that, as $n \rightarrow \infty$,

$$\sqrt{k} L_{n,k}(\theta_0) \xrightarrow{d} \tilde{B}, \quad \text{where } \tilde{B} \sim \mathcal{N}_q(\mathbf{0}, \Gamma(\theta_0)).$$

This follows directly from Einmahl et al. (2012, Proposition 7.3) by replacing $g(\mathbf{x}) d\mathbf{x}$ with $\mu(d\mathbf{x})$. Also, from (C2) and Slutsky's lemma, we have

$$\begin{aligned} \sqrt{k} \nabla f_{n,k,\hat{\Omega}_n}(\theta_0) &= -2\sqrt{k} L_{n,k}(\theta_0)^T \hat{\Omega}_n \dot{\psi}(\theta_0) \\ &\xrightarrow{d} -2\tilde{B}^T \Omega \dot{\psi}(\theta_0) \sim \mathcal{N}_p(\mathbf{0}, 4\dot{\psi}(\theta_0)^T \Omega \Gamma(\theta_0) \Omega \dot{\psi}(\theta_0)). \end{aligned}$$

Since $\hat{\theta}_n$ is a minimizer of $\hat{f}_{k,n}$ we have $\nabla f_{n,k,\hat{\Omega}_n}(\hat{\theta}_n) = 0$. Applying the mean value theorem to the function $t \mapsto \nabla f_{n,k,\hat{\Omega}_n}(\theta_0 + t(\hat{\theta}_n - \theta_0))$ at $t = 0$ and $t = 1$ yields

$$0 = \nabla f_{n,k,\hat{\Omega}_n}(\hat{\theta}_n) = \nabla f_{n,k,\hat{\Omega}_n}(\theta_0) + \mathcal{H}_{n,k,\hat{\Omega}_n}(\tilde{\theta}_n)(\hat{\theta}_n - \theta_0)$$

where $\tilde{\theta}_n$ is a random vector on the segment connecting θ_0 and $\hat{\theta}_n$. As $\hat{\theta}_n \xrightarrow{\mathbb{P}} \theta_0$, we have $\tilde{\theta}_n \xrightarrow{\mathbb{P}} \theta_0$ too. By (A.2) and continuity of $\theta \mapsto \mathcal{H}(\theta; \theta_0)$, it then follows that $\mathcal{H}_{n,k,\hat{\Omega}_n}(\tilde{\theta}_n) \xrightarrow{\mathbb{P}} \mathcal{H}(\theta_0)$. Putting these facts together, we conclude that

$$\sqrt{k}(\hat{\theta}_n - \theta_0) = -(\mathcal{H}_{n,k,\hat{\Omega}_n}(\tilde{\theta}_n))^{-1} \sqrt{k} \nabla f_{n,k,\hat{\Omega}_n}(\theta_0) \xrightarrow{d} \mathcal{N}_p(0, M(\theta_0)),$$

as required. \square

Proof of Corollary 3.3. Assumption (C6) implies that the map $\theta \mapsto \Gamma(\theta)$ is continuous at θ_0 (Einmahl et al., 2008, Lemma 7.2). Further, $\Gamma^{-1}(\hat{\theta}_n^{(0)})$ converges in probability to $\Gamma^{-1}(\theta_0)$, because of the continuous mapping theorem and the fact that $\hat{\theta}_n^{(0)}$ is a consistent estimator of θ_0 . Finally, the choice $\Omega_{\text{opt}} = \Gamma^{-1}(\theta)$ in (3.7) leads to the minimal matrix $M_{\text{opt}}(\theta)$ in (3.8); see for example Abadir and Magnus (2005, page 339). \square

B Calculating the asymptotic variance matrix

Fix the function π describing the pairs, and let for $m, m' \in \{1, \dots, q\}$, $\pi(m) = (u, v)$ and $\pi(m') = (u', v')$. Let $\dot{\ell}_{uv,1}(x_u, x_v)$ and $\dot{\ell}_{uv,2}(x_u, x_v)$ denote the partial derivatives of $\ell_{uv}(x_u, x_v)$ with respect to x_u and x_v respectively and define

$$W_{\ell,uv} = W_{\Lambda}(\{\mathbf{w} \in [0, \infty]^2 \setminus \{(\infty, \infty)\} : w_u \leq x_u \text{ or } w_v \leq x_v\}).$$

Note that

$$B_{uv}(x_u, x_v) = W_{\ell,uv}(x_u, x_v) - \dot{\ell}_{uv,1}(x_u, x_v)W_{\ell,u}(x_u) - \dot{\ell}_{uv,2}(x_u, x_v)W_{\ell,v}(x_v).$$

The (m, m') -th entry of $\Gamma(\theta) \in \mathbb{R}^{q \times q}$ from (3.5) is given by

$$\int_{[0,1]^4} \mathbb{E}[B_{uv}(x_u, x_v)B_{u'v'}(x_{u'}, x_{v'})] d\mathbf{x} = \int_{[0,1]^4} (T_1 - T_2 - T_3 + T_4 + T_5) d\mathbf{x}, \quad (\text{B.1})$$

for $\mathbf{x} = (x_u, x_v, x_{u'}, x_{v'})$ where

$$T_1 = \mathbb{E}[W_{\ell,uv}(x_u, x_v)W_{\ell,u'v'}(x_{u'}, x_{v'})],$$

$$\begin{aligned}
T_2 &= \dot{\ell}_{u'v',1}(x_{u'}, x_{v'}) \mathbb{E}[W_{\ell,uv}(x_u, x_v) W_{\ell,u'}(x_{u'})] + \dot{\ell}_{u'v',2}(x_{u'}, x_{v'}) \mathbb{E}[W_{\ell,uv}(x_u, x_v) W_{\ell,v'}(x_{v'})], \\
T_3 &= \dot{\ell}_{uv,1}(x_u, x_v) \mathbb{E}[W_{\ell,u}(x_u) W_{\ell,u'v'}(x_{u'}, x_{v'})] + \dot{\ell}_{uv,2}(x_u, x_v) \mathbb{E}[W_{\ell,v}(x_v) W_{\ell,u'v'}(x_{u'}, x_{v'})], \\
T_4 &= \dot{\ell}_{uv,1}(x_u, x_v) \dot{\ell}_{u'v',1}(x_{u'}, x_{v'}) \mathbb{E}[W_{\ell,u}(x_u) W_{\ell,u'}(x_{u'})] \\
&\quad + \dot{\ell}_{uv,2}(x_u, x_v) \dot{\ell}_{u'v',2}(x_{u'}, x_{v'}) \mathbb{E}[W_{\ell,v}(x_v) W_{\ell,v'}(x_{v'})], \\
T_5 &= \dot{\ell}_{uv,1}(x_u, x_v) \dot{\ell}_{u'v',2}(x_{u'}, x_{v'}) \mathbb{E}[W_{\ell,u}(x_u) W_{\ell,v'}(x_{v'})] \\
&\quad + \dot{\ell}_{uv,2}(x_u, x_v) \dot{\ell}_{u'v',1}(x_{u'}, x_{v'}) \mathbb{E}[W_{\ell,v}(x_v) W_{\ell,u'}(x_{u'})].
\end{aligned}$$

Suppose (u, v, u', v') are all different and define the sets

$$A_{ij}(z_i, z_j) = \{\mathbf{w} \in [0, \infty]^d \setminus \{(\infty, \dots, \infty)\} : w_i \leq z_i \text{ or } w_j \leq z_j\}.$$

Then

$$\begin{aligned}
\mathbb{E}[W_{\ell,uv}(x_u, x_v) W_{\ell,u'v'}(x_{u'}, x_{v'})] &= \mathbb{E}[W_\Lambda(A_{uv}(x_u, x_v)) W_\Lambda(A_{u'v'}(x_{u'}, x_{v'}))] \\
&= \Lambda(A_{uv}(x_u, x_v) \cap A_{u'v'}(x_{u'}, x_{v'})) \\
&= \Lambda(A_{uv}(x_u, x_v)) + \Lambda(A_{u'v'}(x_{u'}, x_{v'})) \\
&\quad - \Lambda(A_{uv}(x_u, x_v) \cup A_{u'v'}(x_{u'}, x_{v'})) \\
&= \ell_{uv}(x_u, x_v) + \ell_{u'v'}(x_{u'}, x_{v'}) - \ell_{uvu'v'}(x_u, x_v, x_{u'}, x_{v'}).
\end{aligned}$$

Similar calculations for the other terms yield

$$\begin{aligned}
T_1 &= \ell_{uv}(x_u, x_v) + \ell_{u'v'}(x_{u'}, x_{v'}) - \ell_{uvu'v'}(x_u, x_v, x_{u'}, x_{v'}), \\
T_2 &= \dot{\ell}_{u'v',1}(x_{u'}, x_{v'})[\ell_{uv}(x_u, x_v) + x_{u'} - \ell_{uvv'}(x_u, x_v, x_{v'})] \\
&\quad + \dot{\ell}_{u'v',2}(x_{u'}, x_{v'})[\ell_{uv}(x_u, x_v) + x_{v'} - \ell_{uvu'}(x_u, x_v, x_{u'})], \\
T_3 &= \dot{\ell}_{uv,1}(x_u, x_v)[\ell_{u'v'}(x_{u'}, x_{v'}) + x_u - \ell_{uu'v'}(x_u, x_{u'}, x_{v'})] \\
&\quad + \dot{\ell}_{uv,2}(x_u, x_v)[\ell_{u'v'}(x_{u'}, x_{v'}) + x_v - \ell_{vvu'}(x_v, x_{u'}, x_{v'})], \\
T_4 &= \dot{\ell}_{uv,1}(x_u, x_v) \dot{\ell}_{u'v',1}(x_{u'}, x_{v'})[x_u + x_{u'} - \ell_{uu'}(x_u, x_{u'})] \\
&\quad + \dot{\ell}_{uv,2}(x_u, x_v) \dot{\ell}_{u'v',2}(x_{u'}, x_{v'})[x_v + x_{v'} - \ell_{vv'}(x_v, x_{v'})], \\
T_5 &= \dot{\ell}_{uv,1}(x_u, x_v) \dot{\ell}_{u'v',2}(x_{u'}, x_{v'})[x_u + x_{v'} - \ell_{uv'}(x_u, x_{v'})] \\
&\quad + \dot{\ell}_{uv,2}(x_u, x_v) \dot{\ell}_{u'v',1}(x_{u'}, x_{v'})[x_v + x_{u'} - \ell_{vu'}(x_v, x_{u'})].
\end{aligned}$$

Integrating directly over T_1, \dots, T_5 is very slow, so we would like to simplify as many terms as possible. Introduce the notations

$$\begin{aligned}
I(u, v) &:= \int_0^1 \int_0^1 \ell_{uv}(x_u, x_v) dx_u dx_v, \\
I(u, v; x_u) &:= \int_0^1 \ell_{uv}(x_u, x_v) dx_v, \\
I_u(u, v; x_u) &:= \int_0^1 \frac{\partial \ell_{uv}(x_u, x_v)}{\partial x_u} dx_v.
\end{aligned}$$

Now we can write the four-dimensional integrals in (B.1) as:

$$\int_{[0,1]^4} T_1 = I(u, v) + I(u', v') - \int_{[0,1]^4} \ell_{uvu'v'}(x_u, x_v, x_{u'}, x_{v'}) dx_u dx_v dx_{u'} dx_{v'},$$

$$\begin{aligned}
\int_{[0,1]^4} T_2 &= I(u, v)[2I(u', v'; 1) - 1] + 2I(u', v'; 1) - 2I(u', v') \\
&\quad - \int_{[0,1]^3} I_{u'}(u', v'; x_{u'}) \ell_{uvu'}(x_u, x_v, x_{u'}) \, dx_{u'} dx_u dx_v \\
&\quad - \int_{[0,1]^3} I_{v'}(u', v'; x_{v'}) \ell_{uvv'}(x_u, x_v, x_{v'}) \, dx_{v'} dx_u dx_v, \\
\int_{[0,1]^4} T_3 &= I(u', v')[2I(u, v; 1) - 1] + 2I(u, v; 1) - 2I(u, v) \\
&\quad - \int_{[0,1]^3} I_u(u, v; x_u) \ell_{uu'v'}(x_u, x_{u'}, x_{v'}) \, dx_u dx_{u'} dx_{v'} \\
&\quad - \int_{[0,1]^3} I_v(u, v; x_v) \ell_{vv'u'}(x_v, x_{u'}, x_{v'}) \, dx_v dx_{u'} dx_{v'}, \\
\int_{[0,1]^4} T_4 &= [I(u, v) - I(u, v; 1)][1 - 2I(u', v'; 1)] + [I(u', v') - I(u', v'; 1)][1 - 2I(u, v; 1)] \\
&\quad - \int_{[0,1]^2} I_u(u, v; x_u) I_{u'}(u', v'; x_{u'}) \ell_{uu'}(x_u, x_{u'}) \, dx_u dx_{u'} \\
&\quad - \int_{[0,1]^2} I_v(u, v; x_v) I_{v'}(u', v'; x_{v'}) \ell_{vv'}(x_v, x_{v'}) \, dx_v dx_{v'}, \\
\int_{[0,1]^4} T_5 &= [I(u, v) - I(u, v; 1)][1 - 2I(u', v'; 1)] + [I(u', v') - I(u', v'; 1)][1 - 2I(u, v; 1)] \\
&\quad - \int_{[0,1]^2} I_u(u, v; x_u) I_{v'}(u', v'; x_{v'}) \ell_{uv'}(x_u, x_{v'}) \, dx_u dx_{v'} \\
&\quad - \int_{[0,1]^2} I_v(u, v; x_v) I_{u'}(u', v'; x_{u'}) \ell_{vu'}(x_v, x_{u'}) \, dx_v dx_{u'}.
\end{aligned}$$

For the Brown–Resnick process, the integrals $I(u, v)$, $I(u, v; x_u)$ and $I_u(u, v; x_u)$ are analytically computable. To calculate $I(u, v)$, do the change of variables $1/2 \log(x_u/x_v) = z_1$ and $1/2 \log(x_u x_v) = z_2$, so that $dx_u dx_v = 2 \exp(2z_2) dz_1 dz_2$ and the area we integrate over is the area between the lines $z_2 = z_1$ and $z_2 = -z_1$ for $z_2 < 0$. We obtain, for $a = a_{uv} = \sqrt{2\gamma(V(\mathbf{s}_u - \mathbf{s}_v))}$

$$\begin{aligned}
\int_0^1 \int_0^1 \ell_{uv}(x_u, x_v) \, dx_u dx_v &= \int_{z_1=-\infty}^{\infty} \int_{z_2=-\infty}^{-|z_1|} \left[e^{z_2+z_1} \Phi\left(\frac{a}{2} + \frac{2z_1}{a}\right) \right. \\
&\quad \left. + e^{z_2-z_1} \Phi\left(\frac{a}{2} - \frac{2z_1}{a}\right) \right] 2e^{2z_2} \, dz_2 dz_1 \\
&= \Phi(a/2) + \frac{e^{a^2} \Phi(-3a/2)}{3}.
\end{aligned}$$

The other two integrals are given by

$$\begin{aligned}
I(u, v; x_u) &= \frac{1}{2} \Phi\left(\frac{a}{2} - \frac{\log x_u}{a}\right) + x_u \Phi\left(\frac{a}{2} + \frac{\log x_u}{a}\right) + \frac{x_u^2 e^{a^2}}{2} \Phi\left(-\frac{3a}{2} - \frac{\log x_u}{a}\right), \\
I_u(u, v; x_u) &= \Phi\left(\frac{a}{2} + \frac{\log x_u}{a}\right) + x_u e^{a^2} \Phi\left(-\frac{3a}{2} - \frac{\log x_u}{a}\right).
\end{aligned}$$

Acknowledgments

This research is supported by contract “Projet d’Actions de Recherche Concertées” No. 12/17-045 of the “Communauté française de Belgique” and by IAP research network grant nr. P7/06 of the Belgian government (Belgian Science Policy).

REFERENCES

- Abadir, K. M. and J. R. Magnus (2005). *Matrix algebra*, Volume 1. Cambridge University Press.
- Bacro, J.-N. and C. Gaetan (2013). Estimation of spatial max-stable models using threshold exceedances. *Statistics and Computing*, 1–12.
- Beirlant, J., Y. Goegebeur, J. Segers, and J. Teugels (2004). *Statistics of Extremes: Theory and Applications*. Wiley.
- Bienvenüe, A. and C. Y. Robert (2014). Likelihood based inference for high-dimensional extreme value distributions. Available at <http://arxiv.org/abs/1403.0065>.
- Blanchet, J. and A. C. Davison (2011). Spatial modeling of extreme snow depth. *The Annals of Applied Statistics* 5(3), 1699–1724.
- Brown, B. M. and S. I. Resnick (1977). Extreme values of independent stochastic processes. *Journal of Applied Probability* 14(4), 732–739.
- Buishand, T., L. de Haan, and C. Zhou (2008). On spatial extremes: with application to a rainfall problem. *Annals of Applied Statistics* 2(2), 624–642.
- Ceppi, P., P. Della-Marta, and C. Appenzeller (2008). Extreme value analysis of wind speed observations over Switzerland. *Arbeitsberichte der MeteoSchweiz* 219.
- Coles, S. G. (2001). *An Introduction to Statistical Modeling of Extreme Values*. Springer-Verlag Inc.
- Cooley, D., J. Cisewski, R. J. Erhardt, S. Jeon, E. Mannshardt, B. O. Omolo, and Y. Sun (2012). A survey for spatial extremes: measuring spatial dependence and modelling spatial effects. *REVSTAT* 10(1), 135–165.
- Cooley, D., A. C. Davison, and M. Ribatet (2012). Bayesian inference from composite likelihoods, with an application to spatial extremes. *Statistica Sinica* 22(2), 813–845.
- Davis, R. A., C. Klüppelberg, and C. Steinkohl (2013). Statistical inference for max-stable processes in space and time. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 75(5), 791–819.
- Davison, A. C., S. A. Padoan, and M. Ribatet (2012). Statistical modeling of spatial extremes. *Statistical Science* 27(2), 161–186.
- de Haan, L. and A. Ferreira (2006). *Extreme Value Theory: an Introduction*. Springer-Verlag Inc.

- de Haan, L. and T. Lin (2001). On convergence toward an extreme value distribution in $C[0,1]$. *The Annals of Probability* 29(1), 467–483.
- de Haan, L. and T. T. Pereira (2006). Spatial extremes: Models for the stationary case. *The Annals of Statistics* 34(1), 146–168.
- de Haan, L. and S. I. Resnick (1977). Limit theory for multivariate sample extremes. *Zeitschrift fr Wahrscheinlichkeitstheorie und Verwandte Gebiete* 40(4), 317–337.
- Dematteo, A., S. Clemençon, N. Vayatis, and M. Mougeot (2013). Sloshing in the LNG shipping industry: risk modelling through multivariate heavy-tail analysis. Available at <http://arxiv.org/abs/1312.0020>.
- Drees, H. and X. Huang (1998). Best attainable rates of convergence for estimators of the stable tail dependence function. *Journal of Multivariate Analysis* 64(1), 25–47.
- Einmahl, J. H., A. Krajina, and J. Segers (2008). A method of moments estimator of tail dependence. *Bernoulli* 14(4), 1003–1026.
- Einmahl, J. H., A. Krajina, and J. Segers (2012). An M-estimator for tail dependence in arbitrary dimensions. *The Annals of Statistics* 40(3), 1764–1793.
- Engelke, S., A. Malinowski, Z. Kabluchko, and M. Schlather (2014). Estimation of Hüsler-Reiss distributions and Brown-Resnick processes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. To be published.
- Huang, X. (1992). *Statistics of bivariate extreme values*. Ph. D. thesis, Tinbergen Institute Research Series.
- Huser, R. and A. Davison (2013). Composite likelihood estimation for the Brown-Resnick process. *Biometrika* 100(2), 511–518.
- Huser, R. and A. Davison (2014). Space-time modelling of extreme events. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 76(2), 439–461.
- Jeon, S. and R. Smith (2012). Dependence structure of spatial extremes using threshold approach. Available at <http://arxiv.org/abs/1209.6344>.
- Jiang, J. (2010). *Large sample techniques for statistics*. Springer.
- Kabluchko, Z., M. Schlather, and L. de Haan (2009). Stationary max-stable fields associated to negative definite functions. *Annals of Probability* 37(5), 2042–2065.
- Karpa, O. and A. Naess (2013). Extreme value statistics of wind speed data by the ACER method. *Journal of Wind Engineering and Industrial Aerodynamics* 112, 1–10.
- Oesting, M., M. Schlather, and P. Friedrichs (2013). Conditional modelling of extreme wind gusts by bivariate Brown-Resnick processes. Available at <http://arxiv.org/abs/1312.4584>.
- Padoan, S., M. Ribatet, and S. Sisson (2010). Likelihood-based inference for max-stable processes. *Journal of the American Statistical Association (Theory and Methods)* 105(489), 263–277.

- Palutikof, J., B. Brabson, D. Lister, and S. Adcock (1999). A review of methods to calculate extreme wind speeds. *Meteorological Applications* 6(2), 199–132.
- R Core Team (2013). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Reich, B. J. and B. A. Shaby (2012). A hierarchical max-stable spatial model for extreme precipitation. *The Annals of Applied Statistics* 6(4), 1430–1451.
- Resnick, S. I. (1987). *Extreme Values, Regular Variation, and Point Processes*. Springer, New York.
- Ribatet, M. (2013). Spatial extremes: max-stable processes at work. *Journal de la Société Française de Statistique* 154(2), 1–22.
- Ribatet, M., R. Singleton, and R Core team (2013). *SpatialExtremes: Modelling Spatial Extremes*. R package version 2.0-0.
- Schlather, M. (2002). Models for stationary max-stable random fields. *Extremes* 5(1), 33–44.
- Smith, R. L. (1990). Max-stable processes and spatial extremes. Unpublished manuscript.
- Wadsworth, J. L. and J. A. Tawn (2013). Efficient inference for spatial extreme-value processes associated to log-gaussian random functions. To appear in *Biometrika*.
- Wang, Y. and S. Stoev (2011). Conditional sampling for spectrally discrete max-stable random fields. *Advances in Applied Probability* 43, 463–481.
- Yuen, R. A. and S. Stoev (2013). CPRS M-estimation for max-stable models. Available at <http://arxiv.org/abs/1307.7209>.