

Modeling extremal dependence in multivariate and spatial problems: a practical perspective

Boris Béranger¹  and Simone A. Padoan² 

¹*School of Mathematics and Statistics, UNSW Data Science Hub (uDASH), UNSW Sydney*

E-mail: B.Beranger@unsw.edu.au

²*Department of Decision Sciences, Bocconi University*

E-mail: Simone.Padoan@unibocconi.it

April 5, 2026

Abstract

From environmental sciences to finance, there is a growing demand for methods that can assess the risks of extreme events beyond those observed in available data. Extrapolating extreme events beyond the range of the data is not obvious. Risk assessments are often further complicated by the need to account for multiple variables simultaneously. Extreme value theory provides important tools for the analysis of multivariate or spatial extreme events, but these are not easily accessible to professionals without appropriate expertise. This article provides a minimal background on multivariate and spatial extremes and gives simple yet thorough instructions on how to analyse them using the R package **ExtremalDep**. After briefly introducing the statistical methodologies, we focus on road testing the package’s toolbox through several real-world applications.

Keywords: Angular measure, Air pollution, Environmental data analysis, Exchange rates, Extremal coefficient, Extreme sets, Heat Waves, Max-stable processes, Multivariate generalized extreme-value distribution, Pickands dependence function, Quantile regions

1 Introduction

In many applied fields—including environmental sciences, finance, and actuarial sciences—it is essential to assess the risk of extreme events beyond those observed in historical data, with the goal to anticipate the magnitude of hypothetical future catastrophes (e.g., a global financial crisis or large monetary losses due to natural hazards). The aim is to disclose the uncertainty surrounding such extreme events to decision makers in order to carefully plan mitigation strategies. A key challenge arises when extrapolating beyond the range of N available observations, by computing for example quantiles whose probability (p) is smaller than $1/N$, meaning that the largest data observation is not expected to exceed the target event on average. This makes direct empirical estimation impossible. Extreme Value Theory (EVT) addresses this problem by providing a rigorous probabilistic framework and statistical methodology for modeling and extrapolating extreme events. For accessible introductions, see, for instance, [Beirlant et al. \(2004\)](#), [de Haan and Ferreira \(2006\)](#), and [Falk et al. \(2010\)](#).

In the univariate case, the theory and practice is mature with a significant number of packages available on the Comprehensive R Archive Network (CRAN), forming an extensive set of tools. Examples include **ismev** (Heffernan and Stephenson, 2018), **evd** Stephenson (2002), **revdbayes** (Northrop and Grimshaw, 2025), **extRemes** (Gilleland and Katz, 2016) and **ExtremeRisks** (Pacifi et al., 2025) to name a few. In practice, extreme events rarely involve a single variable. Multiple factors are typically observed simultaneously: for example, rainfall, temperature, and wind measured across different locations within a region can jointly drive economic losses from adverse weather. Similarly, high pollution levels often result from the combined effect of several pollutants exceeding safety thresholds.

When it comes to multivariate and spatial problems, statistical methodologies for analyzing multivariate and spatial extremes, and their corresponding software implementations, remain relatively underdeveloped. The dependence structure governing multivariate extreme value distributions, or equivalently, the finite-dimensional distributions of stochastic processes for extremes is termed extremal dependence. It is an infinite-dimensional, nonparametric object subject to some restrictive conditions, which leads to both a complex formulation of the overall multivariate distribution and significant inferential challenges. The literature on multivariate and spatial extremes covers several topics including multivariate extremes based on componentwise maxima (e.g., Beranger and Padoan, 2015) and multiple peaks over threshold (e.g., Rootzén et al., 2018), spatial extremes based on pointwise maxima and peaks over threshold (e.g., Davison et al., 2012; de Fondeville and Davison, 2018), graphical models (Engelke and Hitz, 2020) and time series (e.g., Davis et al., 2011), and remains nowadays a vibrant research field with a number of ongoing new results. These have produced useful R packages such as **mev** (Belzile et al., 2025), **SpatialExtremes** (Ribatet, 2022), **mvPot** (de Fondeville et al., 2025), **RandomFields**, **graphicalExtremes** (Engelke et al., 2025), **extremogram** (Frolova and Cribben, 2016) as well as **ExtremalDep** (Beranger et al., 2025). A review of statistical software for modeling extremes is provided by Belzile et al. (2023), which introduces some of the capabilities of **ExtremalDep**; but we offer here a more in-depth guide to the package, complemented by practical illustrations.

The focus of this work is on modeling the extremal dependence structure arising from the componentwise maxima approach, which enables the assessment of risk associated with multivariate extreme events in a broad sense. Contrary to what might be expected, this framework is not limited to modeling multivariate maxima, but is also well suited for analyzing multivariate threshold exceedances. This is achieved by exploiting approximations of the probabilities of observing at least one exceedance (2.9), joint exceedances (2.10), or an observation falling within an extreme region (2.12), in which at least one component is extreme. Statistical methods for risk analysis based on these probability approximations are presented in Sections 2.3, 2.4.3, and 2.4.4. If the primary interest lies instead in proper multivariate peaks-over-threshold models and methods specifically tailored to their analysis, we refer the reader to Rootzén et al. (2018); de Fondeville and Davison (2018); Kiriliouk et al. (2019). Finally, it is worth noting that several alternative approaches for modeling multivariate extremes also exist, including conditional extremes models (Heffernan and Tawn, 2004), random scale and location mixture models (Huser et al., 2017), and (X-)vine copulae (Kiriliouk et al., 2025). Several of these alternative methods are implemented in the **texmex** package (Southworth et al., 2024).

Most existing packages for the analysis of multivariate extremes rely on parametric representations of extremal dependence (e.g., **mev**, **mvPot**, and **graphicalExtremes**), which simplify the underlying mathematics and inference but impose strong structural assumptions. While

ExtremalDep also accommodates parametric approaches, one of its main contributions lies in extending the available toolkit by providing nonparametric and semiparametric methods that remain closer to the inherently nonparametric nature of extremal dependence. A further distinguishing feature of **ExtremalDep**, compared with existing software, is its emphasis on demonstrating how extremal dependence can be directly exploited for risk assessment related to multivariate extreme events and their occurrence probabilities. In particular, the routines of the package allow one to compute joint return levels and conditional return periods (Section 2.3), conditional probabilities given the occurrence of extreme events (Section 2.4.2), tail probabilities relative to multiple threshold exceedances (Section 2.4.3), and extreme quantile regions (Section 2.4.4). Finally, the package allows for frequentist and Bayesian inference for multivariate extremes and frequentist and Bayesian inference for spatial extremes. This manuscript demonstrates the practical use of the package through several important applications. Because the ordering of points in d -dimensional Euclidean space is arbitrary, we introduce and formalize key concepts such as extreme sets, multivariate extreme quantiles, quantile regions, joint tail and conditional probabilities, and multivariate return levels. In practice, these tools enable a wide range of risk analyses. Examples include risk assessment of future high-pollution episodes in urban areas, identifying regions of a country exposed to heavy precipitation events, or evaluating the risk of future heatwave episodes. The article is divided into two parts: Section 2 discusses multivariate extremes while Section 3 focuses on spatial extremes. In the multivariate part, Section 2.1 gives a brief introduction to the theoretical background while Section 2.3 details how to perform parametric (frequentist and Bayesian) inference for the extremal dependence and for joint and conditional tail probabilities. Section 2.4 explains how to perform nonparametric Bayesian inference for the extremal dependence of a d -dimensional componentwise maxima vector, and once the extremal dependence has been estimated, how to infer joint and conditional tail probabilities in the bivariate case. The focus then successively shifts to the random generation of bivariate extreme values with a flexible semi-parametric dependence structure, the estimation of small probabilities of belonging to certain extreme sets and the estimation extreme quantile regions corresponding to small joint tail probabilities. The paper ends in Section 4 with a summary and a discussion on future developments.

2 Multivariate extremes

2.1 Background

Let $\mathbf{X}_1, \mathbf{X}_2, \dots$ be independent and identically distributed (i.i.d.) random vectors, where $\mathbf{X}_i = (X_{i,1}, \dots, X_{i,d}) \in \mathbb{R}^d$ for some integer $d > 1$, whose joint distribution F is in the domain of attraction of a *multivariate extreme value* distribution G (de Haan and Ferreira, 2006, chap. 6), denoted by $F \in \mathcal{D}(G)$. This implies that for $n = 1, 2, \dots$, there exists norming sequences $\mathbf{a}_n > \mathbf{0}_d = (0, \dots, 0)$ where the inequality is applied element-wise, and $\mathbf{b}_n \in \mathbb{R}^d$ such that

$$\lim_{n \rightarrow \infty} F^n(\mathbf{a}_n \mathbf{x} + \mathbf{b}_n) = G_\gamma(\mathbf{x} \mid D), \quad \mathbf{x} \in \mathbb{R}^d, \quad (2.1)$$

where $\gamma = (\gamma_1, \dots, \gamma_d)^\top \in \mathbb{R}^d$. The limiting distribution in (2.1) takes the form

$$G_\gamma(\mathbf{x} \mid D) = C_{\text{EV}}(G_{\gamma_1}(x_1), \dots, G_{\gamma_d}(x_d) \mid D), \quad \mathbf{x} \in \mathbb{R}^d, \quad (2.2)$$

where C_{EV} is an *extreme-value copula* (Beirlant et al., 2004, chap. 8), representing the so-called *extremal dependence*, which binds the d marginal distributions that are, for all $j \in \{1, \dots, d\}$, univariate GEV distributions of the form

$$G_{\gamma_j}(x_j) = \exp \left\{ - (1 + \gamma_j x_j)^{-\frac{1}{\gamma_j}} \right\}, \quad 1 + \gamma_j x_j > 0, \quad (2.3)$$

where $\gamma_j \in \mathbb{R}$ is the *extreme value index*. This coefficient describes the tail heaviness of the distribution, resulting in a heavy-, light- or short-tailed distribution, depending on whether $\gamma_j > 0$, $\gamma_j = 0$ or $\gamma_j < 0$ (e.g., de Haan and Ferreira, 2006, chap. 1). In (2.2) the extreme-value copula is defined as

$$C_{\text{EV}}(\mathbf{u} \mid D) = \exp \{ -L(-\ln u_1, \dots, -\ln u_d) \}, \quad \mathbf{u} \in (0, 1]^d, \quad (2.4)$$

where $L : [0, \infty)^d \mapsto [0, \infty)$ is a homogeneous function of order 1 named *stable-tail* dependence function (Beirlant et al., 2004, chap. 8). This latter function fully characterizes the extremal dependence structure; however, since it is defined on $[0, \infty)^d$ and takes values in $[0, \infty)$, it is difficult to visualize directly. For this reason, it is common practice to work instead with its restriction to the set $\mathcal{R} := \{\mathbf{t} \in [0, 1]^{d-1} : \|\mathbf{t}\|_1 \leq 1\}$, known as the Pickands dependence function (e.g., Falk et al., 2010, Ch. 4), which is more amenable to graphical representation. While no exhaustive characterization of the Pickands dependence function is available, it must at least satisfy the following convexity and boundary conditions

$$(C1) \quad A(a\mathbf{t}_1 + (1-a)\mathbf{t}_2) \leq aA(\mathbf{t}_1) + (1-a)A(\mathbf{t}_2), \quad a \in [0, 1], \quad \forall \mathbf{t}_1, \mathbf{t}_2 \in \mathcal{R},$$

$$(C2) \quad 1/d \leq \max(t_1, \dots, t_{d-1}, 1 - t_1 - \dots - t_{d-1}) \leq A(\mathbf{t}) \leq 1, \quad \forall \mathbf{t} \in \mathcal{R}.$$

In the bivariate case, these are necessary and sufficient to define a valid Pickands dependence function (for details, see Beirlant et al., 2004, chap. 8). However, there is an alternative way to represent extreme dependence. Indeed, a formal definition of the Pickands dependence function is

$$A(\mathbf{t}) = L(1 - t_1 - \dots - t_{d-1}, \mathbf{t}) = d \int_{\mathcal{S}_d} \max\{(1 - t_1 - \dots - t_{d-1})w_1, \dots, t_{d-1}w_d\} dH(\mathbf{w}), \quad (2.5)$$

where H is a probability measure on the d -dimensional unit simplex $\mathcal{S}_d := \{\mathbf{w} \geq \mathbf{0} : \|\mathbf{w}\|_1 = 1\}$, named *angular measure* (e.g., Falk et al., 2010, Ch. 4). Thus, the level of dependence can be conveyed through H ; however, the same issue arises again: which specific measure should one consider? A valid angular measure is any probability measure satisfying the mean constraint

$$(C3) \quad \int_{\mathcal{S}_d} w_j dH(\mathbf{w}) = 1/d, \quad \forall j \in \{1, \dots, d\}.$$

Since A and H are related through the one-to-one map in formula (2.5), the dependence parameter D of the copula $C_{\text{EV}}(\cdot \mid D)$ can be equivalently meant to be either the angular measure or the Pickands dependence function. Both are possible means to describe the extremal dependence.

The result in (2.1) can be interpreted as follows. Given a sample $(\mathbf{X}_1, \dots, \mathbf{X}_N)$ with joint distribution $F \in \mathcal{D}(G)$, for each of the d variables we compute k maxima are computed over blocks of size n , such that $N = n \cdot k$, where the i th maximum is $M_{i,j}^{(n)} = \max(X_{(i-1)n+1,j}, \dots, X_{in,j})$, $i \in \{1, \dots, k\}$, $j = 1, \dots, d$. We bind the maxima in the vector of componentwise maxima $\mathbf{M}_i^{(n)} = (M_{i,1}^{(n)}, \dots, M_{i,d}^{(n)})^\top$. Now, setting $\boldsymbol{\mu} = \mathbf{b}_n$, $\boldsymbol{\sigma} = \mathbf{a}_n$,




Figure 1: From left to right, the first two panels display angular densities corresponding to strong and weak dependence between two maxima, respectively. In the strong dependence case, the probability mass concentrates around $1/2$, whereas in the weak dependence case it concentrates near the vertices of $(0, 1)$. The third and fourth panels show the corresponding Pickands dependence functions, which are close to the boundary cases $\max(w, 1 - w)$ and 1 , respectively.

$\mathbf{y} = \mathbf{a}_n \mathbf{x} + \mathbf{b}_n$, then thanks to (2.1), the joint distribution of vector of componentwise maxima can be approximated, for large enough n , as $F^n(\mathbf{y}) \approx G_{\boldsymbol{\theta}}(\mathbf{y} \mid D)$, where $\boldsymbol{\theta} = (\boldsymbol{\mu}, \boldsymbol{\sigma}, \boldsymbol{\gamma})^\top$. By construction, the distribution of the univariate maximum $M_{i,j}^{(n)}$ can be approximated by

$$F^n(y_j) \approx G_{\theta_j}(y_j) = \exp \left\{ - \left(1 + \gamma_j \frac{y_j - \mu_j}{\sigma_j} \right)_+^{-\frac{1}{\gamma_j}} \right\}, \quad j \in \{1, \dots, d\}, \quad (2.6)$$

where the dependence among the componentwise maxima $(M_{i,1}^{(n)}, \dots, M_{i,d}^{(n)})$ can be approximated through the extreme-value copula $C_{\text{EV}}(\cdot \mid D)$. When the dependence parameter D is represented by the angular measure H the interpretation is as follows: the more mass H assigns to the center of the simplex $(1/d, \dots, 1/d)$, the stronger the dependence among the maxima; conversely, the dependence weakens as H concentrates its mass near the vertices of the simplex. This behavior is illustrated by the solid lines in the first and second panels of Figure 1. When D is expressed in terms of the Pickands dependence function, dependence strength is reflected by its proximity to the bounds in condition (C2). Specifically, the closer the Pickands dependence function is to the lower bound, the stronger the dependence among the maxima, whereas proximity to the upper bound indicates weaker dependence. These contrasting regimes are illustrated by the solid lines in the third and fourth panels of Figure 1.

A useful summary of the dependence among the maxima is given by the so-called *extremal coefficient* (e.g., Beirlant et al., 2004, Ch. 8), defined as

$$\varrho = dA(1/d, \dots, 1/d), \quad (2.7)$$

which represents the (fractional) number of independent block maxima. The extremal coefficient satisfies $1 \leq \varrho \leq d$, with the lower and upper bounds describing the case of complete dependence and independence.

Even more importantly, beyond these results, the domain of attraction provides also a useful basis to derive small joint tail probabilities corresponding to multivariate extremes relative to the original distribution F such as multiple threshold exceedances and derive

extreme quantile regions corresponding to some small joint tail probabilities. These results are summarized as follows. The limit in (2.1) holds if and only if $n\{1 - F(\mathbf{a}_n\mathbf{x} + \mathbf{b}_n)\} \rightarrow -\log G(\mathbf{x})$ as $n \rightarrow \infty$, for all $\mathbf{x} \in \mathbb{R}^d$, with the convention that $-\log(0) = \infty$. Together with the weak convergence of the copula of the vector of componentwise maxima, namely $\{C(\mathbf{u}^{1/n})\}^n \rightarrow C_{\text{EV}}(\mathbf{u})$ as $n \rightarrow \infty$, for all $\mathbf{u} \in (0, 1]^d$, this result implies that for all $\mathbf{x} \in [0, \infty)^d$, as $n \rightarrow \infty$,

$$n \mathbb{P} \left(\bigcup_{j=1}^d F_j(X_j) > 1 - \frac{x_j}{n} \right) \rightarrow L(\mathbf{x}) = d \int_{\mathcal{S}_d} \max_{j=1, \dots, d} (x_j w_j) dH(\mathbf{w}). \quad (2.8)$$

Recall that the joint distribution function F can be represented in terms of a copula function, C , defined on $[0, 1]^d$ and its marginal distributions $F_i, i = 1, \dots, d$ (assumed to be continuous) through $F(\mathbf{x}) = C\{F_1(x_1), \dots, F_d(x_d)\}$ (see, e.g., Joe, 2014, for details). The limiting result in formula (2.8) implies the following important approximation: let $p_j = x_j/n$, for $j = 1, \dots, d$, and $Q_j(p_j)$ be the $(1 - p)$ quantile of F_j ; then, by homogeneity of the stable tail dependence function, the probability that at least one among the d variables X_j exceeds a high quantile of its own distribution, can be approximated for large n by

$$\mathbb{P}\{X_1 > Q_1(p_1) \text{ or } \dots \text{ or } X_d > Q_d(p_d)\} \approx L(p_1, \dots, p_d). \quad (2.9)$$

This result is particularly relevant in applications because, once the stable tail dependence function has been estimated from the data, for small probabilities p_1, \dots, p_d , it directly provides an estimate of the probability that at least one of the d variables of interest exceeds a high threshold. By the inclusion and exclusion principle, the weak convergence result of the extreme value copula also allows one to obtain

$$n \mathbb{P} \left(\bigcap_{j=1}^d F_j(X_j) > 1 - \frac{x_j}{n} \right) \rightarrow R(\mathbf{x}) = d \int_{\mathcal{S}_d} \min_{j=1, \dots, d} (x_j w_j) dH(\mathbf{w}), \quad n \rightarrow \infty,$$

where R is known as the *tail copula* function, which is also a homogeneous function of order 1, and therefore the probability that all the d variables (X_1, \dots, X_d) exceed simultaneously high quantiles of their own distributions can be approximated, for a sufficiently large n , by

$$\mathbb{P}\{X_1 > Q_1(p_1) \text{ and } \dots \text{ and } X_d > Q_d(p_d)\} \approx R(p_1, \dots, p_d). \quad (2.10)$$

Also this result is very important for practice, as once the tail copula function has been estimated, corresponding to small probabilities p_1, \dots, p_d , it implicitly provides an estimate of the probability that all the d variables of interest exceeds high thresholds.

Finally, an important question is, how to determine an extreme region where two variables, given a small joint probability p of falling in it? We summarize the approach discussed by Cai et al. (2011), Einmahl et al. (2013) and He and Einmahl (2017) as follows with a focus on the bivariate case explored by Einmahl et al. (2013). Let f be the density on \mathbb{R}_+^2 of a bivariate distribution F satisfying $F \in \mathcal{D}(G)$, with $\gamma_1, \gamma_2 > 0$. A possible definition of quantile regions is given by the level sets of f , namely for $\alpha > 0$

$$\mathcal{Q} = \{\mathbf{x} \in \mathbb{R}_+^2 : f(\mathbf{x}) \leq \alpha\}. \quad (2.11)$$

Then, for a small probability $p > 0$, we can derive the region \mathcal{Q} such that $\mathbb{P}(\mathcal{Q}) = p$ and f everywhere on \mathcal{Q} is smaller than f everywhere on \mathcal{Q}^c . Therefore, the latter is the smallest

region satisfying $\mathbb{P}(\mathcal{Q}^c) = 1 - p$. Accordingly, an *extreme quantile region* is defined by a level set \mathcal{Q}_N (with $\alpha = \alpha_N$ in 2.11), which is such that $\mathbb{P}(\mathcal{Q}_N) = p$, where $p = p_N \rightarrow 0$ and the expected number of points falling in it is $Np \rightarrow c > 0$ as $N \rightarrow \infty$. Note that when $c < 1$, the region is so extreme that almost no observations are expected to fall in it. From a practical perspective, given a sample of large size N and a tiny probability p of falling within an extreme quantile region, how can such a region be determined?

The domain of attraction condition implies the existence of a measure ξ , named *exponent measure*, satisfying $\xi(\partial B) = 0$, for all Borel set $B \subset [0, \infty]^2$ that are bounded away from the origin, and such that the probability of falling in extreme region NB can be approximated as

$$\mathbb{P}(\{Q_1(1/N)x_1^{\gamma_1}, Q_2(1/N)x_2^{\gamma_2} : \mathbf{x} \in B\}) \approx \xi(NB), \quad \text{as } N \rightarrow \infty. \quad (2.12)$$

Assume that for some $\mathcal{M} > 0$, f is bounded away from zero on $(0, \mathcal{M}]^2$ and, outside this region, f is decreasing in each coordinate and that the exponent measure admits a density g satisfying

$$\xi(\{\mathbf{v} : v_1 > x_1 \text{ or } v_2 > x_2\}) = \int \int_{\{v_1 > x_1 \text{ or } v_2 > x_2\}} g(\mathbf{v}) d\mathbf{v},$$

and

$$\lim_{N \rightarrow \infty} NQ_1(1/N)Q_2(1/N)f\{Q_1(1/N)x_1^{\gamma_1}, Q_2(1/N)x_2^{\gamma_2}\} = \frac{x_1^{1-\gamma_1}x_2^{1-\gamma_2}}{\gamma_1\gamma_2}g(\mathbf{x}) =: q(\mathbf{x}).$$

Note that g is related to the density h of the angular measure H by $h(w) = 2^{-1}g(w, 1 - w)$, where $w = x_1/r$ and $r = x_1 + x_2$. The main idea behind the derivation of an extreme quantile region is, first to define the so-called *basic set* S defined as

$$S = \{\mathbf{x} : q(\mathbf{x}) \leq 1\} = \{\mathbf{x} : r \geq q_\star^{-1}(w), w \in [0, 1]\},$$

by means of the *angular basic density*

$$q_\star(w) = \left\{ 2\gamma_1^{-1}\gamma_2^{-1}w^{1-\gamma_1}(1-w)^{1-\gamma_2}h(w) \right\}^{-\frac{1}{1+\gamma_1+\gamma_2}}.$$

According to the exponent measure, its size is equal to

$$\xi(S) = 2 \int_{[0,1]} q_\star(w)h(w)dw.$$

Then, the basic set S is suitably inflated into an extreme set $\tilde{\mathcal{Q}}_N$ defined as

$$\tilde{\mathcal{Q}}_N = \left\{ \left(\mu_1 + \sigma_1 \frac{\left(\frac{k\xi(S)x_1}{Np}\right)^{\gamma_1} - 1}{\gamma_1}, \mu_2 + \sigma_2 \frac{\left(\frac{k\xi(S)x_2}{Np}\right)^{\gamma_2} - 1}{\gamma_2} \right) : \mathbf{x} \in S \right\}, \quad (2.13)$$

where $k = k_N$ with $k \rightarrow \infty$ as $N \rightarrow \infty$ and $k = o(N)$, which approximate the target extreme quantile region \mathcal{Q}_N . In particular, $\mu_j = Q_j(k/N)$ and $\sigma_j = a_j(N/k)$ with $j = 1, 2$, where $a_j(\cdot)$ is a suitable positive scale function (de Haan and Ferreira, 2006, Ch. 1-2). Then, the set $\tilde{\mathcal{Q}}_N$ is a good approximation of \mathcal{Q}_N in the sense that $\mathbb{P}(\tilde{\mathcal{Q}}_N) \approx p$ and $\mathbb{P}(\mathcal{Q}_N \Delta \tilde{\mathcal{Q}}_N) \rightarrow 0$ as $N \rightarrow \infty$ (see Einmahl et al., 2013), where $A \Delta B = A \setminus B \cup B \setminus A$.

2.2 Statistical challenges

From a statistical perspective, the multivariate GEV model defines a complex semiparametric family of distributions of the form $G_{\boldsymbol{\theta}}(\mathbf{y} \mid D) = C_{\text{EV}}\{G_{\theta_1}(y_1), \dots, G_{\theta_d}(y_d) \mid D\}$, $\mathbf{y} \in \mathbb{R}^d$, where $\boldsymbol{\theta} \in \Theta \subset (\mathbb{R}, \mathbb{R}_+, \mathbb{R})^d$ is a finite-dimensional vector of marginal parameters and D is an infinite-dimensional dependence parameter of the copula C_{EV} , named extremal dependence, living in a suitable space, depending whether D is the angular measure or the Pickands dependence function. The estimation of such infinite-dimensional parameter, accounting for the corresponding constraints (see conditions (C1)–(C3) of Section 2.1), is far from simple. To reduce complexity, many classes of parametric models were proposed to model D , see Beranger and Padoan (2015) for a review. In order to provide valid extremal dependence structures, the proposed models satisfy the required constraints. Section 2.3 describes how to fit some of the most popular parametric models for D using **ExtremalDep**.

Over time, more sophisticated semi- and non-parametric approaches for modeling and estimating the extremal dependence have been proposed, under both parametrisations. Some examples include: projection estimators (e.g., Fils-Villetard et al., 2008), polynomials (e.g., Marcon et al., 2017) and splines (e.g., Cormier et al., 2014), just to name a few. Section 2.4 details how to use **ExtremalDep** for non-parametric estimation of the extremal dependence, simulation of bivariate extreme values with a flexible semi-parametric dependence structure, estimation of small probabilities to belong to certain extreme sets and estimation extreme quantile regions corresponding to small joint tail probabilities.

2.3 Parametric modeling of the extremal dependence

Theory. The Poisson Point Process (PPP) is a useful characterization for statistical models used to represent the extremal dependence, when the latter is parametrized according to the angular measure and provided that it allows for a density $h(w \mid \boldsymbol{\varphi})$ on \mathcal{S}_d , where $\boldsymbol{\varphi}$ is a suitable vector of parameters (see, e.g., Coles and Tawn, 1991; Engelke et al., 2015). In particular, let $\mathbf{x}_1, \dots, \mathbf{x}_N$ be i.i.d. observations from $F \in \mathcal{D}(G)$ and $\mathbf{y}_1, \dots, \mathbf{y}_N$ be their equivalent with unit Fréchet margins. The pseudopolar transformation T is defined as $T(\mathbf{y}_i) = (r_{(i)}, \mathbf{w}_{(i)})$ where $r_{(i)} = y_{i,1} + \dots + y_{i,d}$ is the radial part and $\mathbf{w}_{(i)} = \mathbf{y}_i/r_{(i)}$ is the angular part of \mathbf{y}_i , $i = 1, \dots, N$. The domain of attraction condition implies that, if r_0 is a large threshold and $W_{r_0} = \{(r, \mathbf{w}) : r > r_0\}$ is a set of points with a radial component larger than r_0 , the number of points falling in W_{r_0} is described approximatively by a Poisson random variable with rate $1/\psi(W_{r_0})$, where $\psi = \xi\{T^{-1}(\cdot)\}$. Thus, conditional on observing k points $\{(r_{(i)}, \mathbf{w}_{(i)}), i = 1, \dots, k\}$ with radial part greater than r_0 , such points are approximately independent with the density function $r^{-2}dr \times H(d\mathbf{w})/\psi(W_{r_0})$. The likelihood function of the exceedances can then be approximated by

$$L(\boldsymbol{\varphi} \mid (r_{(i)}, \mathbf{w}_{(i)}), i = 1, \dots, k) \approx \frac{e^{-\psi(W_{r_0})} \{\psi(W_{r_0})\}^k}{k!} \prod_{i=1}^k \frac{r_{(i)}^{-2} h(\mathbf{w}_{(i)} \mid \boldsymbol{\varphi})}{\psi(W_{r_0})} \\ \propto h(\mathbf{w}_{(i)} \mid \boldsymbol{\varphi}).$$

Estimates of $\boldsymbol{\varphi}$ can be obtained by maximizing the log-likelihood $\ell(\boldsymbol{\varphi}) \propto \sum_{i=1}^k \log h(\mathbf{w}_{(i)} \mid \boldsymbol{\varphi})$. Alternatively, one can appeal to a Bayesian approach (Sabourin et al., 2013), based on the

approximate posterior density of the angular density parameters given by

$$\pi(\boldsymbol{\varphi} \mid \mathbf{w}_{(1)}, \dots, \mathbf{w}_{(k)}) = \frac{\prod_{i=1}^k h(\mathbf{w}_{(i)} \mid \boldsymbol{\varphi}) \phi(\boldsymbol{\varphi})}{\int_{\Psi} \prod_{i=1}^k h(\mathbf{w}_{(i)} \mid \boldsymbol{\varphi}) \phi(\boldsymbol{\varphi}) d\boldsymbol{\varphi}},$$

where $\phi(\boldsymbol{\varphi})$ is a suitable prior distribution on $\boldsymbol{\varphi}$ and Ψ is a suitable parameter space.

Software implementation. A non-exhaustive list of popular parametric models for the extremal dependence includes the asymmetric logistic (AL), pairwise beta (PB), tilted Dirichlet (TD), Hüsler-Reiss (HR), extremal- t (ET) and extremal-skew- t (EST) models (see e.g., [Beranger and Padoan, 2015](#); [Beranger et al., 2017](#)). The routine `fExtDep` with argument `method = "PPP"` allows to estimate $\boldsymbol{\varphi}$ by likelihood maximisation. The model class for h is specified with the argument `model` with the list of available options reported in [Table 2](#). Some classes of models are characterised by an angular density $h(\mathbf{w} \mid \boldsymbol{\varphi})$ that places a continuous mass only in the interior of \mathcal{S}_d (PB, TD, HR), while others also allow to place densities and atoms on subspaces of \mathcal{S}_d . [Beranger et al. \(2017\)](#) describe an extended version of a Maximum Likelihood (ML) method that allows for the estimation of these more complex models up to the three-dimensional case (AL, ET, EST) by exploiting the following approach.

- In the bivariate case, an observation $\mathbf{w} = (w, 1 - w)$ falls at an endpoint of \mathcal{S}_d if $w < c$ or $w > 1 - c$ for some small $c \in (0, 1/2)$ and on the interior of \mathcal{S}_d otherwise.
- A trivariate observation $\mathbf{w} = (w_1, w_2, 1 - w_1 - w_2)$ falls: at the j th corner of \mathcal{S}_d if $\mathbf{w} \in \mathcal{C}_j = \{w_j > 1 - c\}$, on the edge between the j th and l th components if $\mathbf{w} \in \mathcal{E}_{j,l} = \{w_j, w_l < 1 - c, w_j + w_l > 1 - c, w_j > 1 - 2w_l, w_l > 1 - 2w_j\}$, and in the interior of \mathcal{S}_d otherwise (i.e., if $w_j > c$, for all j) (see [Beranger et al., 2017](#), for details).

This approach is implemented when specifying a positive value for the argument `c` of `fExtDep`.

Application. We illustrate the use of `fExtDep` on the analysis of high levels of air pollution recorded in Leeds, UK, over the winter period (November 1st to February 28/29th) between 1994 and 1998 (see [Heffernan and Tawn, 2004](#)). The `pollution` dataset contains the daily maximum of the five air pollutants: particulate matter (PM10), nitrogen oxide (NO), nitrogen dioxide (NO2), ozone (O3), and sulfur dioxide (SO2), on the unit Fréchet scale. For a description of the transformation used see [Beranger and Padoan \(2015\)](#). The datasets `PNS`, `NSN` and `PNN` contains the angular components corresponding to the 100 largest radial components of the triplets of pollutants (PM10, NO, SO2), (NO2, SO2, NO) and (PM10, NO, NO2), respectively.

We run the following commands

```
R> data(pollution, package = "ExtremalDep")
R> f.et05 <- fExtDep(x = PNS, method = "PPP", model = "ET",
+   par.start = c(0.5, 0.5, 0.5, 3), trace = 2, c = 0.05)

R> f.et <- fExtDep(x = PNS, method = "PPP", model = "ET",
+   par.start = c(0.5, 0.5, 0.5, 3), trace = 2 )
```

where the first call to the `fExtDep` routine fits the extremal- t angular density (`model = "ET"`), to the pollution data (`x = PNS`), allowing for point masses at the corners by specifying the argument `c`, while the second call considers a density defined only in the interior of \mathcal{S}_d . The argument `par.start` gives starting values for the parameters and `trace = 2` allows to monitor optimization progress. The optimization method is set by the argument `optim.method` ("BFGS" by default) on which basis the routine `optim` from the `stats` library is used. When `method = "PPP"` or "`composite`", `fExtDep` returns a list object of class `ExtDep.Freq` from which the estimates of the model parameters, their standard errors, the log-likelihood maximum and the Takeuchi Information Criterion (TIC) can be extracted respectively with the `est`, `StdErr`, `logLik` and `tic` functions. The fitted model and estimation method can also be extracted using the `model` and `method` S3 generics. `fExtDep` is first ran for the extremal- t model using different values of c (0, 0.05 and 0.1) to take into account point masses at the corners of \mathcal{S}_d . The TICs indicate that only assuming mass in the interior of \mathcal{S}_d provides the best results (lowest TIC). A performance comparison between the PB, AL, TD, HR, ET and EST models defined on the interior only, is performed with TIC scores reported in Table 1 and as a result the extremal skew- t model fits the data best (lowest TIC).

| model | PB | AL | TD | HR | ET | EST |
|-------|---------|---------|---------|---------|---------|---------|
| TIC | -188.75 | -406.85 | -393.75 | -460.25 | -461.74 | -493.31 |

Table 1: TIC of models fitted to PNS data obtained using `fExtDep` with $c=0$. See Table 2 for a description of the acronyms.

Bayesian inference for the parameter φ is implemented through a model averaging approach (Sabourin et al., 2013) and available in the routine `fExtDep` using `method = "BayesianPPP"`. Table 2 reports the implemented prior densities for each angular density model. Let $\phi(\cdot, a, b)$ denote a univariate Gaussian density with mean a and variance b and $t(\cdot)$ a transformation applied to the parameters to map them to the real line. For each component of φ , the proposal density is Gaussian of the form $\phi\{t(\varphi_j^*); t(\varphi_j^{(i)}), \text{MCpar}\}$, where φ_j^* and $\varphi_j^{(i)}$ are the proposed value and the value at the i th iteration of the algorithm of the j element of φ , respectively, and `MCpar` is the variance term that needs to be specified by the user.

```
R> Hpar.hr <- list(mean.lambda = 0, sd.lambda = 3)
R> PNS.hr <- fExtDep(x = PNS, method = "BayesianPPP", model = "HR",
+   Nsim = 3e3, Nbin = 1e3, Hpar = Hpar.hr, MCpar = 0.05, seed = 14342)
```

The above display illustrates how to perform Bayesian inference for the Hüsler–Reiss model (`model = "HR"`) with 3,000 iterations (`Nsim`), a burn-in period of 1,000 iterations (`Nbin`) while the hyper-parameters are set through `Hpar` and we set the random seed `seed` for reproducibility. Given that we have a Gaussian prior distribution for the model's parameter on the log scale, setting the hyper-parameters `mean.lambda=0` and `sd.lambda=3` provides relatively non-informative prior information about the parameter. When `method = "BayesianPPP"`, `fExtDep` returns an object of class `ExtDep_Bayes` from which the posterior mean, posterior standard deviation and Bayesian Information Criterion can respectively be extracted using the `est`, `StdErr` and `bic` functions. Note that here we have 2,000 posterior draws corresponding to `Nsim - Nbin`. Again, the `model` and `method` function can be used to extract the model name and estimation procedure. The empirical mean (standard deviation) ob-

tained from the approximate posterior distribution are $\hat{\lambda}_{1,2} = 0.65(0.04)$, $\hat{\lambda}_{1,3} = 0.90(0.04)$ and $\hat{\lambda}_{2,3} = 0.98(0.04)$, while the Bayesian Information Criterion (BIC) is -449.65 . Visualization of the fitted angular density and data is obtained by calling the `plot` routine on the object of class `ExtDep_Bayes` with `type = "angular"` as default.

```
R> labs <- c(expression(PM[10]), expression(NO), expression(SO[2]))
R> plot(PNS.hr, labels = labs, cex.lab = 2.4)
```

Figure 2 showcases the fitted dependence structures for the Hüsler-Reiss model as well as the other models listed in Table 2, together with the data (solid black dots). **Note that the package ensures that the covariance matrix involved with the Hüsler-Reiss model is non-negative definite.** From visual inspection, the Hüsler-Reiss, extremal- t and extremal skew- t model appear to follow the data structure well. Given that the former is the simplest model, we focus on the Hüsler-Reiss model for the remainder of this section.

The purpose of estimating the extremal dependence is to estimate small tail probabilities as in (2.10). Similar to Cooley et al. (2010), we define the extreme event $\{PM_{10} > 95, NO > 270, SO_2 > 95\}$ and infer this probability using the routine `pExtDep` with `method = "Parametric"`, `type = "upper"` and specifying the argument `model`. Providing a vector to `par` yields a point estimate of the tail probabilities while a matrix (e.g., the posterior sample) with rows of parameter values derives an approximate posterior distribution for such probabilities. We run the following lines where the object `est` and the routine `transform` are used to transform values to unit Fréchet scale following the steps of Heffernan and Tawn (2004).

```
R> # Function extracting the Generalized Pareto parameter estimates with a
R> # threshold set at the 70% quantile
R< est.fun <- function(x){
+   x <- na.omit(x)
+   unlist(evd::fpot(x, threshold = quantile(x, probs = 0.7)))
+   [c("threshold", "estimate")]
```

| Model | model | φ | $t(\cdot)$ | Prior |
|---------------------|-------|-------------------------------|------------|---|
| pairwise beta | PB | $(\alpha, \beta_{i,j})$ | log | $\phi(\log(\alpha); \text{mean.alpha}, \text{sd.alpha})$ |
| | | | log | $\phi(\log(\beta_{i,j}); \text{mean.beta}, \text{sd.beta})$ |
| asymmetric logistic | AL | $(\alpha_S, \beta_{j,S})$ | log | $\phi(\log(\alpha_S); \text{mean.alpha}, \text{sd.alpha})$ |
| | | | logit | $\phi(\text{logit}(\beta_{j,S}); \text{mean.beta}, \text{sd.beta})$ |
| tilted Dirichlet | TD | α_j | log | $\phi(\log(\alpha_j); \text{mean.alpha}, \text{sd.alpha})$ |
| Hüsler-Reiss | HR | $\lambda_{i,j}$ | log | $\phi(\log(\lambda_{i,j}); \text{mean.lambda}, \text{sd.lambda})$ |
| extremal- t | ET | $(\rho_{i,j}, \nu)$ | atanh | $\phi(\text{atanh}(\rho_{i,j}); \text{mean.rho}, \text{sd.rho})$ |
| | | | log | $\phi(\log(\nu); \text{mean.nu}, \text{sd.nu})$ |
| extremal-skew- t | EST | $(\rho_{i,j}, \alpha_i, \nu)$ | atanh | $\phi(\text{atanh}(\rho_{i,j}); \text{mean.rho}, \text{sd.rho})$ |
| | | | identity | $\phi(\alpha_i; \text{mean.alpha}, \text{sd.alpha})$ |
| | | | log | $\phi(\log(\nu); \text{mean.nu}, \text{sd.nu})$ |

Table 2: List of available angular density model families and prior densities that are considered with them when using `fExtDep(method = "BayesianPPP", ...)`.

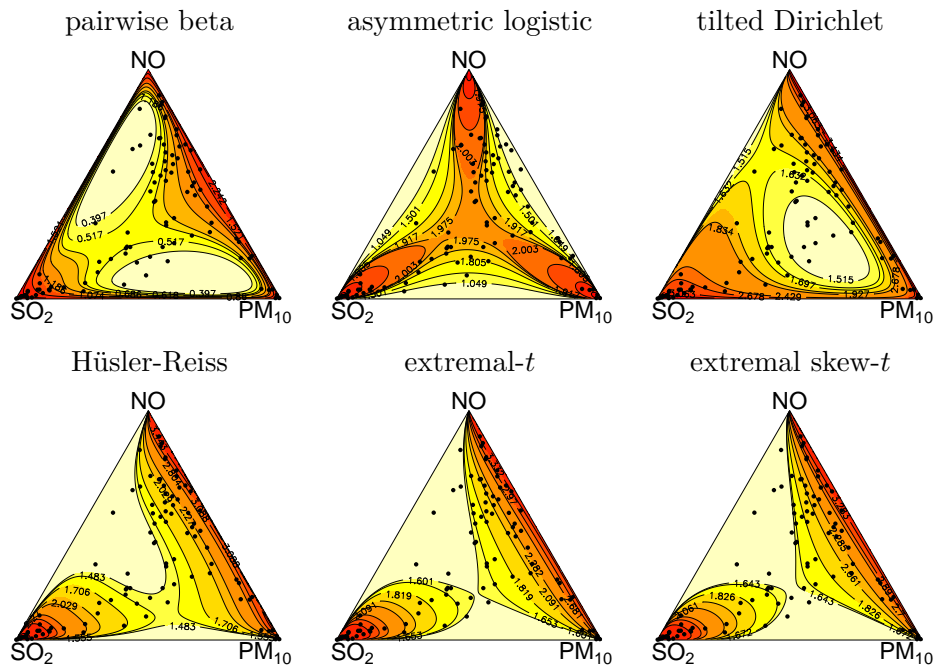


Figure 2: Estimated angular densities of the models from Table 2.

```

+ }

R> est <- apply(winterdat, 2, est.fun)

R> # Transforming univariate data to unit Frechet scale
R> transform <- function (x, data, par){
+   data <- na.omit(data)
+   if(x > par[1]){
+     emp.dist <- mean(data <= par[1])
+     dist <- 1 - (1 - emp.dist) * max(0, 1 + par[3] * (x - par[1]) /
+       par[2])^(-1 / par[3])
+   }else{dist <- mean(data <= x)}
+   return(-1/log(dist))
+ }

R> th <- c(95, 270, 110, NA, 95)
R> Th <- sapply(c(1:3,5), function(x)
+   transform(th[x], data = winterdat[, x], par = est[, x]))

R> names(Th) <- colnames(winterdat[c(1:3, 5)])
R> xl.PNS <- bquote("P(" * PM[10] ~ ">" ~ .(th[1]) * ", NO" ~ ">"
+   ~ .(th[2]) * ", " * SO[2] ~ ">" ~ .(th[5]) * ")")

R> # Compute and plot probabilities of joint exceedances

```

```
R> P.PNS <- pExtDep(q = Th[colnames(PNS)], type = "upper",
+   method = "Parametric", model = "HR", par = PNS.hr$stored.vals,
+   xlab = x1.PNS)
```

The left panel of Figure 3 provides the approximate posterior distribution of the tail probabilities in (2.10) for the extreme events. Note that a smoother output could be obtained by considering 3,000 iterations (`Nsim`) and a burn-in period of 1,000 iterations (`Nburn`). Formula (2.10) can also be used to derive a possible definition of the so-called *joint return level*, see Beranger and Padoan (2015) for details, i.e., the sequence of values $y_{j;p}, j \in J \subset \{1, \dots, d\}$, that satisfies the equation

$$p = \mathbb{P}(Y_j > y_{j;p}, X_i > x_i, j \in J, i \in \{1, \dots, d\} \setminus J),$$

for a given small probability $p \in (0, 1)$, where x_i , with $i \in \{1, \dots, d\} \setminus J$ is a sequence of fixed high thresholds. An example where the plot routine is used on an object of class `ExtDep_Bayes` to infer joint return levels is given by running the following commands.

```
R> Q.fix <- c(NA, Th[c(2, 4)])
R> PM10.range <- seq(from = est[1, 1], to = 400, by = 5)
R> Q.range <- sapply(PM10.range, transform, data = winterdat[, 1],
+   par = est[, 1])
R> set.seed(1)

R> # Compute and plot probabilities of return levels
R> rl.PM10 <- plot(x = PNS.hr, type = "returns", Q.fix = Q.fix,
+   Q.range = Q.range, Q.range0 = PM10.range, labels = expression(PM[10]),
+   main = bquote("Return level for" ~ PM[10] ~ "when NO" ~ ">" ~
+     .(th[2]) ~ "and" ~ SO[2] ~ ">" ~ .(th[5])),
+   cex.lab = 1.4, cex.axis = 1.4, subsamp = 0.05)
```

Only a 5% subset of the posterior is considered (`subsamp = 0.05`) for a faster evaluation but, in practice, we do recommend using the full posterior by leaving `subsamp` unspecified. Returns levels are displayed when `object = "returns"`. In addition, the argument `Q.fix` is required and should take the form of a vector of length equal to the model dimension (2 or 3), where quantile values can be fixed for some components while others (NAs) are left to vary. The `Q.range` argument provides a vector (or matrix) of quantile values on the unit Fréchet scale, for those that aren't fixed. If `Q.fix` contains a single NA then `Q.range` must be a vector or a single column matrix. `Q.range0` provide the same sequences as `Q.range` but on the original scale. This plotting procedure relies on the `pExtDep` routine. The middle panel of Figure 3 reports the univariate joint return level curve for PM_{10} jointly to the extreme event $\{\text{NO} > 270, \text{SO}_2 > 95\}$, corresponding to the return period $1/p$, which has been estimated by the approximate posterior mean and uncertainty is given by the 95% (pointwise) credibility intervals. Similarly, the right panel of Figure 3 depicts the posterior mean and 95% credibility intervals of the estimated contour levels of the bivariate return levels for $(\text{PM}_{10}, \text{NO})$ jointly to extreme event $\text{SO}_2 > 95$. Note that conditional return levels can be obtained by specifying `cond=TRUE`, the conditional event being the fixed event.




Figure 3: Left panel shows the approximate posterior distribution of tail probability with crosses indicating posterior median and 95% (pointwise) credibility interval and a dot indicating the posterior mean. Middle and right panels show the posterior mean (solid line) and posterior 95% credibility interval (dashed lines) of univariate and bivariate return levels associated with return period $1/p$, when respectively fixing two and one components.

2.4 Semi- and non-parametric modeling of the extremal dependence

Common objectives when analyzing extreme values include assessment of the dependence level among extremes, e.g., using the dependence structures in (2.4)–(2.5), simulation of multiple extremes, estimation of small joint tail probability, e.g., (2.9)–(2.10), and estimation of extreme quantile regions. The next sections describe the steps required to accomplish such goals.

2.4.1 Extremal dependence estimation via Bernstein polynomials

Theory. We give here a brief summary of a simple estimation method for the extremal dependence with an arbitrary number of componentwise maxima proposed by Marcon et al. (2017). We refer to the aforementioned paper for further details.

The method consists of two steps: 1) a nonparametric pilot estimation of the Pickands dependence function, 2) a regularization of such estimate by projecting it into a polynomial representation in Bernstein form imposing conditions (C1)–(C2). Given some i.i.d. random vectors $\mathbf{Y}_1, \dots, \mathbf{Y}_k$ (approximately) distributed as multivariate EV distribution in (2.2), the first step is achieved by using the madogram-based estimator of the Pickands dependence function. This is given, for all $\mathbf{t} \in \mathcal{R}$, by

$$\hat{A}_k(\mathbf{t}) = \frac{\hat{\nu}_k(\mathbf{t}) + d^{-1} \sum_{j=1}^d \{t_j / (1 + t_j)\}}{1 - \hat{\nu}_k(\mathbf{t}) - d^{-1} \sum_{j=1}^d \{t_j / (1 + t_j)\}},$$

where

$$\hat{\nu}_k(\mathbf{t}) = \frac{1}{k} \sum_{i=1}^k \left\{ \max_{1 \leq j \leq d} F_{k,j}^{1/t_j}(Y_{i,j}) - \frac{1}{d} \sum_{j=1}^d F_{k,j}^{1/t_j}(Y_{i,j}) \right\},$$

and $F_{k,j}$ denotes the empirical distribution of the j th variable. For the regularization of the pilot estimate, take $\kappa > d$ and let Γ_κ be the set of multi-indices $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_d)^\top \in \{0, 1, \dots, \kappa\}^d$ such that $\alpha_1 + \dots + \alpha_d = \kappa$ and $\alpha_d = \kappa - \alpha_1 - \dots - \alpha_{d-1}$, whose cardinality is

denoted by $C_\kappa = |\Gamma_\kappa|$. Let

$$A_\kappa(\mathbf{t}) = \sum_{\alpha \in \Gamma_\kappa} \beta_\alpha b_\alpha(\mathbf{t}, \kappa), \quad \mathbf{t} \in \mathcal{R}, \quad (2.14)$$

be the Bernstein–Bézier polynomial representation of the Pickands dependence function, where for each $\alpha \in \Gamma_\kappa$,

$$b_\alpha(\mathbf{t}, \kappa) = \frac{\kappa!}{\prod_{j=1}^d \alpha_j!} \prod_{j=1}^{d-1} t_j^{\alpha_j} (1 - t_1 - \dots - t_{d-1})^{\alpha_d}, \quad \mathbf{t} \in \mathcal{R}, \quad (2.15)$$

is the Bernstein polynomial basis function of index α and degree κ .

Now, let \mathcal{A} be the family of functions $f : \mathcal{R} \rightarrow [1/d, 1]$ satisfying conditions (C1)–(C2), and $\mathcal{A}_\kappa = \{\mathbf{t} \mapsto \mathbf{b}_\kappa(\mathbf{t})\boldsymbol{\beta}_\kappa; \boldsymbol{\beta}_\kappa \in [0, 1]^{p_\kappa} \text{ such that } \mathbf{R}_\kappa \boldsymbol{\beta}_\kappa \geq \mathbf{r}_\kappa\}$ be a sequence of families of constrained multivariate Bernstein–Bézier polynomials on \mathcal{R} , where $\mathbf{b}_\kappa(\mathbf{t})$ is the row vector $(b_\alpha(\mathbf{t}, \kappa), \forall \alpha \in \Gamma_\kappa)$, $\boldsymbol{\beta}_\kappa$ is a column vector, \mathbf{R}_κ is a suitable $(q \times p_\kappa)$ matrix of full row rank and \mathbf{r}_κ is a $(q \times 1)$ vector. The constraint $\mathbf{R}_\kappa \boldsymbol{\beta}_\kappa \geq \mathbf{r}_\kappa$ on the coefficient vector $\boldsymbol{\beta}_\kappa$ guarantees that each member of \mathcal{A}_κ satisfies (C1)–(C2). A projection estimator of the Pickands dependence function based on the estimator $\hat{A}_k(\mathbf{t})$ is the solution to the optimization problem

$$\tilde{A}_{k,\kappa} = \arg \min_{f \in \mathcal{A}_\kappa} \|\hat{A}_k - f\|.$$

For a finite set of points $\{\mathbf{t}_u : u = 1, \dots, U\}$, with $U = 1, 2, \dots$ and $\mathbf{t}_u \in \mathcal{R}$ such a solution is obtained by finding the minimizer $\hat{\boldsymbol{\beta}}_\kappa$ of the constrained least-squares problem

$$\begin{aligned} \hat{\boldsymbol{\beta}}_\kappa &= \arg \min_{\boldsymbol{\beta}_\kappa \in [0, 1]^{p_\kappa} : \mathbf{R}_\kappa \boldsymbol{\beta}_\kappa \geq \mathbf{r}_\kappa} \frac{1}{U} \sum_{u=1}^U \left\{ \mathbf{b}_\kappa(\mathbf{t}_u) \boldsymbol{\beta}_\kappa - \hat{A}_k(\mathbf{t}_u) \right\}^2 \\ &= \left(\mathbf{b}_\kappa^\top \mathbf{b}_\kappa \right)^{-1} \mathbf{b}_\kappa^\top \hat{A}_k - \left(\mathbf{b}_\kappa^\top \mathbf{b}_\kappa \right)^{-1} \mathbf{r}_\kappa^\top \boldsymbol{\lambda}, \end{aligned}$$

where $\boldsymbol{\lambda}$ is a vector of Lagrange multipliers.

Software implementation. This method is implemented in the routine `beed` of **ExtremalDep** and its usage is demonstrated through the analysis of heavy rainfall in France. Hydrologists are interested in identifying different geographic regions that differ from each other in that there are clusters of weather stations whose data exhibit substantially different levels of extreme dependence. Within a cluster, climate characteristics are expected to be homogeneous, whereas they can be quite heterogeneous between clusters. The dataset is available through the `PrecipFrance` object which consists of weekly maxima of hourly rainfall (`$precip`) recorded at 92 weather stations in France, during the Fall season between 1993 and 2011, yielding a sample size of $k = 228$ observations. Coordinates of each station are stored in the list elements `$lat` and `$lon`. Note that hourly rainfall has been appropriately pre-processed and with them the weekly maxima have been already computed from [Bernard et al. \(2013\)](#).

Through the `PAMfado` routine, the following code chunk applies the algorithm proposed by [Bernard et al. \(2013\)](#) to the weekly maxima of hourly rainfall such that the weather stations are divided into seven clusters.

```
R> data(PrecipFrance, package = "ExtremalDep")
R> attach(PrecipFrance)
```

```
R> nclust <- 7 # number of clusters
R> PAMmado <- PAMfmado(precip, nclust) # Apply clustering algorithm to data
```

For each cluster, five stations are randomly selected in order to have equal size ($d = 5$) clusters and for each group, the Pickands dependence function is estimated using the Bernstein projection estimator based on the madogram with polynomial degree equal $k = 7$. To summarise, an estimate of the extremal coefficient is computed using (2.7), i.e. $\hat{\varrho}_k = 5\tilde{A}_{k,\kappa}(1/5, \dots, 1/5)$, by running the following commands

```
R> clust <- PAMmado$clustering
R> d <- 5 # dimension = number of stations
R> stationsn <- matrix(NA, nclust, d)
R> xx <- simplex(d = d, n = 15)
R> fit <- list(length = nclust)
R> est <- vector(length = nclust)
R> set.seed(1)
R> for(i in 1:nclust){
+   # Randomly select 5 stations from cluster i
+   stationsn[i, ] <- sample(which(clust == i), 5)
+   data_tmp <- precip[, stationsn[i, ]]
+   # Transform margins to unit Frechet scale
+   data_uf_tmp <- trans2UFrechet(data_tmp, type = "Empirical")
+   # Estimate Pickands dependence function
+   fit[[i]] <- beed(data_uf_tmp, xx, d, "md", "emp", k = 7)
+   # Extract extremal coefficient
+   est[i] <- fit[[i]]$extind}
```

The left panel of Figure 4 indicates the extremal dependence is strongest in the center of the country, away from the coasts, where the conjunction of different densities of air masses produces extreme rain storms. This is consistent with climatologists expectations since extreme precipitation that affects the Mediterranean coast in the fall is caused by the interaction of southern and mountain winds coming from the Pyrénées, Cévennes and Alps regions. In the north, heavy rainfall is produced by mid-latitude perturbations in Brittany (or regions further north) and Paris. Within clusters, extremes are strongly dependent. The right panel of Figure 4 shows the pairwise extremal coefficients from all 92 stations, computed through the estimated Pickands dependence functions using the raw madogram estimator (MD) and its Bernstein projection (MD-BP), versus the geodesic distance between sites. We have $\hat{\varrho}_k \leq 1.5$ for the locations that are less than 200 km apart, meaning that extremes are either strongly or mildly dependent, while for sites more than 200 km apart, we have $\hat{\varrho}_k > 1.5$, meaning that extremes are at most weakly dependent or even independent. The graph also shows the benefits of the projection method: after projection, the extremal coefficients fall within the admissible range $[1, 2]$. The `beed` function is used for the MD-PB estimator whereas `madrogram` implements the raw madogram estimator MD. The following set of commands were used for the estimation.

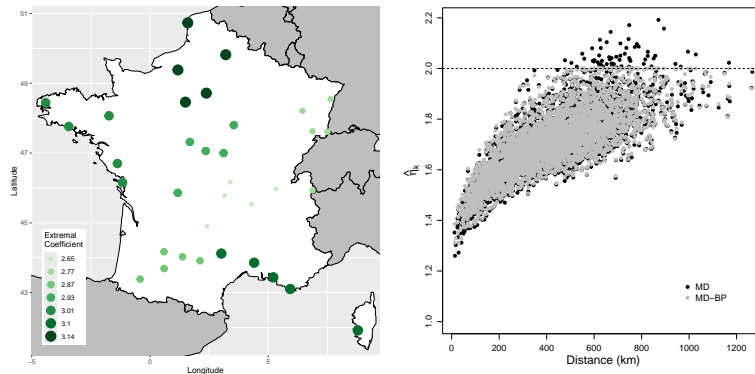


Figure 4: Clusters of 35 weather stations and their estimated extremal coefficients in dimension $d = 5$ using French weekly precipitation maxima (left). Evolution of pairwise extremal coefficients with respect to the distance (right).

```
R> library(geosphere)

R> pairs <- t(combn(92, 2))
R> pairs.LL0 <- cbind(pairs, lon[pairs[, 1]], lat[pairs[, 1]],
+ lon[pairs[, 2]], lat[pairs[, 2]])
R> pairs.dist <- apply(pairs.LL0, 1, function(x) distm(x = x[3:4],
+ y = x[5:6], fun = distGeo)) / 1000

R> pairs.EC <- pairs.ECBP <- vector(length = nrow(pairs))
R> S <- simplex(d = 2, n = 49)
R> for(i in 1:nrow(pairs)){
+ # Transform margins to unit Frechet scale
+ data.tmp <- trans2UFrechet(precip[,pairs[i, ]], type = "Empirical")
+ # Estimate Pickands dependence function
+ fit.tmp <- beed(data.tmp, S, 2, "md", "emp", k = 7)
+ # Extract extremal coefficient
+ pairs.ECBP[i] <- fit.tmp$extind
+ pairs.EC[i] <- 2 * madogram(S, data.tmp, "emp")[which(S[, 1] == 0.5)]}
```

2.4.2 Bernstein polynomials modeling and Bayesian nonparametric inference

Theory. In many applications, it is crucial to estimate joint probabilities of events such as $p = \mathbb{P}(Y_1 > y_1, Y_2 > y_2)$, for pair of extreme values value (y_1, y_2) . These in turn allow for the estimation of related quantities such as small conditional probabilities $\mathbb{P}(Y_j > y_j \mid Y_i > y_i)$ with $i, j = 1, 2$. Section 2.1 describes a theoretical framework that can be used to approximate such probabilities and we now briefly describe a Bayesian nonparametric inference method based on Bernstein polynomials introduced by [Marcon et al. \(2016\)](#).

Let $\mathbf{y}_1, \dots, \mathbf{y}_n$ be n (independent) observations (approximately) distributed as a bivariate GEV distribution. First, a prior distribution on the parameters $\boldsymbol{\theta}_j = (\mu_j, \sigma_j, \gamma_j)^\top, j = 1, 2$, of the marginal GEV distributions is specified as $\Pi(\boldsymbol{\theta}_j) = \Pi(\mu_j)\Pi(\sigma_j)\Pi(\gamma_j) \propto 1/\sigma_j$ with $\sigma_j > 0$, i.e., a product of uniform prior distributions on the real line for $\mu_j, \log(\sigma_j)$ and γ_j .

Note that this improper prior distribution leads to a proper posterior distribution (Northrop and Attalides, 2016). Samples from the posterior distribution are generated using an adaptive (Gaussian) random-walk Metropolis-Hastings (RWMH) algorithm (Haario et al., 2001; Garthwaite et al., 2016). For technical details, see Appendix A.2.

In the bivariate case, the polynomial Pickands dependence function in Bernstein form becomes

$$A_\kappa(t \mid \boldsymbol{\beta}_\kappa) := \sum_{j=0}^{\kappa} \beta_j b_j(t; \kappa), \quad t \in \mathcal{R},$$

for $\kappa = 0, 1, \dots$, where $\boldsymbol{\beta}_\kappa = (\beta_0, \dots, \beta_\kappa)^\top$ and the polynomial basis in (2.15) reduces to

$$b_j(t; \kappa) = \frac{\kappa!}{j!(\kappa-j)!} j^j (1-j)^{\kappa-j}, \quad j = 0, \dots, \kappa.$$

For fixed degree κ , Marcon et al. (2016) derived the restrictions on $\boldsymbol{\beta}_\kappa$ so that A_κ satisfies conditions (C2)–(C3) and is therefore a valid Pickands dependence function. They also demonstrated that the polynomial A_κ implies that the distribution of the corresponding angular measure can be written as a Bernstein polynomial of degree $\kappa - 1$,

$$H_{\kappa-1}([0, w] \mid \boldsymbol{\eta}_\kappa) := \begin{cases} \sum_{j \leq \kappa-1} \eta_j b_j(w; \kappa-1) & \text{if } w \in [0, 1), \\ 1 & \text{if } w = 1, \end{cases} \quad (2.16)$$

where $\boldsymbol{\eta}_\kappa = (\eta_0, \dots, \eta_{\kappa-1})^\top$. Additionally, Marcon et al. (2016) established the conditions on the coefficients $\boldsymbol{\eta}_\kappa$ so that $H_{\kappa-1}$ is the distribution of a valid angular measure and proposed a Bayesian nonparametric procedure for the inference of both A_κ and $H_{\kappa-1}$, which can then be used to compute an approximation of joint tail probabilities. The proposed method consists of three main steps: 1) the specification of a prior distribution for the polynomial order and coefficients $(\kappa, \boldsymbol{\eta}_\kappa)$, 2) the derivation of the likelihood function, 3) the definition of a Markov Chain Monte Carlo (MCMC) algorithm for the posterior distribution computation. In particular, $\Pi(\kappa, \boldsymbol{\eta}_\kappa) = \Pi(\boldsymbol{\eta}_\kappa \mid \kappa) \Pi(\kappa)$, where $\Pi(\kappa)$ is a prior on the polynomial order (e.g., Poisson, negative Binomial, etc.) and $\Pi(\boldsymbol{\eta}_\kappa \mid \kappa) = \Pi(\eta_1, \dots, \eta_{\kappa-2} \mid p_1, p_0, \kappa) \Pi(p_1 \mid \kappa, p_0) \Pi(p_0)$ is a prior on the coefficients $\boldsymbol{\eta}_\kappa$, where $\Pi(p_1 \mid \kappa, p_0)$ and $\Pi(p_0)$ are the priors on the coefficients representative of the atoms $\eta_0 = p_0$ and $\eta_{\kappa-1} = 1 - p_1$ at the edges of the interval $[0, 1]$. Such priors are specified as uniform distributions on suitable intervals, which have been chosen to ensure that the resulting Bernstein polynomial satisfies the constraint (C3). Specification of the prior $\Pi(\boldsymbol{\eta}_\kappa \mid \kappa)$ induces also a prior on the coefficients $\boldsymbol{\beta}_\kappa$ of the corresponding polynomial $A_\kappa(t \mid \boldsymbol{\beta}_\kappa)$, which automatically satisfy constraints (C1)–(C2). To deal with the fact at each MCMC iteration the dimension of $\boldsymbol{\theta}_t$ changes with κ , a trans-dimensional MCMC scheme is considered following Marcon et al. (2016). **At iteration s** , $(\boldsymbol{\eta}_{\kappa^{(s)}}^{(s)}, \kappa^{(s)})$ is updated using the proposal distribution $q(\boldsymbol{\eta}_\kappa, \kappa \mid \boldsymbol{\eta}_{\kappa^{(s)}}^{(s)}, \kappa^{(s)}) = \Pi(\boldsymbol{\eta}_\kappa \mid \kappa) q_\kappa(\kappa \mid \kappa^{(s)})$, where $q_\kappa(\kappa \mid \kappa^{(s)})$ is defined such that if $\kappa^{(s)} = 3$, it places mass on $\kappa = 4$ with probability 1 and if $\kappa^{(s)} > 3$ it places mass on $\kappa^{(s)} - 1$ and $\kappa^{(s)} + 1$ with equal probability. Finally, the likelihood function is defined as $\mathcal{L}(\boldsymbol{\vartheta}) = \prod_{i=1}^m \mathcal{L}(\mathbf{y}_i \mid \boldsymbol{\vartheta})$, where $\boldsymbol{\vartheta} = (\boldsymbol{\theta}_1^\top, \boldsymbol{\theta}_2^\top, \boldsymbol{\beta}_\kappa^\top, \kappa)^\top$ and for any y_j such that $(1 + \gamma_j(y_j - \mu_j)/\sigma_j) > 0$ with $j = 1, 2$,

$$\mathcal{L}(\mathbf{y} \mid \boldsymbol{\vartheta}) = G_{\boldsymbol{\theta}} \{ \mathbf{y} \mid A_\kappa(t \mid \boldsymbol{\beta}_\kappa) \} \frac{1}{\sigma_1 \sigma_2} \left(1 + \gamma_1 \frac{y_1 - \mu_1}{\sigma_1} \right)^{-1/\gamma_1 - 1} \left(1 + \gamma_2 \frac{y_2 - \mu_2}{\sigma_2} \right)^{-1/\gamma_2 - 1}$$

$$\times \left[\{A_\kappa(t | \boldsymbol{\beta}_\kappa) - tA'_\kappa(t | \boldsymbol{\beta}_\kappa)\} \{A_\kappa(t | \boldsymbol{\beta}_\kappa) - (1-t)A'_\kappa(t | \boldsymbol{\beta}_\kappa)\} + \frac{t(1-t)}{r} A''_\kappa(t | \boldsymbol{\beta}_\kappa) \right]$$

and where $\boldsymbol{\theta} = (\boldsymbol{\theta}_1^\top, \boldsymbol{\theta}_2^\top)^\top$, $t = z_2/r$, $r = z_1 + z_2$, see [Marcon et al. \(2016\)](#) and [Beranger et al. \(2021\)](#) for details. The MCMC scheme for the joint inference of marginal distribution and extremal dependence is reported in Algorithm 1 of Appendix A.3. For pairs (y_1^*, y_2^*) of future unobserved yet extremes values, the joint exceeding probability can be estimated through the Bayesian paradigm using the posterior predictive distribution which can be approximated, given a sample $\boldsymbol{\theta}_i$ with $i = 1, \dots, M$, from the posterior distribution as

$$\begin{aligned} \mathbb{P}(Z_1 > z_1^*, Z_2 > z_2^*) &\approx 2 \sum_{i=1}^M \frac{1}{\kappa_i} \sum_{j=0}^{\kappa_i-2} (\eta_{i,j+1} - \eta_{i,j}) \\ &\times \left\{ \frac{(j+1) \text{B}\left(\frac{z_1^*}{z_1^*+z_2^*} \mid j+2, \kappa_i - j - 1\right)}{z_1^*} + \frac{(k_i - j - 1) \text{B}\left(\frac{z_2^*}{z_1^*+z_2^*} \mid \kappa_i - j, j + 1\right)}{z_2^*} \right\}, \end{aligned} \quad (2.17)$$

where (Z_1, Z_2) are distributed as a bivariate GEV with unit-Fréchet margins, $z_j^* = \sum_{i=1}^M \{1 + \gamma_{i,j}(y_j^* - \mu_{i,j})/\sigma_{i,j}\}^{1/\gamma_{i,j}}$ with $j = 1, 2$ and $\text{B}(x \mid a, b)$, $x \in [0, 1]$, denotes the distribution function of a beta random variable with shape parameters $a, b > 0$ (see [Marcon et al., 2016](#), for details).

Application. We show the utility of the methodology by analyzing the joint extremal behavior of log-returns of exchange rates between Great British Pound and U.S. Dollar (GBP/USD), and Great British Pound and Japanese Yen (GBP/JPY). The data are available as `logReturns` and consists of the monthly maxima of daily log-returns exchange rates from March 1991 to December 2014 (286 observations). Exchange rates are `$USD` and `$JPY`, while `$date_USD` and `$date_JPY` are the date when the monthly maxima was attained. The following command lines load the data and remove the trend and seasonality from each maxima series using the `ts()` and `stl()` functions from the `stats` package.

```
R> data(logReturns, package = "ExtremalDep")

R> # Create time-series objects
R> mm_gbp_usd <- ts(logReturns$USD, start = c(1991, 3), end = c(2014, 12),
+   frequency = 12)
R> mm_gbp_jpy <- ts(logReturns$JPY, start = c(1991,3), end = c(2014, 12),
+   frequency = 12)

R> # Decompose time series into seasonal, trend and irregular components
R> seas_usd <- stl(mm_gbp_usd, s.window = "period")
R> seas_jpy <- stl(mm_gbp_jpy, s.window = "period")

R> mm_gbp_usd_filt <- mm_gbp_usd - rowSums(seas_usd$time.series[, -3])
R> mm_gbp_jpy_filt <- mm_gbp_jpy - rowSums(seas_jpy$time.series[, -3])
```

The top-left and bottom-left panels of Figure 5 display the detrended and deseasonalised maxima and the scatterplot in the right panel shows strong dependence between the extremes of both exchanges rates.

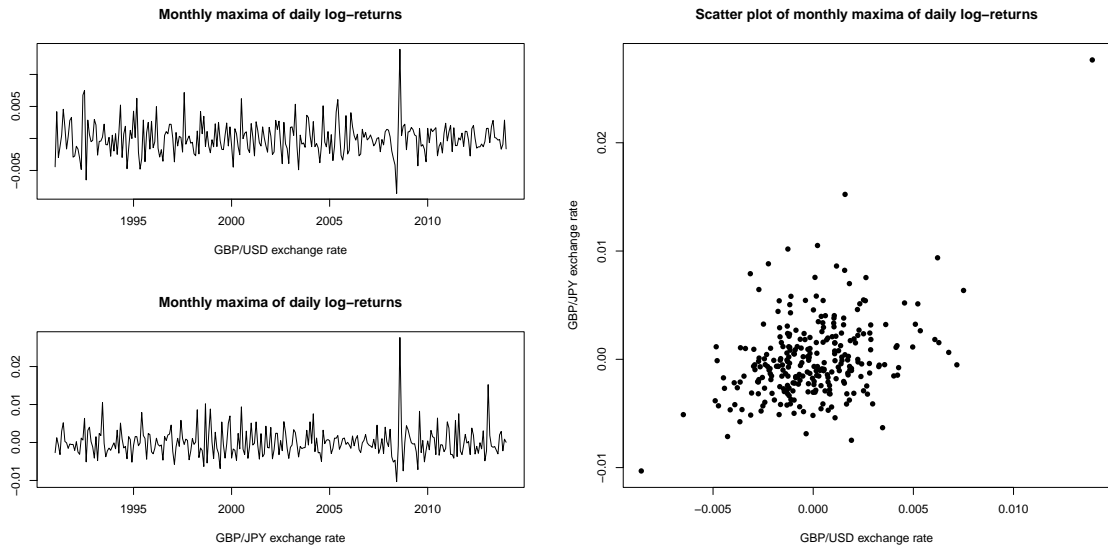


Figure 5: Detrended and deseasonalised times series of monthly-maxima of log-returns of GBP/USD and GBP/JPY exchange rates (left) and corresponding scatterplot (right).

The extremal dependence structure is then estimated using the Bayesian nonparametric framework described above by running the following commands

```
R> hyperparam <- list(mu.nbinom = 3.2, var.nbinom = 4.48)
R> mm_gbp <- cbind(as.vector(mm_gbp_usd_filt), as.vector(mm_gbp_jpy_filt)) * 100

R> set.seed(123)
R> # Estimate dependence structure using Bayesian nonparametric
R> par0 <- c(-0.1, 0.1, -0.1)
R> gbp_mar <- fExtDep.np(x = mm_gbp, method = "Bayesian", par10 = par0,
+   par20 = par0, sig10 = 0.1, sig20 = 0.1, k0 = 5,
+   hyperparam = hyperparam, nsim = 5e+4)

R> diagnostics(gbp_mar)
```

Note that the data has been rescaled by a factor of 100 to avoid numerical instabilities due to the small scale of the data. The `fExtDep.np` routine allows for nonparametric estimation of the extremal dependence and `method = "Bayesian"` specifies that such a dependence is in Bernstein polynomial form and a Bayesian approach is used for the inference. The argument `mar.fit = TRUE` (default) allows for joint estimation of the margins and dependence while `mar.prelim = TRUE` (default) fits the marginal distributions using the RWMH algorithm and the `fGEV` routine, in order to obtain starting values for the marginal parameters. The prior distribution for κ is set to be a negative binomial on $\kappa - 3$ with mean 3.2 and variance 4.48, and for the point masses p_0 and p_1 uniform distributions on $[0, 0.5]$ and $[a, b]$ respectively, where $a = \max\{0, (\kappa - 1)p_0 - \kappa/2 + 1\}$ and $b = (p_0 + \kappa/2 - 1)/(\kappa - 1)$. Note that both the negative binomial and the Poisson distributions are convenient choices for the prior distribution of the polynomial order κ and the selected hyper-parameter values simply allow for a broad

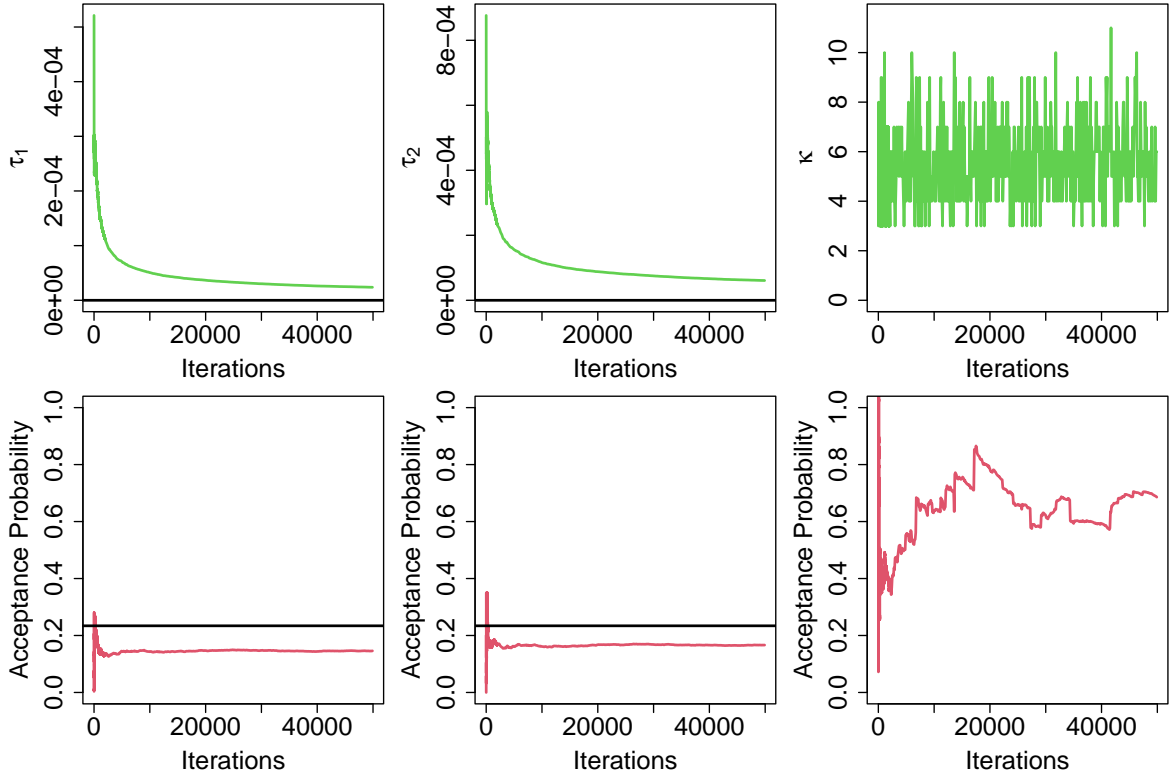


Figure 6: Diagnostic plots for the MCMC algorithm. Left and centre columns focus on the marginal components, illustrating the respective sampler scaling parameter and acceptance probability. Right column focuses on the dependence structure, illustrating the value of the polynomial degree throughout the algorithm and acceptance probability.

range of admissible values. For the prior of the point mass p_1 conditionally on κ and p_0 , the bounds of the uniform distribution are set to fulfill the conditions on the coefficients η_κ . We refer to [Marcon et al. \(2016\)](#) for a comprehensive discussion. In this example, only the hyper-parameters need to be specified since `prior.k = "nbinom"` and `prior.pm = "unif"` by default. Lastly, a two-column object `mm_gbp` representing the data is created.

The `gbp_mar` object is of class `ExtDep_npBayes` and contains the posterior samples for all the parameters: point mass p_0 and p_1 , polynomial coefficients η and degree κ , marginal parameters θ_1 and θ_2 , respectively reported by the arguments `pm`, `eta`, `k`, `mar1` and `mar2`. Its also contains binary vectors indicating the accepted marginal and dependence proposals (`accepted.mar1`, `accepted.mar2` and `accepted`) and the marginal proposals that were rejected right away for not being in the parameter space (`straight.reject1` and `straight.reject2`). It also includes acceptance probabilities at each step (`acc.vec`, `acc.vec.mar1` and `acc.vec.mar2`), the marginal scaling parameters τ_1 and τ_2 (`sig1.vec` and `sig2.vec`), as well as some of the inputs. The `diagnostics` function investigates the convergence of the algorithm, producing Figure 6. The top panels display the scaling parameters τ_1 and τ_2 from the marginal proposals and the polynomial degree κ , as function of the iterations. The bottom panels show the corresponding acceptance probability with the desired acceptance probability of 0.234 (horizontal black solid lines). For both margins, the acceptance rate stabilizes at a level slightly below the nominated level of 0.234 which yields decreasing values of the corresponding τ parameters.

One could be tempted to increasing the number of MCMC iterations or the burn-in period but given the range of values taken, this would have little impact. In addition, the first two rows of the Figure 6 indicate that acceptance probability for the marginal estimates is slightly below the nominated levels. Actual convergence of the marginal chains can be confirmed using traceplots on the columns of the objects `gbp_mar$mar1` and `gbp_mar$mar2`. Overall, Figure 6 suggests a burn-in period of 30,000 iterations and we now run the following commands

```
R> # Compute summary statistics of the MCMC output
R> gbp_mar_sum <- summary(object = gbp_mar, burn = 30000, plot = TRUE)

R> mm_gbp_rg <- apply(mm_gbp, 2, quantile, c(0.9, 0.995))

R> y_gbp_usd <- seq(from = mm_gbp_rg[1, 1], to = mm_gbp_rg[2, 1], length = 20)
R> y_gbp_jpy <- seq(from = mm_gbp_rg[1, 2], to = mm_gbp_rg[2, 2], length = 20)
R> y <- as.matrix(expand.grid(y_gbp_usd, y_gbp_jpy, KEEP.OUT.ATTRS = FALSE))

R> # Compute and plot return levels
R> ret_marg <- returns(x = gbp_mar, summary.mcmc = gbp_mar_sum, y = y,
+   plot = TRUE, data = mm_gbp, xlab = "GBP/USD exchange rate (x100)",
+   ylab = "GBP/JPY exchange rate (x100)")
```

The `summary` routine applied to objects of class `ExtDep_npBayes` computes summary statistics of the MCMC output, including posterior sample, mean and 95% credibility intervals (the latter can be modified using argument `cred`) for the angular density and Pickands dependence function. Posterior mean and credibility intervals are also computed for all parameters. Setting `plot=TRUE` displays the posterior mean and credibility intervals for both angular density and Pickands dependence function, as well as the prior and posterior distribution for the point mass p_0 and the polynomial degree κ . The `returns` routine inputs objects of class `ExtDep_npBayes` and outputs of the `summary` function (via the arguments `x` and `summary.mcmc`) to compute exceeding probabilities as defined in (2.17), for extreme values specified by `y`. The argument `plot` calls the `plot` routine on the object `x` to visualize such probabilities as long as `y` defines a square grid and `data` adds the relevant datapoints. Usage of the `summary` and `returns` functions is provided below with graphical outputs presented in Figure 7. As mentioned at the beginning of this section, computing small conditional probabilities such as $\mathbb{P}(\text{GBP/USD} > q_1 \mid \text{GBP/JPY} > q_2)$ or $\mathbb{P}(\text{GBP/JPY} > q_2 \mid \text{GBP/USD} > q_1)$ can be of interest. We proceed by taking $(q_1, q_2) \approx (0.0069, 0.0102)$ which corresponds to the observed marginal 99% quantiles. The joint probability of exceedance is computed (through `returns`) using the approximation $\mathbb{P}(\text{GBP/USD} > q_1, \text{GBP/JPY} > q_2) \approx \mathbb{P}(Z_1 > z_1^*, Z_2 > z_2^*) \approx 0.0044$, where z_1^* and z_2^* are the transformation of q_1 and q_2 to unit Fréchet scale. Using the `pGEV` function to evaluate the marginal probability of exceedance, we obtain $\mathbb{P}(\text{GBP/USD} > q_1 \mid \text{GBP/JPY} > q_2) \approx 0.3169$ and $\mathbb{P}(\text{GBP/JPY} > q_2 \mid \text{GBP/USD} > q_1) \approx 0.2767$. The computations are presented in the code below.

```
R> qs <- apply(mm_gbp, 2, quantile, c(0.99))

R> # Compute joint probability of exceedance
R> jointP <- returns(x = gbp_mar, summary.mcmc = gbp_mar_sum,
+   y = matrix(qs, ncol = 2))
```

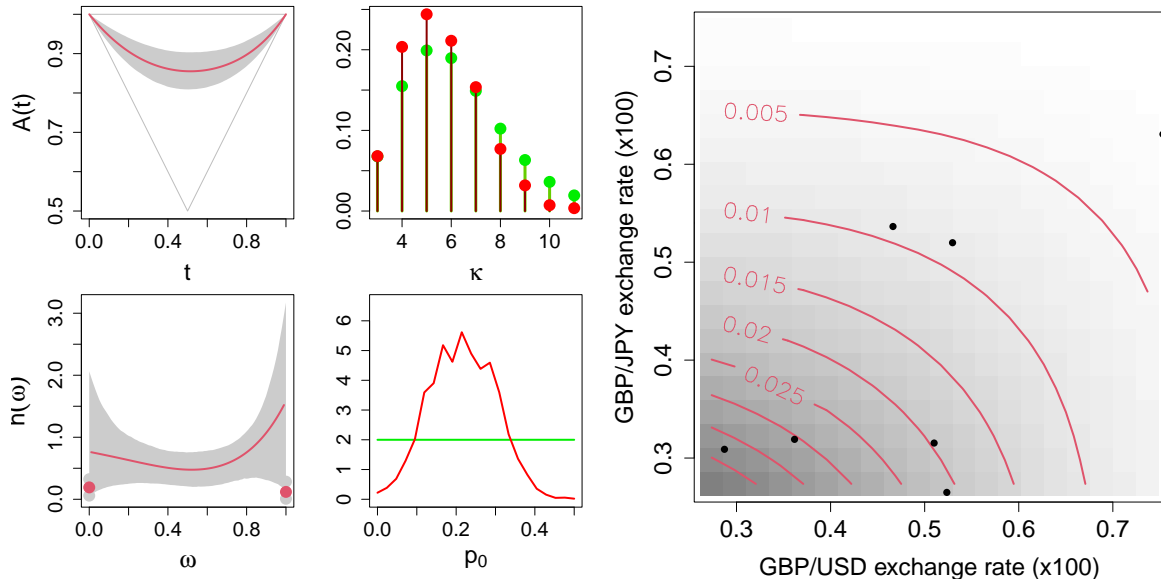


Figure 7: Outputs of the `summary.ExtDep` (left) and `returns` (right) functions.

```
R> # Compute conditional probabilities of exceedance
R> par1 <- gbp_mar_sum$mar1.mean
R> par2 <- gbp_mar_sum$mar2.mean
R> jointP / pGEV(qs[1], loc = par1[1], scale = par1[2], shape = par1[3],
+   lower.tail = FALSE)
R> jointP / pGEV(qs[2], loc = par2[1], scale = par2[2], shape = par2[3],
+   lower.tail = FALSE)
```

The model’s goodness of fit provided can be verified graphically at the marginal and dependence level. First, marginal density histograms can be plotted together with GEV density curves evaluated at the posterior samples of the marginal parameters. Second, after standardizing the data to unit Fréchet marginals, a density histogram of the angular components with the largest radial component can be used as an empirical estimate of the angular density, to which the model estimates (posterior mean, 0.05 and 0.95 quantile) can be added. Figure 14, given in Appendix A.1, provides graphical evidence that the proposed model does provide a good fit to the data.

2.4.3 Simulation of extreme values

Theory. In environmental statistics, there is active research in the field of “stochastic weather generators” which aim at simulating realisations from atmospheric variables according to some stochastic representation. A challenging task is the simulation of multivariate extremes, since their extremal dependence (Pickands dependence function or angular measure) is an infinite dimensional parameter of the multivariate GEV distribution and is therefore challenging to estimate. The ability to simulate multivariate extremes permits to approximate the tail probabilities in (2.9) and (2.10), even though few extremes are available in the dataset.

For simplicity, most of the simulation methods assume that the extremal dependence com-

plies with some parameter model, while few attempt to consider a nonparametric approach. [Marcon et al. \(2017\)](#) proposed a flexible procedure for sampling from bivariate extremes with a semiparametric dependence structure, which is summarized as follows. For every small probabilities p_1 and p_2 such that $L(p_1, p_2) \in [0, 1]$ and $R(p_1, p_2) \in [0, 1]$, the tail probabilities in (2.9) and (2.10) can be approximated by $\mathbb{P}\{X_1 > Q_1(p_1) \text{ or } X_2 > Q_2(p_2)\} \approx L(p_1, p_2)$ and $\mathbb{P}\{X_1 > Q_1(p_1) \text{ and } X_2 > Q_2(p_2)\} \approx R(p_1, p_2)$. For such probabilities p_1 and p_2 , we have that $L(p_1, p_2) = 2\mathbb{E}\{\max(p_1 W, p_2(1 - W))\}$ and $R(p_1, p_2) = 2\mathbb{E}\{\min(p_1 W, p_2(1 - W))\}$, where W is a random variable on $[0, 1]$ distributed according to the angular distribution H . Let $(Z_1, Z_2) = R(W, 1 - W)$, where R is a unit-Pareto random variable, then $\mathbb{P}(Z_1 > z_1 \text{ or } Z_2 > z_2) = 2\mathbb{E}\{\max(W/z_1, (1 - W)/z_2)\}$ and $\mathbb{P}(Z_1 > z_1 \text{ and } Z_2 > z_2) = 2\mathbb{E}\{\min(W/z_1, (1 - W)/z_2)\}$. [Marcon et al. \(2017\)](#) proposed to model H through Bernstein polynomials and demonstrated that H can be written as a finite mixture of beta distributions with weights defined by a suitable transformation of the polynomial coefficients. This leads to a simple algorithm to sample from H (see Algorithm 1) and an algorithm for sampling observations from the tail of a bivariate distribution (see Algorithm 3), which can be used to approximate the corresponding tail probabilities. Briefly, from (2.6) we have that, for sufficiently large n and sufficiently small $p_j = p_{n,j}$ (with $p_j \rightarrow 0$ as $n \rightarrow \infty$), $Q_j(p_j) \approx \mu_j + \gamma_j\{(np_j)^{-\gamma_j} - 1\}/\sigma_j$, for $j = 1, 2$. Let u_1 and u_2 be two high thresholds such that, for every p_1 and p_2 for **which the above approximations hold and for which, in addition** $L(p_1, p_2) \in [0, 1]$ and $R(p_1, p_2) \in [0, 1]$, we have $Q_j(p_j) > u_j$, $j = 1, 2$. For the failure regions $\mathcal{A}_{\mathbf{u}} = \{(v_1, v_2) : v_1 > u_1 \text{ or } v_2 > u_2\}$ and $\mathcal{B}_{\mathbf{u}} = \{(v_1, v_2) : v_1 > u_1 \text{ and } v_2 > u_2\}$ we have $\mathbb{P}(\mathbf{X} \in \mathcal{A}_{\mathbf{u}}) \approx \mathbb{P}(Z_1 > u_1^* \text{ or } Z_2 > u_2^*)$ and $\mathbb{P}(\mathbf{X} \in \mathcal{B}_{\mathbf{u}}) \approx \mathbb{P}(Z_1 > u_1^* \text{ and } Z_2 > u_2^*)$, where $u_j^* = (1 + \gamma_j(u_j - \mu_j)/\sigma_j)^{1/\gamma_j}$, for $j = 1, 2$. By simulating a large sample of angular components (w_1, \dots, w_N) from H and radial components (r_1, \dots, r_N) from a unit-Pareto distribution, we compute $\mathbf{z}_i = 2(r_i w_i, r_i(1 - w_i))$, for $i = 1, \dots, N$ and estimate of the probability of falling in $\mathcal{A}_{\mathbf{u}}$ and $\mathcal{B}_{\mathbf{u}}$ by

$$\hat{p}_{\mathcal{A}_{\mathbf{u}}} = \frac{1}{N} \sum_{i=1}^N \mathbb{1}(z_{i,1} > u_1^* \text{ or } z_{i,2} > u_2^*), \quad \hat{p}_{\mathcal{B}_{\mathbf{u}}} = \frac{1}{N} \sum_{i=1}^N \mathbb{1}(z_{i,1} > u_1^* \text{ and } z_{i,2} > u_2^*). \quad (2.18)$$

Application. The methodology is illustrated on the **Parcay-Meslay** dataset which consists of daily maxima of hourly wind speed (**WS**) and wind gust (**WG**) in meters per second (m/s) and differential of daily range of the hourly air pressure (**DP**) at sea level in millibars. Measurements are taken in the city of Parcay-Meslay, located in the northwest of France, from July 2004 to July 2013. For this analysis, we focus on positive values of DP implying an increase in the daily air pressure level. High air pressure levels are associated with a high content of water vapor in the air which often occurs in stormy weather and leads to strong winds. For brevity, we focus on the relationship between WS and DP. The GEV parameters are first estimated using the point process approach (e.g., [Coles, 2001](#), Ch. 7), implemented in the `fpot` routine of the `evd` package ([Stephenson, 2002](#)), where the observed 0.9 marginal quantiles are set as suitable thresholds. The margins are transformed to unit-Fréchet scale using the `trans2UFrechet` routine with argument `type = "empirical"`, meaning the transformation $y_{i,j} = 1/\{1 - F_{n,j}(x_{i,j})\}$, $i = 1, \dots, n$, $j = 1, 2$ is applied, where $F_{n,j}$ denotes the empirical distribution of the j -th component. These initial steps can be implemented by running the following lines

```
R> data(WindSpeedGust, package = "ExtremalDep")
R> years<- format(ParcayMeslay$time, format = "%Y")
R> attach(ParcayMeslay[which(years %in% c(2004:2013)), ])
```

```

R> WS_th <- quantile(WS, 0.9)
R> DP_th <- quantile(DP, 0.9)

R> # Peaks over threshold modelling using the point processes representation
R> pars.WS <- evd::fpot(WS, WS_th, model = "pp")$estimate
R> pars.DP <- evd::fpot(DP, DP_th, model = "pp")$estimate

R> # Transform margins to unit Frechet scale
R> data_uf <- trans2UFrechet(cbind(WS, DP), type = "Empirical")

```

The angular distribution is then estimated using the approximate likelihood approach (see [Beranger and Padoan, 2015](#)). Observations with radius component greater than their 90% quantile are selected and the routine `fExtDep.np` is then called with arguments `method = "Frequentist"` and `type = "maxima"` to maximize the likelihood of a polynomial angular distribution in Bernstein form through the non-linear optimization routine `nloptr` from the `nloptr` package ([Ypma and Johnson, 2025](#)) subject to constraints established in [Marcon et al. \(2017\)](#). Empirical studies ([Marcon et al., 2017](#)) suggest that the polynomial degree $\kappa = 10$ is enough to represent flexible dependence structures. When `type = "rawdata"`, data are extracted using a threshold on the radial component set by the argument `u` and the likelihood for a sample maxima written as function of the Pickands dependence function in Bernstein form is optimized. In addition, when `mar.fit = TRUE` then marginal empirical transformation to unit Fréchet of the data is applied. The `plot` routine with `type = "summary"` displays graphical summaries (Pickands dependence function and angular density) of the estimated dependence structure (see left and middle panels of [Figure 8](#)). A moderate level of dependence is observed as well as point masses at the vertices. Given that the middle panel of [Figure 8](#) combines empirical information with a summary of the fitted model, it can be leveraged to highlight the good fit of the model. The routine `rExtDep` with arguments `model = "semi.bvevd"` and `angular = TRUE` generates pseudo-angles according to [Algorithm 1](#) of [Marcon et al. \(2017\)](#). The right panel of [Figure 8](#) presents a histogram of 1,000 randomly generated angles and point-masses (black triangles), the red line and dots represent the estimated dependence structure. The corresponding code is

```

R> rdata <- rowSums(data_uf) # Computes radial component
R> r0 <- quantile(rdata, probs = 0.9)
R> extdata <- data_uf[rdata >= r0, ] # Select data with large radial component

R> # Estimate angular density
R> SP_mle <- fExtDep.np(x = extdata, method = "Frequentist", k0 = 10,
+   type = "maxima")
R> plot(x = SP_mle, type = "summary")

R> # Generate angular data
R> SP_wsim <- rExtDep(n = 1000, model = "semi.bvevd",
+   param = SP_mle$Ahat$beta, angular = TRUE)

```

The routine `rExtDep` performs random generation of bivariate maxima (`model = "semi.bvevd"`) and bivariate exceedances (`model = "semi.bvexceed"`), according to [Algorithm 2](#) and [3](#) of




Figure 8: Estimated Pickands dependence function (left) and angular density (middle and right). The middle panel displays the histogram of the data while the right panel displays a histogram of the simulated pseudo-angles as well as the point masses (black triangles).

Marcon et al. (2017). When `angular = TRUE`, only angular components are generated, no matter the `model` argument, meaning that setting `model = "semi.bvexceed"` in the above call of `rExtDep` would produce the same result. The argument `mar` allows for transformations to GEV margins. When `model = "semi.bvexceed"`, one can choose to simulate bivariate observations from the failure regions \mathcal{A}_u (`exceed.type = "or"`) or \mathcal{B}_u (`exceed.type = "and"`), where u is a suitable threshold specified by `threshold`. The command lines below generate 200 observations from both failure regions with threshold $u = (10, 20)$, above the marginal 90% quantiles (see Figure 9 for an illustration).

```
R> set.seed(10)
```

```
R> # Generate exceedances in at least one component
R> SP_exceed_or <- rExtDep(n = 200, model = "semi.bvexceed",
+   param = SP_mle$Ahat$beta, mar = rbind(pars.WS, pars.DP),
+   threshold = c(10, 20), exceed.type = "or")
```

```
R> # Generate exceedances in both components
R> SP_exceed_and <- rExtDep(n = 200, model = "semi.bvexceed",
+   param = SP_mle$Ahat$beta, mar = rbind(pars.WS, pars.DP),
+   threshold = c(10, 20), exceed.type = "and")
```

The routine `pFailure` computes empirical estimates of probabilities of belonging to the failure sets \mathcal{A}_u (`type = "or"`) and \mathcal{B}_u (`type = "and"`) by applying formula (2.18) on data generated by `rExtDep`. Both probabilities are computed when `type = "both"`. The argument `plot` offers the possibility to display contour plots when sequences of thresholds are provided through the arguments `u1` and `u2`. The code provided below generates $N = 50,000$ samples to estimate the probabilities of belonging to each failure sets for thresholds ranging from 19 to 28 for the wind speed and from 40 to 60 for the differential of pressure with outputs given in Figure 10.

```
R> # Estimate probabilities to belong to failure regions "or"/"and"
R> pF <- pFailure(n = 50000, beta = SP_mle$Ahat$beta,
```




Figure 9: Observed (black dots) differential of pressure (mbar) and daily-maximum wind speed (m/s) in Parçay-Meslay, France between 2004 and 2013. Simulation of 200 observations from the failure regions \mathcal{A}_u (red dots) and \mathcal{B}_u (blue dots) with $\mathbf{u} = (10, 20)$ (dashed lines).

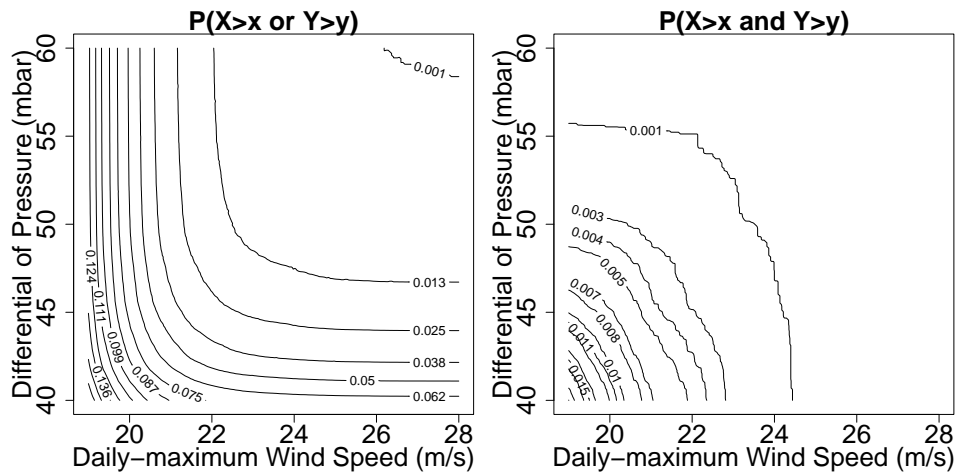


Figure 10: Estimated probabilities to belong to the failure regions \mathcal{A}_u (left) and \mathcal{B}_u (right), for $\mathbf{u} = (u_1, u_2)$, with $18 \leq u_1 \leq 28, 40 \leq u_2 \leq 60$.

```

+           u1 = seq(from = 19, to = 28, length = 200),
+           u2 = seq(from = 40, to = 60, length = 200),
+           mar1 = pars.WS, mar2 = pars.DP, type = "both",
+           plot = TRUE, xlab = "Daily-maximum Wind Speed (m/s)",
+           ylab = "Differential of Pressure (mbar)", nlevels = 15)

```

2.4.4 Estimation of extreme quantile regions

Theory. As highlighted in Section 2.1, an important problem is to define a quantile region Q as in (2.11), given a very small probability p to fall in it. Given a sample of size N , an extreme quantile region can be defined by the level set Q_N such that $\mathbb{P}(Q_N) = p$, where $p = p_N$ satisfies $p \rightarrow 0$ and $Np \rightarrow c > 0$ as $N \rightarrow \infty$. For practical purposes, Q_N can be approximated by the region \tilde{Q}_N given in (2.13) which requires estimating the marginal parameters $\theta_j = (\mu_j, \sigma_j, \gamma_j)^\top$, $j = 1, 2$, the basic set S and the density of the angular measure, allowing, in turn, to obtain an estimate of $\xi(S)$. Beranger et al. (2021) discussed a Bayesian approach for the inference of \tilde{Q}_N based on a polynomial angular measure in Bernstein form and the censored-likelihood approach (e.g., Beirlant et al., 2004, Ch. 8). Since such a methodology is similar to the one described in Section 2.4.2, we refer to Beranger et al. (2021) for a full description.

The methodology is illustrated on the dataset `Milan.winter`, which consists of air pollution levels recorded in Milan, Italy, over the winter period October 31st–February 28/29th, between December 31st 2001 and December 30th 2017. We focus here on the NO_2 and SO_2 pollutants and use the daily maximum temperature (`MaxTemp`) as covariate. The data is prepared before estimating extreme quantile regions using `fExtDep.np`.

```
R> data(MilanPollution, package = "ExtremalDep")
R> data <- Milan.winter[, c("NO2", "SO2", "MaxTemp")]
R> data <- data[complete.cases(data), ]
```

A quadratic relationship between pollutants and maximum temperature is considered, and the j th marginal mean is specified as $\mu_j = \beta_{0,j} + \beta_{1,j}z + \beta_{2,j}t^2$, with $j = 1, 2$, where t is the temperature level. The covariate matrix is defined as

```
R> covar <- cbind(rep(1, nrow(data)), data[, 3], data[, 3]^2 )
```

which will be provided to `fExtDep.np` through the arguments `cov1` and `cov2`. Since a polynomial angular measure in Bernstein form is considered, we specify a prior distribution for the polynomial degree κ as a negative binomial (default) for $\kappa - 3$ with mean 6 and variance 8 and priors for the point masses p_0 and p_1 respectively as uniform distributions (default) on $[0, 0.2]$ and $[a, b]$, where a and b are defined as in the previous two sections. These are the default prior distributions in the `fExtDep.np` routine, therefore the hyper-parameters are specified by

```
R> hyperparam <- list(mu.nbinom = 6, var.nbinom = 8, a.unif = 0, b.unif = 0.2)
```

The prior for p_0 and p_1 are set in such a way to represent the belief that those point masses are likely to be small for these data. The mean and variance of the negative binomial prior for the polynomial degree are set to allow for higher degree polynomial modelling of the angular density if required.

As mentioned above, the estimation is performed through the `fExtDep.np` function where starting values for the marginal parameters $\theta_j^{(0)} = (\beta_{0,j}^{(0)}, \beta_{1,j}^{(0)}, \beta_{2,j}^{(0)}, \sigma_j^{(0)}, \gamma_j^{(0)})^\top$, $j = 1, 2$, the scaling parameter of the sampler on each margin and the polynomial degree are respectively set by the arguments `par10`, `par20`, `sig10` `sig20` and `k0`. The argument `method = "Bayesian"` indicates that a Bernstein polynomial is used to represent the extremal dependence and that the inference part follows the MCMC scheme detailed in Beranger et al.

(2021, Algorithm 1). The argument `u = TRUE` specifies that a censored likelihood approach is applied on raw data with threshold `u` set to the marginal 90% quantile by default. Recall that by default the marginal distributions are fitted jointly with the dependence (`mar.fit = TRUE`) but `mar.prelim = FALSE` indicates that no initial marginal fit is required. The code chunk below provides a full specification of the `fExtDep.np` arguments.

```
R> pollut <- fExtDep.np(x = data[, -3], method = "Bayesian", u = TRUE,
+   cov1 = covar, cov2 = covar, mar.prelim = FALSE,
+   par10 = c(100, 0, 0, 35, 1), par20 = c(20, 0, 0, 20, 1), sig10 = 0.1,
+   sig20 = 0.1, k0 = 5, hyperparam = hyperparam, nsim = 15e+3)
```

Numerical and graphical summaries of the output of the estimation procedure are obtained from the `summary` and `plot` functions by specifying the arguments `object`, `x`, `summary.mcmc` and `type = "Qsets"`. Since covariates are used to model the marginal parameters, we need to provide the covariate levels at which the extreme quantile regions will be computed. In this example, extreme quantile regions are computed at the minimum, median and maximum of the observed daily maximum temperatures. The covariate matrices `QatCov1` and `QatCov2` should not include an intercept term and contain a maximum of three levels as given through

```
R> pollut_sum <- summary_ExtDep(mcmc = pollut, burn = 5e+3)
R> Temp.seq <- c(min(data[, 3]), median(data[, 3]), max(data[, 3]))
R> QatTemp <- cbind(Temp.seq, Temp.seq^2)
```

Representations of quantile regions associated with small probabilities $p = 1/600, 1/1200$ and $1/2400$ are displayed using the `plot` function where we request to include graphical summaries of the extremal dependence as well as displaying the data (`dep=TRUE` and `data`). The output contains a list of dependence quantities in `est.out` and a list of quantile sets in `q.out`. In particular, `est.out` includes the following elements:

- `ghat`: a 3×100 matrix with an estimate of the inverse of the angular basic density q_* evaluated at 100 grid points in $[0, 1]$. Rows gives the posterior 0.05 quantile, mean and 0.95 quantile.
- `Shat_post` and `Shat`: lists where each element is a 100×2 matrix providing an estimate of the basic set \mathcal{S} (obtained through `ghat`). `Shat_post` considers every posterior samples whereas `Shat` takes the pointwise 0.05 quantile, mean and 0.95 quantile.
- `nuShat_post` and `nuShat`: two vectors providing the posterior of the basic set measure $\xi(\mathcal{S})$ and its 0.05 quantile, mean and 0.95 quantile.

The `q.out` list includes:

- `Qset_PA_CovNum_B`: a list of three 100×2 matrices. Each of them is an estimate of the bivariate extreme quantile \tilde{Q}_n . Such regions are computed for the specified probabilities (`P`) and covariate (`QatCov1` and `QatCov2`), e.g., `pl$Qset_P1_CovNum_1` provides an estimate of the (NO_2, SO_2) region corresponding to the probability $1/600$, when the minimum temperature is observed. Each matrix corresponds to the posterior 0.05 quantile, mean and 0.95 quantile, obtained from `Qset_PA_CovNum_B_post`.
- `Qset_PA_CovNum_B_post`: a $2 \times (\text{nsim} - \text{burn})$ matrix providing an estimate of the extreme quantile region for each posterior sample.

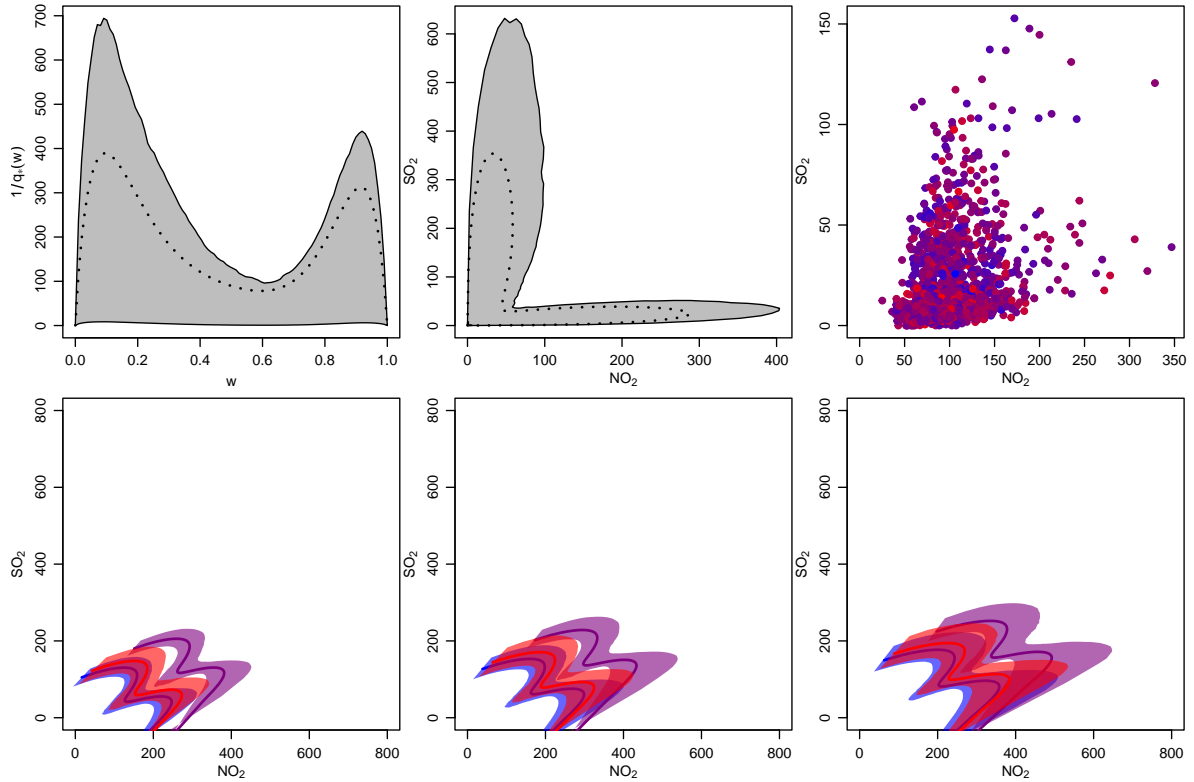


Figure 11: Top row: posterior mean estimate (dotted line) and 90% credible interval (grey) for the inverse of the angular basic density (left), the basic set \mathcal{S} (middle), and observed data (right) with temperature dependent data colouring (from cold = blue to warm = red). Bottom row: posterior mean estimate (dotted line) and 90% credible interval for the extreme bivariate quantiles associated with probabilities $p = 1/600, 1/1200$ and $1/2400$ (left to right) for three maximum daily temperature levels: minimum temperature (blue), median (purple) and maximum (red) temperatures.

An implementation is given below, where the argument `dep=TRUE` produces the top row of Figure 11.

```
R> pl <- plot(x = pollut, type = "Qsets", summary.mcmc = pollut_sum,
+   QatCov1 = QatTemp, P = 1/c(600, 1200, 2400), dep = TRUE,
+   data = data[, -3], xlim = c(0, 800), ylim = c(0, 800),
+   xlab = expression(NO[2]), ylab = expression(SO[2]))
```

The bottom row of Figure 11 shows the extreme quantile regions corresponding to probabilities $1/600$ (left), $1/1,200$ (middle) and $1/2,400$ (right) where the colours indicates different levels of the covariate. Colours of the data (top right), credibility regions and mean from the posterior can be respectively specified by `col.data`, `col.Qfade` and `col.Qfull` (blue to red colour palette by default).

3 Spatial extremes

Max-stable processes (e.g., de Haan and Ferreira, 2006, Ch. 9) are a convenient tool to statistically model extremes and their dependence over a spatial domain \mathcal{D} . Let $\{X_i(s), s \in \mathcal{D}\}$, $i = 1, 2, \dots$, be independent copies of a stochastic process $X(s)$, $s \in \mathcal{D}$ with identical finite-dimensional distribution. If there are functions $a_n(s) > 0$ and $b_n(s)$, for each n , such that the finite-dimensional distributions of a limit process $Y(s)$, given by

$$\{Y(s), s \in \mathcal{D}\} = \lim_{n \rightarrow \infty} \left\{ \max_{1 \leq i \leq n} \left(\frac{X_i(s) - b_n(s)}{a_n(s)} \right), s \in \mathcal{D} \right\},$$

are not degenerate, then they must be multivariate EV as in (2.2), and with the maximum taken pointwise for all $s \in \mathcal{D}$. Refer to Davison et al. (2012) for a detailed review of max-stable processes and the statistical analysis of spatial extremes. The **ExtremalDep** package considers some of the most widely used max-stable models including geometric Gaussian (Davison et al., 2012), Brown–Resnick (Brown and Resnick, 1977), extremal- t (Opitz, 2013) and extremal skew- t processes (Beranger et al., 2017). The routines implemented rely on the results of Dombry et al. (2016) and Beranger et al. (2021) for the exact simulation of the max-stable models and on Stephenson and Tawn (2005) and Beranger et al. (2021) for the inferential procedures.

Software implementation. The routine `rExtDepSpat` extends the `rmaxstab` function from the **SpatialExtremes** package (Ribatet, 2022) by including exact and direct simulation (`method="direct"` or `"exact"`) from the extremal skew- t process, by using the prefix “s” when defining the type of correlation function in the argument `cov.mod` (options are `"whitmat"`, `"cauchy"`, `"powexp"` and `"bessel"`). For the extremal skew- t model, the skewness parameter is represented as

$$\boldsymbol{\alpha} = \alpha_0 + \alpha_1 \mathbf{cov}_1 + \alpha_2 \mathbf{cov}_2 \in \mathbb{R}^d,$$

where \mathbf{cov}_1 and \mathbf{cov}_2 are d -dimensional covariate vectors with d the number of spatial locations $s_j, j = 1, \dots, d$. The argument `alpha` is used to specify $(\alpha_0, \alpha_1, \alpha_2)$ while the covariates are given by `acov1` and `acov2`. Using the command lines provided below, 50 (Ny) replicates of the extremal- t , with $\nu = 1$ degrees of freedom (DoF), are generated at 10 (Ns) spatial locations (`sites`) in the region $[-5, 5]^2$. The correlation function (`cov.mod`) is the power exponential class $\rho(h) = \exp\{- (\|h\|/r)^\eta\}$ with smoothness (`smooth`) parameter $\eta = 1.5$ and range (`range`) parameter $r = 3$.

```
R> set.seed(14342)
R> Ns <- 10; Ny <- 50
R> sites <- matrix(runif(Ns*2) * 10 - 5, nrow = Ns, ncol = 2)
R> for(i in 1:2) sites[, i] <- sites[, i] - mean(sites[, i])

R> # Generate samples from the extremal-t model with power exp corr function
R> z <- rExtDepSpat(Ny, sites, cov.mod = "tpowexp", DoF = 1, range = 3,
+   nugget = 0, smooth = 1.5, control = list(method = "exact"))
```

The routine `rExtDepSpat` returns a list consisting of simulated values at specified locations (`$vals`) and the hitting scenario (`$hits`), both being $Ny \times Ns$ matrices. For a given row of `$hits`, elements with the same value indicate block maxima that occurred at the same

time (for illustrative purposes, one may think that the maxima were obtained from the same “storm”).

The `fExtDepSpat` procedure takes advantage of the availability of the time of occurrence of block maxima to fit the extremal- t (`model = "ET"`) and skew- t (`model = "EST"`) max-stable models using the Stephenson—Tawn likelihood (Stephenson and Tawn, 2005). The correlation function is currently restricted to the power exponential, meaning that the parameters vectors to be estimated are respectively $\boldsymbol{\theta} = (\nu, r, \eta)^\top$ and $\boldsymbol{\theta} = (\nu, r, \eta, \alpha_0, \alpha_1, \alpha_2)^\top$. The corresponding parameters can be fixed through arguments `range`, `smooth`, `DoF` and `alpha` of `fExtDepSpat`, else are estimated. Note that, a vector of length 3 must be provided for `alpha` and therefore a NA value leaves the corresponding parameter free, e.g., `alpha = c(0, NA, NA)` would fit the extremal skew- t model with skewness $\boldsymbol{\alpha} = \alpha_1 \mathbf{cov}_1 + \alpha_2 \mathbf{cov}_2$. Initial values are provided in `par0` in vector form as `c(DoF, range, smooth, alpha0, alpha1, alpha2)`, where fixed or unnecessary parameters are simply omitted. Computations can also be sped-up by considering parallelization (`parallel = TRUE`) and specifying the number of cores (`ncores`) to be used. Arguments `args1` and `args2` are related to specifications of the Monte Carlo simulation scheme to compute the multivariate t cumulative distribution function (cdf). These should take the form of lists including the minimum and maximum number of simulations used (`Nmin` and `Nmax`), the absolute error (`eps`) and whether the error should be controlled on the log-scale (`logeps`). `args1` refers to the $d - 1$ dimensional cdfs terms required to compute the exponent function while `args2` focuses on the $d - m$ ($2 \leq m \leq d - 1$) dimensional cdfs involved in the evaluation of its partial derivatives. When computing the log-likelihood function, the latter terms need to be evaluated on the log-scale, requiring fewer Monte Carlo simulations. In the following command lines, the strategy is to set a higher number of simulations for the partial derivative terms since experiments have shown that they are more important than the terms in the exponent function (see Beranger et al., 2021). The `control` argument offers additional control parameters for the optimization algorithm see `?optim` for more details.

```
R> # Set min and max number of sim to compute CDF terms
R> args1 <- list(Nmax = 50L, Nmin = 5L, eps = 0.001, logeps = FALSE)
R> args2 <- list(Nmax = 500L, Nmin = 50L, eps = 0.001, logeps = TRUE)

R> # Fit true (extremal-t) model using full likelihood
R> fit1 <- fExtDepSpat(x = z$vals, model = "ET", sites = sites, hit = z$hits,
+   par0 = c(3, 1, 1), parallel = TRUE, ncores = 2, args1 = args1,
+   args2 = args2, control = list(trace = 0))
R> est(fit1)

      DoF range smooth
0.991 2.940 1.551

R> method(fit1)

[1] "Full likelihood"
```

The routine `fExtDepSpat` returns an object of class `ExtDep.Spat` from which the estimated parameters, the estimation method, the standard errors computed from the sandwich information matrix, the value of the maximised log-likelihood and the Takeuchi information

criteria can respectively be extracted using the functions `est`, `method`, `StdErr`, `logLik` and `tic`. Here the estimated parameter vector is $\hat{\theta} = (0.991, 2.940, 1.551)^\top$, while the true parameters are $\theta = (1, 3, 1.5)^\top$ using the default full likelihood as indicated by `est1$jw` taking value 20. As proposed in [Beranger et al. \(2021\)](#), composite likelihood estimation ([Padoan et al., 2010](#)) is incorporated when the argument `jw` is specified and is less than the number of locations. Since the number of tuples (pairs if `jw = 2`, triples if `jw = 3`) is increasing with `jw` one can specify a threshold u through the argument `thresh` such that, for a tuple q , the corresponding composite likelihood contribution is weighted according to

$$w_q = \begin{cases} 1 & \text{if } \max_{i,k \in q; i \neq k} \|s_i - s_k\| < u, \\ 0 & \text{otherwise.} \end{cases}$$

In other words, only tuples with maximum pairwise distance less than u are included in the likelihood.

```
R> # Fit true (extremal-t) model using composite likelihood
R> fit2 <- fExtDepSpat(x = z$vals, model = "ET", sites = sites, hit = z$hits,
+   jw = 3, thresh = quantile(dist(sites), 0.25), par0 = c(3, 1, 1),
+   parallel = TRUE, ncores = 2, args1 = args1, args2 = args2,
+   control = list(trace = 0))
R> est(fit2)
```

```
DoF range smooth
1.051 2.975 1.695
```

```
R> method(fit2)
```

```
[1] "3-wise composite likelihood"
```

The above code script indicates that fitting the extremal- t model using the triplewise composite-likelihood approach with threshold u set at the first quartile of pairwise distances yields $\hat{\theta} = (1.051, 2.975, 1.695)^\top$, which is again close to the true one.

Application. We now analyze temperature data around Melbourne, Australia, collected at 90 stations on a 0.15 degree (approximately 13 kilometer) grid in a 9 by 10 formation, over the extended summer period from August to April between 1961 and 2010. Running the command `data(heat, package = "ExtremalDep")` loads several datasets in **R**. The objects `locgrid`, `scalegrid` and `shapegrid` are matrices of the marginal GEV parameters estimated over the grid using unconstrained location and scale while the shape parameter is defined as a linear function of eastings and northings in 100 kilometer units. Below, the **terra** package ([Hijmans, 2026](#)) is leveraged to provide graphical illustrations of these objects.

```
R> library(terra)
```

```
R> data(heat, package = "ExtremalDep")
```

```
R> # Obtain boundary files for the study region
```

```
R> path <- "https://www.abs.gov.au/statistics/standards/australian-statistical-
+   geography-standard-asgs-edition-3/jul2021-jun2026/access-and-downloads/
```




Figure 12: Estimated marginal location (left), scale (middle) and shape (right) parameters.

```
+   digital-boundary-files"
R> file.name <- "SA4_2021_AUST_SHP_GDA2020.zip"
R> download.file(file.path(path, file.name), destfile = file.name)
R> unzip(file.name)
R> geogmel <- vect("SA4_2021_AUST_GDA2020.shp")
R> geogmel <- geogmel[geogmel$GCC_NAME21 == "Greater Melbourne"]
R> unlink(list.files(pattern = "SA4_2021_AUST"))

R> # Create grid and display marginal parameters
R> lat <- seq(from = -38.45, by = 0.15, length = 9)
R> lon <- seq(from = 144.4, by = 0.15, length = 10)
R> xx <- expand.grid(lon = lon, lat = lat)
R> cols <- tim.colors(40)[20:40]
R> for(m in c("locgrid", "scalegrid", "shapegrid")){
+   mat <- mget(m)
+   xx$data <- as.vector(mat[[1]])
+   rd <- rast(xx, crs = crs(geogmel))
+   plot(rd, col = map.pal("gyr", 40), col = cols, plg = list(cex = 2),
+       pax = list(cex.axis = 2))
+   lines(geogmel, lwd = 3)
+ }
```

Note that the code first downloads shape files and extracts the Greater Melbourne region. Also provided in `heat` is the object `heatdata` which is a list with elements:

- `vals`: the 50 yearly maxima (rows) at each of the 90 locations (columns)
- `ufvals`: `vals` marginally transformed to unit-Fréchet scale
- `hits`: the hitting scenarios for each year (row) and location (row). Locations with the same integer value correspond to maxima obtained on the same day (± 3 days).
- `sitesLLO`, `sitesENO`: original location coordinates in latitude-longitude (LL) and easting-northing (EN, in kilometer)
- `sitesLL`, `sitesEN`: same as `sitesLLO`, `sitesENO` but centered.

Below, we fit the extremal skew- t model with $\nu = 5$ to the data. Since the observations were marginalized with GEV shape parameter as function of eastings and northings in 100 kilometre units, the site locations are transformed to be on the same scale. [Beranger et al. \(2021\)](#) used the TIC for model selection between the extremal- t and extremal skew- t with $\nu = 1, 3$ and 5 degrees but, for simplicity, we only consider here the selected model and a subset of the locations (`ind_sites`).

```
R> # Change to 100 kilometer units
R> sites <- heatdata$sitesEN/100

R> # Restrict to sub-region for faster computations
R> ind_sites <- (abs(sites[, 1]) < 0.4 & abs(sites[, 2]) < 0.4)
R> sites <- sites[ind_sites, ]

R> x <- heatdata$ufvals[, ind_sites]
R> hits <- heatdata$hits[, ind_sites]

R> args1 <- list(Nmax = 20L, Nmin = 2L, eps = 0.001, logeps = FALSE)
R> args2 <- list(Nmax = 200L, Nmin = 20L, eps = 0.001, logeps = TRUE)

R> # Fit extremal skew-t model
R> est5 <- fExtDepSpat(x = x, model = "EST", sites = sites, hit = hits,
+   par0 = c(1, 1, 0.5, 0, 0), DoF = 5, acov1 = sites[, 1],
+   acov2 = sites[, 2], parallel = TRUE, ncores = 2, args1 = args1,
+   args2 = args2, control = list(trace = 2))
R> est(est5)

range smooth alpha.0 alpha.1 alpha.2
9.275  1.262  0.496  0.001 -0.009
```

The estimated range and smoothness of the skew- t are respectively $\hat{r} = 9.275$, $\hat{\eta} = 1.262$, while the skewness is estimated as $\hat{\alpha} = 0.496 + 0.001\text{easting} - 0.009\text{northing}$. The largest distance between pairs of locations (in 100 kilometer units) is 1.785, and therefore the smallest correlation is $\exp\left[-(1.785/\hat{r})^{\hat{\eta}}\right] \approx 0.88$, indicating a strong degree of spatial dependence. Finally, the `rExtDepSpat` function is used to simulate from the extremal skew- t conditionally on the hitting scenario and Figure 13 displays an example conditioning on at most two heatwave events causing the annual maxima. These steps are implemented in the following code script.

```
R> set.seed(123)
R> while(TRUE){
+   z.new <- rExtDepSpat(1, sites, cov.mod = "spowexp", DoF = 5, nugget = 0,
+     range = est(est5)[1], smooth = est(est5)[2],
+     alpha = est(est5)[3:5], acov1 = sites[,1], acov2 = sites[,2],
+     control = list(method = "exact"))
+   if(length(unique(as.numeric(z.new$hits))) <= 2) break
+ }
```




Figure 13: Simulation (in $^{\circ}\text{C}$ unit) from the fitted extremal skew- t model with $\nu = 5$, conditioning on at most two heatwave events causing all maxima.

```
R> # Transform back to original margins
R> parmat <- cbind(as.vector(locgrid), as.vector(scalegrid),
+   as.vector(shapegrid))
R> z.new.GEV <- rep(NA, 90)
R> z.new.GEV[ind_sites] <- trans2GEV(data = z.new$vals,
+   pars = parmat[ind_sites, ])

R> xx$data <- z.new.GEV
R> rd <- rast(xx, crs = crs(geogmel))
R> plot(rd, col = map.pal("gyr", 40), col = cols, plg = list(cex = 2),
+   pax = list(cex.axis = 2))
R> lines(geogmel, lwd = 3)
```

4 Conclusions

This paper has presented a comprehensive overview of methodological and practical modeling of multivariate and spatial problems. Building on the componentwise maxima framework, we have shown how extremal dependence, characterized through angular measures and Pickands dependence functions, can be effectively estimated, interpreted, and exploited for risk assessment. A key message throughout the paper is that modeling extremal dependence is not only relevant for block maxima, but also provides a powerful basis for approximating probabilities of multiple threshold exceedances, joint tail events, and extreme quantile regions, thereby supporting a broad range of applied risk analyses.

The **ExtremalDep** package provides a flexible and extensible software that integrates parametric, semiparametric, and nonparametric approaches to modeling extremal dependence. Compared with existing software, it places particular emphasis on nonparametric and semiparametric representations of extremal dependence, to remain close to its nature, while still accommodating parametric models when interpretability or parsimony is required. In addition, it offers a Bayesian inference toolkit that allows for a principled and accessible uncer-

tainty quantification. Through a series of real-world applications, ranging from air pollution and precipitation to financial returns, wind extremes, and heatwaves, the paper has demonstrated how the **R** package can be used to estimate extremal dependence, compute joint and conditional tail probabilities, derive multivariate return levels, simulate extreme events, and construct extreme quantile regions in both multivariate and spatial settings.

We plan the following future developments. An important extension concerns the integration of a more a general peaks-over-threshold inferential framework for multivariate data to better complement the componentwise maxima. Scaling nonparametric and Bayesian approaches to even higher dimensions, for example through sparsity assumptions, graphical representations of extremal dependence, or dimension-reduction techniques. In the spatial domain, incorporating spatio-temporal dynamics and more flexible covariate effects into extremal dependence models remains a challenging avenue for research. Future versions of **ExtremalDep** will aim to improve computational efficiency, for instance by exploiting parallel computing and more efficient Monte Carlo schemes, and to expand the range of implemented models and diagnostic tools and providing additional visualization and decision-oriented outputs.

Funding

BB is supported by the Australian Research Council through the Discovery Project Scheme (DP220103269). SAP is supported by the Bocconi Institute for Data Science and Analytics (BIDSA) and MUR–PRIN Bando 2022–prot. 20229PFAX5, financed by the European Union - Next Generation EU, Mission 4 Component 2 CUP J53D23004260001.

References

- Beirlant, J., Y. Goegebeur, J. Segers, and J. L. Teugels (2004). *Statistics of Extremes: Theory and Applications*. John Wiley & Sons.
- Belzile, L. R., C. Dutang, P. J. Northrop, and T. Opitz (2023). A modeler’s guide to extreme value software. *Extremes* 26(4), 595–638.
- Belzile, L. R., S. D. Grimshaw, R. Huser, P. J. Northrop, and J. L. Wadsworth (2025). *mev: Modelling of Extreme Values*. R package version 2.1.
- Beranger, B. and S. A. Padoan (2015). Extreme dependence models. In D. K. Dey and J. Yan (Eds.), *Extreme Value Modeling and Risk Analysis*, pp. 325–352. Chapman and Hall/CRC.
- Beranger, B., S. A. Padoan, and G. Marcon (2025). *ExtremalDep: Extremal Dependence Models*. R package version 1.0.0.
- Beranger, B., S. A. Padoan, and S. A. Sisson (2017). Models for extremal dependence derived from skew-symmetric families. *Scandinavian Journal of Statistics* 44(1), 21–45.
- Beranger, B., S. A. Padoan, and S. A. Sisson (2021). Estimation and uncertainty quantification for extreme quantile regions. *Extremes* 24(2), 349–375.
- Beranger, B., A. G. Stephenson, and S. A. Sisson (2021). High-dimensional inference using the extremal skew- t process. *Extremes* 24(3), 653–685.

- Bernard, E., P. Naveau, M. Vrac, and O. Mestre (2013). Clustering of maxima: Spatial dependencies among heavy rainfall in France. *Journal of Climate* 26(20), 7929–7937.
- Brown, B. M. and S. I. Resnick (1977). Extreme values of independent stochastic processes. *Journal of Applied Probability* 14(4), 732–739.
- Cai, J.-J., J. H. Einmahl, and L. De Haan (2011). Estimation of extreme risk regions under multivariate regular variation. *The Annals of Statistics* 39(3), 1803–1826.
- Coles, S. (2001). *An Introduction to Statistical Modeling of Extreme Values*. Springer.
- Coles, S. G. and J. A. Tawn (1991). Modelling extreme multivariate events. *Journal of the Royal Statistical Society: Series B (Methodological)* 53(2), 377–392.
- Cooley, D., R. A. Davis, and P. Naveau (2010). The pairwise beta distribution: a flexible parametric multivariate model for extremes. *Journal of Multivariate Analysis* 101(9), 2103–2117.
- Cormier, E., C. Genest, and J. G. Nešlehová (2014). Using B-splines for nonparametric inference on bivariate extreme-value copulas. *Extremes* 17(4), 633–659.
- Davis, R. A., T. Mikosch, and I. Cribben (2011). Estimating extremal dependence in univariate and multivariate time series via the extremogram. <https://arxiv.org/abs/1107.5592>.
- Davison, A. C., S. A. Padoan, and M. Ribatet (2012). Statistical modeling of spatial extremes, with discussion. *Statistical Science* 27(2), 161–186.
- de Fondeville, R., L. R. Belzile, and E. Thibaud (2025). *mvPot: Multivariate Peaks-over-Threshold Modelling for Spatial Extreme Events*. R package version 0.1.7.
- de Fondeville, R. and A. C. Davison (2018). High-dimensional peaks-over-threshold inference. *Biometrika* 105(3), 575–592.
- de Haan, L. and A. Ferreira (2006). *Extreme Value Theory: an Introduction*. Springer.
- Dombry, C., S. Engelke, and M. Oesting (2016). Exact simulation of max-stable processes. *Biometrika* 103(2), 303–317.
- Einmahl, J. H., L. de Haan, and A. Krajina (2013). Estimating extreme bivariate quantile regions. *Extremes* 16(2), 121–145.
- Engelke, S. and A. S. Hitz (2020). Graphical models for extremes. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 82(4), 871–932.
- Engelke, S., A. S. Hitz, N. Gnecco, and M. Hentschel (2025). *graphicalExtremes: Statistical Methodology for Graphical Extreme Value Models*. R package version 0.3.4.
- Engelke, S., A. Malinowski, Z. Kabluchko, and M. Schlather (2015). Estimation of Hüsler–Reiss distributions and Brown–Resnick processes. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 77(1), 239–265.
- Falk, M., J. Hüsler, and R.-D. Reiss (2010). *Laws of Small Numbers: Extremes and Rare Events*. Springer.

- Fils-Villetard, A., A. Guillou, and J. Segers (2008). Projection estimators of Pickands dependence functions. *Canadian Journal of Statistics* 36(3), 369–382.
- Frolova, N. and I. Cribben (2016). *extremogram: Estimation of Extreme Value Dependence for Time Series Data*. R package version 1.0.2.
- Garthwaite, P. H., Y. Fan, and S. A. Sisson (2016). Adaptive optimal scaling of Metropolis–Hastings algorithms using the Robbins–Monro process. *Communications in Statistics–Theory and Methods* 45(17), 5098–5111.
- Gilleland, E. and R. W. Katz (2016). extremes 2.0: An extreme value analysis package in R. *Journal of Statistical Software* 72(8), 1–39.
- Haario, H., E. Saksman, and J. Tamminen (2001). An adaptive Metropolis algorithm. *Bernoulli* 7, 223–242.
- He, Y. and J. H. Einmahl (2017). Estimation of extreme depth-based quantile regions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 79(2), 449–461.
- Heffernan, J. E. and A. G. Stephenson (2018). *ismev: An Introduction to Statistical Modeling of Extreme Values*. R package version 1.4.2.
- Heffernan, J. E. and J. A. Tawn (2004). A conditional approach for multivariate extreme values (with discussion). *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 66(3), 497–546.
- Hijmans, R. J. (2026). *terra: Spatial Data Analysis*. R package version 1.8-93.
- Huser, R., T. Opitz, and E. Thibaud (2017). Bridging asymptotic independence and dependence in spatial extremes using gaussian scale mixtures. *Spatial Statistics* 21, 166–186.
- Joe, H. (2014). *Dependence Modeling with Copulas*. Chapman and Hall/CRC.
- Kiriliouk, A., J. Lee, and J. Segers (2025). X-vine models for multivariate extremes. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 87(3), 579–602.
- Kiriliouk, A., H. Rootzén, J. Segers, and J. L. Wadsworth (2019). Peaks over thresholds modeling with multivariate generalized pareto distributions. *Technometrics* 61(1), 123–135.
- Marcon, G., P. Naveau, and S. Padoan (2017). A semi-parametric stochastic generator for bivariate extreme events. *Stat* 6(1), 184–201.
- Marcon, G., S. A. Padoan, and I. Antoniano-Villalobos (2016). Bayesian inference for the extremal dependence. *Electronic Journal of Statistics* 10(2), 3310–3337.
- Marcon, G., S. A. Padoan, P. Naveau, P. Muliere, and J. Segers (2017). Multivariate non-parametric estimation of the Pickands dependence function using Bernstein polynomials. *Journal of Statistical Planning and Inference* 183, 1–17.
- Northrop, P. J. and N. Attalides (2016). Posterior propriety in Bayesian extreme value analyses using reference priors. *Statistica Sinica* 26(2), 721–743.

- Northrop, P. J. and S. D. Grimshaw (2025). *revdbayes: Ratio-of-Uniforms Sampling for Bayesian Extreme Value Analysis*. R package version 1.5.6.
- Opitz, T. (2013). Extremal t processes: Elliptical domain of attraction and a spectral representation. *Journal of Multivariate Analysis* 122(0), 409 – 413.
- Pacifici, C., S. A. Padoan, and G. Stupfler (2025). *ExtremeRisks: Extreme Risk Measures*. R package version 0.0.5.
- Padoan, S. A., M. Ribatet, and S. A. Sisson (2010). Likelihood-based inference for max-stable processes. *Journal of the American Statistical Association* 105(489), 263–277.
- Ribatet, M. (2022). *SpatialExtremes: Modelling Spatial Extremes*. R package version 2.1-0.
- Roberts, G. O., A. Gelman, and W. R. Gilks (1997). Weak convergence and optimal scaling of random walk Metropolis algorithms. *The Annals of Applied Probability* 7(1), 110–120.
- Rootzén, H., J. Segers, and J. L. Wadsworth (2018). Multivariate peaks over thresholds models. *Extremes* 21(1), 115–145.
- Sabourin, A., P. Naveau, and A.-L. Fougères (2013). Bayesian model averaging for multivariate extremes. *Extremes* 16(3), 325–350.
- Southworth, H., J. E. Heffernan, and P. D. Metcalfe (2024). *texmex: Statistical Modelling of Extreme Values*. R package version 2.4.9.
- Stephenson, A. and J. Tawn (2005). Exploiting occurrence times in likelihood inference for componentwise maxima. *Biometrika* 92(1), 213–227.
- Stephenson, A. G. (2002, June). evd: Extreme value distributions. *R News* 2(2), 31–32.
- Ypma, J. and S. G. Johnson (2025). *The NLOpt nonlinear-optimization package*. R package version 2.2.1.

A Appendix

A.1 Modelling of log-returns of exchange rates between GBP/USD and GPB/JPY: Goodness of fit

A.2 Adaptive Gaussian random-walk Metropolis–Hastings (RWMH) algorithm

At the current state s of the chain, $\boldsymbol{\theta}_j^{(s)}$ is updated by the proposal $\boldsymbol{\theta}'_j \sim \phi_3(\boldsymbol{\theta}^{(s)}, \tau^{(s)}\Sigma^{(s)})$, where $\phi_d(\mathbf{a}, A)$ is a d -dimensional Gaussian density with mean \mathbf{a} and covariance A . The proposal covariance matrix $\Sigma^{(s)}$ is specified as

$$\Sigma^{(s+1)} = \begin{cases} (1 + [\tau^{(s)}]^2/s) \mathbb{I}_3, & s \leq 100 \\ \frac{1}{s-1} \sum_{l=1}^s (\boldsymbol{\theta}^{(l)} - \bar{\boldsymbol{\theta}}^{(s)}) (\boldsymbol{\theta}^{(l)} - \bar{\boldsymbol{\theta}}^{(s)})^\top + \left\{ (\tau^{(s)})^2 / s \right\} \mathbb{I}_3, & s > 100, \end{cases} \quad (\text{A.1})$$




Figure 14: Left and middle panels: marginal density histograms (light blue) for GBP/USD and GBP/JPY with GEV density curves (solid blue) evaluated at each posterior sample for its corresponding parameter. Right panel: histogram of the empirical angular density $h(\omega)$ (light blue) with posterior mean of the estimated angular density (dotted blue) and 90% credibility envelope.

where \mathbb{I}_d is the d -dimensional identity matrix, $\bar{\boldsymbol{\theta}}^{(s)} = s^{-1} (\boldsymbol{\theta}^{(1)} + \dots + \boldsymbol{\theta}^{(s)})$, and $\tau^{(s)} > 0$ is a scaling parameter that affects the acceptance rate of proposal parameter values (Haario et al., 2001) and, τ is updated using a Robbins–Monro process so that

$$\log \tau^{(s+1)} = \log \tau^{(s)} + c (\pi^{(s)} - \pi^*), \quad (\text{A.2})$$

where $c = (2\pi)^{1/2} \exp(\zeta_0^2/2) / (2\zeta_0)$ is a steplength constant, $\zeta_0 = -1/\Phi(\pi^*/2)$, and Φ is the univariate standard Gaussian distribution function (Garthwaite et al., 2016). To control the desired sampler acceptance probability, the parameter $\pi^* = 0.234$ is specified as in Roberts et al. (1997). We note that some improvements of the adaptive MCMC could be implemented to gain in efficiency, for example by taking as initial covariance a multiple of the inverse Fisher information. **Furthermore, alternative strategies to guarantee diminishing adaptation could be implemented, e.g. by specifying a long burn-in period during which to adapt, then stop and start sampling.**

A.3 MCMC Scheme for the joint inference of marginal distribution and extremal dependence

Algorithm 1: Trans-dimensional MCMC scheme

- 1 **Initialize:** Set $\boldsymbol{\vartheta}^{(0)} = (\boldsymbol{\theta}_1^{(0)}, \boldsymbol{\theta}_2^{(0)}, \kappa^{(0)}, \boldsymbol{\eta}_{\kappa^{(0)}}^{(0)})$, $\boldsymbol{\eta}_{\kappa^{(0)}}^{(0)}$, $\tau_i^{(0)}$ and $\Sigma_i^{(0)}$ for $j = 1, 2$.
 - 2 **for** $s = 0$ **to** M **do**
 - 3 **Step 1:** *Marginal component 1:*
 - 4 Draw proposal $\boldsymbol{\theta}'_1 \sim \text{MVN}(\boldsymbol{\theta}_1^{(s)}, \tau_1^{(s)} \Sigma_1^{(s)})$.
 - 5 Compute acceptance probability $\pi_1 = \min \left\{ \frac{\mathcal{L}(\boldsymbol{\theta}'_1, \boldsymbol{\theta}_2^{(s)}, \kappa^{(s)}, \boldsymbol{\beta}_{\kappa^{(s)}}^{(s)}) \Pi(\boldsymbol{\theta}'_1)}{\mathcal{L}(\boldsymbol{\theta}_1^{(s)}, \boldsymbol{\theta}_2^{(s)}, \kappa^{(s)}, \boldsymbol{\beta}_{\kappa^{(s)}}^{(s)}) \Pi(\boldsymbol{\theta}_1^{(s)})}, 1 \right\}$.
 - 6 Draw $U_1 \sim \mathcal{U}(0, 1)$. If $\pi_1 < U_1$ then set $\boldsymbol{\theta}_1^{(s+1)} = \boldsymbol{\theta}'_1$ else set $\boldsymbol{\theta}_1^{(s+1)} = \boldsymbol{\theta}_1^{(s)}$.
 - 7 Update $\Sigma_1^{(s)}$ according to (A.1).
 - 8 Update $\tau_1^{(s)}$ according to (A.2).
 - 9 **Step 2:** *Marginal component 2:*
 - 10 Draw proposal $\boldsymbol{\theta}'_2 \sim \text{MVN}(\boldsymbol{\theta}_2^{(s)}, \tau_2^{(s)} \Sigma_2^{(s)})$.
 - 11 Compute acceptance probability $\pi_2 = \min \left\{ \frac{\mathcal{L}(\boldsymbol{\theta}_1^{(s+1)}, \boldsymbol{\theta}'_2, \kappa^{(s)}, \boldsymbol{\beta}_{\kappa^{(s)}}^{(s)}) \Pi(\boldsymbol{\theta}'_2)}{\mathcal{L}(\boldsymbol{\theta}_1^{(s+1)}, \boldsymbol{\theta}_2^{(s)}, \kappa^{(s)}, \boldsymbol{\beta}_{\kappa^{(s)}}^{(s)}) \Pi(\boldsymbol{\theta}_2^{(s)})}, 1 \right\}$.
 - 12 Draw $U_2 \sim \mathcal{U}(0, 1)$. If $\pi_2 < U_2$ then set $\boldsymbol{\theta}_2^{(s+1)} = \boldsymbol{\theta}'_2$ else set $\boldsymbol{\theta}_2^{(s+1)} = \boldsymbol{\theta}_2^{(s)}$.
 - 13 Update $\Sigma_2^{(s)}$ according to (A.1).
 - 14 Update $\tau_2^{(s)}$ according to (A.2).
 - 15 **Step 3:** *Dependence structure:*
 - 16 Draw proposal $\kappa' \sim q_\kappa(\kappa \mid \kappa^{(s)})$ and $\boldsymbol{\eta}'_{\kappa'} \sim q_\eta(\boldsymbol{\eta}_\kappa \mid \kappa')$, and compute $\boldsymbol{\eta}'_{\kappa'}$.
 - 17 Set $c = 1/2$ if $\kappa^{(s)} = 3$ or $c = 1$ if $\kappa^{(s)} > 3$.
 - 18 Compute acceptance probability

$$\pi_3 = \min \left\{ c \frac{\Pi(\kappa')}{\Pi(\kappa^{(s)})} \frac{\mathcal{L}(\boldsymbol{\theta}_1^{(s+1)}, \boldsymbol{\theta}_2^{(s+1)}, \kappa', \boldsymbol{\beta}'_{\kappa'})}{\mathcal{L}(\boldsymbol{\theta}_1^{(s+1)}, \boldsymbol{\theta}_2^{(s+1)}, \kappa^{(s)}, \boldsymbol{\beta}_{\kappa^{(s)}}^{(s)})}, 1 \right\}.$$
 - 19 Draw $U_3 \sim \mathcal{U}(0, 1)$. If $\pi_3 < U_3$ then set $\kappa^{(s+1)} = \kappa'$, $\boldsymbol{\eta}_{\kappa^{(s+1)}}^{(s+1)} = \boldsymbol{\eta}'_{\kappa'}$, $\boldsymbol{\beta}_{\kappa^{(s+1)}}^{(s+1)} = \boldsymbol{\beta}'_{\kappa'}$ else set $\kappa^{(s+1)} = \kappa^{(s)}$, $\boldsymbol{\eta}_{\kappa^{(s+1)}}^{(s+1)} = \boldsymbol{\eta}_{\kappa^{(s)}}^{(s)}$, $\boldsymbol{\beta}_{\kappa^{(s+1)}}^{(s+1)} = \boldsymbol{\beta}_{\kappa^{(s)}}^{(s)}$.
-