



HAL
open science

Efficient inference and simulation for elliptical Pareto processes

Emeric Thibaud, Thomas Opitz

► **To cite this version:**

Emeric Thibaud, Thomas Opitz. Efficient inference and simulation for elliptical Pareto processes. *Biometrika*, 2015, 102 (3), 10.1093/biomet/asv045 . hal-02641446

HAL Id: hal-02641446

<https://hal.inrae.fr/hal-02641446>

Submitted on 31 Jul 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Efficient inference and simulation for elliptical Pareto processes

Emeric Thibaud^{*1} and Thomas Opitz^{†2}

¹Department of Statistics, Colorado State University, Fort Collins, Colorado 80523-1877, U.S.A.

²Biostatistics and Spatial Processes Unit, French National Institute for Agronomic Research, 84914 Avignon, France

Abstract

Recent advances in extreme value theory have established ℓ -Pareto processes as the natural limits for extreme events defined in terms of exceedances of a risk functional. Here we provide methods for the practical modelling of data based on a tractable yet flexible dependence model. We introduce the class of elliptical ℓ -Pareto processes, which arise as the limit of threshold exceedances of certain elliptical processes characterized by a correlation function and a shape parameter. An efficient inference method based on maximizing a full likelihood with partial censoring is developed. Novel procedures for exact conditional and unconditional simulation are proposed. These ideas are illustrated using precipitation extremes in Switzerland.

Keywords: Censored likelihood; Elliptical extremes; Extremal- t process; Pareto process; Simulation.

1 Introduction

There has recently been increasing development of methodologies for modelling spatial extremes, motivated by numerous applications in climatology and environmental sciences. Classical extreme value theory relies on max-stable processes, which extend the univariate generalized extreme-value distribution to stochastic processes. Such processes are the only possible nondegenerate limits for rescaled maxima of spatial processes (de Haan and Ferreira, 2006, Ch. 9) and provide a natural modelling framework for asymptotically dependent extremes.

Inference for spatial extremes has been based on various max-stable models (e.g., Davison et al., 2012; Ribatet, 2013). The extremal Gaussian model (Schlather, 2002) or the Brown–Resnick model (Kabluchko et al., 2009) have proven to be well-suited for modelling extremal dependence of environmental data (Davison et al., 2012; Ribatet, 2013). The extremal- t process, which can be seen as generalizing these two models, is the max-stable limit of all asymptotically dependent elliptical processes (Opitz, 2013). Its distribution depends on a correlation function and a shape parameter, providing a flexible dependence structure for spatial extremes. Because of the complicated form of the distribution of a max-stable process, composite likelihood methods have been used to fit such models (Padoan et al., 2010), leading to a loss in efficiency. More recently, efficient full likelihood inference methods were developed in a point process framework. Engelke et al. (2015) developed full likelihood methods based either on the distribution of increments with respect to a fixed extreme component, or on the multivariate spectral measure. Wadsworth and Tawn (2014) calculated

^{*}Emeric.Thibaud@colostate.edu

[†]Thomas.Opitz@paca.inra.fr

a full likelihood for exceedances of a thresholding field while censoring the part of the observation vector falling below this threshold.

In this paper, we propose the use of ℓ -Pareto processes (Dombry and Ribatet, 2015) for modelling extremes of spatial processes defined in terms of the exceedance of a risk functional. Ferreira and de Haan (2014) and Dombry and Ribatet (2015) showed that Pareto processes are the only possible asymptotic limits for threshold exceedances of spatial processes. Inference based on these processes is currently limited to nonparametric estimation (Dombry and Ribatet, 2015). We introduce the elliptical ℓ -Pareto process, which is the limiting process for threshold exceedances of all asymptotically dependent elliptical processes, and propose an efficient inference approach for it based on a full likelihood with partial censoring. The resulting inferential procedures, potentially more efficient than composite likelihood methods, are discussed, and efficiency gains over a pairwise likelihood are assessed in a simulation study. In addition, we propose a new approach to exact simulation from extremal- t and elliptical Pareto processes, and we show how conditional simulations can be obtained very easily for the latter. Finally, we illustrate the use of elliptical Pareto processes in an application to extreme precipitation in Switzerland.

We develop our results for processes with continuous sample paths defined on a nonempty compact domain $K \subset \mathbb{R}^m$, $m \geq 1$. The assumption of continuity is natural in applications and ensures that Pareto processes are well-defined. We focus here on the practical use of Pareto processes; for more technical details on the definitions of these processes and related convergence in functional spaces we refer to Ferreira and de Haan (2014) and Dombry and Ribatet (2015).

2 Functional extreme value theory

2.1 ℓ -Pareto processes

We let $C(K)$ denote the space of continuous functions over K , endowed with the supremum norm $\|f\|_\infty = \sup_{s \in K} |f(s)|$. The restriction of $C(K)$ to non-negative functions is denoted by $C_+(K)$. In univariate and multivariate theory, a generalized Pareto limit is obtained by conditioning on the exceedance of a high threshold in at least one component (Rootzén and Tajvidi, 2006). Ferreira and de Haan (2014) extended this idea to infinite-dimensional spaces by conditioning on exceedances of the supremum of the process over the space, leading to the notion of a generalized Pareto process. This definition was generalized by Dombry and Ribatet (2015) to so-called ℓ -Pareto processes by considering exceedances defined in terms of a linear risk functional.

A functional $\ell : C_+(K) \rightarrow [0, \infty)$ is called a risk functional, or cost functional, if it is continuous and homogeneous, i.e., $\ell(tf) = t\ell(f)$ for $t \geq 0$. In what follows, the unit sphere in $C_+(K)$ with respect to ℓ is written $S_\ell = \{f \in C_+(K) : \ell(f) = 1\}$. Suppose that we are given a risk functional ℓ and a probability measure ρ on S_ℓ . We call any process $Y^* = \{Y^*(s)\}_{s \in K}$ a standard ℓ -Pareto process with ℓ -spectral distribution ρ if it can be represented as

$$Y^*(s) = Rf_0(s), \quad R \sim \text{Par}(1), \quad \{f_0(s)\} \sim \rho, \quad (1)$$

i.e., $\Pr(R > y) = 1/y$ ($y \geq 1$), with R independent of the spectral function $f_0 = \{f_0(s)\}$. For continuous real functions $\sigma(s) > 0$, $\mu(s)$, $\xi(s)$ defined over K , the process

$$\begin{cases} \mu(s) + \sigma(s) \{Y^*(s)^{\xi(s)} - 1\} / \xi(s), & \xi(s) \neq 0, \\ \mu(s) + \sigma(s) \log Y^*(s), & \xi(s) = 0, \end{cases} \quad s \in K, \quad (2)$$

is termed a generalized ℓ -Pareto process. To avoid confusion between processes associated to different risk functionals ℓ , we will write Y_ℓ^* for Y^* . Such processes Y_ℓ^* generalize the peaks-over-threshold stability of multivariate generalized Pareto distributions to infinite dimensions: for any $u \geq 1$, the renormalized

threshold-exceeding process $\{u^{-1}Y_\ell^* \mid \ell(Y_\ell^*) \geq u\}$ is equal to Y_ℓ^* in distribution (Dombry and Ribatet, 2015). The interpretation of the construction (1) is that R characterizes the intensity of an extreme event in terms of the risk functional ℓ , whereas f_0 describes the corresponding spatial profile.

In applications, different choices of ℓ may be used to answer different questions. If $\ell(f) = \max_{j=1, \dots, D} f(s_j)/u_j$ for certain sites $s_j \in K$ ($j = 1, \dots, D$), we focus on processes with at least one exceedance of the thresholds $u_j > 0$. By contrast, $\ell(f) = \min_{j=1, \dots, D} f(s_j)/u_j$ requires exceedances at each of the D sites. The original definition of a Pareto process (Ferreira and de Haan, 2014) uses $\ell(f) = \sup_{s \in K} f(s)$, but conditioning on another $\ell(f)$ is desirable in applications where data are only observed at a finite number of sites.

2.2 Limiting processes of extremes

We recall the different forms of convergence of extremes of continuous processes in terms of block maxima, threshold exceedances and point processes. Throughout, the symbol \implies indicates weak convergence of random elements from the univariate, multivariate or functional domain. For independent and identically distributed copies X_1, X_2, \dots of a stochastic process $X = \{X(s)\}_{s \in K}$ with continuous sample paths, we say that X is in the maximum domain of attraction of a max-stable process $Z = \{Z(s)\}_{s \in K}$ (de Haan and Ferreira, 2006, Ch. 9) if there exist sequences of normalizing continuous functions $a_n(s) > 0$ and $b_n(s)$ such that

$$\left\{ \max_{i=1, \dots, n} a_n(s)^{-1} \{X_i(s) - b_n(s)\} \right\} \implies \{Z(s)\}, \quad n \rightarrow \infty, \quad (3)$$

in $C(K)$, with the limit process Z having nondegenerate univariate distributions. Convergence of the dependence structure and of marginal distributions can be viewed separately in (3) (de Haan and Ferreira, 2006, §9.2). Therefore, we define a normalized process X^* by $X^*(s) = 1/[1 - F_{X(s)}\{X(s)\}]$ where $F_{X(s)}$ denotes the distribution of $X(s)$. If we assume that X has continuous marginal distributions, then X^* has marginal standard Pareto distributions. For $a_n(s) \equiv n$ and $b_n(s) \equiv 0$ the max-stable limit for X^* in (3) is a standard max-stable process Z^* with univariate unit Fréchet distributions.

The distribution of the process Z^* is fully characterized by a so-called exponent measure Λ on $C(K) \setminus \{0\}$ through the relation (Giné et al., 1990)

$$\Lambda \left[\bigcup_{j=1, \dots, D} \left\{ f \in C(K) : \sup_{s \in K_j} f(s) \geq z_j \right\} \right] = -\log \Pr \left\{ \sup_{s \in K_1} Z^*(s) \leq z_1, \dots, \sup_{s \in K_D} Z^*(s) \leq z_D \right\} \quad (4)$$

for any collection of nonempty compact sets $K_j \subset K$ and $z_j > 0$ ($j = 1, \dots, D$). The measure Λ is uniquely defined if we impose the constraint $\Lambda\{C(K) \setminus C_+(K)\} = 0$. We obtain the unique version as $\Lambda_+(B) = \Lambda[\{f \in C(K) : f_+ \in B\}]$, for measurable $B \subset C_+(K) \setminus \{0\}$, with $f_+(s) = \max\{f(s), 0\}$. When the sets $K_j = \{s_j\}$ are singletons, expression (4) is called the exponent function and denoted by $V(z_1, \dots, z_D)$, where $V(1, \dots, 1)$ is known as the extremal coefficient of the sites s_1, \dots, s_D . Finite-dimensional marginal measures of Λ relative to D sites $s = (s_1, \dots, s_D)$ are written as Λ_s , i.e., $\Lambda_s(\times_{j=1, \dots, D} [a_j, b_j]) = \Lambda[\bigcap_{j=1, \dots, D} \{f \in C(K) : f(s_j) \in [a_j, b_j]\}]$ for $0 < a_j < b_j$ ($j = 1, \dots, D$). In particular, $V(z_1, \dots, z_D) = \Lambda_s \{(\times_{j=1, \dots, D} [-\infty, z_j])^C\}$.

Max-stability of Z^* implies that the measure Λ_+ is homogeneous of order -1 , i.e., $\Lambda_+(tB) = t^{-1}\Lambda_+(B)$, $t > 0$. For some risk functional ℓ and $f \in C_+(K)$ with $\ell(f) > 0$, consider the pseudo-polar coordinates (r, f_0) with $r = \ell(f)$ and $f_0 = f/\ell(f)$. If $\kappa_\ell(K) = \Lambda_+\{\ell(f) \geq 1\} > 0$, arguments similar to de Haan and Ferreira (2006, §9.4) imply the factorization

$$\Lambda_+(df) = \kappa_\ell(K) r^{-2} dr \rho_\ell(df_0), \quad r > 0, \quad (5)$$

with ρ_ℓ an ℓ -spectral distribution on S_ℓ .

Assuming X^* is in the maximum domain of attraction of a standard max-stable process Z^* , [Dombry and Ribatet \(2015, Theorem 3\)](#) proved the convergence of standard ℓ -exceedances

$$\{n^{-1}X^*(s) : \ell(X^*) > n\} \implies \{Y_\ell^*(s)\}, \quad n \rightarrow \infty, \quad (6)$$

where Y_ℓ^* is a standard ℓ -Pareto process with ℓ -spectral distribution ρ_ℓ related to the exponent measure Λ_+ of Z^* through (5).

Convergence (6) establishes a basis for threshold-based inference as follows. From definition (1), the distribution of a standard ℓ -Pareto process is $r^{-2}dr\rho_\ell(df_0)$ on $[1, \infty) \times S_\ell$, which is also equal to $\Lambda_+(df)/\kappa_\ell(K)$ from (5) with $f = rf_0$. Hence the convergence in (6) conveys that, for large n ,

$$\Pr\{X^* \in B \mid \ell(X^*) > n\} \approx n\Lambda_+(B)/\kappa_\ell(K),$$

for $B \subset \{f \in C_+(K) : \ell(f) > n\}$.

A point process framework links max-stable limits for maxima and ℓ -Pareto limits for threshold exceedances. Convergence (3) for the normalized process X^* is equivalently expressed in terms of the standard point process convergence ([de Haan and Ferreira, 2006, Theorem 9.3.9](#))

$$\{n^{-1}X_i^*(s), i = 1, \dots, n\} \implies \mathcal{P}, \quad n \rightarrow \infty, \quad (7)$$

where $\mathcal{P} = \{P_i(s), i = 1, 2, \dots\}$ is a Poisson process with intensity measure Λ_+ . Then, $Z^*(s) = \max_{i=1, \dots, n} P_i(s)$ and, from Poisson process theory ([Daley and Vere-Jones, 2007, Ch. 9](#)), the points P_i with $\ell(P_i) \geq 1$ are independent and have distribution $\Lambda_+(df)/\kappa_\ell(K)$; they are realisations of the ℓ -Pareto process Y_ℓ^* .

Max-stable convergence in (3) implies the convergence of marginal pointwise maxima to generalized extreme-value distributions, for which standard theory ([Beirlant et al., 2004, Ch. 5](#)) provides alternative characterizations in terms of univariate threshold exceedances or point process convergence. From a Pareto process perspective, it is convenient to fix a high threshold function $u(s)$ and to assume that

$$\Pr\{X(s) > x\} = [1 + \xi(s)\{x - \mu(s)\}/\sigma(s)]_+^{-1/\xi(s)}, \quad x > u(s), \quad (8)$$

corresponding to the univariate tail probabilities of the generalized Pareto process in (2), with real parameters $\mu(s) < u(s)$, $\sigma(s) > 0$ and $\xi(s)$, such that the right-hand side of (8) is less than unity.

2.3 Elliptical extremes

A random vector $X \in \mathbb{R}^D$ is said to follow an elliptical distribution if it can be written as

$$X = RAU + \mu, \quad (9)$$

with R a nonnegative random variable, A a $D \times D$ deterministic nonsingular matrix defining the dispersion matrix $\Sigma = AA'$, U a random vector independent of R and distributed uniformly on the Euclidean unit sphere $\{x \in \mathbb{R}^D : x'x = 1\}$ and $\mu \in \mathbb{R}^D$ a deterministic shift vector. The restriction to nonsingular square matrices A excludes some special cases of minor practical importance. Examples of elliptical distributions are the multivariate Gaussian and the multivariate t distributions. As an extension of (9), a random process X is called elliptical if all its finite-dimensional distributions are elliptical with dispersion matrices Σ defined through a covariance function. The max-stable limits in (3) for elliptical processes are either processes with independent univariate marginal distributions in the case of asymptotic independence, as for instance the limits of Gaussian processes, or are extremal- t processes in all other cases. In terms of unit Fréchet margins, extremal- t processes can be represented as

$$Z^*(s) = m_\alpha \max_{i=1,2,\dots} W_i(s)_+^\alpha / Q_i, \quad m_\alpha = \pi^{1/2} 2^{1-\alpha/2} \Gamma\{(\alpha+1)/2\}^{-1}, \quad (10)$$

where $0 < Q_1 < Q_2 < \dots$ are the points of a unit-rate Poisson process on the positive half-line, and $W_i = \{W_i(s)\}$ are independent replicates of a standard Gaussian process with continuous sample paths and correlation function ϱ (Opitz, 2013). In particular, $\alpha = 1$ yields the extremal Gaussian process (Schlather, 2002). By interpreting the processes W_i as independent marks of the points of the Poisson process $\{Q_i\}$, we see that the point process $\{P_i\} = \{m_\alpha(W_i)_+^\alpha/Q_i\}$ is Poisson with intensity measure Λ_+ . We use this for simulation from the corresponding ℓ -Pareto process, see Section 4. We use the term elliptical ℓ -Pareto process since the tails of its finite-dimensional distributions correspond to elliptical distributions with a Pareto-distributed radial variable R in (9). The finite-dimensional dependence structure associated to D sites $s = (s_1, \dots, s_D)$ is characterized by the exponent function (Nikoloulopoulos et al., 2009)

$$\begin{aligned} V(z) &= -\log \text{pr}\{Z^*(s_1) \leq z_1, \dots, Z^*(s_D) \leq z_D\} \\ &= \sum_{j=1}^D z_j^{-1} t_{\alpha+1} \left\{ (z_{-j}/z_j)^{1/\alpha}, \Sigma_{-j,j}, (\alpha+1)^{-1} (\Sigma_{-j,-j} - \Sigma_{-j,j} \Sigma'_{-j,j}) \right\}, \end{aligned} \quad (11)$$

with the correlation matrix $\Sigma = \{\varrho(s_{j_1}, s_{j_2})\}_{1 \leq j_1, j_2 \leq D}$ related to the correlation function ϱ , and with $t_\alpha(\cdot; \mu, \Sigma)$ the cumulative distribution function of a multivariate t distribution with α degrees of freedom and parameters μ and Σ .

Dependence structures of Brown–Resnick type arise as a special case of extremal- t dependence when $\alpha \rightarrow \infty$. By analogy with (10), a standard Brown–Resnick process is constructed as $Z_{\text{BR}}^*(s) = \max_{i=1,2,\dots} \exp\{\tilde{W}_i(s) - \gamma(s)\}/Q_i$ (Kablichko et al., 2009), where \tilde{W}_i are independent and identically distributed copies of an intrinsically stationary centered Gaussian process characterized by its variogram $2\gamma(s) = \text{E}\{\tilde{W}_1(s)^2\}$ and with $\tilde{W}_1(0) = 0$ almost surely. For processes W_1 whose correlation function ϱ_α depends on α such that the limit $\gamma(s_2 - s_1) = \lim_{\alpha \rightarrow \infty} \alpha\{1 - \varrho_\alpha(s_1, s_2)\}$ exists and satisfies $0 < \gamma(s_2 - s_1) < \infty$ for all sites s_1, s_2 with $s_1 \neq s_2$, the extremal- t process Z^* in (10) converges to Z_{BR}^* as $\alpha \rightarrow \infty$ (Nikoloulopoulos et al., 2009). For instance, the correlation function $\varrho_\alpha(s_1, s_2) = \exp[-\{\|s_1 - s_2\|/(\alpha^{1/\kappa}\lambda)\}^\kappa]$ with $\kappa \in (0, 2]$ and $\lambda > 0$ yields the variogram $2(\|s_1 - s_2\|/\lambda)^\kappa$.

The truncation of W_i at zero in (10) implies that the measure Λ_+ of an extremal- t process has positive mass on the set $\{f \in C_+(K) \setminus \{0\} : \min_{s \in K} f(s) = 0\}$, which is not the case for Brown–Resnick processes. We later discuss the implications for inference on elliptical ℓ -Pareto processes.

3 Inference

3.1 Likelihoods for ℓ -Pareto processes

We now consider a collection $s = (s_1, \dots, s_D)$ of sites in K , and $X_{s,1}, \dots, X_{s,n}$ independent replicates of a finite-dimensional observation vector $X_s = \{X(s_1), \dots, X(s_D)\}$, which is embedded in a process X . We suppose that X is in the maximum domain of attraction of a max-stable process Z . We assume the marginal parameters $\mu(s)$, $\sigma(s)$ and $\xi(s)$ in (8) were estimated in a first step, and we consider the standardized process X^* , whose finite-dimensional vectors relative to s are denoted by X_s^* . Here we describe the estimation of Λ_+ based on ℓ -exceedances of X_s^* with a suitably chosen risk functional ℓ . We use the elliptical Pareto process with a parametric correlation function defining a parametric model for the measure Λ_+ .

Different choices of ℓ yield different approaches to inference, but it is crucial that $\ell(X^*)$ can be determined from X_s^* , so we need $\ell(X^*) = \ell(X_s^*)$. Without loss of generality, we define the exceedance observations in terms of $\ell(X_s^*) \geq 1$. We approximate the distribution of the points X_s^* with $\ell(X_s^*) \geq 1$ by the distribution of the elliptical ℓ -Pareto process Y_ℓ^* .

For a standard ℓ -Pareto process Y_ℓ^* , the density of the vector $Y_s^* = \{Y_\ell^*(s_1), \dots, Y_\ell^*(s_D)\}$ on $\{y \in \mathbb{R}_+^D \setminus \{0\} : \ell(y) \geq 1\}$ is $\lambda_{+,s}(y)/\kappa_\ell(K)$, where $\lambda_{+,s}$ is the density of $\Lambda_{+,s}$, the finite-dimensional marginal

measure of Λ_+ relative to the sites s . When $\Lambda_{+,s}$ is absolutely continuous with respect to Lebesgue measure, $\lambda_{+,s}$ is the full derivative $-V_{1:D}(y)$ of the negated exponent function V . Otherwise, when $\Lambda_{+,s}$ puts positive mass on lower-dimensional subspaces of \mathbb{R}_+^D , we get slightly different expressions for $\lambda_{+,s}$ on those subspaces (Coles and Tawn, 1991). In the extremal elliptical model, we find positive mass on $\{y \in \mathbb{R}_+^D \setminus \{0\} : \|y\|_\infty > 0\}$, see §3.2.

Based on the sample of ℓ -exceedances $X_{s,k}^*$ ($k = 1, \dots, N_u$) satisfying $\ell(X_{s,k}^*) \geq 1$, and assuming a parametric ℓ -Pareto model with parameter vector ψ , we obtain the full likelihood

$$\tilde{L}_\ell(\psi) = \prod_{k=1}^{N_u} \frac{\lambda_{+,s}(X_{s,k}^*)}{\kappa_\ell(K)}. \quad (12)$$

When $\kappa_\ell(K)$ cannot be calculated explicitly, Monte Carlo approximations are required to evaluate the likelihood function (12). For a choice of ℓ that is both tractable and useful in practice, we focus on $\ell(f) = \max_{j=1, \dots, D} f(s_j)/u_j$ with a high multivariate threshold $u = (u_1, \dots, u_D) > 0$, which select observations for which at least one component exceed its marginal threshold. Then, $\kappa_\ell(K) = V(u)$ and (Ferreira and de Haan, 2014)

$$\text{pr}(Y_s^* \leq y) = \frac{V\{\min(y, u)\} - V(y)}{V(u)}, \quad y \not\leq u,$$

which is the multivariate Pareto distribution defined by Rootzén and Tajvidi (2006). Specifying $\kappa_\ell(K) = V(u)$ in (12) yields the corresponding likelihood

$$\tilde{L}_1(\psi) = \prod_{k=1}^{N_u} \frac{\lambda_{+,s}(X_{s,k}^*)}{V(u)}.$$

Inference based on \tilde{L}_1 might be compromised in practice: first, using the full information from an observation $X_{s,k}^*$ with $\ell(X_{s,k}^*) \geq 1$ might be inefficient since the asymptotic distribution might model the non-exceeding components badly and thus induce bias in the estimators. Second, positive mass on the boundary of $\mathbb{R}_+^D \setminus \{0\}$ creates a discontinuity due to the weak convergence of the data process to the ℓ -Pareto process in (6), as is the case for the elliptical model. The margins of $X_{s,k}^*$ are standard Pareto, and so are strictly positive, which is incoherent with the possible mass on the axis for $\Lambda_{+,s}$. To overcome these two issues, we propose the use of a censoring scheme. We consider the censored observations $X_{s,k}^c = \max(X_{s,k}^*, u)$, where the maximum is taken componentwise. The corresponding likelihood is

$$\tilde{L}_2(\psi) = \prod_{k=1}^{N_u} \frac{-V_{I_k}(X_{s,k}^c)}{V(u)},$$

where V_{I_k} denotes the partial derivative of V with respect to the indices $I_k \subset \{1, \dots, D\}$ associated to the components that exceed their corresponding marginal thresholds.

When both n and N_u are observed, we propose to incorporate the information provided by the binomial variable $n - N_u$, that represents the number of fully-censored observations. We use the approximation $\text{Pr}\{\ell(X^*) \geq 1\} = \text{Pr}\{\max_{j=1, \dots, D} X^*(s_j)/u_j \geq 1\} \approx V(u)$, which follows from (3) for a high threshold vector u (de Haan and Ferreira, 2006, Theorem 9.3.1), and define the likelihoods

$$L_m = \{1 - V(u)\}^{n - N_u} V(u)^{N_u} \times \tilde{L}_m \quad (m = 1, 2).$$

The threshold vector u must be high enough to yield $V(u) \leq 1$.

Full likelihood inference based on L_1 or L_2 is possible if $\lambda_{+,s}$, the function V and its partial derivatives, are known. We derive these expressions for elliptical Pareto processes in Section 3.2; expressions for Brown–Resnick processes were derived by [Wadsworth and Tawn \(2014\)](#). By contrast, inference for max-stable extremal- t and Brown–Resnick processes is typically based on composite likelihoods ([Padoan et al., 2010](#)). In Section 5, we use simulation to investigate the gain in efficiency from the use of full likelihoods.

We relate our approach to [Wadsworth and Tawn \(2014\)](#) and [Engelke et al. \(2015\)](#), who proposed full likelihood inference based on the finite-dimensional convergence to the Poisson process in (7). [Wadsworth and Tawn \(2014\)](#) proposed a censored approach with the likelihood

$$\exp\{-nV(u)\}V(u)^{N_u} \times \tilde{L}_2(\psi),$$

which differs from L_2 only through the distribution assumed for the number of exceedance N_u : binomial for the Pareto approach, and Poisson for the point process approach. Since n is large and $V(u)$ is small in practice, these two approaches give very similar results. By contrast, [Engelke et al. \(2015\)](#) considered $X_{s,k}^*$ as an exceedance when $\sum_{j=1}^D X_{s_j,k}^* > u$, leading to inference based on the multivariate sum spectral measure ([Coles and Tawn, 1991](#)). An equivalent approach in the framework of Pareto processes is obtained for $\ell(f) = \sum_{j=1}^D f(s_j)/u$ for a threshold $u > 0$, where $\kappa_\ell(K) = D/u$ in the likelihood (12). Although this approach seems to perform well for Brown–Resnick processes, it would be inefficient for elliptical Pareto processes due to the singularities in $\Lambda_{+,s}$, just like the uncensored likelihood L_1 ; see Section 5. [Engelke et al. \(2015\)](#) further considered the use of extremal increments, corresponding to $\ell(f) = f(s_0)/u_0$ for fixed $s_0 \in K$ and threshold $u_0 > 0$, but this approach has the same disadvantages in the case of elliptical Pareto processes.

3.2 Densities and partial derivatives of the exponent function for extremal- t processes

We derive the density $\lambda_{+,s}$ of the finite-dimensional exponent measure $\Lambda_{+,s}$ and the partial derivatives V_{I_k} through calculations similar to those of [Wadsworth and Tawn \(2014\)](#) for Brown–Resnick processes. A complication for extremal- t processes arises from the singularities of $\Lambda_{+,s}$ on the boundary of $\mathbb{R}_+^D \setminus \{0\}$. To resolve this, we observe that the extremal- t process Z^* in (10) arises as the pointwise maximum of a Poisson process with points $\tilde{P}_i = m_\alpha T_\alpha(W_i)/Q_i$, where $T_a(x) = \text{sign}(x)|x|^a$ for $a > 0$. The truncation of W_i at zero in (10) is irrelevant because Z^* is constituted from pointwise maxima that are positive almost surely. If the point process of the \tilde{P}_i has intensity measure Λ , the unique measure Λ_+ of the extremal- t process is obtained by projecting the negative values to zero. Therefore, we first calculate the intensity $\lambda_s(y)$ for $y \in \mathbb{R}^D \setminus \{0\}$. To derive $\lambda_{+,s}(y)$ when some components of y are zero, say $y = (\tilde{y}, 0)$ with $\tilde{y} > 0$, we can integrate λ_s over all negative values of the zero-components in y such that $\lambda_{+,s}(y) = \int_{-\infty}^0 \lambda_s(\tilde{y}, z) dz$.

[Ribatet \(2013\)](#) gives the density λ_s of Λ_s ,

$$\begin{aligned} \lambda_s(y) &= \alpha^{1-D} \pi^{(1-D)/2} |\Sigma_s|^{-1/2} \Gamma\{(\alpha+1)/2\}^{-1} \Gamma\{(\alpha+D)/2\} \\ &\quad \times \prod_{j=1}^D |y_j|^{1/\alpha-1} \{T_{1/\alpha}(y)' \Sigma_s^{-1} T_{1/\alpha}(y)\}^{-(\alpha+D)/2}, \quad y \in \mathbb{R}^D, \end{aligned}$$

where $\Sigma_s = \{\varrho(s_{j_1}, s_{j_2})\}_{1 \leq j_1, j_2 \leq D}$ denotes the finite-dimensional correlation matrix stemming from the correlation function ϱ of the extremal- t dependence structure relative to the sites $s = (s_1, \dots, s_D)$. The density $\lambda_{+,s}$ of $\Lambda_{+,s}$ on $(0, \infty)^D$ equals λ_s . The partial derivatives V_{I_k} of the exponent function V are calculated by integrating λ_s with respect to the components in the set complementary to I_k . The integration is carried out using conditional intensities. Given a collection $s_0 = (s_{0,1}, \dots, s_{0,d})$ of d conditioning locations

with values y_0 , the conditional intensity $\lambda_{s|s_0,y_0}(y) = \lambda_{(s,s_0)}(y, y_0)/\lambda_{s_0}(y_0)$ equals (Ribatet, 2013)

$$\lambda_{s|s_0,y_0}(y) = \alpha^{-D} \pi^{-D/2} (d + \alpha)^{-D/2} |\tilde{\Sigma}|^{-1/2} \Gamma\{(\alpha + d)/2\}^{-1} \Gamma\{(\alpha + D + d)/2\} \\ \times \prod_{j=1}^D |y_j|^{1/\alpha-1} \left[1 + \frac{\{T_{1/\alpha}(y) - \tilde{\mu}\}' \tilde{\Sigma}^{-1} \{T_{1/\alpha}(y) - \tilde{\mu}\}}{d + \alpha} \right]^{-(\alpha+D+d)/2}, \quad (13)$$

with

$$\tilde{\mu} = \Sigma_{s:s_0} \Sigma_{s_0}^{-1} T_{1/\alpha}(y_0), \quad \tilde{\Sigma} = \frac{T_{1/\alpha}(y_0)' \Sigma_{s_0}^{-1} T_{1/\alpha}(y_0)}{d + \alpha} (\Sigma_s - \Sigma_{s:s_0} \Sigma_{s_0}^{-1} \Sigma_{s_0:s}),$$

where $\Sigma_{s:s_0}$ denotes the matrix of covariances between the random vectors corresponding to the location vectors s and s_0 . Expression (13) is the density of a random vector $T_\alpha(X)$, where X follows a D -dimensional t distribution with $d + \alpha$ degrees of freedom and parameters $\tilde{\mu}$ and $\tilde{\Sigma}$.

Without loss of generality, we consider the partial derivative $V_{1:d}(y)$ of V with respect to the indices 1 to d such that $I_k = \{1, \dots, d\}$, obtained by calculating the integral of $\lambda_{s_{(d+1):D}|s_{1:d},y_{1:d}}(y_{(d+1):D})$ and by multiplying the resulting expression by $\lambda_{s_{1:d}}(y_{1:d})$. The required integral of the conditional density is $t_{d+\alpha}(y_{(d+1):D}^{1/\alpha}; \tilde{\mu}, \tilde{\Sigma})$. We get

$$-V_{1:d}(y) = t_{d+\alpha} \left(y_{(d+1):D}^{1/\alpha}; \tilde{\mu}, \tilde{\Sigma} \right) \alpha^{1-d} \pi^{(1-d)/2} |\Sigma_{1:d}|^{-1/2} \Gamma\{(\alpha + 1)/2\}^{-1} \\ \times \Gamma\{(\alpha + d)/2\} \left(\prod_{j=1}^d |y_j| \right)^{1/\alpha-1} \left\{ (y'_{1:d})^{1/\alpha} \Sigma_{1:d}^{-1} y_{1:d}^{1/\alpha} \right\}^{-(\alpha+d)/2}, \quad (14)$$

with $\tilde{\mu} = \Sigma_{(d+1):D,1:d} \Sigma_{1:d}^{-1} y_{1:d}^{1/\alpha}$ and $\tilde{\Sigma} = (d + \alpha)^{-1} (y'_{1:d})^{1/\alpha} \Sigma_{1:d}^{-1} y_{1:d}^{1/\alpha} (\Sigma_{(d+1):D} - \Sigma_{(d+1):D,1:d} \Sigma_{1:d}^{-1} \Sigma_{1:d,(d+1):D})$. Equation (14) also gives the densities $\lambda_{+,s}$ for a point y on the boundary of $\mathbb{R}_+^D \setminus \{0\}$: if $y_{1:d} > 0$ and $y_{(d+1):D} = 0$, then the density on the corresponding subset of $\mathbb{R}_+^D \setminus \{0\}$ is $-V_{1:d}(y)$, see Coles and Tawn (1991, §3.1).

3.3 Maximum likelihood inference

Numerical maximization of L_1 or L_2 yields the maximum likelihood estimate $\hat{\psi}$ for the vector of parameters ψ of an elliptical Pareto process. Assuming that the data come from the limiting model, standard regularity conditions ensure consistency and asymptotic normality of $\hat{\psi}$, with an asymptotic covariance matrix that equals the inverse Fisher information matrix (van der Vaart, 2000, Ch. 5). In practice, the asymptotic covariance matrix can be estimated by the Hessian matrix of the negated log-likelihood evaluated at $\hat{\psi}$.

A practical inconvenience for maximum likelihood inference based on L_1 or L_2 is the need to calculate the t probabilities in (11) and (14). They can be calculated using Monte Carlo approximations (Genz and Bretz, 2009), but the use of full likelihood inference might be too slow if $D > 50$. In larger dimensions, one could partition the sample sites into moderately large groups and use a composite likelihood based on the full likelihood contribution from each group.

4 Exact simulation procedures

We now describe exact finite-dimensional simulation procedures for extremal- t and elliptical ℓ -Pareto processes. Due to the elliptical structure of the points $P_i^{1/\alpha}$ from the point process $\{P_i\}$ in (10), an equivalent representation of the finite-dimensional projection of an extremal- t process relative to D sites $s = (s_1, \dots, s_D)$ is obtained by setting

$$P_{s,i} = \{E(U_{1,1})_+^\alpha\}^{-1} (A_s U_i)_+^\alpha / Q_i, \quad (15)$$




Fig. 1: Left: simulation (black line) from an extremal- t process with $\alpha = 1$ and $\varrho(h) = \exp(-\|h\|)$. The grey lines show the points $P_{s,i}$ in the spectral decomposition (15). Right: independent simulations from the corresponding elliptical ℓ -Pareto process with $\ell(f) = \sup_{s \in [0,5]} f(s)$ are given by the points $P_{s,i}$ with $\ell(P_{s,i}) \geq 1$.

with A_s the Cholesky root of the correlation matrix $\Sigma_s = A_s A_s'$ and $U_i = (U_{i,1}, \dots, U_{i,D})'$ independent and identically distributed copies of a vector U uniformly distributed on the Euclidean unit sphere (Opitz, 2013, Theorem 3.2). This allows exact simulation of both max-stable and Pareto processes due to the boundedness $\|(A_s U_i)_+^\alpha\|_\infty \leq 1$.

In practice, max-stable processes are simulated using only a finite number of $P_{s,i}$; see Fig. 1. When a finite boundary $b < \infty$ exists for the components of $Q_i P_{s,i}$ such that $\text{pr}\{\max_{i=1,2,\dots} Q_i P_{s,i}(s_j) \leq b\} = 1$ ($j = 1, \dots, D$), exact simulation of Z_s^* can be achieved from a finite number of points $P_{s,i}$ (Schlather, 2002, Theorem 4). Since the components of $\{E(U_{1,1})_+^\alpha\}^{-1} (A_s U_i)_+^\alpha$ in (15) are always bounded by $b = \{E(U_{1,1})_+^\alpha\}^{-1}$, exact simulation of extremal- t processes is possible. For $i \geq 1$, $\|P_{s,i}\|_\infty \leq b/Q_i$ with an increasing sequence $\{Q_i\}$. If $\|\max_{i=1,\dots,\tau_b} P_{s,i}\|_\infty \geq b/Q_{\tau_b}$ for some $\tau_b > 1$, then the points $P_{s,i}$ for $i > \tau_b$ cannot contribute to the maximum in (10) and we have $Z_s^* = \max_{i=1,\dots,\tau_b} P_{s,i}$. Two numerical limitations may restrict the applicability of this simulation approach: first, standard algorithms for determining the Cholesky root A_s of Σ_s require $O(D^3)$ basic operations; second, b may be large if α or D are large, requiring the simulation of a very large number of points $P_{s,i}$. More precisely,

$$b = 2\pi^{1/2} \frac{\Gamma\{(D + \alpha)/2\}}{\Gamma\{(\alpha + 1)/2\}\Gamma(D/2)} \approx 2^{1-\alpha/2} \pi^{1/2} \frac{(D + \alpha - 2)^{\alpha/2}}{\Gamma\{(\alpha + 1)/2\}}, \quad D \rightarrow \infty,$$

using Stirling's formula. In certain situations, notably when D indexes a fine spatial grid of points, these limitations are too restrictive. Then the conventional approach for approximate simulation can be used. Since the tails of $W(s_j)_+^\alpha$ become heavier when α increases, the approximation error in the simulated max-stable process also increases.

The simulation of the points $P_{s,i}$ in (15) yields an algorithm for the simulation of elliptical ℓ -Pareto processes: as mentioned in §2.3, the points $P_{s,i}$ with $\ell(P_{s,i}) \geq 1$ are independent realisations from the standard ℓ -Pareto process; see Fig. 1. Moreover, for $u_0 > 0$, the homogeneity of $\Lambda_{+,s}$ implies that the points $u_0^{-1} P_{s,i}$ with $\ell(P_{s,i}) \geq u_0$ are also realisations from the standard ℓ -standard process. The existence of the upper bound b allows us to simulate all the points $P_{s,i}$ in a set $A = ([0, u]^D)^C \subset \mathbb{R}_+^D$ for $u > 0$. Since the set $\ell(y) \geq u_0$ is a subset of A for suitably chosen u , we can obtain exact simulations from every elliptical ℓ -Pareto process.

Instead of simulating the points of the Poisson process, it is possible to use an acceptance-rejection algorithm to generate realisations of $Y_\ell^*(s)$ without dealing with a random number of realisations. First, we consider the simulation of a standard ℓ -Pareto process with $\ell(f) = \max_{j=1,\dots,D} f(s_j)$. From the previous

paragraph, the points $b^{-1}P_{s,i}$ with $\ell(P_{s,i}) \geq b$ are realisations from the standard ℓ -Pareto process. The condition $\ell(P_{s,i}) \geq b$ implies $1/Q_i \geq 1$, hence only the points $P_{s,i}$ with $1/Q_i \geq 1$ need to be simulated. Since the points $1/Q_i$ satisfying $1/Q_i \geq 1$ are distributed according to the standard Pareto distribution, any vector $R(A_s U)_+^\alpha$ with U independent of a standard Pareto distributed random variable R is a standard ℓ -Pareto process if $\ell\{R(A_s U)_+^\alpha\} \geq 1$. When ℓ is different from the componentwise maximum, we proceed as before and fix $u_0 > 0$ such that $\max_{j=1,\dots,D} f(s_j) \geq 1$ whenever $\ell(f) \geq u_0$. Then the vector $u_0^{-1}R(A_s U)_+^\alpha$, given that $\ell\{R(A_s U)_+^\alpha\} \geq u_0$, is a realisation of $Y_\ell^*(s)$. We can get a sample of $Y_\ell^*(s)$ by repeatedly simulating random vectors $u_0^{-1}R(A_s U)_+^\alpha$ and retaining only those vectors fulfilling the condition $\ell\{R(A_s U)_+^\alpha\} \geq u_0$. To minimise the frequency of rejections, u_0 should be chosen as small as possible.

Whereas conditioning Λ_+ on exceedances of $\ell(f)$ over unity yields the distribution of the ℓ -Pareto process, one might instead be interested in the conditional distribution when values $y_0 > 0$ for a collection of sites $s_0 = (s_{0,1}, \dots, s_{0,d})$ are fixed. The finite-dimensional conditional distribution for the sites $s = (s_1, \dots, s_D)$, disjoint with s_0 , has density (13). The conditional process defined on $K \setminus \{s_0\}$ corresponds to a transformed t process that can easily be simulated.

5 Simulation study

We used simulation to investigate the efficiency of the full likelihoods L_1 and L_2 for estimating the parameters of elliptical Pareto processes. For comparison, we also report results from the pairwise censored likelihood approach based on all pairs, which represents a standard approach to fitting max-stable models. Based on the exact simulation procedure introduced in Section 4, we generated samples of 1000 elliptical ℓ -Pareto processes with $\ell(f) = \max_{s \in K} f(s)$ at 16 locations given by $K = \{(s_1/3, s_2/3)\}_{s_1, s_2 \in \{0,1,2,3\}}$. We chose the stable correlation function $\varrho(h) = \exp\{-(\|h\|/\lambda)^\kappa\}$, where h is the lag vector between two locations, $\lambda > 0$ is a range parameter and $0 < \kappa \leq 2$ is a smoothness parameter. Different combinations of values for the degrees of freedom $\alpha > 0$ and for λ and κ were considered, covering small to strong dependence with different degrees of smoothness. We estimated the vector of parameters $\psi = (\log \lambda, \kappa, \alpha)$ using the three approaches, each of them based on marginal thresholds equal to the 95% quantiles. The mean squared error of $\hat{\psi}$ can be decomposed into a sum of bias and variance terms: $\text{MSE}(\hat{\psi}) = \|\text{E}(\hat{\psi}) - \psi\|^2 + \text{tr}(V)$, where V is the covariance matrix of $\hat{\psi}$. For each parameter configuration, estimates for 1000 samples were calculated to obtain the bias and covariance matrix of each estimator. Table 1 shows the relative efficiency of the three estimators, here defined as the ratio of the trace of their covariance matrices. Unreported results showed that all estimators have only little or no bias. Throughout, the full uncensored likelihood estimator L_1 was found to be more efficient than the full censored estimator L_2 , owing to the loss of information from censoring. The difference is larger when dependence is weak, that is, when more components are censored in exceedances, though more exceedances are observed. The L_2 -based estimator is more efficient than its pairwise equivalent, and efficiency gains are larger for smooth processes with weak dependence. Overall, the relative reduction in variance is around 60%. Other simulations indicated that the efficiency improvements of the full likelihood over the pairwise likelihood become larger when the number of sites is larger: in a similar estimation framework, we found a reduction of variance of around 35% for nine locations and of around 10% for four locations.

To investigate the impact of the convergence to a limiting elliptical Pareto process, we further simulated samples $X_{s,1}, \dots, X_{s,1000}$ of t processes with α degrees of freedom on the same grid as before. Marginal distributions of $X_{s,k}$ were transformed to the standard Pareto scale by the transformation $X_{s,k}^* = 1/\{1 - t_\alpha(X_{s,k})\}$ ($k = 1, \dots, 1000$), where t_α denotes the cumulative distribution function of a univariate t variable with α degrees of freedom. We fitted elliptical ℓ -Pareto processes to threshold exceedances of the simulated $X_{s,k}^*$ over the marginal 95%, 98% and 99% quantiles using the two full likelihoods L_1 and L_2 and the pairwise censored likelihood. We then considered the bias, variance and mean squared error of

Table 1: Relative efficiency (in %) of full and pairwise likelihood estimators for the parameters of elliptical Pareto processes with the stable correlation function. For each combination of κ and α , three values of λ were chosen to give pairwise extremal coefficients $\theta \in \{1.2, 1.4, 1.6\}$ at distance 0.5. Each cell gives the ratio of the covariance matrix traces for the uncensored and censored full likelihood estimators and for the censored full and pairwise likelihood estimators, separated by /.

θ/α	$\kappa = 0.5$				$\kappa = 1$				$\kappa = 1.5$			
	1	2	5	10	1	2	5	10	1	2	5	10
1.2	61/62	51/59	48/60	45/58	50/45	39/47	31/42	36/41	43/29	27/28	21/27	27/21
1.4	48/55	29/51	16/52	15/53	41/43	22/37	8/35	7/44	34/27	18/23	5/19	5/21
1.6	30/50	14/39	5/45	4/39	34/36	13/35	3/31	2/37	34/25	16/22	3/19	1/25

these estimators. As opposed to the simulations discussed in the previous paragraph, the elliptical Pareto model is only valid asymptotically and so the estimators are biased. Table 2 reports the bias and the empirical covariance matrix trace of $\hat{\psi}$ calculated from 1000 estimates and the relative efficiencies of the two full likelihood and the two censored likelihood estimators, here defined as the ratio of their mean squared error. For all thresholds, the uncensored estimator L_1 has the largest mean squared error because of its very large bias. The two censored estimators have small bias when $\alpha = 1$, but the bias increases as α increases. This may be explained by the slower convergence to the limiting dependence structure for larger α . The bias is reduced by increasing the thresholds such that the exceedance distribution is closer to the asymptotic model; variances increase accordingly. Variances are always smaller for the full likelihood estimator than for the pairwise one, but the bias of the full likelihood estimator is often larger. In terms of mean squared error for the 95% threshold, the full likelihood estimator outperforms the pairwise one for $\alpha < 6$, but not otherwise. The bias of the full likelihood estimator decreases for higher thresholds, and the full likelihood estimator generally has a smaller mean squared error than the pairwise estimator owing to its smaller variance. Hence for large values of α , very high thresholds are needed for the full likelihood estimator to outperform the pairwise estimator in terms of mean squared error.

The results of these simulations suggest that censored approaches are the best in practice when the model is misspecified. Moreover, full likelihood inference improves estimation efficiency when the distribution of extremes is close to the limiting model, but the pairwise approach appears more robust to certain kinds of model misspecification.

6 Application

We illustrate the use of ℓ -Pareto processes for modelling precipitation extremes in the region of Zürich, Switzerland. Daily cumulative rainfall data at 44 locations were provided by MétéoSuisse; see Fig. 2. Elevations vary from 327 to 718 m for these stations. Our analysis is based on summer data recorded from 1 June to 31 August for the years 1962–2012. A preliminary study showed no signs of non-stationarity in the time series and only weak day-to-day dependence in exceedances over the 95% quantiles, leading us to model the daily data as independent and identically distributed. The data seem coherent with the assumption of asymptotic dependence, which suggests modelling threshold exceedances using Pareto processes; see the Supplementary Material. We selected 25 stations for the fit of the spatial model, see Fig. 2; the other stations are kept for validation. First, we fitted a spatially varying model for the univariate marginal distributions (8) over marginal thresholds taken to be the 95% percentiles at each of the 25 stations. We used a Bayesian hierarchical model to capture spatial random effects in $\mu(s)$ and $\sigma(s)$, similar to the latent variable model of Davison et al. (2012), see the Supplementary Material. The shape parameter ξ was assumed to be constant over the region. Its estimate and 95% credible interval is 0.11 (0.08, 0.14) corresponding to heavy-tailed marginal distributions. We then transformed the original data at each location to the standard Pareto scale by

Table 2: Estimation of elliptical ℓ -Pareto processes based on exceedances of t processes for the stable correlation function with $\kappa = 1$. For each α in $1, \dots, 10$, the values of λ were chosen to yield the pairwise extremal coefficient $\theta = 1.4$ at distance 0.5. For each of the thresholds chosen at the 95%, 98% and 99% quantiles and for each of the three estimators based on L_1 , L_2 or the pairwise censored distributions, the bias/variance terms of $\hat{\psi}$ are reported. For each threshold, the last row reports the ratio δ of the mean squared error for the censored and uncensored full likelihood estimators and for the censored full and pairwise likelihood estimators, separated by /. All numbers have been multiplied by 100.

	α	1	2	3	4	5	6	7	8	9	10
95%	L_1	303/9	241/6	223/5	215/6	210/7	208/8	205/10	203/10	202/10	201/11
	L_2	2/3	12/6	34/9	53/11	69/14	82/16	92/21	100/21	108/23	115/25
	pw	3/7	15/14	28/20	43/27	55/32	67/37	71/42	77/44	79/51	82/48
	δ	0/44	1/45	4/72	8/87	14/98	19/101	24/114	29/117	34/124	38/135
98%	L_1	257/16	187/10	171/8	170/9	172/12	173/13	176/18	179/20	182/24	186/26
	L_2	7/7	2/13	13/22	30/28	41/34	57/44	68/53	79/59	90/70	91/70
	pw	6/18	16/38	26/55	42/70	55/86	64/98	66/107	71/103	76/114	76/108
	δ	1/42	4/32	8/38	13/42	17/44	25/55	30/66	35/78	42/88	41/93
99%	L_1	240/25	160/14	141/12	142/14	145/18	151/25	157/30	162/33	169/43	174/47
	L_2	12/15	14/26	6/40	15/57	23/75	43/97	50/114	62/115	79/171	87/205
	pw	8/39	19/70	30/111	48/140	54/175	72/208	75/234	80/220	84/277	80/237
	δ	3/41	10/38	19/33	27/36	35/39	46/44	50/48	52/54	71/67	80/93

using the fitted marginal distributions above the thresholds and the empirical distributions below them.

In a second step, we used the likelihood L_2 to model the dependence in the standardized data with ℓ -Pareto processes for $\ell(f) = \max_{j=1, \dots, 25} f(s_j)/20$; 20 is the 95% quantile of a standard Pareto distribution. We fitted an elliptical Pareto process with a stable correlation function $\varrho(h) = \exp\{- (\|h\|/\lambda)^\kappa\}$ and shape $\alpha > 0$. For comparison, we also fitted the Pareto model of a Brown–Resnick process with $\gamma(h) = (\|h\|/\lambda)^\kappa$, corresponding to the limiting model when $\alpha \rightarrow \infty$; see §2.3. We used the Akaike information criterion to select the best model: for the elliptical Pareto model, it is 864 less than that of the Brown–Resnick model. The parameter estimates and standard errors for the elliptical Pareto model are $\hat{\lambda} = 520$ (73) km, $\hat{\kappa} = 0.63$ (0.02) and $\hat{\alpha} = 6.3$ (0.4), which yields a process with realisations that are continuous but not differentiable. These results are consistent with those found by Davison et al. (2012) who identified an extremal- t model with $\hat{\alpha} = 5.5$ (2.1) as the best max-stable model for yearly maxima of daily cumulative rainfall on the same region and found a similar estimate for the smoothness parameter of the correlation function.

We validated the accuracy of our model for modelling spatial extremes using the data from the other 19 stations. The right panel of Fig. 2 shows estimates of pairwise extremal coefficients related to these stations. The estimates for validation stations are only slightly more variable. Overall, the extremal dependence for validation data is adequately represented by the model.

Using the conditional distribution (13), we considered simulation at the 19 validation stations conditional on the values observed at the other stations when at least one of the 25 components exceeded its marginal 95% threshold. To compare the observed extreme values at the 19 validation locations with those predicted by the model, we simulated 30000 conditional realisations for each day, and we measured the proportion of true values falling between the 2.5% and 97.5% quantiles of the simulations. Over the 19 locations and the 986 days with at least one exceedance at the other 25 locations, approximately 85% of the observed values were in the 95% simulated intervals. We do not attain 95% corresponding to a perfect prediction, which can at least partly be explained by the fact that the simulations used conditional data below the thresholds and also ignored the uncertainty of the estimates of the fitted model; simulations taking these two aspects into account should be more variable and thus have a higher coverage probability.

Finally, we illustrate the ability of the Pareto process approach to easily simulate conditional rainfall given observed extremes at some set of locations. For a particular day and given the observed data at the 25

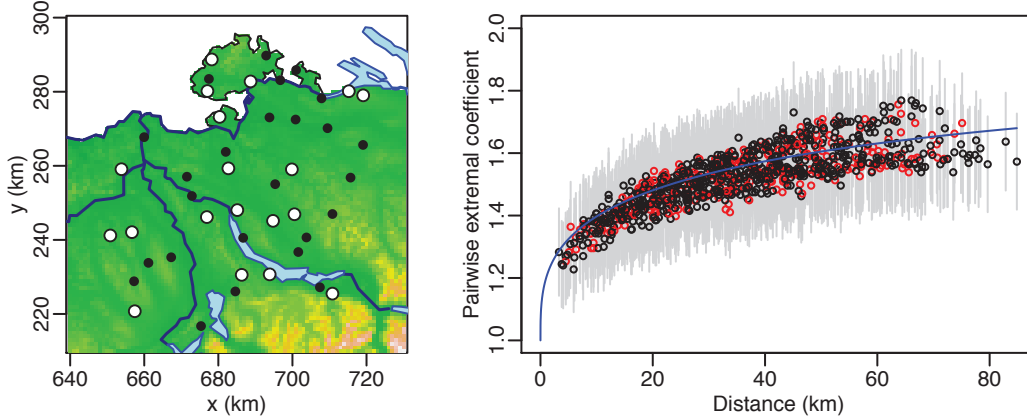


Fig. 2: Modelling extreme precipitation around Zürich. Left: region of Zürich. Black dots correspond to the stations used for the fit, white dots to the stations used for validation. Right: empirical pairwise extremal coefficients (with 95% confidence intervals in grey) for the data used to fit the model (in red) and for the validation data (in black), with the fitted extremal coefficient curve for the best elliptical Pareto model in blue.

locations used for the fit, we simulated conditional values of rainfall over the region using the transformed t process characterized in (13); the left panel of Fig. 3 shows the mean of these simulated rainfall fields and the right panel its standard deviation.

7 Discussion

Pareto models are appealing because they generalize peaks-over-threshold stability to the spatial context and they appropriately exploit the regularity of the exponent measure in pseudo-polar coordinates for extrapolation on extremes. In this paper we have introduced inference and simulation procedures for the elliptical Pareto model. Numerical results suggest that the censored approach based on marginal thresholds gives more reliable estimates for this process, but the choice of a sufficiently high threshold is crucial to guarantee that the limiting model provides an adequate approximation to the tail of the data. The pairwise likelihood method was found to be more robust to model misspecification.

We modelled rainfall extremes using a two-step approach that combines a latent variable model for the margins with the fit of an elliptical Pareto process. The Bayesian approach enables the modelling of complex spatial trends in univariate marginal distributions and provides a more flexible alternative to regression models. Ideally we would like to use a full Bayesian model based on generalized elliptical Pareto processes, but it is prone to computational difficulties. In our application, one evaluation of the censored likelihood function takes several minutes, preventing the use of this likelihood in Markov chain Monte Carlo algorithms.

The distribution of an elliptical Pareto process is fully specified by its exponent measure Λ_+ which depends on a correlation function and a shape parameter. In practice, data are observed on a finite set of sites s , so inference focuses on the estimation of the finite-dimensional exponent measure $\Lambda_{+,s}$, based on the choice of a risk functional that can be evaluated at s . Estimation is based on finite-dimensional distributions from which parameters can be identified, which induces the unique dependence structure on $C_+(K) \setminus \{0\}$. It thus enables projection of extremes on $C_+(K) \setminus \{0\}$, eventually using different risk functionals than those used for the estimation.

The properties of the maximum likelihood estimators for elliptical ℓ -Pareto processes under the limiting model are well-known but further work to investigate the theoretical properties of such estimators under domain of attraction assumptions would be valuable.




Fig. 3: Conditional simulation of extreme rainfall over the region of Zürich for 1 June 1962. Left: mean of the conditional simulation of daily cumulative rainfall (in mm). Black dots correspond to the locations on which we conditioned, white dots to the data used for validation with their observed values (written in parentheses under the thresholds). Right: standard deviation for the conditional simulation.

Acknowledgments

We thank A. C. Davison, R. Huser and two anonymous referees for their comments on the manuscript. We thank the Swiss National Science Foundation for financial support. T. Opitz is grateful for financial support from two projects: McSim, funded by the French National Research Agency, and MIRACCLE, funded by the French Ministry for ecology.

Supplementary material

Supplementary material available at Biometrika online includes details on the latent variable model used in Section 6 and on a new estimator for extremal coefficients, and discussion on asymptotic independence for the rainfall data.

References

- Beirlant, J., Goegebeur, Y., Segers, J., and Teugels, J. *Statistics of Extremes: Theory and Applications*. Wiley, New York, 2004.
- Coles, S. G. and Tawn, J. A. Modelling extreme multivariate events. *Journal of the Royal Statistical Society: Series B*, 53: 377–392, 1991.
- Daley, D. J. and Vere-Jones, D. *An Introduction to the Theory of Point Processes: General Theory and Structure*, volume 2. Springer, New York, second edition, 2007.
- Davison, A. C., Padoan, S. A., and Ribatet, M. Statistical modelling of spatial extremes (with Discussion). *Statistical Science*, 27: 161–186, 2012.
- de Haan, L. and Ferreira, A. *Extreme Value Theory: An Introduction*. Springer, New York, 2006.
- Dombry, C. and Ribatet, M. Functional regular variations, Pareto processes and peaks over threshold. *Statistics and Its Interface*, 8: 9–17, 2015.

- Engelke, S., Malinowski, A., Kabluchko, Z., and Schlather, M. Estimation of Hüsler–Reiss distributions and Brown–Resnick processes. *Journal of the Royal Statistical Society: Series B*, 77: 239–265, 2015.
- Ferreira, A. and de Haan, L. The generalized Pareto process; with a view towards application and simulation. *Bernoulli*, 20: 1717–1737, 2014.
- Genz, A. and Bretz, F. *Computation of Multivariate Normal and t Probabilities*. Springer, New York, 2009.
- Giné, E., Hahn, M. G., and Vatan, P. Max-infinitely divisible and max-stable sample continuous processes. *Probability Theory and Related Fields*, 87: 139–165, 1990.
- Kabluchko, Z., Schlather, M., and de Haan, L. Stationary max-stable fields associated to negative definite functions. *The Annals of Probability*, 37: 2042–2065, 2009.
- Nikoloulopoulos, A., Joe, H., and Li, H. Extreme value properties of multivariate t copulas. *Extremes*, 12: 129–148, 2009.
- Opitz, T. Extremal t processes: Elliptical domain of attraction and a spectral representation. *Journal of Multivariate Analysis*, 122: 409–413, 2013.
- Padoan, S. A., Ribatet, M., and Sisson, S. A. Likelihood-based inference for max-stable processes. *Journal of the American Statistical Association*, 105: 263–277, 2010.
- Ribatet, M. Spatial extremes: Max-stable processes at work. *Journal de la Société Française de Statistique*, 154: 156–177, 2013.
- Rootzén, H. and Tajvidi, N. Multivariate generalized Pareto distributions. *Bernoulli*, 12: 917–930, 2006.
- Schlather, M. Models for stationary max-stable random fields. *Extremes*, 5: 33–44, 2002.
- van der Vaart, A. W. *Asymptotic Statistics*. Cambridge University Press, Cambridge, UK, 2000.
- Wadsworth, J. L. and Tawn, J. A. Efficient inference for spatial extreme value processes associated to log-Gaussian random functions. *Biometrika*, 101: 1–15, 2014.