Fast and flexible inference for spatial extremes

Peng Zhong, Scott A. Sisson, **Boris Béranger**

CMStatistics 2025,

Birkbeck, University of London, 13 December 2025

Introduction

Broad context

• Interest in the extremes of a stochastic process $X(s), s \in \mathcal{S}$.
E.g. $X(\cdot)$ *measures the amount rainfall at locations over Florida*

• **Goal:** Model the dependence structure of the extremes of $X(s)$

• What characterises an extreme event? → Tailored approach

• Focus on asymptotic dependent processes: max-stable and r-Pareto.

In this talk

1. Establish theoretical conditions for max-stable and r-Pareto models to have a continuous exponent measure
2. Derive two new max-stable and r-Pareto models
3. Provide a fast inference methodology using spectral likelihoods

Modelling framework for max-stable and r-Pareto processes

Theoretical & methodological results

Simulation experiements

Analysis of extreme rainfall over Florida

Max-stable processes

Definition (Schlather, 2002)

A max-stable process with unit Fréchet margins can be characterized as

$$Z(\mathbf{s}) = \sup_{i=1}^{\infty} R_i W_i(\mathbf{s}), \ \mathbf{s} \in \mathcal{S},$$

where $R_1, R_2, \ldots$, are the points of a PPP on $(0, \infty)$ and $W_1(\mathbf{s}), W_2(\mathbf{s}), \ldots$, are independent copies of a stochastic processes $W(\mathbf{s})$ on $\mathcal{S}$ with unit mean.

The distribution function can be expressed as

$$G(\boldsymbol{x}) = \exp\left\{-V(\boldsymbol{x})\right\} = \exp\left\{-\kappa([0, \boldsymbol{x}]^c)\right\}, \quad \mathbf{x} \in \Omega = \mathbb{R}_+^D \setminus \{\mathbf{0}\},$$

where

$$\kappa\left([\mathbf{0}, \boldsymbol{x}]^c\right) = \int_0^\infty 1 - \Pr(\boldsymbol{W} \in [\mathbf{0}, \boldsymbol{x}r])\mathrm{d}r,$$

denotes the exponent measure restricted onto $\mathbb{R}_+^D$, and $\boldsymbol{W} = (W(\mathbf{s}_1), \ldots, W(\mathbf{s}_D))^\top$.

Max-stable processes

Let $B_D = \{1, \dots, D\}$ and $B_k = \{b_1, \dots, b_k\} \subset B_D$, where $b_1 < \cdots < b_k$.

Let $\Omega_{B_k} = \{\boldsymbol{x} \in \Omega : x_j = 0 \text{ if } j \notin B_k\}$ such that:

- $\partial\Omega = \{\Omega_{B_k}, \forall B_k$ and $k = 1, \dots, D-1\}$ represents the boundaries of Ω,
- $\Omega^\circ = \Omega \setminus \partial\Omega$ denotes the Interior of Ω.

Important

Depending on the choice of W, the exponent measure κ **can put mass on both $\partial\Omega$ and Ω°** with the intensity function on each subspace Ω_{B_k}

$$\lim_{x_i \to 0, i \notin B_k} -V_{B_k}(\boldsymbol{x}), \quad V_{B_k} = \frac{\partial^k V}{\partial x_{b_1} \dots \partial x_{b_k}}.$$

On Ω°, it can be expressed as $\kappa(\boldsymbol{x}) = -V_{B_D}(\boldsymbol{x})$, where the function κ is referred to as the intensity function of the max-stable process.

Max-stable processes - Inference

Full likelihood: intractable!

Composite likelihood: Popular but still limited.

Stephenson-Tawn likelihood: Can be biased, moderate dimensions.

Spectral likelihood (Coles & Tawn, 1991)

If data $\in$ MDA($\boldsymbol{Z}$) then can be approximately treated as points of a PPP with measure $\kappa(\cdot)$. For a model with parameter $\boldsymbol{\theta}$, the log-likelihood is

$$\ell_A(\boldsymbol{\theta}; \mathbf{x}_1, \ldots, \mathbf{x}_n) \propto \sum_{i \in \{m : \|\mathbf{x}_m\|_1 > u\}} \log \kappa\left(\boldsymbol{x}_i; \boldsymbol{\theta}\right).$$

for some high enough threshold u.

This requires convergence of:

a) X to the max-stable process Z by taking pointwise maxima.

b) X to the Poisson point process.

The fact that κ can put mass on $\partial\Omega$ hinders the convergence of $X \Longrightarrow$ bias.

r-Pareto processes

Definition (Dombry & Ribatet, 2015)

Assuming the process X with unit Pareto margins satisfying $\lim_{u\to\infty} u \Pr(X/u \in B) = \kappa(B), \forall B \subset C^+(\mathcal{S})$, then the limiting process

$$\tilde{Z}(\mathbf{s}) = \lim_{u\to\infty} \frac{X(\mathbf{s})}{u} | r\left(\{X(\mathbf{s}), \mathbf{s} \in \mathcal{S}\}\right) > u,$$

defines a simple r-Pareto process on $\mathcal{A}_r = \{f \in C^+(\mathcal{S}) : r(f) > 1\}$ with probability measure $\kappa(\cdot \cap \mathcal{A}_r)/\kappa(\mathcal{A}_r)$.

The finite dimensional density is therefore

$$\frac{\kappa(\mathbf{x})}{\kappa(\mathcal{A}_r^D)}, \quad \mathbf{x} \in \mathcal{A}_r^D,$$

where κ is the intensity function and $\mathcal{A}_r^D$ is the set $\mathcal{A}_r$ restricted to D dimensions.

r-Pareto processes - Inference

The log-likelihood is thus

$$\ell_{rP}(\boldsymbol{\theta}; \mathbf{x}_1, \ldots, \mathbf{x}_n) = \sum_{i \in \{m : r(\mathbf{x}_m) > u\}} \log\left(\frac{\kappa(\mathbf{z}_i; \boldsymbol{\theta})}{\mathsf{K}(\mathcal{A}_r; \boldsymbol{\theta})}\right),$$

where $\mathbf{z}_i = \mathbf{x}_i / u$ represent the realizations of the r-Pareto process.

Important

- $\mathsf{K}(\mathcal{A}_r; \boldsymbol{\theta})$ involves integration over $\mathbb{R}_+^D$, $\Longrightarrow$ intractability
- de Fondeville & Davison (2018):
  - ⋆ Simplifications for specific choices of $r(\cdot)$.
  - ⋆ Score matching.
- $r(\mathbf{x}) = \|\mathbf{x}\|_1 \Longrightarrow$ spectral likelihood.
- If the exponent measure K has discontinuities (presence of mass on $\partial\mathcal{A}_r^D$),

$\Longrightarrow$ Inference requires evaluation of $-V_{B_k}(\boldsymbol{x})$.

  - ⋆ Restriction to the Brown-Resnick models

Modelling framework for max-stable and r-Pareto processes

Theoretical & methodological results

Simulation experiements

Analysis of extreme rainfall over Florida

Ensuring continuous exponent measures

Theorem 1 (Zhong, Sisson & Béranger, 2025)

Consider the max-stable process $\{Z(\mathbf{s}), \mathbf{s} \in \mathcal{S}\}$ defined at D locations and assume the partial derivatives of the function V exist.
The intensity function on $\partial\Omega$ is <u>zero almost everywhere</u> if and only if the conditional probability of $\boldsymbol{W}$ satisfies

$$\Pr(\boldsymbol{W}_{\bar{B}_k} = \mathbf{0}_{D-k} \mid \boldsymbol{W}_{B_k} = \boldsymbol{x}_{B_k}) = 0, \; \forall\, k \in \{1, \ldots, D-1\}, \; \boldsymbol{x}_{B_k} > \mathbf{0}_k.$$

Brown-Resnick: $W = \exp\left(\tilde{W} - \frac{\sigma^2}{2}\right)$, with $\tilde{W}$ a centered Gaussian process
$\Longrightarrow$ Condition satisfied

skew extremal-t : $W = \max(\tilde{W}^{\nu}, 0)$ with $\tilde{W}$ a skew-normal process, $\nu > 0$.
$\Longrightarrow$ Condition NOT satisfied

Extending current classes of max-stable and r-Pareto models

Theorem 2 (Zhong, Sisson & Béranger, 2025)

Assume $Y(\mathbf{s})$ is a process whose finite-dimensional distribution is skew-normal with scale matrix Σ.

a) skewed Brown-Resnick: Let $W(\mathbf{s}) = \exp\{Y(\mathbf{s}) - a(\mathbf{s})\}$ with slant parameter $\boldsymbol{\alpha}$, and $a(\mathbf{s}) = \log \mathbb{E}\left[\exp\{Y(\mathbf{s})\}\right]$.

b) truncated extremal-t: Let $W(\mathbf{s}) = \tilde{Y}(\mathbf{s})^{\nu}/a(\mathbf{s})$, with $\nu > 0$, $\tilde{Y}(\mathbf{s}) = Y(\mathbf{s}) | Y(\mathbf{s}) > 0$, $Y(\mathbf{s})$ has unit variances and $a(\mathbf{s}) = \mathbb{E}\left[\tilde{Y}(\mathbf{s})^{\nu}\right]$.

$\Longrightarrow$ Both models have no mass on $\partial\Omega$.

Comments:

- The sBR model has a non-stationary dependence structure.
- The intensity of the truncated extremal-t is somewhat difficult to compute...
- Removal of the mass on $\partial\Omega$ increases the dependence strength

Improved inference for r-Pareto models

Where does the idea come from?

[Dombry, Legrand & Opitz (2024)] Using rejection sampling, one can generate samples from a r-Pareto process with risk functional r_2 from samples of a r-Pareto process associated with risk functional r_1 as long as $Mr_1(\cdot) \geq r_2(\cdot), M > 0$.

Focus: Observations $i \in \{m : r(\mathbf{x}_m) > u\}$

Proposal: use the likelihood of the L_1-Pareto process to make inference about any r-Pareto process with a different risk functional by choosing a high threshold $u > M$.

This particularly applies to L_p norms, $p > 1$, since L_p bounds L_1 for finite p.

$\implies \| \cdot \|_1 \leq D^{1-1/p} \| \cdot \|_p, \ p > 1$

$\implies$ choose $u > D^{1-1/p}, p > 1$, to infer the L_p-Pareto process.

Benefit: Avoids to compute the normalising constant!!

Modelling framework for max-stable and r-Pareto processes

Theoretical & methodological results

Simulation experiements

Analysis of extreme rainfall over Florida

Spectral likelihoods vs score matching

Setup:

- Generate $n = 2,000$ obs from the skewed Brown-Resnick model on a 15×15 grid ($D = 225$).
- Power-law semivariogram $\gamma(h) = (h/\lambda)^{\vartheta}$ with range $\lambda = 5, 10$ and smoothness $\vartheta = 1, 1.5$.
- Skewness represented through spline functions with 2 Gaussian kernel basis functions $(b_1, b_2) = (0, 0), (-1, -2), (-1, 1)$.
- L_1 and L_3 risk functionals.
- An observation is considered extreme when exceeding the 95% empirical quantile of $r(\boldsymbol{X}_1), \ldots, r(\boldsymbol{X}_n)$.
- 300 replicates.

Spectral likelihoods vs score matching

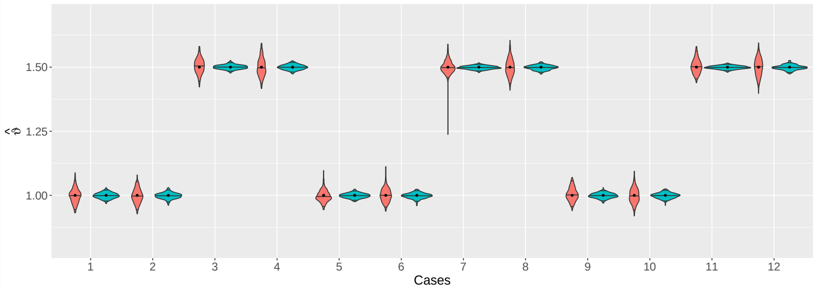

Figure 1: Violin plots for score matching (red) and spectral likelihood (blue) estimates of ϑ for the skewed Brown-Resnick r-Pareto process with L_3 norm risk functional. Black dots indicate the parameter true values.

- The spectral likelihood provides unbiased, low variability estimates.
- The score matching produces unbiased but more variable estimates.

Spectral likelihoods vs score matching

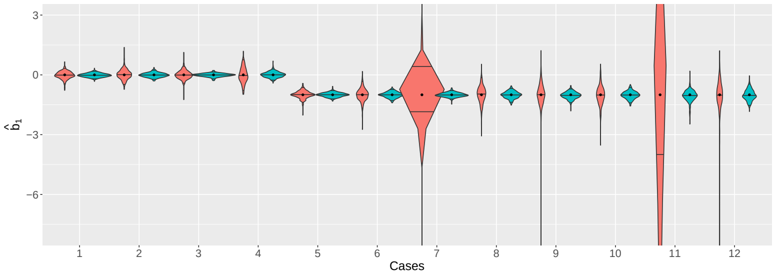

Figure 2: Violin plots for score matching (red) and spectral likelihood (blue) estimates of b_1 for the skewed Brown-Resnick r-Pareto process with L_3 norm risk functional. Black dots indicate the parameter true values.

- Score matching estimates can become numerically unstable (cases 7–12).
- Spectral likelihood is ~ 5 times faster than the score matching approach (141 versus 704 seconds on average using 3 CPU cores).

Modelling framework for max-stable and r-Pareto processes

Theoretical & methodological results

Simulation experiements

Analysis of extreme rainfall over Florida

Analysis of extreme rainfall over Florida

Data:

- Location: Tampa Bay area, Florida. Regular 2km grid with $4,449$ spatial observations.

- Measurements: radar images recorded at 15 minute intervals between 1995–2019 during the wet season (June–September). Total $n = 139,881$ images.

- Smaller version of the dataset analysed in de Fondeville & Davison (2018).

- Risk functions:

  $\rightarrow$ L_∞ norm: defines extremes events as locally intense rainfall events at any location within the region

  $\rightarrow$ L_1 norm selects events with high cumulative rainfall over the whole region.

Analysis of extreme rainfall over Florida

Modelling:

- Brown-Resnick (BR) and skewed Brown-Resnick (sBR) with anisotropic semivariogram.
- Skewness of sBR expressed using 4 kernels.
- Fitting using score matching and spectral likelihood.

Outcomes:

- Brown-Resnick:

  → Spectral likelihood and score matching provide consistent estimates.

  → Spectral likelihood is 80% (L_1 norm) and 18% (L_∞ norm) faster.

- Brown-Resnick vs skewed Brown-Resnick:

  → AIC favours the skewed Brown-Resnick for both L_1 and L_∞ norms.

Analysis of extreme rainfall over Florida

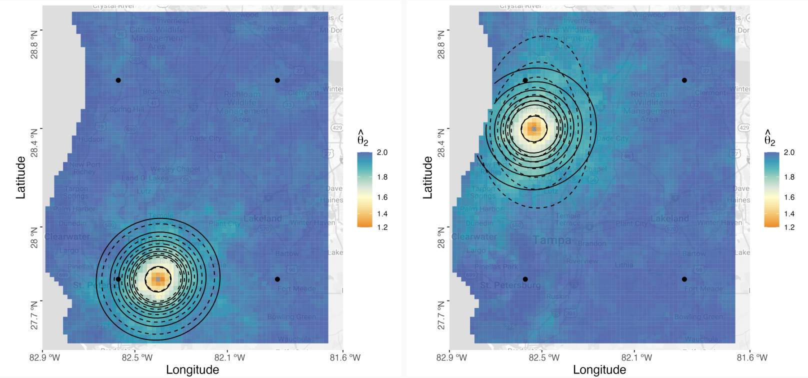

Figure 3: Maps of bivariate empirical extremal coefficients (shading) with respect to two different reference points, and contours of the extremal coefficient of the fitted sBR (dashed line) and BR (solid line) r-Pareto models with L_∞ norm risk functional. Black dots denote the kernel centres used in the sBR model.

Conclusion

- Established condition ensuring the intensity function of a max-stable process only places mass on Ω°;

  $\rightarrow$ No discontinuities in the associated exponent measure;

  $\rightarrow$ Simplifying the evaluation of the density of the r-Pareto process.

- Likelihood-based inference can be successfully implemented via the spectral likelihood.

- Two new models: skewed Brown-Resnick and truncated Extremal-t.

- Not presented: improved rejection sampling algorithm for r-Pareto processes.

THANK YOU

https://arxiv.org/pdf/2407.13958

B.Beranger@unsw.edu.au