

Graphical models for extremes

Sebastian Engelke

University of Geneva, Switzerland

and Adrien S. Hitz

University of Oxford and Materialize.X, Enterprise Laboratory, London, UK

[Read before The Royal Statistical Society at a meeting organized by the Research Section on Wednesday, February 12th, 2020, Professor G. P. Nason in the Chair]

Summary. Conditional independence, graphical models and sparsity are key notions for parsimonious statistical models and for understanding the structural relationships in the data. The theory of multivariate and spatial extremes describes the risk of rare events through asymptotically justified limit models such as max-stable and multivariate Pareto distributions. Statistical modelling in this field has been limited to moderate dimensions so far, partly owing to complicated likelihoods and a lack of understanding of the underlying probabilistic structures. We introduce a general theory of conditional independence for multivariate Pareto distributions that enables the definition of graphical models and sparsity for extremes. A Hammersley–Clifford theorem links this new notion to the factorization of densities of extreme value models on graphs. For the popular class of Hüsler–Reiss distributions we show that, similarly to the Gaussian case, the sparsity pattern of a general extremal graphical model can be read off from suitable inverse covariance matrices. New parametric models can be built in a modular way and statistical inference can be simplified to lower dimensional marginals. We discuss learning of minimum spanning trees and model selection for extremal graph structures, and we illustrate their use with an application to flood risk assessment on the Danube river.

Keywords: Conditional independence; Extreme value theory; Graphical models; Multivariate Pareto distribution; Sparsity

1. Introduction

Evaluation of the risk that is related to heat waves, extreme flooding, financial crises or other rare events requires the quantification of their small occurrence probabilities. Empirical estimates are unreliable since the regions of interest are in the tail of the distribution and typically contain few or no data points. Extreme value theory provides the theoretical foundation for extrapolations to the distributional tail of a d -dimensional random vector \mathbf{X} . The univariate case $d = 1$ is well studied and the generalized extreme value and Pareto distributions are widely applied in areas such as hydrology (Katz *et al.*, 2002), climate science (Min *et al.*, 2011) and finance (McNeil *et al.*, 2015); see also Embrechts *et al.* (1997) and Beirlant *et al.* (2004).

In the multivariate setting, $d \geq 2$, the result of the extrapolation strongly depends on the strength of extremal dependence between the components of \mathbf{X} . Most current statistical models assume multivariate regular variation for \mathbf{X} (Resnick, 2008) since this entails mathematically elegant descriptions of the asymptotic tail distribution. Similarly to the univariate setting, two different but closely related approaches exist. Max-stable distributions arise as limits of nor-

Address for correspondence: Sebastian Engelke, Research Center for Statistics, University of Geneva, Boulevard du Pont d’Arve 40, 1205 Geneva, Switzerland.
E-mail: sebastian.engelke@unige.ch

malized maxima of independent copies of \mathbf{X} and have been extensively studied and applied in multivariate and spatial risk problems (see de Haan (1984), Gudendorf and Segers (2010) and Davison *et al.* (2012)). In contrast, multivariate Pareto distributions describe the random vector \mathbf{X} conditioned on the event that at least one component exceeds a high threshold; see Rootzén and Tajvidi (2006), Rootzén *et al.* (2018) and Kiriliouk *et al.* (2018) for their construction, stability properties and statistical inference.

A drawback of the current multivariate models is their limitation to rather moderate dimensions d , and the construction of tractable parametric models in higher dimensions is challenging, both for max-stable and for multivariate Pareto distributions. Sparse multivariate models require the notion of conditional independence (Dawid, 1979), which is not easy to define for tail distributions. In fact, Papastathopoulos and Strokorb (2016) have shown that, if (Z_1, Z_2, Z_3) is a max-stable random vector with positive continuous density, then the conditional independence of $Z_1 \perp\!\!\!\perp Z_3 | Z_2$ already implies the independence $Z_1 \perp\!\!\!\perp Z_3$; see also Dombry and Éyi-Minko (2014). Meaningful conditional independence structures can thus only be obtained for max-stable distributions with discrete spectral measure (Gissibl and Klüppelberg, 2018). Since these models do not admit densities, this excludes most of the currently used parametric families.

In this work we take the perspective of threshold exceedances and introduce a new notion of conditional independence for a multivariate Pareto distribution $\mathbf{Y} = (Y_1, \dots, Y_d)$, which we denote by ‘ \perp_e ’ to stress that it is designed for extremes. It is different from classical conditional independence since the support of \mathbf{Y} is not a product space, but the homogeneity property of \mathbf{Y} can be used to show that it is well defined. Conditional independence is tightly linked to graphical models. For an undirected graph $\mathcal{G} = (V, E)$ with nodes $V = \{1, \dots, d\}$ and edge set E , we say that \mathbf{Y} is an extremal graphical model if it satisfies the pairwise Markov property

$$Y_i \perp_e Y_j | \mathbf{Y}_{\setminus \{i,j\}}, \quad (i, j) \notin E. \quad (1)$$

The main advantage of conditional independence and graphical models is that they imply a simple probabilistic structure and possibly sparse patterns in multivariate random vectors (Lauritzen, 1996; Wainwright and Jordan, 2008). For extremal graphical models on decomposable graphs, we prove a Hammersley–Clifford-type theorem stating that property (1) is equivalent to the factorization of the density $f_{\mathbf{Y}}$ of \mathbf{Y} into lower dimensional marginals. This underlines that our notion of conditional independence is in fact natural for multivariate Pareto distributions.

Applications of this result are manifold. From a probabilistic perspective, we analyse models in the literature regarding their graphical properties in the sense of our definition (1). Extremal graphical models whose underlying graph is a tree have a particularly simple multiplicative stochastic representation in terms of extremal functions: a notion that is known from the simulation of max-stable processes (Dombry *et al.*, 2016). In multivariate extremes, one may argue that the family of Hüsler–Reiss distributions (Hüsler and Reiss, 1989) takes a similar role to that of Gaussian distributions in the non-extreme world. Instead of covariance matrices, they are parameterized by a variogram matrix Γ . We show that the extremal graphical structure of a Hüsler–Reiss distribution can be identified by zero patterns on matrices derived from Γ .

Extremal graphical models enable the construction of parsimonious models for multivariate Pareto distributions \mathbf{Y} , which further enjoy the advantage of interpretability in terms of the underlying graph. Thanks to the factorization of the densities, statistical inference can be efficiently carried out on lower dimensional marginals. For decomposable graphs with singleton separator sets, so-called block graphs, this allows the use of multivariate Pareto models in fairly high dimensions. In many cases the underlying graphical structure is unknown and must be learned from data. We discuss how a maximum likelihood tree can be obtained by using

1 standard algorithms by Kruskal (1956) or Prim (1957), and how the best model can be selected
2 among different extremal graphical models.

3 There is previous work on the construction of parsimonious extreme value models. A large
4 body of literature studies spatial max-stable random fields (Schlather, 2002; Kabluchko *et al.*,
5 2009; Opitz, 2013). Such models have small parameter dimension but they rely on strong as-
6 sumptions on stationarity and cannot be applied to multivariate, non-spatial data without infor-
7 mation on an underlying space. Other approaches include constructions through factor copulas
8 (Lee and Joe, 2018), ensembles of trees combining bivariate copulas (Yu *et al.*, 2017), graphical
9 models for large censored observations (Hitz and Evans, 2016) and eigendecompositions (Coo-
10 ley and Thibaud, 2018). Closely related to our concept of conditional independence is the work
11 of Coles and Tawn (1991) and Smith *et al.* (1997) who proposed a Markov chain model where
12 all bivariate marginals are extreme value distributions. This can be seen as a special case of our
13 approach when the graph has the simple structure of a chain. Similar limiting objects also arise
14 as the tail chains in the theory of extremes of stationary Markov chains with regularly varying
15 marginals (Smith, 1992; Basrak and Segers, 2009; Janssen and Segers, 2014). This theory has re-
16 cently been extended to regularly varying Markov trees (Segers, 2019). Gissibl and Klüppelberg
17 (2018) and Gissibl *et al.* (2018) studied the causal structure of directed acyclic graphs for max-
18 linear models, and they developed methods for model identification based on tail dependence
19 coefficients. Their work is in some sense complementary to ours, since their models do not have
20 densities whereas we shall explicitly assume the existence of densities.

21 To the best of our knowledge, our work is the first principled attempt to define conditional
22 independence for general multivariate extreme value models that naturally extends to the factor-
23 ization of densities, sparsity and graphical models. Section 2 introduces background on extreme
24 value theory and graphical models that is needed throughout the paper. The new notion of condi-
25 tional independence is defined in Section 3 and equivalent properties are derived. Section 4
26 contains the main probabilistic results on extremal graphical models, the representation of trees
27 and the characterization for Hüsler–Reiss distributions. Statistical models on block graphs and
28 their estimation, simulation and model selection are discussed in Section 5. In these graphical
29 models the dependence is modelled directly between lower dimensional subsets of variables,
30 whereas the global dependence is implicitly implied by the conditional independence structure
31 of the graph. There are many potential applications of extremal graphical models. In Section 6,
32 we illustrate the advantages of this structured approach compared with classical extreme value
33 models on a data set related to flooding on a river network in the upper Danube basin (see Asadi
34 *et al.* (2015)). The interpretation of the graphical structures that are obtained in this application
35 is particularly interesting since there is a seemingly natural underlying tree associated with the
36 flow connections. Our conditional independence is formulated for multivariate Pareto distribu-
37 tions, but the results in this paper have implications for max-stable distributions. This point and
38 further research directions will be addressed in the discussion in Section 7. Appendices A–F
39 contain proofs and some additional results.

40 An implementation for R (R Core Team, 2019) is available in the package `graphicalExtre`
41 `mes` (Engelke *et al.*, 2019). The code for the simulation study and application is available
42 from
43 <http://www.rss.org.uk/preprints>.

44 2. Background

45 2.1. Notation

46 We introduce some standard notation that is used throughout the paper. Symbols in bold such
47
48

as $\mathbf{x} \in \mathbb{R}^d$ are column vectors with components denoted by x_i , $i \in \{1, \dots, d\}$, and operations and relationships involving such vectors are meant componentwise. The vectors $\mathbf{0} = (0, \dots, 0)$ and $\mathbf{1} = (1, \dots, 1)$ are used as generic vectors with suitable dimension. Denoting the index set by $V = \{1, \dots, d\}$, for a non-empty subset $I \subset V$, we write for the subvectors $\mathbf{x}_I = (x_i)_{i \in I}$ and $\mathbf{x}_{\setminus I} = (x_i)_{i \in V \setminus I}$. Similar notation is used for random vectors $\mathbf{X} = (X_i)_{i \in V}$ with values in \mathbb{R}^d . For a matrix $A = (A_{ij})_{i,j \in V} \in \mathbb{R}^{d \times d}$ with entries indexed by V , and subsets $I, J \subset V$, we let $A_{IJ} = (A_{ij})_{i \in I, j \in J}$ denote the $|I| \times |J|$ submatrix of A , and we abbreviate $A_I = A_{II}$. For $\mathbf{a}, \mathbf{b} \in \mathbb{R}^d$ with $\mathbf{a} \leq \mathbf{b}$, a multivariate interval is denoted by $[\mathbf{a}, \mathbf{b}] = [a_1, b_1] \times \dots \times [a_d, b_d]$. The l_p -norm of a vector $\mathbf{x} \in \mathbb{R}^d$ for $p \geq 1$ is $\|\mathbf{x}\|_p = (\sum_{i \in V} |x_i|^p)^{1/p}$, and its l_∞ -norm is $\|\mathbf{x}\|_\infty = \max_{i \in V} |x_i|$. The density of a random vector \mathbf{X} , if it exists, is denoted by $f_{\mathbf{X}}$. The density of the marginal \mathbf{X}_I for a non-empty $I \subset V$ is denoted by f_I , if there is no ambiguity regarding the random vector.

2.2. Multivariate extreme value theory

The tail behaviour of the random vector $\mathbf{X} = (X_1, \dots, X_d)$ can be described through two different approaches: one based on componentwise maxima and the other on threshold exceedances. We briefly discuss both approaches and the close link between them.

Let $\mathbf{X}_i = (X_{i1}, \dots, X_{id})$, $i = 1, \dots, n$, be independent copies of \mathbf{X} and denote the componentwise maximum by $\mathbf{M}_n = (M_{1n}, \dots, M_{dn}) = (\max_{i=1}^n X_{i1}, \dots, \max_{i=1}^n X_{id})$. Under mild conditions on the marginal distribution of X_j there are sequences of normalizing constants $b_{jn} \in \mathbb{R}$, $a_{jn} > 0$, $j = 1, \dots, d$, such that

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(\frac{M_{jn} - b_{jn}}{a_{jn}} \leq x \right) = G_j(x) = \exp\{-(1 + \xi_j x)_+^{-1/\xi_j}\}, \quad x \in \mathbb{R}, \quad (2)$$

where $z_+ = \max(z, 0)$, and G_j is the generalized extreme value distribution whose shape parameter $\xi_j \in \mathbb{R}$ determines the heaviness of the tail of X_j ; see de Haan and Ferreira (2006), Embrechts *et al.* (1997) and Beirlant *et al.* (2004) for details. For analysis of the dependence structure, the marginal distributions F_j of X_j are typically estimated first to normalize the data by $1/\{1 - F_j(X_j)\}$ to standard Pareto distributions. For simplicity, we assume in what follows that the F_j are known and the vector \mathbf{X} has been normalized to standard Pareto marginals. Joint estimation of marginals and dependence is discussed in Section 5.2.

The standardized vector \mathbf{X} is said to be in the max-domain of attraction of the random vector $\mathbf{Z} = (Z_1, \dots, Z_d)$ if for any $\mathbf{z} = (z_1, \dots, z_d)$

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(\max_{i=1, \dots, n} X_{i1} \leq nz_1, \dots, \max_{i=1, \dots, n} X_{id} \leq nz_d \right) = \mathbb{P}(\mathbf{Z} \leq \mathbf{z}). \quad (3)$$

In this case, \mathbf{Z} is max-stable with standard Fréchet marginals $\mathbb{P}(Z_j \leq z) = \exp(-1/z)$, $z \geq 0$, and we may write

$$\mathbb{P}(\mathbf{Z} \leq \mathbf{z}) = \exp\{-\Lambda(\mathbf{z})\}, \quad \mathbf{z} \in \mathcal{E}, \quad (4)$$

where the exponent measure Λ is a Radon measure on the cone $\mathcal{E} = [0, \infty)^d \setminus \{\mathbf{0}\}$, and $\Lambda(\mathbf{z})$ is shorthand for $\Lambda(\mathcal{E} \setminus [\mathbf{0}, \mathbf{z}])$. If Λ is absolutely continuous with respect to Lebesgue measure on \mathcal{E} , its Radon–Nikodym derivative, denoted by λ , has the following properties:

- (a) homogeneity of order $-(d+1)$, i.e. $\lambda(t\mathbf{y}) = t^{-(d+1)}\lambda(\mathbf{y})$ for any $t > 0$ and $\mathbf{y} \in \mathcal{E}$;
- (b) normalized marginals, i.e. for any $i = 1, \dots, d$,

$$\int_{\mathbf{y} \in \mathcal{E}: y_i > 1} \lambda(\mathbf{y}) d\mathbf{y} = 1.$$

The two properties follow from the max-stability and the standard Fréchet marginals of \mathbf{Z} , respectively. For a non-empty subset $I \subset \{1, \dots, d\}$, we define the marginal of λ by

$$\lambda_I(\mathbf{y}_I) = \int_{[0, \infty)^{d-|I|}} \lambda(\mathbf{y}) d\mathbf{y}_{\setminus I}, \tag{5}$$

and note that it is homogeneous of order $-(|I| + 1)$. In particular, if $I = \{i\}$ for some $i = 1, \dots, d$, then $\lambda_{\{i\}}(y_i) = 1/y_i^2$ as a consequence of properties (a) and (b). Conversely, any positive and continuous function λ satisfying properties (a) and (b) defines a valid density of an exponent measure $\Lambda(\mathbf{z})$ by integration over $\mathcal{E} \setminus \{\mathbf{0}, \mathbf{z}\}$, $\mathbf{z} \in \mathcal{E}$, that satisfies similar homogeneity and normalization properties as λ . By expression (4) this also defines a max-stable distribution.

Another perspective on multivariate extremes is through threshold exceedances. By proposition 5.17 in Resnick (2008), the convergence in equation (3) is equivalent to

$$\lim_{u \rightarrow \infty} u \{1 - \mathbb{P}(\mathbf{X} \leq u\mathbf{z})\} = \Lambda(\mathbf{z}), \quad \mathbf{z} \in \mathcal{E}.$$

Consequently, the multivariate distribution of the threshold exceedances of \mathbf{X} satisfies

$$\mathbb{P}(\mathbf{Y} \leq \mathbf{z}) = \lim_{u \rightarrow \infty} \mathbb{P}\left(\frac{\mathbf{X}}{u \leq \mathbf{z}} \|\mathbf{X}\|_\infty > u\right) = \frac{\Lambda(\mathbf{z} \wedge \mathbf{1}) - \Lambda(\mathbf{z})}{\Lambda(\mathbf{1})}, \quad \mathbf{z} \in \mathcal{E}. \tag{6}$$

The distribution of the limiting random vector \mathbf{Y} is called a multivariate Pareto distribution (see Rootzén and Tajvidi, (2006)). It is defined through the exponent measure Λ of the max-stable distribution \mathbf{Z} , with support on the L -shaped space $\mathcal{L} = \{\mathbf{x} \in \mathcal{E} : \|\mathbf{x}\|_\infty > 1\}$. We say that \mathbf{Z} and \mathbf{Y} are associated, since their distributions mutually determine each other.

Multivariate Pareto distributions are the only possible limits in expression (6) and, owing to the homogeneity of the exponent measure, they enjoy certain stability properties (see Rootzén *et al.* (2018)). The exponent measure Λ , and hence the distribution of \mathbf{Y} , may place mass on some lower dimensional faces of the space \mathcal{E} . For the remainder of this paper we exclude this case to avoid technical difficulties. We further assume that the distribution of \mathbf{Y} admits a positive and continuous density $f_{\mathbf{Y}}$ on \mathcal{L} , which is

$$f_{\mathbf{Y}}(\mathbf{y}) = \frac{\partial^d}{\partial y_1 \dots \partial y_d} \mathbb{P}(\mathbf{Y} \leq \mathbf{y}) = \frac{\lambda(\mathbf{y})}{\Lambda(\mathbf{1})}, \quad \mathbf{y} \in \mathcal{L},$$

since $\Lambda(\mathbf{y} \wedge \mathbf{1})$ is always constant along at least one co-ordinate for $\mathbf{y} \in \mathcal{L}$. The density $f_{\mathbf{Y}}$ is thus proportional to the density λ of the exponent measure Λ . By the homogeneity of λ , $f_{\mathbf{Y}}$ is also homogeneous of order $-(d + 1)$. The normalization constant $\Lambda(\mathbf{1}) \in [1, d]$ is known as the d -variate extremal coefficient (see Schlather and Tawn (2003)). The assumption of a positive and continuous density $f_{\mathbf{Y}}$ implies that the multivariate Pareto distributions that we study are models for asymptotic extremal dependence, and all p -variate extremal coefficients, $1 \leq p \leq d$, are strictly smaller than their upper limit p .

For some non-empty subset $I \subset \{1, \dots, d\}$, the subvector $\mathbf{X}_I = (X_j)_{j \in I}$, properly normalized, given that its l_∞ -norm is large converge in the sense of expression (6) to the marginal $\mathbf{Y}_I = (Y_j)_{j \in I}$ of \mathbf{Y} defined on $\mathcal{L}_I = \{\mathbf{x}_I \in [0, \infty)^{|I|} \setminus \{\mathbf{0}\} : \|\mathbf{x}_I\|_\infty > 1\}$ with homogeneous density of order $-(|I| + 1)$ given by

$$f_I(\mathbf{y}_I) = \frac{\Lambda(\mathbf{1})}{\Lambda_I(\mathbf{1})} \int_{[0, \infty)^{d-|I|}} f_{\mathbf{Y}}(\mathbf{y}) d\mathbf{y}_{\setminus I} = \frac{\lambda_I(\mathbf{y}_I)}{\Lambda_I(\mathbf{1})}, \quad \mathbf{y}_I \in \mathcal{L}_I, \tag{7}$$

where Λ_I is the exponent measure of \mathbf{Z}_I , and λ_I is the density of Λ_I .

2.2.1. *Example 1 (logistic distribution)*

The extremal logistic distribution with parameter $\theta \in (0, 1)$ induces a multivariate Pareto distribution with density

$$f_{\mathbf{Y}}(\mathbf{y}) = \frac{1}{d^\theta} (y_1^{-1/\theta} + \dots + y_d^{-1/\theta})^{\theta-d} \prod_{i=1}^{d-1} \left(\frac{i}{\theta} - 1 \right) \prod_{i=1}^d y_i^{-1/\theta-1}, \quad \mathbf{y} \in \mathcal{L}. \quad (8)$$

2.2.2. *Example 2 (Hüsler–Reiss distribution)*

The Hüsler–Reiss distribution (Hüsler and Reiss, 1989) is parameterized by a symmetric, strictly conditionally negative definite matrix $\Gamma = \{\Gamma_{ij}\}_{1 \leq i, j \leq d}$ with $\text{diag}(\Gamma) = \mathbf{0}$ and non-negative entries, i.e. $\mathbf{a}^\top \Gamma \mathbf{a} < 0$ for all non-zero vectors $\mathbf{a} \in \mathbb{R}^d$ with $\sum_{i=1}^d a_i = 0$. The corresponding density of the exponent measure can be written for any $k \in \{1, \dots, d\}$ as (see Engelke *et al.* (2015))

$$\lambda(\mathbf{y}) = y_k^{-2} \prod_{i \neq k} y_i^{-1} \phi_{d-1}(\tilde{\mathbf{y}}_{\setminus k}; \Sigma^{(k)}), \quad \mathbf{y} \in \mathcal{E}, \quad (9)$$

where $\phi_p(\cdot; \Sigma)$ is the density of a centred p -dimensional normal distribution with covariance matrix Σ , $\tilde{\mathbf{y}} = \{\log(y_i/y_k) + \Gamma_{ik}/2\}_{i=1, \dots, d}$ and

$$\Sigma^{(k)} = \frac{1}{2} \{\Gamma_{ik} + \Gamma_{jk} - \Gamma_{ij}\}_{i, j \neq k} \in \mathbb{R}^{(d-1) \times (d-1)}. \quad (10)$$

The matrix $\Sigma^{(k)}$ is strictly positive definite; see Appendix B for details. The representation of the density in expression (9) seems to depend on the choice of k , but, in fact, the value of the right-hand side of this equation is independent of k . The Hüsler–Reiss multivariate Pareto distribution has density $f_{\mathbf{Y}}(\mathbf{y}) = \lambda(\mathbf{y})/\Lambda(\mathbf{1})$ and the strength of dependence between the i th and j th component is parameterized by Γ_{ij} , ranging from complete dependence for $\Gamma_{ij} = 0$ and independence for $\Gamma_{ij} = \infty$. In the bivariate case $d = 2$ we have

$$\lambda(y_1, y_2) = \frac{y_1^{-2} y_2^{-1}}{\sqrt{2\pi\Gamma_{12}}} \exp\left[-\frac{\{\log(y_2/y_1) + \Gamma_{12}/2\}^2}{2\Gamma_{12}}\right], \quad (y_1, y_2) \in \mathcal{E}, \quad (11)$$

and $\Lambda(1, 1) = 2\Phi(\sqrt{\Gamma_{12}/2})$, where Φ is the standard normal distribution function. The extension of Hüsler–Reiss distributions to random fields are Brown–Resnick processes (Brown and Resnick, 1977; Kabluchko *et al.*, 2009), which are widely used models for spatial extremes.

2.2.3. *Example 3 (bivariate Pareto distribution)*

In the general bivariate case $d = 2$, because of homogeneity, the density λ of the exponent measure can be characterized by a univariate distribution. Indeed, for any positive random variable U_2^1 with density $f_{U_2^1}$ and $\mathbb{E}(U_2^1) = 1$,

$$\lambda(y_1, y_2) = y_1^{-3} f_{U_2^1}(y_2/y_1), \quad (y_1, y_2) \in \mathcal{E}, \quad (12)$$

satisfies conditions (a) and (b) in Section 2.2 and thus defines a valid exponent measure density. We call U_2^1 the extremal function at co-ordinate 2, relative to co-ordinate 1 (see Dombry *et al.* (2013, 2016)). Equivalently, we can write the density in terms of the extremal function U_1^2 at co-ordinate 1, relative to co-ordinate 2, as $\lambda(y_1, y_2) = y_2^{-3} f_{U_1^2}(y_1/y_2)$, $(y_1, y_2) \in \mathcal{E}$, and U_1^2 is related to U_2^1 via the measure change $\mathbb{P}(U_1^2 \leq z) = \mathbb{E}(\mathbf{1}\{1/U_2^1 \leq z\} U_2^1)$, $z > 0$.

The above principle is a general construction principle, since every valid exponent measure density can be obtained in this way. The bivariate Hüsler–Reiss distribution in expression

(11) corresponds to the case of log-normal U_2^1 and U_1^2 , but many other parametric and non-parametric examples are available (e.g. Boldi and Davison (2007), Cooley *et al.* (2010), Ballani and Schlather (2011) and de Carvalho and Davison (2014)).

2.3. Graphical models

A graph $\mathcal{G} = (V, E)$ is defined as a set of nodes $V = \{1, \dots, d\}$, also called vertices, together with a set of edges $E \subset V \times V$ of pairs of distinct nodes. The graph is called undirected if, for two nodes $i, j \in V$, $(i, j) \in E$ if and only if $(j, i) \in E$. For notational convenience, for undirected graphs we sometimes represent edges as unordered pairs $\{i, j\} \in E$. When counting the number of edges, we count $\{i, j\} \in E$ such that each edge is considered only once. A subset $C \subset V$ of nodes is called complete if it is fully connected in the sense that $(i, j) \in E$ for all $i, j \in C$. We denote by \mathcal{C} the set of all cliques, i.e. the complete subsets that are not properly contained within any other complete subset.

To each node $i \in V$ we associate a random variable X_i with continuous state space $\mathcal{X}_i \subset \mathbb{R}$. The resulting random vector $\mathbf{X} = (X_i)_{i \in V}$ takes values in the Cartesian product $\mathcal{X} = \times_{i \in V} \mathcal{X}_i$. Suppose that \mathbf{X} has a positive and continuous Lebesgue density $f_{\mathbf{X}}$ on \mathcal{X} . For three disjoint subsets $A, B, C \subset V$ whose union is V , we say that \mathbf{X}_A is conditionally independent of \mathbf{X}_C given \mathbf{X}_B if the density factorizes as

$$f_{\mathbf{X}}(\mathbf{x}) = \frac{f_{A|B}(\mathbf{x}_{A|B}) f_{B|C}(\mathbf{x}_{B|C})}{f_B(\mathbf{x}_B)}, \tag{13}$$

and we write $\mathbf{X}_A \perp\!\!\!\perp \mathbf{X}_C | \mathbf{X}_B$. If $B = \emptyset$, then equation (13) amounts to independence of \mathbf{X}_A and \mathbf{X}_C .

The random vector \mathbf{X} is said to be a probabilistic graphical model on the graph $\mathcal{G} = (V, E)$ if its distribution satisfies the pairwise Markov property relative to \mathcal{G} , i.e. $X_i \perp\!\!\!\perp X_j | \mathbf{X}_{\setminus\{i,j\}}$ for all $(i, j) \notin E$. If in addition, for any disjoint subsets $A, B, C \subset V$ such that B separates A from C in \mathcal{G} , $\mathbf{X}_A \perp\!\!\!\perp \mathbf{X}_C | \mathbf{X}_B$, then \mathbf{X} is said to obey the global Markov property relative to \mathcal{G} . Since $f_{\mathbf{X}}$ is positive and continuous, it follows from the Hammersley–Clifford theorem (see Lauritzen (1996), theorem 3.9) that the two Markov properties are equivalent, and they are further equivalent to the factorization of the density

$$f_{\mathbf{X}}(\mathbf{x}) = \prod_{C \in \mathcal{C}} \psi_C(\mathbf{x}_C), \quad \mathbf{x} \in \mathcal{X}, \tag{14}$$

for suitable functions ψ_C on $\times_{i \in C} \mathcal{X}_i$. If the graph \mathcal{G} is decomposable, then this factorization can be rewritten in terms of marginal densities

$$f_{\mathbf{X}}(\mathbf{x}) = \frac{\prod_{C \in \mathcal{C}} f_C(\mathbf{x}_C)}{\prod_{D \in \mathcal{D}} f_D(\mathbf{x}_D)}, \quad \mathbf{x} \in \mathcal{X}, \tag{15}$$

where \mathcal{D} is a multiset containing intersections between the cliques called separator sets; see Lauritzen (1996) and Appendix A for the definition of decomposability and separator sets.

2.3.1. Example 4

We recall that, for a normal distribution $\mathbf{W} = (W_i)_{i \in V}$ with invertible covariance matrix Σ , the precision matrix Σ^{-1} contains the conditional independences, or equivalently the graph structure, since, for $i, j \in V$,

$$W_i \perp\!\!\!\perp W_j | \mathbf{W}_{\setminus\{i,j\}} \Leftrightarrow \Sigma_{ij}^{-1} = 0.$$

3. Conditional independence for threshold exceedances

The notion of conditional independence has not been exploited in extreme value theory. In fact, for max-stable distributions it leads only to trivial probabilistic structures (Papastathopoulos and Stokorb, 2016). An exception is directed acyclic graphs for max-linear models studied in Gissibl and Klüppelberg (2018) and Gissibl *et al.* (2018), which do, however, not admit densities.

We therefore approach the problem from the perspective of threshold exceedances. Since the notion of independence is defined only on product spaces, the meaning of conditional independence is not straightforward for a multivariate Pareto distribution $\mathbf{Y} = (Y_i)_{i \in V}$, $V = \{1, \dots, d\}$, with support on the L -shaped space $\mathcal{L} = \{\mathbf{x} \in \mathcal{E} : \|\mathbf{x}\|_\infty > 1\}$. In this section we show that there is nevertheless a natural definition of conditional independence for \mathbf{Y} . For this, we restrict \mathbf{Y} to product spaces. For any $k \in V$, we consider the random vector \mathbf{Y}^k defined as \mathbf{Y} conditioned on the event that $\{Y_k > 1\}$. Clearly, \mathbf{Y}^k has support on the product space $\mathcal{L}^k = \{\mathbf{x} \in \mathcal{L} : x_k > 1\}$ (Fig. 1) and it admits the density

$$f^k(\mathbf{y}) = \frac{f_{\mathbf{Y}}(\mathbf{y})}{\int_{\mathcal{L}^k} f_{\mathbf{Y}}(\mathbf{y}) d\mathbf{y}} = \lambda(\mathbf{y}), \quad \mathbf{y} \in \mathcal{L}^k, \quad (16)$$

since $\int_{\mathcal{L}^k} f_{\mathbf{Y}}(\mathbf{y}) d\mathbf{y} = 1/\Lambda(\mathbf{1})$ because of property (b) in Section 2.2. From expression (16) we see that the densities f^1, \dots, f^d coincide with λ on the intersection of their supports. Therefore there are no problems with lack of self-consistency as for instance in the model of Heffernan and Tawn (2004).

For any set $I \subset V$ with $k \in I$, the marginal \mathbf{Y}_I^k has density

$$f_I^k(\mathbf{y}_I) = \int_{[0, \infty)^{d-|I|}} \lambda(\mathbf{y}) d\mathbf{y}_{\setminus I} = \lambda_I(\mathbf{y}_I), \quad \mathbf{y}_I \in \mathcal{L}_I^k,$$

which is homogeneous of order $-(|I| + 1)$ on $\mathcal{L}_I^k = \{\mathbf{x}_I \in \mathcal{L}_I : x_k > 1\}$; see expression (5). This is, however, not so if $k \notin I$, since integration over $\mathbf{y}_{\setminus I}$ then includes y_k whose domain is $(1, \infty)$ in \mathcal{L}^k , and thus in general $f_I^k(\mathbf{y}_I) \neq \lambda_I(\mathbf{y}_I)$, $\mathbf{y}_I \in [0, \infty)^{|I|}$.

Definition 1. Suppose that \mathbf{Y} is multivariate Pareto and admits a positive and continuous density $f_{\mathbf{Y}}$ on \mathcal{L} , and let $A, B, C \subset V$ be non-empty disjoint subsets whose union is $V = \{1, \dots, d\}$. We say that \mathbf{Y}_A is conditionally independent of \mathbf{Y}_C given \mathbf{Y}_B if




Fig. 1. The blue shaded areas are (a) the support \mathcal{L} of \mathbf{Y} , and the supports (b) \mathcal{L}^1 of \mathbf{Y}^1 and (c) \mathcal{L}^2 of \mathbf{Y}^2 :
 •, samples of \mathbf{Y}

$$\forall k \in \{1, \dots, d\}: \mathbf{Y}_A^k \perp\!\!\!\perp \mathbf{Y}_C^k | \mathbf{Y}_B^k. \quad (17)$$

In this case we write $\mathbf{Y}_A \perp_e \mathbf{Y}_C | \mathbf{Y}_B$.

In fact, this definition can be shown to be equivalent to a slightly weaker condition, and to a factorization of the exponent measure density λ .

Proposition 1. Let $f_{\mathbf{Y}}$ and the sets A, B and C be as in the above definition; then $\mathbf{Y}_A \perp_e \mathbf{Y}_C | \mathbf{Y}_B$ is equivalent to any of the following two conditions.

(a)

$$\exists k \in B: \mathbf{Y}_A^k \perp\!\!\!\perp \mathbf{Y}_C^k | \mathbf{Y}_B^k. \quad (18)$$

(b) The density of the exponent measure factorizes as

$$\lambda(\mathbf{y}) = \frac{\lambda_{A \cup B}(\mathbf{y}_{A \cup B}) \lambda_{B \cup C}(\mathbf{y}_{B \cup C})}{\lambda_B(\mathbf{y}_B)}, \quad \mathbf{y} \in \mathcal{L}. \quad (19)$$

A natural question is whether we can extend the definition of $\mathbf{Y}_A \perp_e \mathbf{Y}_C | \mathbf{Y}_B$ to the case where $B = \emptyset$, meaning that \mathbf{Y}_A and \mathbf{Y}_C are independent on \mathcal{L} . In terms of the original definition, that would mean that, for any $k \in V$, $f^k(\mathbf{y}) = f_A^k(\mathbf{y}_A) f_C^k(\mathbf{y}_C)$ for all $\mathbf{y} \in \mathcal{L}^k$. Without losing generality, suppose that $k \in A$; then $f_C^k(\mathbf{y}_C) = \lambda(\mathbf{y}_A, \mathbf{y}_C) / \lambda_A(\mathbf{y}_A)$ for any $\mathbf{y}_A \in \mathcal{L}_A^k$ and $\mathbf{y}_C \in [0, \infty)^{|C|}$. Therefore f_C^k would be homogeneous of order $-|C|$ and thus not integrable on $[0, \infty)^{|C|}$, which is a contradiction. In the next section we show that this property implies that all graphical models that are defined in terms of the conditional independence \perp_e must be connected.

4. Graphical models for threshold exceedances

The notion of conditional independence enables us to define graphical models for threshold exceedances. As before, let $f_{\mathbf{Y}}$ be the positive and continuous density on \mathcal{L} of a multivariate Pareto distribution \mathbf{Y} , proportional to the density λ of the exponent measure Λ , and homogeneous of order $-(d+1)$. Let $\mathcal{G} = (V, E)$ be an undirected graph with nodes $V = \{1, \dots, d\}$ and edge set E . Similarly to the classical probabilistic graphical models that were described in Section 2.3, we say that \mathbf{Y} satisfies the pairwise Markov property on \mathcal{L} relative to \mathcal{G} if

$$Y_i \perp_e Y_j | \mathbf{Y}_{\setminus \{i, j\}}, \quad (i, j) \notin E, \quad (20)$$

i.e. Y_i and Y_j are conditionally independent in the sense of definition 1 given all other nodes whenever there is no edge between i and j in \mathcal{G} . By definition, this is equivalent to saying that \mathbf{Y}^k satisfies the usual pairwise Markov property on \mathcal{L}^k relative to \mathcal{G} for all $k \in V$. The global Markov property for \mathbf{Y} is defined similarly.

Definition 2. Let $\mathcal{G} = (V, E)$ be an undirected graph. If the multivariate Pareto distribution \mathbf{Y} with positive and continuous density $f_{\mathbf{Y}}$ satisfies the pairwise Markov property (20) relative to \mathcal{G} , we call the distribution of \mathbf{Y} an extremal graphical model with respect to \mathcal{G} .

For a decomposable graph \mathcal{G} we obtain a factorization of the density $f_{\mathbf{Y}}$ that is similar to the classical Hammersley–Clifford theorem, showing that definition 1 of conditional independence is natural for multivariate Pareto distributions. Let \mathcal{C} and \mathcal{D} be the sequences of cliques and separators of \mathcal{G} respectively, satisfying the running intersection property (44) in Appendix A.

Theorem 1. Let $\mathcal{G} = (V, E)$ be a decomposable graph and suppose that \mathbf{Y} has a multivariate Pareto distribution with positive and continuous density $f_{\mathbf{Y}}$ on \mathcal{L} . Denote the corresponding

exponent measure and its density by Λ and λ respectively. Then the following results are equivalent.

- (a) The distribution of \mathbf{Y} satisfies the pairwise Markov property relative to \mathcal{G} .
- (b) The distribution of \mathbf{Y} satisfies the global Markov property relative to \mathcal{G} .
- (c) The density $f_{\mathbf{Y}}$ factorizes according to \mathcal{G} , i.e.

$$f_{\mathbf{Y}}(\mathbf{y}) = \frac{1}{\Lambda(\mathbf{1})} \frac{\prod_{C \in \mathcal{C}} \lambda_C(\mathbf{y}_C)}{\prod_{D \in \mathcal{D}} \lambda_D(\mathbf{y}_D)}, \quad \mathbf{y} \in \mathcal{L}, \quad (21)$$

where the marginals λ_I are positive, continuous and homogeneous of order $-(|I| + 1)$ for any $I \subset V$.

In all cases, the graph \mathcal{G} is necessarily connected.

Remark 1. Theorem 1 shows that only connected extremal graphical models can arise. This is related to the assumption of multivariate regular variation in expression (3) that implies asymptotic dependence between all components. Loosely speaking, unconnected components would correspond to asymptotically independent components.

Remark 2. If the graph \mathcal{G} in theorem 1 is non-decomposable, it is expected that the density $f_{\mathbf{Y}}$ still factorizes into factors on the cliques of the graph. These factors can, however, no longer be identified with marginal densities, and it is an open problem whether they can be chosen to be homogeneous.

Since \mathcal{L} is not a product space, unlike in the classical Hammersley–Clifford theorem for decomposable graphs in expression (15), the factors in the factorization of the density $f_{\mathbf{Y}}$ in expression (21) are not the marginals f_I but the marginals of the exponent measure density λ_I . It holds, however, that $f_I(\mathbf{y}_I) = \lambda_I(\mathbf{y}_I) / \Lambda_I(\mathbf{1})$ for all $\mathbf{y}_I \in \mathcal{L}_I \subset \{\mathbf{x}_I : \mathbf{x} \in \mathcal{L}\}$.

As a first application, theorem 1 enables us to analyse formally the conditional independences and graphical structures of models in the multivariate extreme value literature.

4.1. Example 5

One of the simplest examples of a graph is a chain, i.e.

$$E = \{\{1, 2\}, \{2, 3\}, \dots, \{d-1, d\}\}.$$

Coles and Tawn (1991) proposed a model that factorizes with respect to this chain where all bivariate marginals are logistic (see example 1). This was extended to general bivariate marginals in Smith *et al.* (1997). More generally, in the study of extremes of stationary Markov chains the limiting objects are so-called tail chains. The latter induce multivariate Pareto distributions that can readily be seen to factorize with respect to a chain; see Smith (1992), Basrak and Segers (2009) and Janssen and Segers (2014).

4.2. Example 6

It turns out that many of the multivariate models in the literature do not have any conditional independences, i.e. their underlying graph is fully connected. For instance, this holds for the logistic multivariate Pareto distribution in example 1, the Dirichlet mixture model in Boldi and Davison (2007) and the pairwise beta distribution in Cooley *et al.* (2010).

4.3. Example 7

Similarly to Gaussian distributions, an appealing property of the Hüsler–Reiss model is its stability under taking marginals. Indeed, for any $I \subset V$ and $k \in I$ the marginal density of the exponent measure is

$$\lambda_I(\mathbf{y}_I) = \int_{[0, \infty)^{d-|I|}} \lambda_I(\mathbf{y}) d\mathbf{y}_{\setminus I} = y_k^{-2} \prod_{i \in I \setminus \{k\}} y_i^{-1} \phi_{|I|-1} \{\tilde{\mathbf{y}}_{I \setminus \{k\}}; \Sigma_I^{(k)}\},$$

with the notation of example 2, where $\Sigma_I^{(k)}$ is the matrix in expression (10) induced by the submatrix Γ_I . Thus, $f_I(\mathbf{y}_I) = \lambda_I(\mathbf{y}_I) / \Lambda_I(\mathbf{1})$ is the density of the $|I|$ -dimensional Hüsler–Reiss Pareto distribution with parameter matrix Γ_I .

By theorem 1, the density of a Hüsler–Reiss distribution that satisfies the pairwise Markov property relative to some decomposable graph \mathcal{G} factorizes into lower dimensional Hüsler–Reiss distributions. The explicit formula is given in corollary 2 in Appendix C.

Theorem 1 can also be seen as a construction principle to build new classes of extreme value distributions in a modular way by combining lower dimensional marginals. The following corollary shows how a multivariate Pareto distribution can be defined that factorizes according to a desired underlying graphical structure. This is particularly useful in high dimensional problems to ensure model sparsity.

Corollary 1. Let \mathcal{G} be a decomposable and connected graph and suppose that $\{\lambda_I : I \in \mathcal{C} \cup \mathcal{D}\}$ is a set of valid, positive and continuous exponent measure densities in the sense of properties (a) and (b) in section 2.2. For $D \subset C$, $D \in \mathcal{D}$ and $C \in \mathcal{C}$, assume that they satisfy the consistency constraint

$$\lambda_D(\mathbf{y}_D) = \int_{[0, \infty)^{C \setminus D}} \lambda_C(\mathbf{y}_C) d\mathbf{y}_{C \setminus D}. \quad (22)$$

The density of a valid d -dimensional exponent measure Λ is then given by

$$\lambda(\mathbf{y}) = \frac{\prod_{C \in \mathcal{C}} \lambda_C(\mathbf{y}_C)}{\prod_{D \in \mathcal{D}} \lambda_D(\mathbf{y}_D)}, \quad \mathbf{y} \in \mathcal{L},$$

and the function $f_{\mathbf{Y}}(\mathbf{y}) = \lambda(\mathbf{y}) / \Lambda(\mathbf{1})$, $\mathbf{y} \in \mathcal{L}$, is the density of a multivariate Pareto distribution satisfying the pairwise Markov property relative to \mathcal{G} .

4.4. Tree graphical models

A tree is a special case of a decomposable graphical model that is connected and has no cycles. All cliques are then of size 2 and the separators contain only one node. Let $\mathcal{T} = (V, E)$ be an undirected tree with nodes $V = \{1, \dots, d\}$ and edge set $E \subset V \times V$. Suppose that $\mathbf{Y} = (Y_i)_{i \in V}$ follows a multivariate Pareto distribution on \mathcal{L} with density $f_{\mathbf{Y}}$ that is an extremal graphical model with respect to the tree \mathcal{T} . Theorem 1 yields the factorization

$$f_{\mathbf{Y}}(\mathbf{y}) = \frac{1}{\Lambda(\mathbf{1})} \prod_{\{i, j\} \in E} \frac{\lambda_{ij}(y_i, y_j)}{y_i^{-2} y_j^{-2}} \prod_{i \in V} y_i^{-2}, \quad (23)$$

where $\lambda_{ij} = \lambda_{\{i, j\}}$ are the bivariate marginals of the exponent measure density λ corresponding to \mathbf{Y} . Formula (23) enables the extension of the modelling approach by Smith *et al.* (1997) that is




Fig. 2. A tree \mathcal{T}^2 rooted at node $k=2$ with the extremal functions in proposition 2 on the edges

described in example 5 from time series to general tree structures. Such tree models can represent more complex dependences and, moreover, are suitable beyond temporal data for multivariate or spatial applications.

Thanks to the relatively simple structure of a tree, more explicit results can be derived than for general graphical models. For this, we define a new directed tree $\mathcal{T}^k = (V, E^k)$ rooted at an arbitrary but fixed node $k \in V$. The edge set E^k consist of all edges $e \in E$ of the tree \mathcal{T} pointing away from node k ; see Fig. 2 for an example with $k=2$. For the resulting directed tree we define a set $(U_e)_{e \in E^k}$ of independent random variables, where for $e = (i, j)$, the distribution of $U_e = U_j^i$ is the extremal function of λ_{ij} at co-ordinate j , relative to co-ordinate i ; see expression (12) in example 3 for the definition of extremal functions. The following stochastic representation of the random vectors $\mathbf{Y}^k, k \in V$, provides a better understanding of the stochastic structure of multivariate Pareto distributions factorizing on a tree, and it is the main ingredient for efficient simulation in Section 5.4.

Proposition 2. Let \mathbf{Y} be a multivariate Pareto distribution with positive and continuous density on \mathcal{L} that factorizes with respect to the tree \mathcal{T} . With the notation above, and for a standard Pareto distribution P , we have the joint stochastic representation for \mathbf{Y}^k on \mathcal{L}^k

$$Y_i^k \stackrel{d}{=} \begin{cases} P, & \text{for } i = k, \\ P \times \prod_{e \in \text{ph}(ki)} U_e, & \text{for } i \in V \setminus \{k\}, \end{cases} \quad (24)$$

where $\text{ph}(ki)$ denotes the set of edges on the unique path from node k to node i on the tree \mathcal{T}^k .

Remark 3. The same object as in representation (24) has been obtained in Segers (2019) as the limit of regularly varying random vectors that satisfy a Markov condition on a tree. In analogy to the tail chains in example 5, they term it a tail tree.

4.5. Example 8

Suppose that all bivariate marginals λ_{ij} for $\{i, j\} \in E$ of a tree Pareto model are of logistic type with parameter $\theta_{ij} \in (0, 1)$ as defined in example 1. This tree logistic model is a generalization of the chain logistic model that was considered in Coles and Tawn (1991). In this

symmetric case, the extremal functions U_j^i and U_i^j have the same distribution with stochastic representation F/G , where F follows a Fréchet $(1/\theta, c_\theta)$ distribution with scale parameter $c_\theta = \Gamma(1 - \theta)^{-1}$ and $(G/c_\theta)^{-1/\theta}$ follows a gamma $(1 - \theta, 1)$ distribution, where we abbreviated $\theta = \theta_{ij}$ and Γ is the gamma function.

Similarly we can define a Hüsler–Reiss tree model, or use asymmetric models for λ_{ij} such as the Dirichlet model in Boldi and Davison (2007) for some or all of the edges $\{i, j\} \in E$. In asymmetric models, the extremal functions U_j^i and U_i^j in general have different distributions. We

refer to Section 4 in Dombry *et al.* (2016) for explicit formulae for extremal function distributions of commonly used model classes.

4.6. Hüsler–Reiss graphical models

In many ways, the class of Hüsler–Reiss distributions that were introduced in example 2 can be seen as the natural analogue of Gaussian distributions in the world of asymptotically dependent extremes. They are parameterized by the variogram of Gaussian distributions, and their statistical inference (Wadsworth and Tawn, 2014; Engelke *et al.*, 2015) and exact simulation (Dombry *et al.*, 2016) involves tools that are closely related to the corresponding methods for normal models.

Despite the similarities to Gaussian distributions, there are subtle but important differences that render analysis and statistical inference of Hüsler–Reiss distributions more difficult. To characterize conditional independence and graphical structures in these models, we first recall some results that are related to the original construction. The max-stable Hüsler–Reiss distribution has a stochastic representation as componentwise maxima

$$\mathbf{Z} = \max_{l \in \mathbb{N}} U_l \exp\{\mathbf{W}_l - \text{diag}(\Sigma)/2\}, \quad (25)$$

where $\{U_l : l \in \mathbb{N}\}$ is a Poisson point process on $[0, \infty)$ with intensity $u^{-2}du$, and \mathbf{W}_l are independent copies of a d -dimensional normal distribution \mathbf{W} with zero mean and covariance matrix Σ . Subtracting $\mathbb{E}\{\exp \mathbf{W}\} = \text{diag}(\Sigma)/2$ in the exponent normalizes the marginals of \mathbf{Z} to be standard Fréchet. Kabluchko *et al.* (2009) showed that the distribution of \mathbf{Z} depends on only the strictly conditionally negative definite variogram matrix of \mathbf{W} ,

$$\Gamma_{ij} = \mathbb{E}(W_i - W_j)^2, \quad i, j \in V.$$

Importantly, this implies that the representation in equation (25) is not unique since any centred, possibly degenerate normal distribution \mathbf{W} with variogram matrix Γ leads to the same max-stable Hüsler–Reiss distribution. Let

$$\mathcal{S}_\Gamma = \{\Sigma \in \mathbb{R}^{d \times d} \text{ covariance matrix} : \mathbf{1}\text{diag}(\Sigma)^T + \text{diag}(\Sigma)\mathbf{1}^T - 2\Sigma = \Gamma\}, \quad (26)$$

be the set of all covariance matrices that correspond to the same variogram matrix Γ ; see Appendix B. The Hüsler–Reiss Pareto distribution \mathbf{Y} that is associated with \mathbf{Z} is defined by its density in example 2, which is also parameterized by Γ . We recall that, for a normal distribution \mathbf{W} with invertible covariance matrix Σ , the precision matrix Σ^{-1} contains the conditional independences; see example 4. A first naive guess would be that the graph structure of \mathbf{W} that is used in the construction of \mathbf{Z} directly translates into the extremal graph structure of the Hüsler–Reiss Pareto distribution \mathbf{Y} . This is, however, not so.

4.7. Example 9

We consider three examples for \mathbf{W} in the representation (25) with $d = 4$.

- (a) Let W_i , $i = 1, \dots, 4$, be independent standard normal distributions; then $\Sigma^{-1} = \text{diag}(1, \dots, 1)$ and Γ_{ij} equals 2 if $i \neq j$ and 0 otherwise. The graph underlying the distribution of \mathbf{W} is the graph with four unconnected nodes. The graph of the corresponding Hüsler–Reiss Pareto distribution \mathbf{Y} turns out to be the fully connected graph in Fig. 3(a).
- (b) Consider the centred normal distribution \mathbf{W} with precision matrix and variogram matrix




Fig. 3. Hüsler–Reiss graphical models described in example 9

$$\Sigma^{-1} = \begin{pmatrix} 12 & -4 & -4 & -1 \\ -4 & 2 & 1 & 0 \\ -4 & 1 & 2 & 0 \\ -1 & 0 & 0 & 1 \end{pmatrix},$$

$$\Gamma = \begin{pmatrix} 0 & 1 & 1 & 1 \\ 1 & 0 & 2 & 2 \\ 1 & 2 & 0 & 2 \\ 1 & 2 & 2 & 0 \end{pmatrix}$$

respectively. The Gaussian graphical model is the graph in Fig. 3(b) with an additional edge between nodes 2 and 3. In contrast, the corresponding Hüsler–Reiss model factorizes according to the graph in Fig. 3(b).

- (c) Consider the centred normal distribution \mathbf{W} with precision matrix and variogram matrix

$$\Sigma^{-1} = \begin{pmatrix} 2 & -0.5 & -0.5 & 0 \\ -0.5 & 1 & 0 & -0.5 \\ -0.5 & 0 & 1 & -0.5 \\ 0 & -0.5 & -0.5 & 1 \end{pmatrix},$$

$$\Gamma = \begin{pmatrix} 0 & 1.5 & 1.5 & 2 \\ 1.5 & 0 & 2 & 1.5 \\ 1.5 & 2 & 0 & 1.5 \\ 2 & 1.5 & 1.5 & 0 \end{pmatrix}$$

respectively. It can be checked that both the Gaussian and the Hüsler–Reiss graphical model are as in Fig. 3(c). Also note that this graph is not decomposable.

The above examples show that it is not possible simply to transfer the Gaussian graphical model of the covariance matrix Σ in the construction (25) to the extremal graphical structure of the corresponding Hüsler–Reiss Pareto distribution. This is not surprising since the covariance matrices in the set \mathcal{S}_Γ can have very different graph structures, but all lead to the same Hüsler–Reiss graphical model. There is, however, a set of particular matrices that enable us to identify conditional independences and thus the graphical structure of a Hüsler–Reiss Pareto distribution. Recall the definition of $\Sigma^{(k)} \in \mathbb{R}^{(d-1) \times (d-1)}$ in expression (10). The same matrix including the k th row and column

$$\tilde{\Sigma}^{(k)} = \frac{1}{2} \{ \Gamma_{ik} + \Gamma_{jk} - \Gamma_{ij} \}_{i,j \in V} \in \mathbb{R}^{d \times d}, \tag{27}$$

is degenerate since the k th component has zero variance, but it is a valid choice in the construction (25), i.e. $\tilde{\Sigma}^{(k)} \in \mathcal{S}_\Gamma$, for any $k \in V$. Let \mathbf{W}^k be a centred normal distribution with covariance matrix $\tilde{\Sigma}^{(k)}$ and note that $W_k^k = 0$ almost surely. For a random variable P with standard Pareto distribution, independent of \mathbf{W}^k , it can be seen that

$$\mathbf{Y}^k \stackrel{d}{=} P \exp(\mathbf{W}^k - \Gamma_{\cdot k}/2), \quad (28)$$

by comparing the density of the right-hand side with expression (9) and noting that $\text{diag}(\tilde{\Sigma}^{(k)}) = \Gamma_{\cdot k}$. This together with the original definition of conditional independence in expression (17) suggests that the matrices $\Sigma^{(k)}$ contain the graphical structure of a Hüsler–Reiss Pareto distribution.

We denote the precision matrix of $\Sigma^{(k)}$ by $\Theta^{(k)} = (\Sigma^{(k)})^{-1}$. For notational convenience, the indices of the matrices $\Sigma^{(k)}$ and $\Theta^{(k)}$ range in $\{1, \dots, d\} \setminus \{k\}$ instead of $\{1, \dots, d-1\}$.

Lemma 1. For $k, k' \in V, k \neq k'$, the precision matrices $\Theta^{(k)}$ and $\Theta^{(k')}$ satisfy for $i, j \in V \setminus \{k'\}$

$$\begin{aligned} \Theta_{ij}^{(k')} &= \Theta_{ij}^{(k)}, & \text{if } i, j \neq k, \\ \Theta_{ik}^{(k')} &= - \sum_{l \neq k} \Theta_{il}^{(k)}, & \text{if } i \neq k, j = k, \\ \Theta_{kk}^{(k')} &= \sum_{l, m \neq k} \Theta_{lm}^{(k)}, & \text{if } i, j = k. \end{aligned}$$

Lemma 1 is of independent interest since it explains the link between the precision matrices $\Theta^{(k)}$ for different $k \in V$. The proof uses blockwise inversion of the precision matrices. This result is the crucial ingredient to characterize conditional independence in Hüsler–Reiss models.

Proposition 3. For a Hüsler–Reiss Pareto distribution \mathbf{Y} with parameter matrix Γ , it holds for $i, j \in V$ with $i \neq j$, and for any $k \in V$, that

$$Y_i \perp_e Y_j | \mathbf{Y}_{\setminus \{i, j\}} \Leftrightarrow \begin{cases} \Theta_{ij}^{(k)} = 0, & \text{if } i, j \neq k, \\ \sum_{l \neq k} \Theta_{lj}^{(k)} = 0, & \text{if } i = k, j \neq k, \\ \sum_{l \neq k} \Theta_{il}^{(k)} = 0, & \text{if } j = k, i \neq k. \end{cases} \quad (29)$$

For any $k \in V$, the single matrix $\Theta^{(k)}$ contains all information on conditional independence of \mathbf{Y} . Conditional independence concerning the k th component is encoded in the row and column sums of $\Theta^{(k)}$, and it might sometimes be easier to switch to another representation $\Theta^{(k')}, k' \neq k$, where it simply figures as a zero entry. In example 9 we can now easily determine the graphical model $\mathcal{G} = (V, E)$ for each of the three Hüsler–Reiss Pareto distributions. For a given Σ we first compute the matrix Γ as in expression (26), then we transform it by expression (10) to obtain $\Sigma^{(k)}$ for any $k \in V$ and use proposition 3 to decide whether $(i, j) \in E$ for all $i, j \in V$. These transformations are implemented in our R-package `graphicalExtremes` (Engelke *et al.*, 2019).

4.8. Example 10

In spatial extreme value statistics, the finite dimensional distributions of the Brown–Resnick process (Kablichko *et al.*, 2009) at locations $t_1, \dots, t_d \in \mathbb{R}^D$ are Hüsler–Reiss distributed with variogram matrix $\Gamma_{ij} = \gamma(t_i - t_j), i, j \in \{1, \dots, d\}$, where γ is a variogram function on \mathbb{R}^D . The most commonly used model is the fractal variogram family $\gamma_\alpha(h) = \|h\|_2^\alpha$, for some $\alpha \in (0, 2]$. The corresponding d -variate Hüsler–Reiss distribution does not have conditional independences and its graph is thus fully connected. The only exception is the case of the original Brown–Resnick process in Brown and Resnick (1977) with $\alpha = 1$ and $D = 1$, where the corresponding graph is a chain as in example 5.

In this section, we have so far not required that the underlying graph \mathcal{G} is decomposable. If this is so then, as shown in example 7, theorem 1 implies that the density of the Hüsler–Reiss

graphical model factorizes into lower dimensional Hüsler–Reiss densities; see corollary 2 in appendix C.

5. Statistical inference for block graphs

5.1. Model construction

The notion of conditional independence and graphical models for multivariate Pareto distributions enables the construction of new statistical models with two major advantages. First, sparsity can be imposed on the model, which is a crucial ingredient for tractable and parsimonious models in higher dimensions. Second, under certain graphical structures, the model parameters can be estimated separately on lower dimensional subsets of the data.

We consider here, and throughout the rest of the paper, decomposable and connected graphs $\mathcal{G} = (V, E)$ with clique set \mathcal{C} and separator set \mathcal{D} , where all separators in \mathcal{D} are single nodes. Such graph structures with singleton separator sets are known as block graphs (see Harary (1963)) and have already been seen to have appealing properties for discrete data (Loh and Wainwright, 2013). In our case, they are a convenient way of restricting the model complexity to obtain a tractable class of extremal graphical models. In fact, corollary 1 provides a simple construction principle for multivariate Pareto distributions that factorize with respect to the block graph \mathcal{G} .

- (a) For each clique $C \in \mathcal{C}$, choose possibly different parametric families of valid exponent measure densities $\{\lambda_C(\cdot; \theta_C) : \theta_C \in \Omega_C\}$ for suitable parameter spaces Ω_C . If \mathcal{G} is a tree \mathcal{T} , then this reduces to choosing $d - 1$ bivariate exponent measure densities λ_{ij} , for each $\{i, j\} \in E$; see example 3 for a general representation of such densities.
- (b) Since all separator sets consist of a single node, the consistency constraint (22) is trivially fulfilled as a consequence of properties (a) and (b) in Section 2.2 and the fact that $\lambda_D(y_D) = y_D^{-2}$ for all $D \in \mathcal{D}$.
- (c) For any fixed combination of parameters $\theta = (\theta_C)_{C \in \mathcal{C}} \in \Omega = \times_{C \in \mathcal{C}} \Omega_C$, the product of the normalized lower dimensional exponent measure densities,

$$f_{\mathbf{Y}}(\mathbf{y}; \theta) = \frac{1}{\Lambda(\mathbf{1}; \theta)} \prod_{C \in \mathcal{C}} \frac{\lambda_C(\mathbf{y}_C; \theta_C)}{\prod_{j \in C} y_j^{-2}} \prod_{i \in V} y_i^{-2}, \quad \mathbf{y} \in \mathcal{L}, \quad (30)$$

defines a valid d -variate Pareto distribution factorizing according to the graph \mathcal{G} , which is a member of the parametric family parameterized by $\theta \in \Omega$. For a tree \mathcal{T} , this reduces to the density in equation (23).

Concrete examples for this construction are tree logistic or tree Hüsler–Reiss models as described in example 8, where all cliques have the same type of distributions. The above construction is much more flexible, as it allows us to use different distribution families for the different cliques. Moreover, some, or even all, of the cliques may be modelled by non-parametric methods; see Lafferty *et al.* (2012) for non-parametric tree models in the non-extreme case. In this direction, there is a line of research on kernel-based estimation of exponent measure densities (see de Carvalho and Davison (2014), Marcon *et al.* (2017) and Kiriliouk *et al.* (2018)) that could be used as clique models. We shall not follow this approach here.

In the graphical models above, the dependence inside each clique is modelled directly, whereas dependence between components from different cliques is implicitly implied by the conditional independence structure of the graph. Even if all cliques are modelled with the same type of parametric family, the joint distribution (30) is typically not of this distribution type. For a tree logistic distribution, for instance, this can easily be seen by comparing its density (23) with that

of the d -variate logistic distribution in example 1. The latter has only one parameter governing the whole d -dimensional dependence structure, whereas the tree has $d - 1$ logistic parameters $\{\theta_{ij}; \{i, j\} \in E\}$ and thus much higher flexibility.

An important exception is the family of Hüsler–Reiss distributions, which is stable under taking marginal distributions; see example 7. The following proposition shows that for a given graphical structure as above, if all cliques have Hüsler–Reiss distributions, then so has the full d -dimensional model. This is the converse of corollary 2 in Appendix C.

Proposition 4. Let $\mathcal{G} = (V, E)$ be a block graph as above, and suppose that, on each clique $C \in \mathcal{C}$, \mathbf{Y} has a $|C|$ -variate Hüsler–Reiss distribution with exponent measure density $\lambda_C(\cdot; \Gamma^{(C)})$ parameterized by a $(|C| \times |C|)$ -dimensional variogram matrix $\Gamma^{(C)}$. Then there is a unique solution to the problem

$$\begin{aligned} & \text{find a } (d \times d)\text{-dimensional variogram matrix } \Gamma, \\ \text{subject to } & \begin{cases} \Gamma_{ij} = \Gamma_{ij}^{(C)}, & \text{for } i, j \in C \text{ and all } C \in \mathcal{C}, \\ \Theta_{ij}^{(k)} = 0, & \text{for all } k \in V, i, j \neq k \text{ and } (i, j) \notin E, \end{cases} \end{aligned} \quad (31)$$

with the notation from proposition 3. The corresponding d -variate Hüsler–Reiss distribution factorizes according to the graph \mathcal{G} into the lower dimensional Hüsler–Reiss densities on the cliques.

This is a matrix completion problem for variograms similar to what Dempster (1972) introduced for covariance matrices. In our case, the graph is decomposable and the above result relates to the marginal problem that was studied in Kellerer (1964) and Dawid and Lauritzen (1993). For Hüsler–Reiss marginals on block graphs we even see that the implied d -dimensional distribution is again Hüsler–Reiss. We give a direct constructive proof in Appendix F. This provides a method to construct high dimensional Hüsler–Reiss distributions out of many low dimensional distributions. The full d -variate Hüsler–Reiss model without any conditional independences has $d(d - 1)/2$ parameters. A Hüsler–Reiss distribution as in proposition 4 that factorizes on a block graph with click set \mathcal{C} has only

$$\frac{1}{2} \sum_{C \in \mathcal{C}} |C|(|C| - 1)$$

parameters, which can be much smaller than $d(d - 1)/2$.

5.2. Estimation

Extremal graphical models can be used to build parsimonious statistical models for the tail of a multivariate random vector. In this section we discuss how the model parameters can be estimated efficiently by considering each clique distribution separately.

Let $\mathbf{X} = (X_j)_{j \in V}$, $V = \{1, \dots, d\}$, be a random vector in the max-domain of attraction of the max-stable random vector \mathbf{Z} as in expression (3), with marginal distribution X_j in the max-domain of attraction of a generalized extreme value distribution with shape parameter ξ_j , $j \in V$. Equivalently, there are a sequence of high thresholds $\mathbf{t}_u = (t_{u1}, \dots, t_{ud})$ with t_{uj} tending to the upper end point of X_j as $u \rightarrow \infty$, and positive normalizing functions $\sigma_u = (\sigma_{u1}, \dots, \sigma_{ud})$, such that the distribution of exceedances converges weakly:

$$\frac{\mathbf{X} - \mathbf{t}_u}{\sigma_u} \left\| \left\| \frac{\mathbf{X}}{\mathbf{t}_u} \right\|_{\infty} \right\| > 1 \longrightarrow \frac{\mathbf{Y}^{\xi} - 1}{\xi}, \quad u \rightarrow \infty, \quad (32)$$

where \mathbf{Y} is the multivariate Pareto distribution that is associated with \mathbf{Z} . We assume that \mathbf{Y} is in the model class of the previous section with density (30), and for now we suppose that the underlying graph $\mathcal{G} = (V, E)$ is known and fixed. The conditional density of $\mathbf{X} - \mathbf{t}_u$ given that $\|\mathbf{X}/\mathbf{t}_u\|_\infty > 1$ is then approximated by

$$f_{\mathbf{Y}} \left\{ \left(1 + \xi \frac{\mathbf{x}}{\sigma_u} \right)^{1/\xi}; \theta \right\} \prod_{j \in V} \frac{1}{\sigma_{uj}} \left(1 + \xi_j \frac{x_j}{\sigma_{uj}} \right)^{1/\xi_j - 1}. \quad (33)$$

This density can be used to estimate jointly the marginal parameters (σ_{uj}, ξ_j) , $j \in V$, and the dependence parameter vector $\theta = (\theta_C)_{C \in \mathcal{C}}$ of $f_{\mathbf{Y}}$.

In what follows we concentrate on estimation of the dependence, and we therefore assume that the marginal parameters are known or have been estimated separately. As described in Section 2.2, we can then normalize \mathbf{X} to standard Pareto marginals, in which case $\xi_j = 1$, $t_{uj} = u$ and $\sigma_{uj} = u$ for all $j \in V$. We recover the standardized setting of expression (6) considered throughout the paper, where \mathbf{X}/u given that $\|\mathbf{X}\|_\infty > u$ converges to \mathbf{Y} , whose likelihood is proportional as a function of θ to

$$f_{\mathbf{Y}}(\mathbf{y}; \theta) \propto \frac{1}{Z_\theta} \prod_{C \in \mathcal{C}} \frac{\lambda_C(\mathbf{y}_C; \theta_C)}{\Lambda_C(\mathbf{1}; \theta_C)}, \quad Z_\theta = \frac{\Lambda(\mathbf{1}; \theta)}{\prod_{C \in \mathcal{C}} \Lambda_C(\mathbf{1}; \theta_C)}. \quad (34)$$

Direct maximization of the likelihood with contributions (34) for each data point is tedious since the normalizing constant Z_θ contains all parameters and does not factorize. Fortunately the class of block graphs has the property that we can estimate the parameters θ_C of each λ_C separately, without having to enforce the consistency constraints at the separator sets. In fact, we use the following observation. If \mathbf{X} is in the domain of attraction of the family of multivariate Pareto distributions $\{f_{\mathbf{Y}}(\cdot; \theta) : \theta \in \Omega\}$, then, for a fixed clique $C \in \mathcal{C}$, the subvector \mathbf{X}_C is in the domain of attraction of $\{f_C(\cdot; \theta_C) : \theta_C \in \Omega_C\}$, and the distribution of the normalized exceedance $\mathbf{X}_C/u \|\mathbf{X}_C\|_\infty > u$ is approximated for large u by \mathbf{Y}_C with density

$$f_C(\mathbf{y}_C; \theta_C) = \frac{\lambda_C(\mathbf{y}_C; \theta_C)}{\Lambda_C(\mathbf{1}; \theta_C)}, \quad \mathbf{y}_C \in \mathcal{L}_C; \quad (35)$$

see expression (7) in Section 2.2. We can therefore obtain an estimate of θ_C based only on data of the components in C , whose dimension is typically much smaller than the dimension d of the full graph. Estimating the cliques separately might in principle result in a loss of estimation efficiency compared with using the joint likelihood (34). The normalizing constant Z_θ does, however, not contain much information on the parameter θ and the maximum likelihood estimate by using $f_{\mathbf{Y}}(\mathbf{y}; \theta)$ is generally very close to the estimate that is obtained by maximizing separate likelihoods based on expression (35). We discuss this point in the simulation study in Section 5.5.

In practice, some components of \mathbf{X} might not have converged to the limiting distribution \mathbf{Y} . To avoid biased estimates of the dependence parameters θ_C , it has become a standard approach to apply censoring to the data; see Ledford and Tawn (1997) and Smith *et al.* (1997). For a data point \mathbf{X}_C with $\|\mathbf{X}_C\|_\infty > u$ for a high threshold $u > 0$, define J to be the set of indices $j \in C$ such that $Y_j < 1$, i.e. $X_j < u$. For this data point we use the censored likelihood contribution

$$f_C^{\text{cens}}(\mathbf{y}_C; \theta_C) = \int_{[0, 1]^{|J|}} f_C(\mathbf{y}_C; \theta_C) d\mathbf{y}_J, \quad \mathbf{y}_C \in \mathcal{L}_C, \quad (36)$$

which uses for all $j \in J$ only the information that this component of \mathbf{Y}_C is smaller than 1, but not its exact value. For explicit forms of the censored likelihoods for many parametric models see Dombry *et al.* (2017) and Kiriliouk *et al.* (2018).

For n independent data $\mathbf{y}^{(h)} \in \mathcal{L}$, $h = 1, \dots, n$, of $\mathbf{X}/u\|\mathbf{X}\|_\infty > u$, for each clique C we define $\hat{\theta}_C$ as the maximizer of the censored log-likelihood

$$L(\theta_C; \mathbf{y}^{(1)}, \dots, \mathbf{y}^{(n)}) = \sum_{\mathbf{y}^{(h)} \in \mathcal{L}_C} \log\{f_C^{\text{cens}}(\mathbf{y}_C^{(h)}; \theta_C)\}, \quad (37)$$

where $\mathcal{L}_C = \{\mathbf{y} \in \mathcal{L} : \exists j \in C \text{ subject to } y_j > 1\}$, and each $\mathbf{y}_C^{(h)}$ has its own censoring set $J^{(h)} \subset C$.

Maximum likelihood estimation is only one possibility to infer the parameters θ_C on the basis of exceedances of \mathbf{X}_C and the limiting distribution (35). Alternative methods use M -estimators (Einmahl *et al.*, 2012, 2016) or proper scoring rules (de Fondeville and Davison, 2018).

5.3. Model selection

Up to now we have assumed that a graphical structure \mathcal{G} was *a priori* given and we analysed models that factorize with respect to this structure. In many applications the underlying graph structure is unknown and should be learned in a data-driven way. Theorem 1 implies that all extremal graphical structures are connected, and a simple and flexible class of connected graphs is trees; see Section 4.1. It is thus natural first to build a suitable tree as a baseline model, and then to extend the tree by adding additional edges to obtain more complex graphs.

Since trees are a special case of general graphical models, there are specific methods to learn these simpler structures. The notion of a minimum spanning tree is crucial (Kruskal, 1956). Let $\mathcal{G}_0 = (V, E_0)$ be the fully connected graph on $V = \{1, \dots, d\}$ with edge set $E_0 = \{(i, j) : i, j \in V\}$. Suppose that a positive weight $w_{ij} > 0$ is attached to each edge $(i, j) \in E_0$ of \mathcal{G}_0 . This number can be seen as the length of the edge (i, j) or the distance between nodes i and j , and it is assumed that $w_{ij} = w_{ji}$ and $w_{ii} = 0$, $i, j \in V$. The minimum spanning tree is the tree $\mathcal{T}_{\text{mst}} = (V, E_{\text{mst}})$, with $E_{\text{mst}} \subset E_0$, that minimizes the sum of weights on that tree, i.e.

$$\mathcal{T}_{\text{mst}} = \arg \min_{T=(V,E)} \sum_{(i,j) \in E} w_{ij}. \quad (38)$$

If all edges of \mathcal{G}_0 have distinct lengths, then \mathcal{T}_{mst} is unique. This minimization problem can be solved efficiently by the greedy algorithms that were proposed in Kruskal (1956) or Prim (1957).

The weights w_{ij} determine the tree structure and should be chosen carefully. A common approach in graphical modelling is to search the conditional independence structure that maximizes the likelihood (see Cowell *et al.* (2006), chapter 11). Such a tree is also called a Chow–Liu tree (Chow and Liu, 1968). We fix a parametric family of bivariate Pareto distributions that is used for all pairs of nodes $\{f(\cdot; \theta_{ij}) : \theta_{ij} \in \Omega\}$. For n independent data $\mathbf{y}^{(h)}$, $h = 1, \dots, n$, the maximal log-likelihood of a fixed tree within this parametric class is essentially the sum over the maximized clique log-likelihoods in equation (37) over all edges of this tree. To find the tree that maximizes the log-likelihood over all trees and all distributions in this parametric family, we therefore find the minimum spanning tree in equation (38) with weights

$$w_{ij} = -L(\hat{\theta}_{ij}; \mathbf{y}^{(1)}, \dots, \mathbf{y}^{(n)}) - 2 \sum_{y_i^{(h)} > 1} \log(y_i^{(h)}) - 2 \sum_{y_j^{(h)} > 1} \log(y_j^{(h)}), \quad (39)$$

where we include the censored marginal densities y_i^{-2} and y_j^{-2} in distribution (30) for the clique $\{i, j\}$, since now the edges are no longer fixed but parameters of the optimization. The resulting tree \mathcal{T}_{mst} is the baseline model for the data. If the model fit is not satisfactory, it is possible to extend this tree to graphs with more complex structures by adding additional edges. The family of Hüsler–Reiss distributions is particularly appealing since the bivariate marginals remain in the same class. We illustrate this model extension through a greedy forward selection in Section 5.5.

The different multivariate Pareto models can then be compared by the Akaike information criterion (Kiriliouk *et al.*, 2018):

$$\text{AIC} = 2p - 2L(\hat{\theta}; \mathbf{y}^{(1)}, \dots, \mathbf{y}^{(n)}), \quad (40)$$

where p is the number of parameters in the respective model, and the second term is twice the negative log-likelihood based on the censored version of expression (34), evaluated at the optimized parameters of each clique.

5.4. Exact simulation

Exact simulation of a max-stable random vector \mathbf{Z} relies on the notion of extremal functions (Dombry and Éyi-Minko, 2013). The extremal function of \mathbf{Z} , or of its associated multivariate Pareto distribution \mathbf{Y} , relative to co-ordinate $k \in V$ is the d -dimensional random vector \mathbf{U}^k with $U_k^k = 1$ such that the exponent measure density of \mathbf{Z} can be written as

$$\lambda(\mathbf{y}) = y_k^{-(d+1)} f_{\mathbf{U}^k}(\mathbf{y}_{\setminus k}/y_k). \quad (41)$$

The distributions of the extremal functions \mathbf{U}^k , $k \in V$, for most commonly used models have explicit forms and are derived in section 4 of Dombry *et al.* (2016). Theorem 2 in Dombry *et al.* (2016) relates the distribution of the so-called spectral measure to these extremal functions. Together with the following representation of \mathbf{Y} , this enables simulation of multivariate Pareto distributions by rejection sampling. Recall that, for any $k \in V$, the random vector \mathbf{Y}^k is defined as \mathbf{Y} conditioned on the event that $\{Y_k > 1\}$.

Lemma 2. The distribution of the extremal function \mathbf{U}^k of \mathbf{Y} relative to co-ordinate $k \in V$ is given by the distribution of \mathbf{Y}^k/Y_k^k . Independently, let P be a standard Pareto random variable and T uniformly distributed on $\{1, \dots, d\}$. We then have for any Borel set $A \subset \mathcal{L}$

$$\mathbb{P}(\mathbf{Y} \in A) = \mathbb{P}\left(\frac{P\mathbf{Y}^T}{\|\mathbf{Y}^T\|_1} \in A \mid \frac{P\|\mathbf{Y}^T\|_\infty}{\|\mathbf{Y}^T\|_1} > 1\right). \quad (42)$$

This representation yields a simple algorithm for exact simulation of \mathbf{Y} ; see also de Fondeville and Davison (2018).

5.4.1. Algorithm 1 (exact simulation of a multivariate Pareto distribution \mathbf{Y})

Step 1: simulate a standard Pareto random variable P .

Step 2: simulate T uniformly on $\{1, \dots, d\}$ and sample a realization of the extremal function \mathbf{U}^T relative to co-ordinate T .

Step 3: if $\max\{P\|\mathbf{U}^T\|_\infty/\|\mathbf{U}^T\|_1\} > 1$, return $\mathbf{Y} = P\mathbf{U}^T/\|\mathbf{U}^T\|_1$ as a realization of the multivariate Pareto distribution.

Step 4: otherwise, reject the simulation and go back to step 1.

The complexity of this simulation algorithm as a function of the dimension d of the vector \mathbf{Y} is driven by the number of times that we must sample from one of the extremal functions $\mathbf{U}^1, \dots, \mathbf{U}^d$, since simulation of the variables P and T requires much less computational effort. Let $C_{\mathbf{Y}}(d)$ denote the number of extremal functions that must be simulated in algorithm 1. The random variable $C_{\mathbf{Y}}(d)$ follows a geometric distribution and from equation (50) in the proof of lemma 2 in Appendix F its expectation is

$$\mathbb{E}\{C_{\mathbf{Y}}(d)\} = d/\Lambda(\mathbf{1}) \in [1, d].$$

The expected complexity therefore depends on both the dimension and the strength of extremal dependence in \mathbf{Y} . Weak dependence implies a large coefficient $\Lambda(\mathbf{1})$ closer to d and therefore reduces the computational effort that is required for exact simulation. The simulation of multivariate Pareto distributions is in general computationally easier than for the associated max-stable distribution \mathbf{Z} . Indeed, exact simulation of the latter is also based on samples from a mixture of the $\mathbf{U}^1, \dots, \mathbf{U}^d$, and the fastest algorithm in Dombry *et al.* (2016) has expected complexity $\mathbb{E}\{C_{\mathbf{Z}}(d)\} = d$; see also Dieker and Mikosch (2015) and Oesting *et al.* (2018) for other exact simulation methods.

The complexity measures $C_{\mathbf{Y}}(d)$ and $C_{\mathbf{Z}}(d)$ consider only the number of extremal functions that are required for one exact simulation of \mathbf{Y} and \mathbf{Z} respectively. The computational effort of sampling \mathbf{U}^k can, however, be significantly lower if \mathbf{Y} has a sparse structure. If \mathbf{Y} factorizes according to a graph, then, by the definition 1 of conditional independence, the $\mathbf{Y}^1, \dots, \mathbf{Y}^d$ inherit the sparsity of this graph structure. This is particularly important in the case of trees and for Hüsler–Reiss distributions, as shown in the examples below. It is important to note that more efficient simulation of the extremal functions speeds up exact simulation of the multivariate Pareto distribution \mathbf{Y} , but also of the max-stable distribution \mathbf{Z} .

5.4.2. Example 11

Suppose that \mathbf{Y} factorizes according to a tree $\mathcal{T} = (V, E)$. It follows from proposition 2 and lemma 2 that the extremal function \mathbf{U}^k relative to co-ordinate $k \in V$ is

$$\mathbf{Y}^k / Y_k^k \stackrel{d}{=} \left(\prod_{e \in \text{ph}(ki)} U_e \right)_{i \in V}.$$

For exact simulation of \mathbf{Y} it therefore suffices to simulate the univariate random variables U_e . This is feasible even in very large dimensions.

5.4.3. Example 12

If \mathbf{Y} has a Hüsler–Reiss distribution that factorizes on the graph $\mathcal{G} = (V, E)$, then it follows from expression (28) that the extremal function \mathbf{U}^k relative to co-ordinate $k \in V$ is

$$\mathbf{Y}^k / Y_k^k \stackrel{d}{=} \exp(\mathbf{W}^k - \Gamma_{\cdot k} / 2),$$

where \mathbf{W}^k is a centred normal distribution with covariance matrix $\tilde{\Sigma}^{(k)}$ in expression (27); see also proposition 4 in Dombry *et al.* (2016). The normal distribution $\mathbf{W}^k_{\setminus k}$ factorizes in the classical sense on the subgraph $\mathcal{G}_{\setminus k}$, and efficient simulation algorithms exist if the graph is sparse (e.g. Rue and Held (2005)).

The exact simulation algorithms for both multivariate Pareto and max-stable distributions are implemented in our R package `graphicalExtremes` (Engelke *et al.*, 2019).

5.5. Simulation study

We assess the efficiency of parameter estimation and model selection in the framework of graphical models for extremes described in the previous sections. We fix a dimension d of variables or nodes $V = \{1, \dots, d\}$ and a block graph $\mathcal{G} = (V, E)$ as in Section 5.1. In this study we simulate samples directly from the limiting distribution \mathbf{Y} by using the exact algorithm 1, but we use the censored estimator since this is common practice in applications.

We first choose $d = 5$ and let \mathcal{G} be the undirected version of the tree in Fig. 2. We simulate $n \in \{100, 200\}$ samples $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(n)}$ of a Hüsler–Reiss distribution with parameter matrix Γ that

factorizes according to \mathcal{G} . The entries of Γ need to be specified only on the submatrices $\Gamma^{(C)}$ for all cliques $C \in \mathcal{C}$ of \mathcal{G} , since the solution to the matrix completion problem in proposition 4 then yields the unique variogram matrix Γ . In this simulation we set

$$\Gamma = \begin{pmatrix} 0 & \mathit{1} & \mathit{2} & \mathit{2} & \mathit{3} \\ \mathit{1} & 0 & 3 & 1 & 2 \\ \mathit{2} & 3 & 0 & 4 & 5 \\ \mathit{2} & 1 & 4 & 0 & 3 \\ \mathit{3} & \mathit{2} & 5 & 3 & 0 \end{pmatrix}, \quad (43)$$

where we specified only the four parameters Γ_{ij} for $(i, j) \in E, i < j$, to the values in italics, and the rest of the matrix is implied by the graph structure.

In this dimension we can still maximize the censored version of the joint likelihood (34) to obtain an estimate $\hat{\Gamma}_{ij}^{\text{joint}}, \{i, j\} \in E$, of the parameters corresponding to the four edges of the tree. We also obtain estimates $\hat{\Gamma}_{ij}, \{i, j\} \in E$, of the parameters of each clique separately by maximizing the censored clique likelihood (37). In both cases, the four estimated parameters yield estimates $\hat{\Gamma}^{\text{joint}}$ and $\hat{\Gamma}$ of the whole variogram matrix Γ through the graph structure. We repeat the simulation and estimation 200 times and compare the efficiency of both approaches in Fig. 4, displaying only the four free parameters that have actually been estimated.

The difference in efficiency between the joint and clique likelihoods seems to be small or even negligible. This is due to two reasons. For non-censored points the two likelihoods differ by only the normalizing constant Z_θ . Since this constant measures only the global strength of dependence and does not depend on the data, it seems not very sensitive to changes in the parameter θ . The second difference between the two approaches is that they use slightly different data. Consider a clique $C \in \mathcal{C}$ and the corresponding model parameter θ_C . The joint likelihood uses all data \mathbf{Y} in the space $\mathcal{L} = \{\mathbf{y} \in \mathcal{E} : \exists j \in V \text{ subject to } y_j > 1\}$, but censors all components with $y_j \leq 1$. In contrast, the clique likelihood uses the marginals \mathbf{Y}_C of all data \mathbf{Y} in $\mathcal{L}_C = \{\mathbf{y} \in \mathcal{L} : \exists j \in C \text{ subject to } y_j > 1\}$. Consequently, the additional data that are used in the joint likelihood are in $\mathcal{L} \setminus \mathcal{L}_C = \{\mathbf{y} \in \mathcal{L} : y_j \leq 1 \text{ for all } j \in C\}$. But the contribution to the joint likelihood of data in this set with regard to the parameter θ_C is completely censored and does therefore not add significant additional information. These two considerations underline that estimating the parameters for each clique separately does not result in significant losses of efficiency. This is one of the main advantages of graphical models, namely that the distribution is defined locally by the cliques

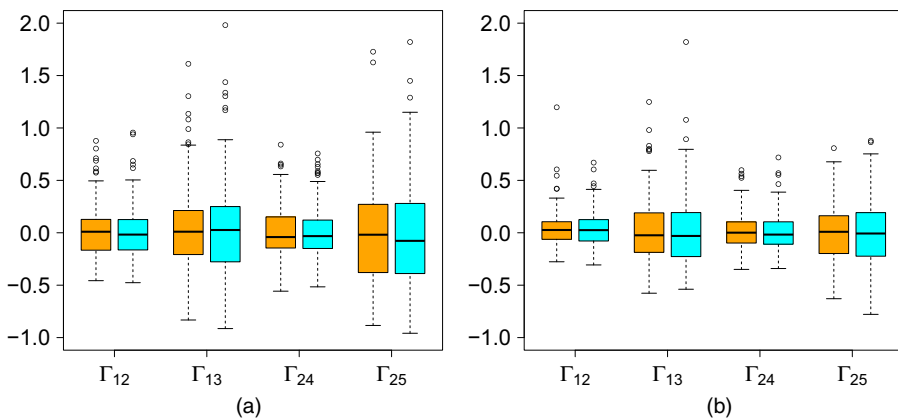


Fig. 4. Boxplots of errors of the four parameter estimates of the Hüsler–Reiss tree model in expression (43) based on joint (orange) and clique likelihood (cyan) with sample size (a) $n = 100$ and (b) $n = 200$




Fig. 5. (a) True underlying graphical structure and (b) the estimated structure in the second experiment, where the line thickness indicates the number of times that the edge has been selected

and extends globally by the conditional independence structure. In terms of computational aspects, the joint likelihood becomes infeasible even in moderate dimensions, whereas the clique likelihood is applicable in high dimensions as long as the cliques have sufficiently small sizes. Moreover, the computations for different cliques can be easily parallelized.

For the second experiment we take $d = 16$ and let \mathcal{G} be the graph in Fig. 5(a), which is not a tree. We simulate $n = 100$ samples of a Hüsler–Reiss distribution with parameter matrix Γ that factorizes according to \mathcal{G} . The parameters of the $p = 18$ edges are independently sampled from a uniform distribution on $(0.5, 1)$, under the constraint that Γ is conditionally negative definite on cliques with three nodes. We illustrate how we can choose the best graphical model, where we restrict to block graphs as in Section 5.1 with cliques of sizes 2 and 3. We first construct the minimum spanning tree as described in Section 5.3 within the class of Hüsler–Reiss distributions. The estimated edge set of this tree is denoted by E_1 . The 15 parameter estimates $\hat{\Gamma}_{ij}, \{i, j\} \in E_1$ that are obtained by fitting the clique likelihoods of each clique of the tree yield a unique estimate $\hat{\Gamma}$ of the $(d \times d)$ -dimensional variogram matrix; see proposition 4. This tree model does not contain all edges of the true underlying graph. We therefore perform a greedy forward selection to add additional edges and to improve the model. In each step, we define an enlarged edge set $E_{m+1} = E_m \cup \{i, j\}, m = 1, 2, \dots$, restricting to those edges $\{i, j\}, i, j \in V$, that still yield a block graph with cliques of maximal size 3. We continue this process until no more edges can be added in this way. For the same parameter matrix Γ , we repeat the simulation and model selection 100 times. Fig. 5(b) shows the graph with the selected edges, where the line width of each edge indicates the number of times that it has been selected among the first 18 edges. It can be seen that the graph structure is generally very well identified. For each model and each repetition we also compute the resulting AIC according to expression (40). The proportion of times that the model with $\{15, \dots, 20\}$ edges has the smallest AIC are $\{0.01, 0.11, 0.23, 0.39, 0.23, 0.03\}$. Even though AIC is a criterion that was built for model estimation and not for identification (see Arlot and Celisse (2010)), it seems to be well suited to select the correct degree of sparsity for this extremal graphical model.

6. Application

We illustrate the applicability of extremal graphical models by the example of river discharges in the upper Danube basin: a region that is prone to serious flooding. The data are provided by the Bavarian Environmental Agency (<http://www.gkd.bayern.de>) and we use $d = 31$ gauging stations with 50 years of common daily data from 1960–2009. The tree that is induced by the

physical flow connections at these stations is shown in Fig. 6(a), where the path $10 \rightarrow 9 \rightarrow \dots \rightarrow 1$ is on the Danube and the other branches are tributaries. The spatial extremal dependence structure of this data set has been studied in Asadi *et al.* (2015) and we follow their preprocessing steps to make the results comparable. Out of all daily data only the three months June, July and August are considered since the most severe floods occur in this period and are caused by heavy summer rain (Böhm and Wetzel, 2006). The $50 \times 92 = 4600$ observations in these months are declustered in time to remove temporal dependence and to match slightly shifted peak flows at different locations. We refer to Asadi *et al.* (2015) for more details on the data, the declustering method and exploratory analysis concerning stationarity and asymptotic dependence; see also Keef *et al.* (2009, 2013) for other approaches to flood risk assessment.

The declustering yields $N = 428$ supposedly independent events $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)} \in \mathbb{R}^d$. The univariate marginal distributions of these data were estimated in Asadi *et al.* (2015) by a regionalized extreme value model. We focus on estimation of the extremal dependence and normalize the data empirically to standard Pareto marginals. This still guarantees consistent inference of the dependence parameters (e.g. Genest *et al.* (1995) and Joe (2015)). We obtain $n = 117$ approximate samples of \mathbf{Y} by $\mathbf{y}^{(h)} = \mathbf{x}^{(h)}/u$ for all observations with $\|\mathbf{x}^{(h)}\|_\infty > u$, where we choose the threshold u as the 90%-quantile of the marginal Pareto distribution.

The max-stable Brown–Resnick model in Asadi *et al.* (2015) corresponds to a parametric family of Hüsler–Reiss Pareto distributions $\{f_{\mathbf{Y}}(\cdot; \theta) : \theta \in \Omega\}$ at the 31 gauging stations. The dependence model is tailor made for this particular application to river extremes and uses several covariates such as distance on the river network, catchment sizes and altitudes. In terms of our new notion of extremal graphical models it is readily checked by using the results of proposition 3 that for any parameter value $\theta \in \Omega$ their model does not exhibit conditional independences.

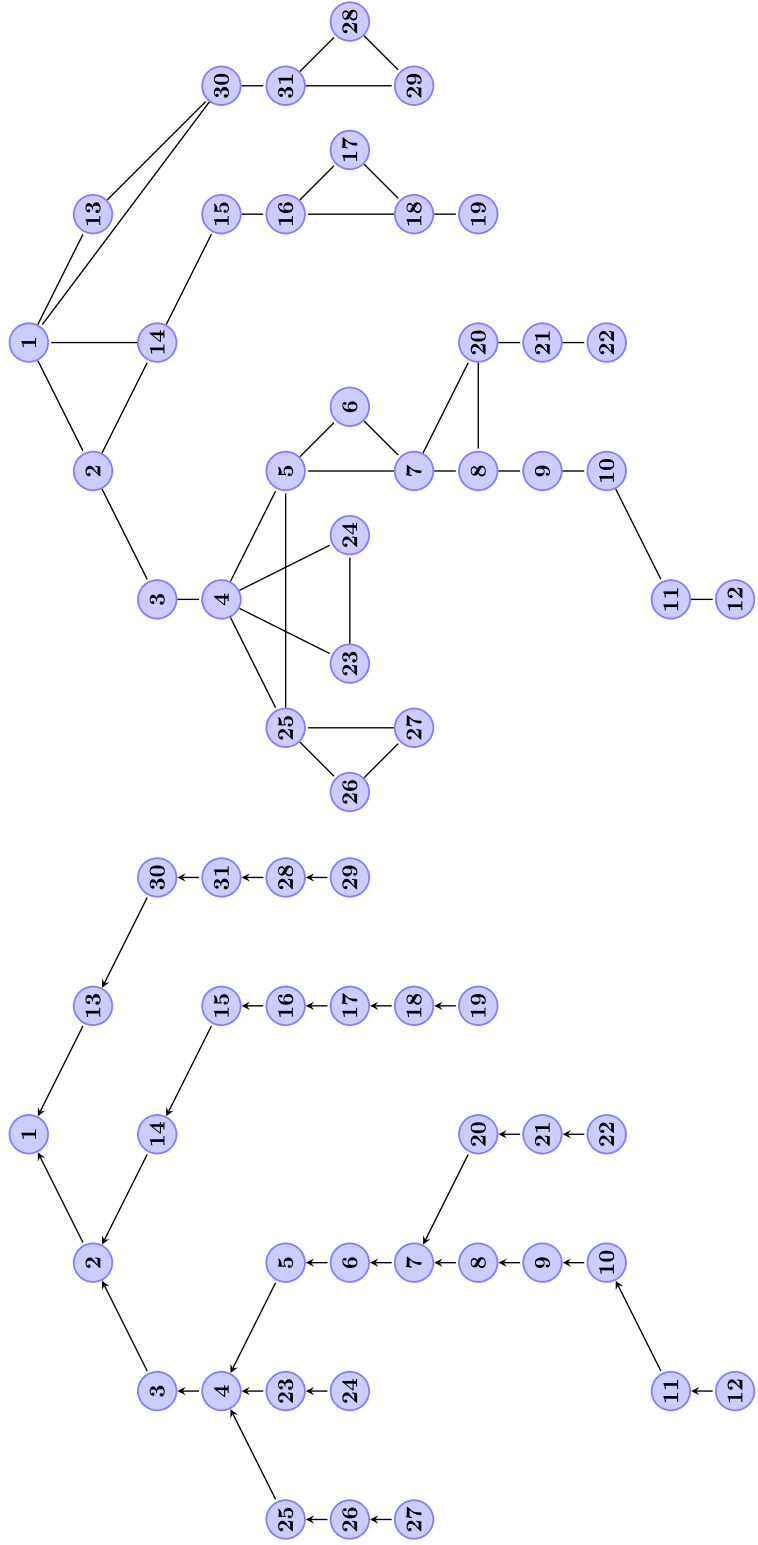
We propose a different Hüsler–Reiss model that factorizes according to a sparse graph and does not require any domain knowledge or additional covariates. In fact, we propose a sequence of models

$$M^{(l)} = \{f_{\mathbf{Y}}(\cdot; \theta^{(l)}) : \theta^{(l)} \in \Omega^{(l)}\}, \quad l = 1, \dots, L,$$

where $\theta^{(l)} = (\theta_C^{(l)})_{C \in \mathcal{C}^{(l)}}$, and $\mathcal{C}^{(l)}$ is the set of all cliques of the l th extremal graphical model $\mathcal{G}^{(l)}$ according to which the model family $M^{(l)}$ factorizes. As the simplest model we take $\mathcal{G}^{(1)}$ to be the minimum spanning tree within the family of Hüsler–Reiss models as described in Section 5.3. Similarly to the simulation study in Section 5.5, we obtain $\mathcal{G}^{(2)}, \dots, \mathcal{G}^{(L)}$ by successively adding edges to the tree $\mathcal{G}^{(1)}$ in a greedy way while restricting the model class to block graphs with cliques of size at most 3. The estimated tree $\mathcal{G}^{(1)}$ is shown in Fig. 9(a) in Appendix D. It is very similar to the tree in Fig. 6 that corresponds to the tree that is induced by the flow connections of the river network. There are, however, differences, and it is important to note that the flow connection tree is not necessarily the optimal tree structure in terms of extreme river discharges. Appendix D also contains a sensitivity analysis of the tree structure for various thresholds u , and a comparison with a Gaussian tree model fitted to non-extremal data.

Fig. 7 shows the AIC-values for the models $M^{(1)}, \dots, M^{(L)}$. Forward selection is a greedy approach and it does not guarantee finding the optimal graph. We therefore also initialize the forward selection with the simplest model $\mathcal{G}^{(1)}$ being the flow connection tree in Fig. 6(a). This tree must have a larger AIC than the minimum spanning tree but, interestingly, Fig. 7(a) shows that by adding additional edges the optimal AIC is better than the previous optimal AIC. In this particular case, we thus choose the graph that is initiated with the flow connection tree with nine additional edges. In general, a tree structure appears to be too simple for this application. The reason is that only part of the extremal dependence of discharges at different locations can be explained by flow connections. Additional dependence may arise even between

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48



(a) Tree induced by flow connections for the 31 stations in the upper Danube basin and (b) the estimated graph with the optimal AIC




Fig. 7. (a) AIC-values for extremal graphical models for the Danube data set with an increasing number of edges, starting from the minimum spanning tree (●) and the flow connection tree (●) (---, AIC of the spatial model in Asadi *et al.* (2015)), and empirically estimated χ -coefficients against those implied by (b) the fitted spatial and (c) graphical model minimizing the AIC: ●, flow connected stations

flow-unconnected locations due to proximity of their catchments that are affected by the same spatial precipitation events. Asadi *et al.* (2015) modelled this explicitly through a variogram with two parts: one for the dependence on the river network and one for the spatial meteorological dependence. The nine additional edges of the graphical model in Fig. 6(b), which minimizes AIC, partly improve the model in terms of this spatial dependence between flow-unconnected stations but also strengthen it between some flow-connected locations. This best graphical model has 39 edges and an AIC of 5269.43. It significantly outperforms the simpler tree models with 30 edges and the spatial model of Asadi *et al.* (2015), which has only six parameters but an AIC of 5291.34, which is indicated by the broken orange line in Fig. 7(a).

A popular summary statistic for extremal dependence between Y_i and Y_j , $i, j \in V$, is the tail correlation (see Coles *et al.* (1999)), which can be expressed as $\chi_{ij} = 2 - \Lambda_{ij}(1, 1)$. Figs 7(b) and 7(c) compare empirical estimates of these statistics for all pairs of stations with those implied by the fitted models. In terms of this bivariate summary, both models seem to fit the data well, even though the graphical model seems to be slightly less biased than the spatial model. There are also versions of χ that assess how a model captures the higher order extremal dependence structure. In Fig. 11 in Appendix E we compare the trivariate empirical χ -coefficients with those implied from the fitted spatial and graphical model. Both models fit well the trivariate dependence, again with a slightly lower bias of the graphical model.

In this application we have considered only block graphs, which are particularly convenient in terms of statistical inference as seen in the previous sections. In general it should be assessed whether this sparse model class is justified for the data. In our case, the bivariate and trivariate χ -coefficients indicate that block graphs are sufficiently flexible to capture the extremal dependence structure of the river data. This is further supported by the fact that the AIC-curve in Fig. 7 attains its minimum even before the maximal number of edges has been added in this model class. It is an important question for future research how extremal graphical models with more complicated structures can be estimated.

7. Discussion

The conditional independence relationship \perp_c that is introduced in this paper is natural for a multivariate Pareto distribution \mathbf{Y} as it explains the factorization of its density \mathbf{f}_Y into lower dimensional marginals (see theorem 1). This establishes a link of extreme value statistics to the broad field of graphical models, and it opens the door to define sparsity and to perform structure learning for tail distributions. In this work we have studied the probabilistic structure and statistical inference for some important models, with the main purpose of modelling the extremal dependence structure. Many subsequent research directions are possible. Directed acyclic graphs as in Gissibl and Klüppelberg (2018) for max-linear models may be formulated in our setting and would yield factorizations that are different from those for undirected graphs, and this would form the basis for extending work on causal inference for extremes (Naveau *et al.*, 2018; Mhalla *et al.*, 2019; Gnecco *et al.*, 2019) to continuous extreme value distributions. The models in this paper are well suited for asymptotic dependence. Another line of research focuses on multivariate tail models under asymptotic independence (Ledford and Tawn, 1997; Heffernan and Tawn, 2004; Wadsworth *et al.*, 2017). Conditional independence and graphical models have not been studied in this framework, except for the special case of Markov chains (Kulik and Soulier, 2015; Papastathopoulos *et al.*, 2017).

Conditional independence for \mathbf{Y} does not carry over to factorization of the density of the associated max-stable distribution \mathbf{Z} . By proposition 1, the conditional independence relationship \perp_c does, however, imply the factorization of the exponent measure density λ of \mathbf{Z} , which is the key object in simulation (Dombry *et al.*, 2016) and full likelihood estimation (Thibaud *et al.*, 2016; Dombry *et al.*, 2017; Huser *et al.*, 2019) of max-stable processes. Thus, sparsity in our notion for multivariate Pareto distributions also facilitates inferential tasks for max-stable distributions: a fact that has been briefly discussed for simulation in Section 5.4 but which deserves further investigation.

The application to flood risk assessment is just one illustrative example. Unlike spatial models, extremal graphical models can be applied to multivariate problems without domain knowledge, as for instance in financial or insurance applications. The ability to learn underlying structures in a data-driven way has also great practical potential for exploratory analysis and data visualization. In on-going research we investigate efficient learning of extremal tree structures and, in the case of Hüsler–Reiss distributions, of more general graphs based on l_1 -regularization.

Acknowledgements

We thank Robin J. Evans and Nicola Gnecco for helpful discussions. We are grateful to the editorial team and the referees for knowledgeable comments that improved the paper. Financial support by the Swiss National Science Foundation (S. Engelke) and by the Berrow Foundation

(A. S. Hitz) is gratefully acknowledged. The paper was completed while S. Engelke was a visitor at the Department of Statistical Sciences, University of Toronto.

Appendix A: Definitions for graphical models

Let $\mathcal{G} = (V, E)$ be an undirected graph with node set $V = \{1, \dots, d\}$ and edge set $E \subset V \times V$; see Section 2.3. We define the notion decompositions and decomposability for the graph \mathcal{G} (see Lauritzen (1996), definition 2.1).

Definition 3. A triplet (A, B, C) of disjoint subsets of V is said to form a decomposition of \mathcal{G} into the components $\mathcal{G}_{A \cup B}$ and $\mathcal{G}_{B \cup C}$ if $V = A \cup B \cup C$ and

- (a) B separates A from C (i.e. every path from A to C intersects B) and
- (b) B is a complete subset.

The decomposition is called proper if A and C are both non-empty. A graph \mathcal{G} is decomposable if it is complete or if there is a proper decomposition (A, B, C) into decomposable subgraphs $\mathcal{G}_{A \cup B}$ and $\mathcal{G}_{B \cup C}$. Decomposable graphs are also known as triangulated or chordal graphs.

For instance, $(\{1, 2, 3, 4, 5\}, \{4, 5\}, \{4, 5, 6\})$ is a proper decomposition of the decomposable graph in Fig. 8.

For a connected, decomposable graph \mathcal{G} , we can order the set of the cliques $\mathcal{C} = \{C_1, \dots, C_m\}$ such that, for all $i = 2, \dots, m$,

$$D_i := C_i \cap \bigcup_{j=1}^{i-1} C_j \subset C_k \quad \text{for some } k < i, \tag{44}$$

which is a condition called the running intersection property; see Lauritzen (1996), chapter 2, and Green and Thomas (2013). The sets D_i , $i = 2, \dots, m$, are called separators of the graph, and both \mathcal{C} and the collection of separators $\mathcal{D} = \{D_2, \dots, D_m\}$ are uniquely determined up to different orderings. The separators may not all be distinct, and we say that \mathcal{D} is a multiset. A possible enumeration of cliques and separators for the graph in Fig. 8 that satisfies the running intersection property is

$$\mathcal{C} = (\{1, 2\}, \{2, 3, 4, 5\}, \{4, 5, 6\}), \quad \mathcal{D} = (\{2\}, \{4, 5\}).$$

From expression (44) we note that the clique C_m intersects the other cliques only in D_m . Consider the connected decomposable subgraph \mathcal{G}_{m-1} of \mathcal{G} with node set $V_{m-1} = V \setminus (C_m \setminus D_m)$ and corresponding induced edge set. Property (44) then holds for \mathcal{G}_{m-1} , which has one clique fewer. Continuing this process, we note that each C_j intersects the subgraph \mathcal{G}_j only in D_j , $j = 2, \dots, m$, and \mathcal{G}_1 with nodes $V_1 = C_1$ is complete.

Appendix B: Link between variogram and covariance matrices

For $k \in V = \{1, \dots, d\}$, we denote by \mathcal{P}_{d-1}^k the set of all strictly positive definite covariance matrices $\Sigma^{(k)} \subset \mathbb{R}^{(d-1) \times (d-1)}$ indexed by $V \setminus \{k\}$. In contrast, the space of strictly conditionally negative definite $d \times d$ matrices is denoted by




Fig. 8. Decomposable graph with set of nodes $V = \{1, \dots, 6\}$: the cliques of the graph are $\{1, 2\}$, $\{2, 3, 4, 5\}$ and $\{4, 5, 6\}$; the separators are $\{2\}$ and $\{4, 5\}$

$$\mathcal{D}_d = \left\{ \Gamma \in [0, \infty)^{d \times d} : \mathbf{a}^\top \Gamma \mathbf{a} < 0 \text{ for all } \mathbf{a} \in \mathbb{R}^d \setminus \{0\} \text{ with } \sum_{i \in V} a_i = 0, \Gamma_{ii} = 0, \Gamma_{ij} = \Gamma_{ji} \text{ for all } i, j \in V \right\}.$$

Lemma 3. For any $k \in V$, there is a bijection $\varphi_k : \mathcal{D}_d \rightarrow \mathcal{P}_{d-1}^k$ given by

$$\begin{aligned} \varphi_k : \Gamma &\mapsto \frac{1}{2} \{ \Gamma_{ik} + \Gamma_{jk} - \Gamma_{ij} \}_{i, j \neq k}, \\ \varphi_k^{-1} : \Sigma^{(k)} &\mapsto \mathbf{1} \text{diag}(\tilde{\Sigma}^{(k)})^\top + \text{diag}(\tilde{\Sigma}^{(k)}) \mathbf{1}^\top - 2\tilde{\Sigma}^{(k)}, \end{aligned} \quad (45)$$

where $\tilde{\Sigma}^{(k)}$ is the $d \times d$ matrix that coincides with $\Sigma^{(k)}$ for $i, j \neq k$ and that has 0s in the k th column and row.

Proof. It is easy to check that the mappings are their mutual inverses. To see that the strict positive definiteness of $\Sigma^{(k)}$ is equivalent to the strict conditionally negative definiteness of Γ , we observe for any $\mathbf{a}_{\setminus k} \in \mathbb{R}^{d-1} \setminus \{0\}$ and $a_k = -\sum_{i \neq k} a_i$

$$\mathbf{a}_{\setminus k}^\top \Sigma^{(k)} \mathbf{a}_{\setminus k} = \frac{1}{2} \sum_{i, j \neq k} a_i a_j (\Gamma_{ik} + \Gamma_{jk} - \Gamma_{ij}) = - \sum_{i \neq k} a_i a_k \Gamma_{ik} - \frac{1}{2} \sum_{i, j \neq k} a_i a_j \Gamma_{ij} = -\mathbf{a}^\top \Gamma \mathbf{a},$$

using the fact that Γ is symmetric and $\Gamma_{ii} = 0$ for all $i \in V$. The assertion then follows; see also the proof of lemma 3.2.1 in Berg *et al.* (1984).

Appendix C: Hüsler–Reiss densities on decomposable graphs

Corollary 2. Let $\mathcal{G} = (V, E)$ be a decomposable and connected graph, and suppose that \mathbf{Y} is a Hüsler–Reiss Pareto distribution that satisfies the pairwise Markov property

$$Y_i \perp_{\mathbf{e}} Y_j | \mathbf{Y}_{\setminus \{i, j\}} \quad \text{if } (i, j) \notin E.$$

Then the density of \mathbf{Y} factorizes according to \mathcal{G} into lower dimensional Hüsler–Reiss densities, i.e.

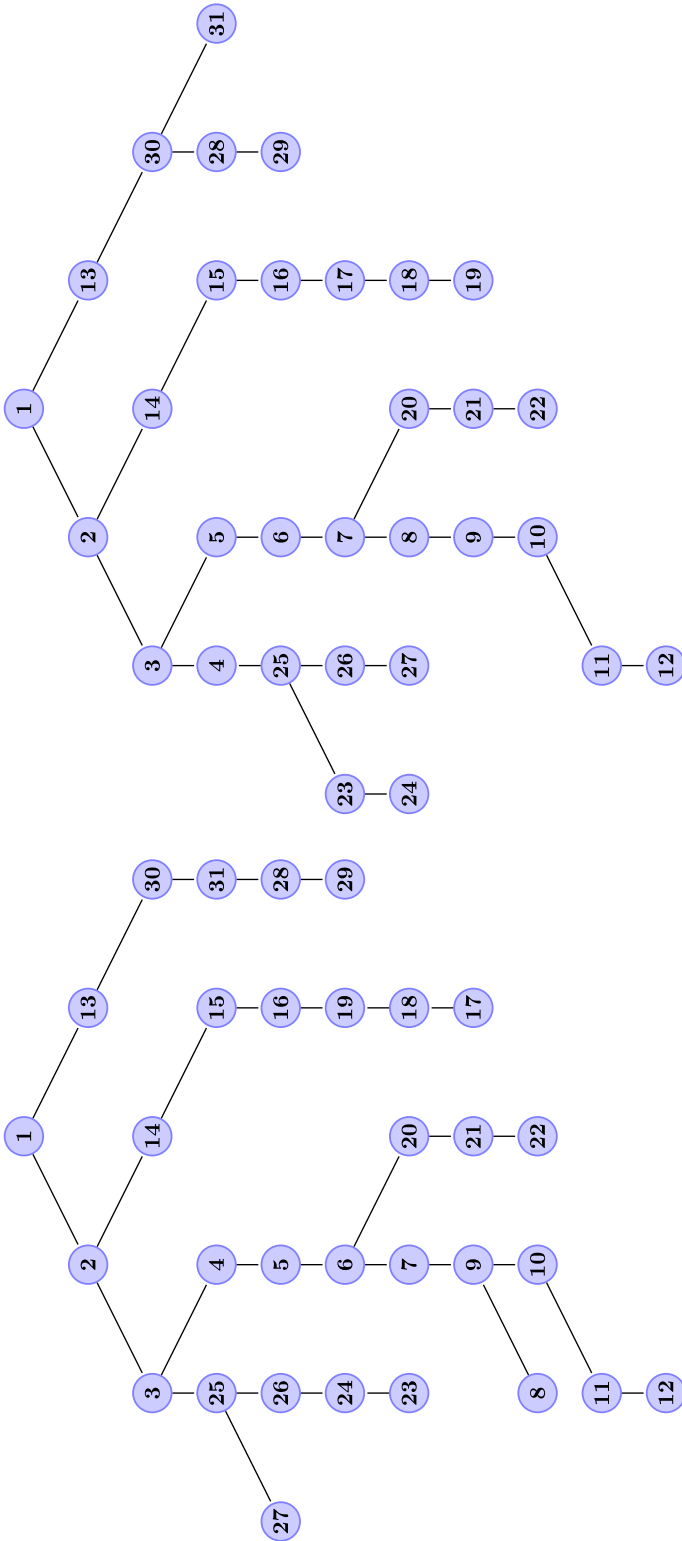
$$f_{\mathbf{Y}}(\mathbf{y}) = \frac{y_{k_1}^{-2} \prod_{j \neq k_1} y_j^{-1} \prod_{i=1}^m \phi_{|C_i|-1} \{ \log(\mathbf{y}_{C_i \setminus \{k_i\}} / y_{k_i}) + \Gamma_{C_i \setminus \{k_i\}, k_i; \Sigma_{C_i}^{(k_i)} \}}}{\Lambda(\mathbf{1}) \prod_{i=1}^{m-1} \phi_{|D_i|-1} \{ \log(\mathbf{y}_{D_i \setminus \{k_i\}} / y_{k_i}) + \Gamma_{D_i \setminus \{k_i\}, k_i; \Sigma_{D_i}^{(k_i)} \}}}, \quad \mathbf{y} \in \mathcal{L},$$

where the sequences of cliques $\{C_1, \dots, C_m\}$ and separator sets $\{D_2, \dots, D_m\}$ have the running intersection property (44), and $k_i \in D_i$, $i = 2, \dots, m$, $k_1 \in C_1$.

Proof. Theorem 1 and proposition 3 yield the factorization. It remains to show that the factors in front of the normal densities simplify to $y_{k_m-1}^{-2} \prod_{i \neq k_m-1} y_i^{-1}$. Indeed, since we choose $k_i \in D_i \subset C_i$, $i = 2, \dots, m$, the ratio $\lambda_{C_i}(\mathbf{y}_{C_i}) / \lambda_{D_i}(\mathbf{y}_{D_i})$ contributes the factor y_j^{-1} for all $j \in C_i \setminus D_i$, and each such j appears exactly once. For $i = 1$, the contribution of $\lambda_{C_1}(\mathbf{y}_{C_1})$ is $y_{k_1}^{-2} \prod_{i \in C_1 \setminus \{k_1\}} y_i^{-1}$.

Appendix D: Minimum spanning tree for the Danube river

Fig. 9(a) shows the estimated Hüsler–Reiss minimum spanning tree for the Danube data in Section 6 for a threshold u chosen as the 90% quantile of the marginal Pareto distribution. To assess the sensitivity of the tree structure with respect to the choice of threshold, we estimate the minimum spanning tree for thresholds u corresponding to a range of quantiles. The similarities of these trees in terms of the number of identical edges compared with the 90% quantile tree are shown in Fig. 10. We can see that there is some variation of the tree structure for different thresholds, but that most of the 30 edges are fairly stable throughout a wide range of thresholds. As a comparison, Fig. 9(b) shows the Gaussian minimum spanning tree fitted to all log-transformed data, using $\log(1 - \rho_{ij}^2)$ as distances in expression (38), where ρ_{ij} is the correlation coefficient between nodes $i, j \in V$. The Gaussian tree, which is a model for non-extremal data, is similar to the Hüsler–Reiss tree, which is a model for extreme flooding, but there are also some differences. For instance, for the extremal data the ordering of stations 16–19 seems to be less important since large discharges affect all at the same time. This is confirmed by the fact that, when the Hüsler–Reiss tree is extended to a block graph, then additional edges are introduced between these stations.



(a)

(b)

Fig. 9. (a) Estimated Hüsler–Reiss minimum spanning tree for the Danube data with 90% quantile threshold and (b) Gaussian minimum spanning tree using all log-transformed data

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48

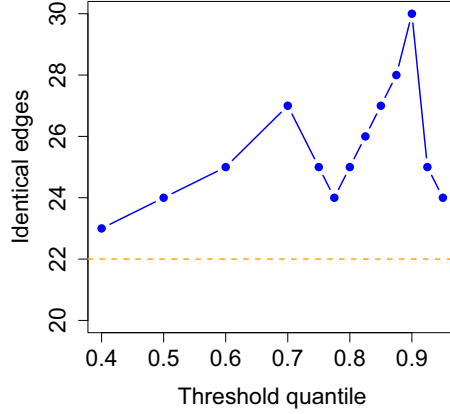


Fig. 10. For estimated minimum spanning trees corresponding to different threshold quantiles, the number of edges that are identical to the 90% quantile tree (●) and (---) the number of identical edges for a Gaussian minimum spanning tree by using all log-transformed data

Appendix E: Trivariate χ -coefficients

Fig. 11 shows the empirical estimates of the trivariate coefficients

$$\chi_{ijk} = 3 - \Lambda_{ij}(1, 1) - \Lambda_{ik}(1, 1) - \Lambda_{jk}(1, 1) + \Lambda_{ijk}(1, 1, 1), \quad i, j, k \in V,$$

against those implied by the fitted spatial model in Asadi *et al.* (2015) and our graphical model minimizing AIC.

Appendix F: Proofs

F.1. Proof (of proposition 1)

The implication that expression (17) implies part (a) is trivial. To show that part (a) implies part (b) let $k \in B$ and suppose that result (18) holds, i.e.

$$f^k(\mathbf{y}) = \frac{f_{A \cup B}^k(\mathbf{y}_{A \cup B}) f_{B \cup C}^k(\mathbf{y}_{B \cup C})}{f_B^k(\mathbf{y}_B)}, \quad \mathbf{y} \in \mathcal{L}^k.$$

For any $\mathbf{y} \in \mathcal{L}$ choose $0 < t < \min(y_k, 1)$, i.e. $\mathbf{y}/t \in \mathcal{L}^k$, and observe that

$$\begin{aligned} \lambda(\mathbf{y}) &= t^{-(d+1)} f^k(\mathbf{y}/t) \\ &= t^{-(d+1)} \frac{f_{A \cup B}^k(\mathbf{y}_{A \cup B}/t) f_{B \cup C}^k(\mathbf{y}_{B \cup C}/t)}{f_B^k(\mathbf{y}_B/t)} \\ &= t^{-(d+1)} \frac{\lambda_{A \cup B}(\mathbf{y}_{A \cup B}/t) \lambda_{B \cup C}(\mathbf{y}_{B \cup C}/t)}{\lambda_B(\mathbf{y}_B/t)} \\ &= \frac{\lambda_{A \cup B}(\mathbf{y}_{A \cup B}) \lambda_{B \cup C}(\mathbf{y}_{B \cup C})}{\lambda_B(\mathbf{y}_B)}, \end{aligned}$$

using the homogeneity of the λ_I , and the fact that $f_I^k(\mathbf{y}_I/t) = \lambda_I(\mathbf{y}_I/t)$ for any $I \subset V$ with $k \in I$. For this argument it is crucial that k is in an element of all three sets B , $A \cup B$ and $B \cup C$.

To show that part (b) implies expression (17) suppose that the factorization (19) of λ holds, and let $k \in V$. For all $\mathbf{y} \in \mathcal{L}^k$

$$f^k(\mathbf{y}) = \frac{\lambda_{A \cup B}(\mathbf{y}_{A \cup B}) \lambda_{B \cup C}(\mathbf{y}_{B \cup C})}{\lambda_B(\mathbf{y}_B)} = g(\mathbf{y}_{A \cup B}) h(\mathbf{y}_{B \cup C}),$$

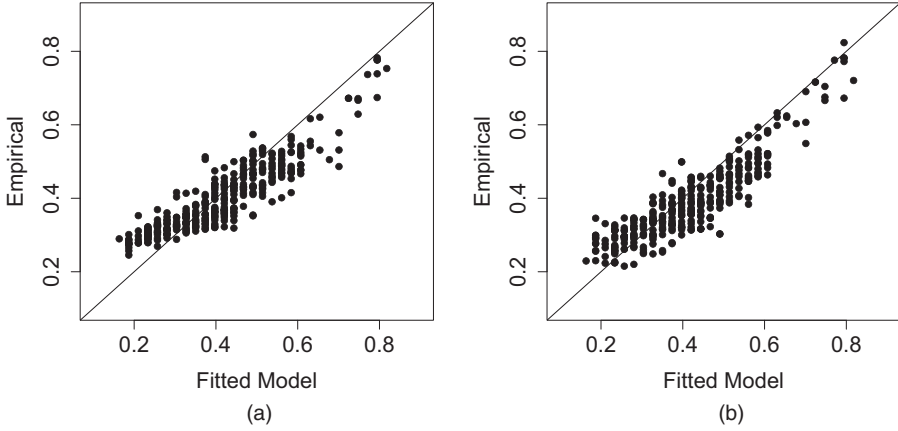


Fig. 11. Empirical estimates of the trivariate coefficients χ_{ijk} , $i, j, k \in V$, against those implied by (a) the fitted spatial model in Asadi *et al.* (2015) and (b) our graphical model minimizing AIC: only coefficients for 400 randomly selected triplets are shown

for suitable functions g and h , implying the required conditional independence of f^k (see Lauritzen (1996), chapter 3). This shows that condition (17) indeed holds and thus $\mathbf{Y}_A \perp_e \mathbf{Y}_C | \mathbf{Y}_B$.

F.2. Proof (of theorem 1)

We start by proving that, if \mathbf{Y} satisfies the pairwise Markov property relative to \mathcal{G} , then the graph \mathcal{G} is necessarily connected. Indeed, suppose that V can be split into non-empty disjoint subsets $V_1, V_2 \subset V$ such that for $(i, j) \in E$ it holds that either $i, j \in V_1$ or $i, j \in V_2$. For an arbitrary $k \in V$, by assumption, the pairwise Markov property relative to \mathcal{G} is satisfied for f^k on \mathcal{L}^k and the classical Hammersley–Clifford theorem implies the global Markov property for f^k , and in particular

$$f^k(\mathbf{y}) = f_{V_1}^k(\mathbf{y}_{V_1}) f_{V_2}^k(\mathbf{y}_{V_2}), \quad \mathbf{y} \in \mathcal{L}^k.$$

The discussion after proposition 1 shows that such a factorization contradicts integrability of the multivariate Pareto density, and therefore the graph must be connected.

We now show that part (a) implies part (c). The pairwise Markov property of f^k relative to \mathcal{G} implies by the classical Hammersley–Clifford theorem that

$$f^k(\mathbf{y}) = \frac{\prod_{C \in \mathcal{C}} f_C^k(\mathbf{y}_C)}{\prod_{D \in \mathcal{D}} f_D^k(\mathbf{y}_D)}, \quad \mathbf{y} \in \mathcal{L}^k.$$

This representation is not of direct use since it cannot be extended to $f_{\mathbf{Y}}$ on the whole space \mathcal{L} , since all f_i^k with $k \notin I$ are not homogeneous. The result, however, tells us that \mathbf{Y}^k also satisfies the global Markov property on \mathcal{L}^k relative to \mathcal{G} , as defined in Section 2.3. The running intersection property implies that D_m separates $C_m \setminus D_m$ from $(C_1 \cup \dots \cup C_{m-1}) \setminus D_m$. Choose $k \in D_m$; then the global Markov property for \mathbf{Y}^k yields

$$f^k(\mathbf{y}) = \frac{f_{C_m}^k(\mathbf{y}_{C_m}) f_{C_1 \cup \dots \cup C_{m-1}}^k(\mathbf{y}_{C_1 \cup \dots \cup C_{m-1}})}{f_{D_m}^k(\mathbf{y}_{D_m})} = \frac{\lambda_{C_m}(\mathbf{y}_{C_m}) \lambda_{C_1 \cup \dots \cup C_{m-1}}(\mathbf{y}_{C_1 \cup \dots \cup C_{m-1}})}{\lambda_{D_m}(\mathbf{y}_{D_m})}, \quad \mathbf{y} \in \mathcal{L}^k,$$

where the second equality holds since $k \in D_m$, and D_m is a subset of both C_m and $C_1 \cup \dots \cup C_{m-1}$. By a homogeneity argument similar to the proof of proposition 1, this factorization extends to λ on the whole space \mathcal{L} , i.e.

$$\lambda(\mathbf{y}) = \frac{\lambda_{C_m}(\mathbf{y}_{C_m}) \lambda_{C_1 \cup \dots \cup C_{m-1}}(\mathbf{y}_{C_1 \cup \dots \cup C_{m-1}})}{\lambda_{D_m}(\mathbf{y}_{D_m})}, \quad \mathbf{y} \in \mathcal{L}.$$

It remains to decompose $\lambda_{C_1 \cup \dots \cup C_{m-1}}$ in the same manner. For this, choose a new $k \in D_{m-1}$ and note that

$$f_{C_1 \cup \dots \cup C_{m-1}}^k(\mathbf{y}_{C_1 \cup \dots \cup C_{m-1}}) = \int_{[0, \infty)^{C_m \setminus D_m}} \frac{\prod_{C \in \mathcal{C}} f_C^k(\mathbf{y}_C)}{\prod_{D \in \mathcal{D}} f_D^k(\mathbf{y}_D)} d\mathbf{y}_{C_m \setminus D_m} = \frac{\prod_{C \neq C_m} f_C^k(\mathbf{y}_C)}{\prod_{D \neq D_m} f_D^k(\mathbf{y}_D)},$$

and therefore satisfies the global Markov property relative to the subgraph induced on $C_1 \cup \dots \cup C_{m-1}$. Since $f_{C_1 \cup \dots \cup C_{m-1}}^k = \lambda_{C_1 \cup \dots \cup C_{m-1}}$ on \mathcal{L}^k , applying successively the same reasoning as before yields the factorization of λ that directly implies representation (21) for f_Y .

To show that part (c) implies part (b), we need only to verify that \mathbf{Y}^k satisfies the global Markov property on \mathcal{L}^k for any $k \in V$. For disjoint sets $A, B, C \subset V$ such that B separates A from C , factorization (21) entails that

$$f^k(\mathbf{y}) = \Lambda(\mathbf{1}) f_Y(\mathbf{y}) = g(\mathbf{y}_{A \cup B}) h(\mathbf{y}_{B \cup C}),$$

for suitable functions g and h , and thus $\mathbf{Y}_A^k \perp\!\!\!\perp \mathbf{Y}_C^k \mid \mathbf{Y}_B^k$.

The implication that part (b) implies part (a) holds trivially.

F.3. Proof (of corollary 1)

It is easy to check that λ and f_Y are homogeneous of order $-(d+1)$ on \mathcal{L} . Let $\{C_1, \dots, C_m\}$ and $\{D_2, \dots, D_m\}$ be the sequences of cliques and separators with the running intersection property (44). Sequential integration of the function f_Y on $C_m \setminus D_m, \dots, C_2 \setminus D_2$, together with the consistency constraint, yields that it defines in fact a probability density. Theorem 1 implies that the corresponding distribution on \mathcal{L} satisfies the Markov property relative to \mathcal{G} .

F.4. Proof (of proposition 2)

The density of the random vector on the right-hand side of expression (24) is

$$y_k^{-2} \prod_{e=(i,j) \in E^k} y_i^{-1} f_{U_e}(y_j/y_i) = y_k^{-2} \frac{\prod_{(i,j) \in E^k} \lambda_{ij}(y_i, y_j)}{\prod_{(i,j) \in E^k} y_i^{-2}} = \prod_{\{i,j\} \in E} \frac{\lambda_{ij}(y_i, y_j)}{y_i^{-2} y_j^{-2}} \prod_{i \in V} y_i^{-2},$$

where we used expression (12) for the first equation, and the fact that each node $i \in V \setminus \{k\}$ has exactly one incoming arrow, and the k th node has no incoming arrows. However, we recall that the density of \mathbf{Y}^k is $\lambda(\mathbf{y}) = \Lambda(\mathbf{1}) f_Y(\mathbf{y})$, which factorizes with respect to the tree \mathcal{T} . Comparing the above density with expression (23) yields the result.

F.5. Proof (of lemma 1)

Without losing generality, we may and do assume that $k' = 1$ and $k = 2$. Let the vector $\mathbf{W}^1 = (0, W_2^1, \dots, W_d^1)$ have a centred normal distribution with covariance matrix $\Sigma = \{\sigma_{ij}\} = \tilde{\Sigma}^{(1)}$, such that

$$\Sigma^{(1)} = \Sigma_{\setminus \{1\}} = \begin{pmatrix} \sigma_{22} & \Sigma_{2, \setminus \{1,2\}} \\ \Sigma_{\setminus \{1,2\}, 2} & \Sigma_{\setminus \{1,2\}} \end{pmatrix}.$$

The precision matrix is obtained by blockwise inversion as

$$\Theta^{(1)} = \begin{pmatrix} \sigma_{22}^{-1} + \sigma_{22}^{-2} \Sigma_{2, \setminus \{1,2\}} S^{-1} \Sigma_{\setminus \{1,2\}, 2} & -\sigma_{22}^{-1} \Sigma_{2, \setminus \{1,2\}} S^{-1} \\ -\sigma_{22}^{-1} S^{-1} \Sigma_{\setminus \{1,2\}, 2} & S^{-1} \end{pmatrix},$$

where $S = \Sigma_{\setminus \{1,2\}} - \sigma_{22}^{-1} \Sigma_{\setminus \{1,2\}, 2} \Sigma_{2, \setminus \{1,2\}}$ is the Schur complement of upper left-hand block σ_{22} in the matrix $\Sigma^{(1)}$. The random vector \mathbf{W}^1 can be transformed into

$$\mathbf{W}^2 = (-W_2^1, 0, W_3^1 - W_2^1, \dots, W_d^1 - W_2^1),$$

which is readily verified to have a centred normal distribution with covariance matrix $\tilde{\Sigma}^{(2)}$. In contrast, we may write the covariance matrix $\Sigma^{(2)}$ of $(-W_2^1, W_3^1 - W_2^1, \dots, W_d^1 - W_2^1)$ in terms of Σ as

$$\Sigma^{(2)} = \begin{pmatrix} \sigma_{22} & \sigma_{22}\mathbf{1}^T - \Sigma_{2, \setminus \{1,2\}} \\ \sigma_{22}\mathbf{1} - \Sigma_{\setminus \{1,2\}, 2} & \Sigma_{\setminus \{1,2\}} + \sigma_{22}\mathbf{1}\mathbf{1}^T - \Sigma_{\setminus \{1,2\}, 2}\mathbf{1}^T - \mathbf{1}\Sigma_{2, \setminus \{1,2\}} \end{pmatrix}.$$

It can be checked that the Schur complement of the upper left-hand block σ_{22} in the matrix $\Sigma^{(2)}$ is again S . Thus, blockwise inversion yields

$$\Theta^{(2)} = \begin{pmatrix} \sigma_{22}^{-1} + \sigma_{22}^{-2}(\sigma_{22}\mathbf{1}^T - \Sigma_{2, \setminus \{1,2\}})S^{-1}(\sigma_{22}\mathbf{1} - \Sigma_{\setminus \{1,2\}, 2}) & -\sigma_{22}^{-1}(\sigma_{22}\mathbf{1}^T - \Sigma_{2, \setminus \{1,2\}})S^{-1} \\ -\sigma_{22}^{-1}S^{-1}(\sigma_{22}\mathbf{1} - \Sigma_{\setminus \{1,2\}, 2}) & S^{-1} \end{pmatrix}.$$

Comparing these representations of $\Theta^{(1)}$ and $\Theta^{(2)}$ yields the assertion for $i, j \in V \setminus \{1, 2\}$. For $i \neq 2$ and $j = 2$, we observe that

$$\sum_{l \neq 2} \Theta_{il}^{(2)} = - \sum_{m \neq 1, 2} S_{im}^{-1} + \sigma_{22}^{-1} \sum_{m \neq 1, 2} S_{im}^{-1} \sigma_{m2} + \sum_{m \neq 1, 2} S_{im}^{-1} = -\Theta_{i2}^{(1)}.$$

The case $i, j = 2$ follows similarly.

F.6. Proof (of proposition 3)

Let $i, j \in V$ with $i \neq j$ be fixed and choose a $k \neq i, j$. Let P and \mathbf{W} be as in representation (28). Since $Y_k^k = P$ and because of the independence of P and \mathbf{W} we obtain

$$\begin{aligned} Y_i^k \perp\!\!\!\perp Y_j^k | \mathbf{Y}_{\setminus \{i,j\}}^k &\Leftrightarrow P \times e^{W_i^k - \Gamma_{ik}/2} \perp\!\!\!\perp P \times e^{W_j^k - \Gamma_{jk}/2} | P, \mathbf{W}_{\setminus \{i,j,k\}}^k \\ &\Leftrightarrow W_i^k \perp\!\!\!\perp W_j^k | \mathbf{W}_{\setminus \{i,j,k\}}^k \\ &\Leftrightarrow \Theta_{ij}^{(k)} = 0, \end{aligned}$$

where the variable W_k^k can be deleted from the conditioning since it is deterministic given P , and therefore the reduced precision matrix $\Theta^{(k)}$ of the vector $\mathbf{W}_{\setminus k}^k$ appears. The last equivalence follows from the well-known fact that conditional independence in multivariate normal models corresponds to 0s in the precision matrix (see example 4).

Let now $k = i \neq j$ and choose a $k' \notin \{i, j\}$. Lemma 1 implies that

$$-\sum_{l \neq k} \Theta_{jl}^{(k)} = \Theta_{jk}^{(k')}. \quad (46)$$

Since $k' \in V \setminus \{i, j\}$, by proposition 1, $Y_i \perp_e Y_j | \mathbf{Y}_{\setminus \{i,j\}}^{k'}$ is equivalent to $Y_i^{k'} \perp\!\!\!\perp Y_j^{k'} | \mathbf{Y}_{\setminus \{k,j\}}^{k'}$. The latter, by the first part of the proof, is then equivalent to $\Theta_{jk}^{(k')} = 0$, which, together with equation (46), yields the assertion. The case $k = j \neq i$ is analogous by symmetry.

F.7. Proof (of proposition 4)

Let C_1, \dots, C_m be an enumeration of the cliques of the decomposable connected graph $\mathcal{G} = (V, E)$. Recall that, by assumption, all intersections between pairs of cliques are either empty or contain a single node. We show how to obtain the unique $(d \times d)$ -dimensional variogram matrix Γ that solves the completion problem (31) by adding one clique after another. We first set

$$\Gamma_{ij} = \Gamma_{ij}^{(C_1)}, \quad \text{for } i, j \in C_1. \quad (47)$$

Let $I_{p-1} = C_1 \cup \dots \cup C_{p-1}$ be the union of the first $p-1$ cliques, $2 \leq p \leq m$ cliques that have been chosen in an order such that \mathcal{G} restricted to I_{p-1} forms a connected graph. Suppose that we have already constructed a unique $(|I_{p-1}| \times |I_{p-1}|)$ -dimensional variogram matrix $\Gamma^{(I_{p-1})}$ that satisfies

$$\begin{aligned} \Gamma_{ij}^{(I_{p-1})} &= \Gamma_{ij}^{(C_l)}, & \text{for } i, j \in C_l \text{ and all } l = 1, \dots, p-1, \\ \Theta_{ij}^{(I_{p-1}, k)} &= 0, & \text{for all } i, j, k \in I_{p-1}, i, j \neq k \text{ and } (i, j) \notin E, \end{aligned} \quad (48)$$

where here and in what follows we use the notation $\Theta^{(J,k)}$ as the inverse of $\Sigma^{(J,k)} = \varphi_k(\Gamma^{(J)})$ for a variogram matrix $\Gamma^{(J)}$ on some index set $J \subset V$ and $k \in J$. We next choose a clique, say C_p , that intersects I_{p-1} , and this intersection must be a single node, say $k_0 \in V$. Let $I_p = I_{p-1} \cup C_p$ and define the matrix

$$\Theta^{(I_p,k_0)} = \begin{pmatrix} \Theta^{(I_{p-1},k_0)} & 0 \\ 0 & \Theta^{(C_p,k_0)} \end{pmatrix}. \tag{49}$$

This matrix is an invertible covariance matrix since its blocks are invertible covariance matrices, and its inverse $\Sigma^{(I_p,k_0)}$ has the same property with blocks $\Sigma^{(I_{p-1},k_0)}$ and $\Sigma^{(C_p,k_0)}$. This yields an $(|I_p| \times |I_p|)$ -dimensional variogram matrix $\Gamma^{(I_p)}$ through the mapping $\varphi_{k_0}^{-1}$, which has the form

$$\Gamma_{ij}^{(I_p)} = \begin{cases} \Gamma_{ij}^{(I_{p-1})}, & \text{for } i, j \in I_{p-1}, \\ \Gamma_{ij}^{(C_p)}, & \text{for } i, j \in C_p, \\ \Gamma_{ik_0}^{(I_{p-1})} + \Gamma_{jk_0}^{(I_{p-1})}, & \text{for } i \in I_{p-1}, j \in C_p \text{ or } j \in I_{p-1}, i \in C_p. \end{cases}$$

This variogram matrix clearly solves problem (48) with I_{p-1} replaced by I_p . It is unique by construction and the fact that φ_{k_0} and $\varphi_{k_0}^{-1}$ are bijections.

Starting with expression (47) and then adding all cliques for $p = 2, \dots, m$ according to the above procedure, we obtain a unique $(d \times d)$ -dimensional variogram $\Gamma = \Gamma^{(I_m)}$ matrix that satisfies all constraints in expression (31). Comparing with corollary 2 it follows that the corresponding density in expression (30) is d -variate Hüsler–Reiss with parameter matrix Γ .

F.8. Proof (of lemma 2)

The general formula for extremal functions in proposition 1 in Dombry *et al.* (2016) can be written in terms of the exponent measure density λ as

$$\begin{aligned} \mathbb{P}(\mathbf{U}^k \in A) &= \int_{\mathcal{E}} \mathbf{1}\{\mathbf{y}/y_k \in A\} \mathbf{1}\{y_k > 1\} \lambda(\mathbf{y}) d\mathbf{y} \\ &= \int_{\mathcal{L}^k} \mathbf{1}\{\mathbf{y}/y_k \in A\} f^k(\mathbf{y}) d\mathbf{y} \\ &= \mathbb{P}(\mathbf{Y}^k/Y_k^k \in A). \end{aligned}$$

Since the density of $\mathbf{U}_{\setminus k}^k = \mathbf{Y}_{\setminus k}^k/Y_k^k$ is readily seen to be $\lambda(\mathbf{y})$ for $\mathbf{y}_{\setminus k} \in [0, \infty)^{d-1}$ and $y_k = 1$, it follows with

$$\lambda(\mathbf{y}) = y_k^{-(d+1)} \lambda(\mathbf{y}/y_k) = y_k^{-(d+1)} f_{\mathbf{U}_{\setminus k}^k}(\mathbf{y}_{\setminus k}/y_k), \quad \mathbf{y} \in \mathcal{E},$$

that expression (41) is an equivalent definition of extremal functions.

It follows from theorem 2 in Dombry *et al.* (2016) that, for a uniform distribution T on $\{1, \dots, d\}$, the random vector $\mathbf{Y}^T/\|\mathbf{Y}^T\|_1$ follows the distribution of the spectral measure H on $S_{d-1} = \{\mathbf{x} \in \mathcal{E} : \|\mathbf{x}\|_1 = 1\}$ associated with the max-stable distribution \mathbf{Z} , i.e.

$$\Lambda(A) = d \int_{S_{d-1}} \int_0^\infty u^{-2} \mathbf{1}\{u\mathbf{w} \in A\} du H(d\mathbf{w}), \quad A \subset \mathcal{E}.$$

If $A \subset \mathcal{L}$, then $u\mathbf{w} \in A$ implies that $u \geq 1$, and therefore

$$\begin{aligned} \mathbb{P}\left(\frac{P\mathbf{Y}^T}{\|\mathbf{Y}^T\|_1} \in A\right) &= \int_{S_{d-1}} \int_1^\infty f_P(u) \mathbf{1}\{u\mathbf{w} \in A\} du H(d\mathbf{w}) \\ &= \frac{1}{d} \int_A \lambda(\mathbf{y}) d\mathbf{y}, \end{aligned}$$

since $f_P(u) = 1/u^2$, $u \geq 1$. For $A = \mathcal{L} = \mathcal{E} \setminus [\mathbf{0}, \mathbf{1}]$ this yields for the conditioning event in expression (42)

$$\mathbb{P}\left(\frac{P\|\mathbf{Y}^T\|_\infty}{\|\mathbf{Y}^T\|_1} > 1\right) = \frac{\Lambda(\mathcal{L})}{d} = \frac{\Lambda(\mathbf{1})}{d}. \tag{50}$$

Since \mathbf{Y} has density $\lambda(\mathbf{y})/\Lambda(\mathbf{1})$, this concludes the proof.

References

- Arlot, S. and Celisse, A. (2010) A survey of cross-validation procedures for model selection. *Statist. Surv.*, **4**, 40–79.
- Asadi, P., Davison, A. C. and Engelke, S. (2015) Extremes on river networks. *Ann. Appl. Statist.*, **9**, 2023–2050.
- Ballani, F. and Schlather, M. (2011) A construction principle for multivariate extreme value distributions. *Biometrika*, **98**, 633–645.
- Basrak, B. and Segers, J. (2009) Regularly varying multivariate time series. *Stoch. Processes Appl.*, **119**, 1055–1080.
- Beirlant, J., Goegebeur, Y., Teugels, J. and Segers, J. (2004) *Statistics of Extremes*. Chichester: Wiley.
- Berg, C., Christensen, J. P. R. and Ressel, P. (1984) Theory of positive definite and related functions. In *Harmonic Analysis on Semigroups*. New York: Springer.
- Böhm, O. and Wetzel, K.-F. (2006) Flood history of the Danube tributaries Lech and Isar in the alpine foreland of Germany. *Hydrol. Sci. J.*, **51**, 784–798.
- Boldi, M.-O. and Davison, A. C. (2007) A mixture model for multivariate extremes. *J. R. Statist. Soc. B*, **69**, 217–229.
- Brown, B. M. and Resnick, S. I. (1977) Extreme values of independent stochastic processes. *J. Appl. Probab.*, **14**, 732–739.
- de Carvalho, M. and Davison, A. C. (2014) Spectral density ratio models for multivariate extremes. *J. Am. Statist. Ass.*, **109**, 764–776.
- Chow, C. and Liu, C. (1968) Approximating discrete probability distributions with dependence trees. *IEEE Trans. Inform. Theory*, **14**, 462–467.
- Coles, S., Heffernan, J. and Tawn, J. (1999) Dependence measures for extreme value analyses. *Extremes*, **2**, 339–365.
- Coles, S. G. and Tawn, J. A. (1991) Modelling extreme multivariate events. *J. R. Statist. Soc. B*, **53**, 377–392.
- Cooley, D., Davis, R. A. and Naveau, P. (2010) The pairwise beta distribution: a flexible parametric multivariate model for extremes. *J. Multiv. Anal.*, **101**, 2103–2117.
- Cooley, D. and Thibaud, E. (2018) Decompositions of dependence for high-dimensional extremes. *Preprint* (Available from <http://arxiv.org/abs/1612.07190>.)
- Cowell, R. G., Dawid, P., Lauritzen, S. L. and Spiegelhalter, D. J. (2006) *Probabilistic Networks and Expert Systems: Exact Computational Methods for Bayesian Networks*. Berlin: Springer.
- Davison, A. C., Padoan, S. A. and Ribatet, M. (2012) Statistical modeling of spatial extremes. *Statist. Sci.*, **27**, 161–186.
- Dawid, A. P. (1979) Conditional independence in statistical theory (with discussion). *J. R. Statist. Soc. B*, **41**, 1–31.
- Dawid, A. P. and Lauritzen, S. L. (1993) Hyper-Markov laws in the statistical analysis of decomposable graphical models. *Ann. Statist.*, **21**, 1272–1317.
- Dempster, A. P. (1972) Covariance selection. *Biometrics*, **28**, 157–175.
- Dieker, A. B. and Mikosch, T. (2015) Exact simulation of Brown–Resnick random fields at a finite number of locations. *Extremes*, **18**, 301–314.
- Dombry, C., Engelke, S. and Oesting, M. (2016) Exact simulation of max-stable processes. *Biometrika*, **103**, 303–317.
- Dombry, C., Engelke, S. and Oesting, M. (2017) Bayesian inference for multivariate extreme value distributions. *Electron. J. Statist.*, **11**, 4813–4844.
- Dombry, C. and Éyi-Minko, F. (2013) Regular conditional distributions of continuous max-infinitely divisible random fields. *Electron. J. Probab.*, **18**.
- Dombry, C. and Éyi-Minko, F. (2014) Stationary max-stable processes with the Markov property. *Stoch. Processes Appl.*, **124**, 2266–2279.
- Dombry, C., Éyi-Minko, F. Ribatet, M. (2013) Conditional simulation of max-stable processes. *Biometrika*, **100**, 111–124.
- Einmahl, J. H. J., Kiriliouk, A., Krajina, A. and Segers, J. (2016) An M -estimator of spatial tail dependence. *J. R. Statist. Soc. B*, **78**, 275–298.
- Einmahl, J. H. J., Krajina, A. and Segers, J. (2012) An M -estimator for tail dependence in arbitrary dimensions. *Ann. Statist.*, **40**, 1764–1793.
- Embrechts, P., Klüppelberg, C. and Mikosch, T. (1997) *Modelling Extremal Events: for Insurance and Finance*. London: Springer.
- Engelke, S., Hitz, S. A. and Gnecco, N. (2019) graphicalExtremes: statistical methodology for graphical extreme value models. *R Package Version 0.1.0*. (Available from <https://CRAN.R-project.org/package=graphicalExtremes>.)
- Engelke, S., Malinowski, A., Kabluchko, Z. and Schlather, M. (2015) Estimation of Hüsler–Reiss distributions and Brown–Resnick processes. *J. R. Statist. Soc. B*, **77**, 239–265.
- de Fondeville, R. and Davison, A. C. (2018) High-dimensional peaks-over-threshold inference. *Biometrika*, **105**, 575–592.
- Genet, C., Ghoudi, K. and Rivest, L.-P. (1995) A semiparametric estimation procedure of dependence parameters in multivariate families of distributions. *Biometrika*, **82**, 543–552.

- 1 Gissibl, N. and Klüppelberg, C. (2018) Max-linear models on directed acyclic graphs. *Bernoulli*, **24**, 2693–2720.
- 2 Gissibl, N., Klüppelberg, C. and Otto, M. (2018) Tail dependence of recursive max-linear models with regularly
 3 varying noise variables. *Econometr. Statist.*, **6**, 149–167.
- 4 Gnecco, N., Meinshausen, N., Peters, J. and Engelke, S. (2019) Causal discovery in heavy-tailed models. *Preprint*.
 5 (Available from <https://arxiv.org/abs/1908.05097>.)
- 6 Green, P. J. and Thomas, A. (2013) Sampling decomposable graphs using a Markov chain on junction trees.
 7 *Biometrika*, **100**, 91–110.
- 8 Gudendorf, G. and Segers, J. (2010) Extreme-value copulas. In *Copula Theory and Its Applications*, pp. 127–145.
 9 Berlin: Springer.
- 10 de Haan, L. (1984) A spectral representation for max-stable processes. *Ann. Probab.*, **12**, 1194–1204.
- 11 de Haan, L. and Ferreira, A. (2006) *Extreme Value Theory*. New York: Springer.
- 12 Harary, F. (1963) A characterization of block-graphs. *Can. Math. Bull.*, **6**, 1–6.
- 13 Heffernan, J. E. and Tawn, J. A. (2004) A conditional approach for multivariate extreme values (with discussion).
 14 *J. R. Statist. Soc. B*, **66**, 497–546.
- 15 Hitz, S. A. and Evans, J. R. (2016) One-component regular variation and graphical modeling of extremes. *J. Appl.*
 16 *Probab.*, **53**, 733–746.
- 17 Huser, R., Dombry, C., Ribatet, M. and Genton, M. G. (2019) Full likelihood inference for max-stable data. *Stat.*,
 18 **8**, no. 1, article e218.
- 19 Hüsler, J. and Reiss, R.-D. (1989) Maxima of normal random vectors: between independence and complete
 20 dependence. *Statist. Probab. Lett.*, **7**, 283–286.
- 21 Janssen, A. and Segers, J. (2014) Markov tail chains. *J. Appl. Probab.*, **51**, 1133–1153.
- 22 Joe, H. (2015) *Dependence Modeling with Copulas*. Boca Raton: CRC Press.
- 23 Kabluchko, Z., Schlather, M. and de Haan, L. (2009) Stationary max-stable fields associated to negative definite
 24 functions. *Ann. Probab.*, **37**, 2042–2065.
- 25 Katz, R. W., Parlange, M. B. and Naveau, P. (2002) Statistics of extremes in hydrology. *Adv. Wat. Resour.*, **25**,
 26 1287–1304.
- 27 Keef, C., Tawn, J. and Svensson, C. (2009) Spatial risk assessment for extreme river flows. *Appl. Statist.*, **58**,
 28 601–618.
- 29 Keef, C., Tawn, J. A. and Lamb, R. (2013) Estimating the probability of widespread flood events. *Environmetrics*,
 30 **24**, 13–21.
- 31 Kellerer, H. G. (1964) Verteilungsfunktionen mit gegebenen Marginalverteilungen. *Z. Wahrsch. Ver. Geb.*, **3**,
 32 247–270.
- 33 Kiriliouk, A., Rootzén, H., Segers, J. and Wadsworth, J. L. (2018) Peaks over thresholds modeling with multivariate
 34 generalized Pareto distributions. *Technometrics*, 1–13.
- 35 Kiriliouk, A., Segers, J. and Tafakori, L. (2018) An estimator of the stable tail dependence function based on the
 36 empirical beta copula. *Extremes*, **21**, 581–600.
- 37 Kruskal, Jr, J. B. (1956) On the shortest spanning subtree of a graph and the traveling salesman problem. *Proc.*
 38 *Am. Math. Soc.*, **7**, 48–50.
- 39 Kulik, R. and Soulier, P. (2015) Heavy tailed time series with extremal independence. *Extremes*, **18**, 273–299.
- 40 Lafferty, J., Liu, H. and Wasserman, L. (2012) Sparse nonparametric graphical models. *Statist. Sci.*, **27**, 519–537.
- 41 Lauritzen, S. L. (1996) *Graphical Models*. Oxford: Oxford University Press.
- 42 Ledford, A. W. and Tawn, J. A. (1997) Modelling dependence within joint tail regions. *J. R. Statist. Soc. B*, **59**,
 43 475–499.
- 44 Lee, D. and Joe, H. (2018) Multivariate extreme value copulas with factor and tree dependence structures. *Ex-*
 45 *tremes*, **21**, 147–176.
- 46 Loh, P.-L. and Wainwright, M. (2013) Structure estimation for discrete graphical models: generalized covariance
 47 matrices and their inverses. *Ann. Statist.*, **41**, 3022–3049.
- 48 Marcon, G., Padoan, S., Naveau, P., Muliere, P. and Segers, J. (2017) Multivariate nonparametric estimation of
 the Pickands dependence function using Bernstein polynomials. *J. Statist. Planng Inf.*, **183**, 1–17.
- McNeil, A. J., Frey, R. and Embrechts, P. (2015) *Quantitative Risk Management: Concepts, Techniques and Tools*.
 Princeton: Princeton University Press.
- Mhalla, L., Chavez-Demoulin, V. and Dupuis, D. J. (2019) Causal mechanism of extreme river discharges in the
 upper Danube basin network. *Preprint*. (Available from <https://arxiv.org/abs/1907.03555>.)
- Min, S.-K., Zhang, X., Zwiers, F. and Hegerl, G. (2011) Human contribution to more-intense precipitation
 extremes. *Nature*, **470**, 378–381.
- Naveau, P., Ribes, A., Zwiers, F., Hannart, A., Tuel, A. and Yiou, P. (2018) Revising return periods for record
 events in a climate event attribution context. *J. Clim.*, **31**, 3411–3422.
- Oesting, M., Schlather, M. and Zhou, C. (2018) Exact and fast simulation of max-stable processes on a compact
 set using the normalized spectral representation. *Bernoulli*, **24**, 1497–1530.
- Opitz, T. (2013) Extremal t processes: elliptical domain of attraction and a spectral representation. *J. Multiv.*
Anal., **122**, 409–413.
- Papastathopoulos, I. and Strokorb, K. (2016) Conditional independence among max-stable laws. *Statist. Probab.*
Let., **108**, 9–15.

- 1 Papastathopoulos, I., Strokorb, K., Tawn, J. A. and Butler, A. (2017) Extreme events of Markov chains. *Adv.*
2 *Appl. Probab.*, **49**, 134–161.
- 3 Prim, R. C. (1957) Shortest connection networks and some generalizations. *Bell Sys. Tech. J.*, **36**, 1389–1401.
- 4 R Core Team (2019) *R: a Language and Environment for Statistical Computing*. Vienna: R Foundation for Statistical
5 Computing.
- 6 Resnick, S. I. (2008) *Extreme Values, Regular Variation and Point Processes*. New York: Springer.
- 7 Rootzén, H., Segers, J. and Wadsworth, J. L. (2018) Multivariate peaks over thresholds models. *Extremes*, **21**,
8 115–145.
- 9 Rootzén, H. and Tajvidi, N. (2006) Multivariate generalized Pareto distributions. *Bernoulli*, **12**, 917–930.
- 10 Rue, H. and Held, L. (2005) Theory and applications. In *Gaussian Markov Random Fields*. Boca Raton: Chapman
11 and Hall–CRC.
- 12 Schlather, M. (2002) Models for stationary max-stable random fields. *Extremes*, **5**, 33–44.
- 13 Schlather, M. and Tawn, J. A. (2003) A dependence measure for multivariate and spatial extreme values: properties
14 and inference. *Biometrika*, **90**, 139–156.
- 15 Segers, J. (2019) One- versus multi-component regular variation and extremes of Markov trees. *Preprint*. (Available
16 from <https://arxiv.org/abs/1902.02226>.)
- 17 Smith, R. L. (1992) The extremal index for a Markov chain. *J. Appl. Probab.*, **29**, 37–45.
- 18 Smith, R., Tawn, J. and Coles, S. (1997) Markov chain models for threshold exceedances. *Biometrika*, **84**, 249–268.
- 19 Thibaud, E., Aalto, J., Cooley, D. S., Davison, A. C. and Heikkinen, J. (2016) Bayesian inference for the Brown–
20 Resnick process, with an application to extreme low temperatures. *Ann. Appl. Statist.*, **10**, 2303–2324.
- 21 Wadsworth, J. L. and Tawn, J. A. (2014) Efficient inference for spatial extreme value processes associated to
22 log-Gaussian random functions. *Biometrika*, **101**, 1–15.
- 23 Wadsworth, J. L., Tawn, J. A., Davison, A. C. and Elton, D. M. (2017) Modelling across extremal dependence
24 classes. *J. R. Statist. Soc. B*, **79**, 149–175.
- 25 Wainwright, M. J. and Jordan, M. I. (2008) Graphical models, exponential families, and variational inference.
26 *Foundns Trends Mach. Learn.*, **1**, 1–305.
- 27 Yu, H., Uy, W. and Dauwels, J. (2017) Modeling spatial extremes via ensemble-of-trees of pairwise copulas. *IEEE*
28 *Trans. Signal Process.*, **65**, 571–586.
- 29
- 30
- 31
- 32
- 33
- 34
- 35
- 36
- 37
- 38
- 39
- 40
- 41
- 42
- 43
- 44
- 45
- 46
- 47
- 48