

Spatial wildfire risk modeling using mixtures of tree-based multivariate Pareto distributions

Daniela Cisneros¹, Arnab Hazra², and Raphaël Huser¹

August 7, 2023

Abstract. Wildfires pose a severe threat to the ecosystem and economy, and risk assessment is typically based on fire danger indices such as the McArthur Forest Fire Danger Index (FFDI) used in Australia. Studying the joint tail dependence structure of high-resolution spatial FFDI data is thus crucial for estimating current and future extreme wildfire risk. However, existing likelihood-based inference approaches are computationally prohibitive in high dimensions due to the need to censor observations in the bulk of the distribution. To address this, we construct models for spatial FFDI extremes by leveraging the sparse conditional independence structure of Hüsler–Reiss-type generalized Pareto processes defined on trees. These models allow for a simplified likelihood function that is computationally efficient. Our framework involves a mixture of tree-based multivariate Pareto distributions with randomly generated tree structures, resulting in a flexible model that can capture nonstationary spatial dependence structures. We fit the model to summer FFDI data from different spatial clusters in Mainland Australia and 14 decadal windows between 1999–2022 to study local spatiotemporal variability with respect to the magnitude and extent of extreme wildfires. Our results demonstrate that our proposed method fits the margins and spatial tail dependence structure adequately, and is helpful to provide extreme wildfire risk measures.

Keywords: Climate change; Graphical model; Generalized Pareto process; Hüsler–Reiss distribution; McArthur Forest Fire Danger Index; Spatial threshold exceedance; Wildfire risk assessment.

¹Statistics Program, Computer, Electrical and Mathematical Sciences and Engineering (CEMSE) Division, King Abdullah University of Science and Technology (KAUST), Thuwal 23955-6900, Saudi Arabia. E-mails: daniela.cisnerosarce@kaust.edu.sa; raphael.huser@kaust.edu.sa

²Department of Mathematics and Statistics, Indian Institute of Technology Kanpur, Kanpur 208016, India. E-mail: ahazra@iitk.ac.in

1 Introduction

Wildfires pose severe threats to human lives, the ecosystem, and properties, resulting in major economic costs (Beyers [et al.](#), 2005; Thomas [et al.](#), 2017). Smoke from wildfires can also lead to poor air quality and high-temperature events (Sandberg, 2003; Jhariya and Raj, 2014). With global warming and other triggering factors, wildfires are becoming more frequent, particularly in countries like Australia and the United States (Westerling, 2016; Stephenson [et al.](#), 2013; Marlon [et al.](#), 2012). The Black Summer bushfires in Australia in 2019–2020 were particularly devastating, scorching millions of hectares and causing significant loss of wildlife (Godfree [et al.](#), 2021).

To assess wildfire risk over space and time, a variety of fire danger indices have been proposed, such as the McArthur Forest Fire Danger Index (FFDI) used in Australia. These fire indices integrate meteorological and fuel information into a single measure, which is then used by various governmental agencies (e.g., the Australian Bureau of Meteorology) for issuing forecasts and warnings and estimating fire suppression difficulty (Van Wagner [et al.](#), 1987; Sirca [et al.](#), 2006; Fiorucci [et al.](#), 2008). A number of studies examine the reliability (Andrews [et al.](#), 2003) and sensitivity (Dowdy [et al.](#), 2010) of fire danger indices using high-resolution datasets, while others attempt to interpolate them spatially (Sanabria [et al.](#), 2013), through various statistical machine learning algorithms, such as logistic regression, random forests, and inverse-distance weighting. However, these classical methods ignore the fact that the most devastating and impactful wildfires often lie in the tail of the distribution. Therefore, for accurate risk assessment and effective wildfire management, it is crucial to develop flexible yet resilient statistical approaches that are specifically designed to estimate the risk of extreme (rather than medium) wildfires, under both current and future conditions. While most studies neglect the tail behavior, a recent exception is the paper by Hazra [et al.](#) (2018), in which a flexible semiparametric Bayesian model is developed to analyze both the bulk and tail properties of the Fosberg Fire Weather Index (FFWI), with the ultimate goal of producing high-resolution fire risk maps in southern California. However, the spatial dependence model developed therein, which is based on a Dirichlet mixture of skew- t processes, is fairly computationally

intensive to fit in high dimensions and is also fitted to the entire dataset (extreme and non-extreme wildfires combined), thus favoring model calibration for the more prevalent non-extreme wildfire events. Another related contribution is [Cisneros et al. \(2023\)](#), where a multi-stage multivariate spatial model is proposed to jointly analyze wildfire counts and burnt areas over the entire United States. However, their model again does not explicitly rely on extreme-value theory.

In this work, our primary goal is to develop a new statistical methodology that relies on extreme-value theory and that is computationally efficient with large spatial data, in order to help improve the assessment of extreme wildfires and estimate how their marginal distribution and dependence structure varies over both space and time. In our real data application, we focus on estimating the spatial tail dependence structure of FFDI data over Mainland Australia in a spatio-temporal moving window fashion. Our local estimation approach not only allows us to reduce the effective dimensionality of the problem (thus fitting extreme-value models faster), but more importantly, it also allows us to account for the strong spatio-temporal non-stationarity displayed in the spatial tail dependence structure of FFDI data, and to quantify temporal trends and the effect of climate change on the spatially-aggregated risk of extreme wildfires over the past decades. Furthermore, we note that even if the FFDI only serves as a proxy for wildfire risk, our spatial tail dependence analysis also sheds some light on the (local) severity and spatial extent of extreme wildfires.

To tackle this challenge and estimate spatial wildfire risk in each spatial cluster and time window, we develop a spatial model for high-dimensional data that combines ideas from extreme-value theory and graphical models. The theory of graphical models provides a probabilistic framework for studying conditional independence relationships in a wide variety of contexts ([Lauritzen, 1996](#); [Koller and Friedman, 2009](#)); one benefit of this framework is that it allows one to construct models with a sparse probabilistic structure, thus facilitating both inference and interpretation. Despite the widespread use of graphical models in various fields and applications, they have been mostly limited to the multivariate Gaussian distribution ([Ren et al., 2015](#); [Epskamp et al., 2018](#)). However, Gaussian models are unsuitable for studying extreme events in risk analysis due to their thin tails

and restrictive dependence structure, which can lead to heavily underestimated probabilities of concurrent extremes ([Davison et al., 2013](#)). Therefore, developing multivariate or spatial graphical models specifically designed for capturing tail dependence in high-dimensional data is crucial. In the framework of max-stable distributions, [Gissibl and Klüppelberg \(2018\)](#) and [Klüppelberg and Lauritzen \(2019\)](#) developed max-linear models on directed acyclic graphs. However, such models have a discrete underlying angular probability measure, which is typically unrealistic for spatial data, and [Papastathopoulos and Strokorb \(2016\)](#) showed that density factorization of the most popular max-stable distributions (with continuous angular density) on graphs is not possible, except in trivial cases. Another perspective on multivariate extremes, however, is based on threshold exceedances—instead of maxima—, leading to the multivariate Pareto distribution (MPD)—instead of max-stable distributions—occurring in the limit as the threshold grows arbitrarily; see, e.g., [Rootzén and Tajvidi \(2006\)](#) and [Rootzén et al. \(2018\)](#). For MPDs, [Engelke and Hitz \(2020\)](#) introduced new notions of conditional independence and extremal graphical models, which enable data density factorization through a Hammersley–Clifford-type theorem; see also [Segers \(2020\)](#) and [Asenova et al. \(2021\)](#) for related work on limits of regularly-varying Markov trees. While providing a flexible framework for multivariate extremes, the single tree-based extremal graphical model of [Engelke and Hitz \(2020\)](#) is often unrealistic for spatial data; [Engelke and Hitz \(2020\)](#) also considered block graphs as simple extensions to tree structures, but these remain overly restrictive in realistic spatial settings. In some related work, [Yu et al. \(2016\)](#) developed a copula model for station-wise block maxima constructed from an ensemble of trees; however, the linking copulas used in this paper are not motivated by extreme-value theory. Additionally, the analytical expression of the extremal dependence χ -measure ([Coles et al., 1999](#)) is intractable for their proposed model; only a lower bound depending on the underlying tree structure is derived.

In this paper, we model high FFDI threshold exceedances over space using a tree-based Hüsler–Reiss-type generalized Pareto process mixture. For simplicity, we estimate the marginal distribution parameters and the spatial dependence parameters separately in two stages. For marginal estima-

tion (first stage), we borrow information from the neighboring locations to obtain spatially-smooth estimates of the parameters of the location-specific generalized Pareto distributions; using these estimates, we then transform the margins to the unit Pareto scale. For dependence structure estimation (second stage), we build upon [Engelke and Hitz \(2020\)](#) and [Yu et al. \(2016\)](#) and fit a tree-based Hüsler–Reiss-type Pareto process model to the transformed data. While a single tree is unable to properly capture the spatial dependence, a mixture of (appropriately constructed) trees has considerably higher flexibility and is, therefore, better suited for modeling spatial data; moreover, this mixture modeling approach keeps at the same time the computational benefit that comes with the use of tree structures. Thus, our approach is not only justified by multivariate extreme-value theory, but it is also amenable to high spatial dimensions. The resulting model enjoys high flexibility and can in particular capture nonstationary spatial extremal dependence structures by adapting the tree mixture weights in a data-driven way. Furthermore, unlike [Yu et al. \(2016\)](#), we can here obtain the analytical expression of the extremal dependence χ -measure. A simulation study showcases the flexibility of our model in capturing the extremal dependence and the upper tail behavior; see the supplementary material of this paper. For our data application, we split Mainland Australia into 25, 50, and 100 spatial clusters and fit our model locally to summer (Nov., Dec., Jan., Feb.) FFDI data available across 23 summer seasons between 1999–2022. To understand the temporal trends in marginal and dependence structures across Australia, we split the 23 seasons into 14 overlapping decadal windows and fit our model separately for each cluster-window combination. Because high spatial aggregates indicate a higher probability of large-scale extreme fire events, we assess wildfire risk across the spatial clusters and study their temporal evolution.

The rest of the paper is organized as follows. In [Section 2](#), an exploratory analysis of the Australian FFDI dataset is presented. [Section 3](#) describes our proposed mixture model based on tree-based Hüsler–Reiss-type multivariate Pareto distributions. In [Section 4](#), we discuss model fitting. In [Section 5](#), we apply the tree-based mixture model to analyze the Australian FFDI dataset and discuss the results. [Section 6](#) concludes and points out some future research directions.

2 Australian McArthur Forest Fire Danger Index data

The Australian McArthur Forest Fire Danger Index (FFDI) dataset that we analyze in this paper contains daily fire index observations obtained using weather forecasts from historical simulations. The FFDI data are made publicly available by the European Forest Fire Information System (EFFIS, <https://effis.jrc.ec.europa.eu/>) of the European Centre for Medium-Range Weather Forecasts (ECMWF) Re-Analysis, 5th generation (ERA5). The FFDI is a standard measure used across Australia for examining the influence of near-surface weather conditions on fire behavior, with the Australian Bureau of Meteorology routinely issuing FFDI forecasts for use by fire management authorities throughout Australia. Mathematically, FFDI is defined as a nonlinear function of temperature (T [$^{\circ}\text{C}$]), relative humidity (RH [%]), and wind speed (ν [km/hr]) on a given day, as well as a dimensionless number representing fuel availability, called the drought factor (DF; Griffiths, 1999); precisely, it is calculated as

$$\text{FFDI} = 2 \exp\{-0.45 + 0.987 \log(\text{DF}) - 0.0345 \text{RH} + 0.0338 T + 0.0234 \nu\}. \quad (1)$$

This formula is a computationally efficient rearranged definition of the commonly used formulation due to Noble *et al.* (1980). The Australian Bureau of Meteorology considers an FFDI value in the interval 0–11 to be low-to-moderate, 12–31 to be high, 32–49 to be very high, 50–74 to be severe, 75–99 to be extreme, and above 100 to be catastrophic.

We limit our study to the Australian mainland that has been seriously affected by wildfires in recent years. We perform some pre-processing steps, which include thinning the grid cells by keeping every second grid cell in each direction to avoid nearly-identical values in neighboring locations (due to wildfires being typically smooth large-scale phenomena), which can cause numerical instabilities. The daily FFDI data are then extracted over the years 1999–2022 across Mainland Australia at a spatial resolution of $0.5^{\circ} \times 0.5^{\circ}$, yielding a total of 2749 grid cells across the study domain. We focus on the summer months (Nov., Dec., Jan., and Feb.), which are more prone to wildfires; thus, we analyze data across 23 summer seasons. Figure 1 displays the mean and maximum FFDI

values, calculated at each grid cell across the entire time period. While the mean FFDI profile is highest near the mid-western region of Australia, the maximum profile shows higher values near the southern coastal areas; these different spatial patterns indicate that the FFDI data behave differently in the bulk and the tail, suggesting that extreme wildfire risk should be assessed using a specialized tail model fitted to the largest observations only. Interestingly, maximum FFDI values fall in the extreme or catastrophic category for a major part of Mainland Australia, except for coastal regions in the Eastern part of the country, as well as a small region in the South-West.

To explore temporal trends in the tails of the FFDI distribution, we compute, for each grid cell and summer season, the 98%-empirical quantile and we then fit a simple linear regression model with the season number as a covariate. We present the estimated annual rate of change at each grid cell in the left panel of Figure 2. While the estimated slope coefficient (i.e., annual rate of change) varies between -0.39 and 0.78 , it is positive at 2600 grid cells and negative at 149 grid cells only. Based on grid cell-wise t -tests, the slope coefficients are found to be significant at 536 grid cells (i.e., over about 20% of the country), and all of them are significantly positive. We obtain similar results for other high quantiles. This exploratory analysis shows that there has been a significant positive shift in FFDI extremes over the past two decades, and this motivates us to split the 23 summer seasons into 14 overlapping decadal windows and to perform window-specific analyses.

Finally, we also investigate the local spatial extremal dependence structure of FFDI data across Australia and assess the extent of spatial nonstationarity. Mathematically, extremal dependence between two random variables $Y_1 \sim F_1$ and $Y_2 \sim F_2$ can be described by the χ -measure, $\chi_{12} = \lim_{u \rightarrow 1} \chi_{12}(u)$, where $\chi_{12}(u) = \Pr \{Y_1 > F_1^{-1}(u) \mid Y_2 > F_2^{-1}(u)\}$. In the spatial context, the χ -measure is often used to discriminate between processes that are asymptotically dependent (with $\chi_{12} > 0$ for all pairs of sites) and asymptotically independent (with $\chi_{12} = 0$ for all pairs of sites); see [Huser and Wadsworth \(2020\)](#) for an in-depth discussion. Since wildfires are large-scale events, we assume hereafter that the local dependence structure is asymptotically dependent. In case replications of the variables Y_1 and Y_2 are available, we can easily estimate χ_{12} by setting a high

threshold $u \approx 1$ and calculating the empirical counterpart of $\chi_{12}(u)$. Since we have spatial data with temporal replicates, we perform the above-mentioned procedure for every pair of sites $\{\mathbf{s}_1, \mathbf{s}_2\}$ and then study the pairwise empirical χ_{12} as a function of the distance $h = \|\mathbf{s}_1 - \mathbf{s}_2\|$ between sites. Here, we split Mainland Australia into 50 different non-overlapping spatial clusters obtained by k -means clustering and estimate the local cluster-specific χ_{12} (assuming local stationarity and isotropy) for each decadal window. The right panel of Figure 2 displays smoothed χ_{12} function estimates plotted with respect to geodesic distance h (in km) for two clusters—one near the eastern coastal region and the other one near the southwestern coastal region, centered at the coordinates (151.56°E, 30.21°S) and (117.67°E, 33.98°S), respectively—for both the first and the last decadal time windows. For the first window (1999–2009), the χ_{12} values are larger for the cluster near the southwestern coast, while for the last window (2012–2022), the χ_{12} profiles for the two clusters cross each other. This indicates a nonstationary dependence structure as a whole across Australia and it motivates us to model the data using a local cluster-specific approach. For the eastern cluster, we notice that the empirical χ_{12} profile is higher for the last decadal window than for the first window. The opposite pattern can be noticed for the other southwestern cluster, though the observed change is much weaker than that for the eastern cluster. Since the χ_{12} profile increases more rapidly near the eastern coast, where the majority of the largest cities of Australia are located, studying these temporal changes based on a theoretically-justified spatial statistical model is of utmost importance.

3 Methodology

Spatial extremes modeling is typically done in two subsequent stages. First, parameters of marginal distributions are estimated by fitting extreme-value theory-justified limit distributions to sample extremes—here, high threshold exceedances. Then, the data are transformed to the same scale using the probability-integral transform. Second, the dependence structure is inferred by fitting a suitable multivariate or spatial extreme-value model. We adopt such a two-stage inference approach to study the behavior of spatial FFDI extremes within each spatial cluster and decadal window.

Section 3.1 describes marginal modeling, while Section 3.2 presents our proposed spatial dependence model based on tree-based MPD mixtures and discusses its properties.

3.1 Marginal modeling

Here we describe the estimation of the marginal distribution at each site based on classical extreme-value theory (Davison and Huser, 2015). Suppose $Z_t(\cdot)$, $t = 1, \dots, T$, denote the T replicates of a spatial stochastic process observed at spatial locations \mathbf{s}_i , $i = 1, \dots, N$. When the process is observed only over a discretized domain, we write $Z_{i,t} \equiv Z_t(\mathbf{s}_i)$ for simplicity, and we write the set of available observations at site \mathbf{s}_i as $\mathcal{Z}_i = \{Z_{i,1}, \dots, Z_{i,T}\}$. Here the elements of \mathcal{Z}_i form a random sample from some marginal distribution $F_i(\cdot)$ and suppose the Pickands–Balkema–de Haan Theorem (Balkema and De Haan, 1974; Pickands III, 1975) holds; that is, we assume that there exists a normalizing function $a(u_i) > 0$ such that

$$\mathrm{P}\left(\frac{Z_{i,t} - u_i}{a(u_i)} > z \mid Z_{i,t} > u_i\right) \rightarrow 1 - H_i(z), \quad \text{as } u_i \rightarrow z_{i\star}, \quad (2)$$

where $z_{i\star}$ is the upper endpoint of the support of $F_i(\cdot)$, and $H_i(\cdot)$ is a non-degenerate distribution. In that case, the limit $H_i(\cdot)$ must be a generalized Pareto (GP) distribution function (Davison and Smith, 1990) and has the form

$$H_i(z) = \begin{cases} 1 - (1 + \xi_i z / \sigma_i)_+^{-1/\xi_i}, & \text{if } \xi_i \neq 0, \\ 1 - \exp(-z/\sigma_i), & \text{if } \xi_i = 0, \end{cases} \quad (3)$$

where $(a)_+ = \max\{0, a\}$, and $\sigma_i > 0$ and $\xi_i \in \mathbb{R}$ are scale and shape parameters, respectively; we denote this distribution by $\mathrm{GP}(\sigma_i, \xi_i)$. The shape parameter ξ_i determines the behavior of the GP distribution: if $\xi_i < 0$, $H_i(\cdot)$ has an upper bound at $-\sigma_i/\xi_i$, and if $\xi_i \geq 0$, the upper tail of $H_i(\cdot)$ is unbounded and can be either light ($\xi_i = 0$) or heavy ($\xi_i > 0$). By replacing the limit in (2) with a equality in distribution for some high site-specific threshold u_i , and letting the term $a(u_i)$ be absorbed into the GP scale parameter σ_i , we then obtain the approximation

$$\mathrm{P}(Z_{i,t} > z) \approx \zeta_{u_i} \left(1 + \xi_i \frac{z - u_i}{\sigma_i}\right)^{-1/\xi_i}, \quad z > u_i, \quad \xi_i \neq 0, \quad (4)$$

where $\zeta_{u_i} = \text{P}(Z_{i,t} > u_i)$, and thus, the level $z_i^{(m)}$ that is exceeded on an average once every m observations is given by

$$z_i^{(m)} = \begin{cases} u_i + \frac{\sigma_i}{\xi_i} \{(m\zeta_{u_i})^{\xi_i} - 1\}, & \text{if } \xi_i \neq 0, \\ u_i + \sigma_i \log(m\zeta_{u_i}), & \text{if } \xi_i = 0, \end{cases} \quad (5)$$

provided m is sufficiently large to ensure that $z_i^{(m)} > u_i$. The high quantile $z_i^{(m)}$ is commonly called the m -observation return level, though its meaning is less clear under temporal non-stationarity.

To estimate marginal parameters at each site \mathbf{s}_i , we adopt a local inference approach, whereby the GP model is fitted by maximum likelihood by pooling information from a set of sites $\{\mathbf{s}_j; j \in \mathcal{N}_i\}$ containing \mathbf{s}_i and its four nearest neighbors. Using data from neighboring sites allows us to increase the effective sample size, thus reducing estimation uncertainty, and to smooth marginal parameters spatially. For each site \mathbf{s}_i , $i = 1, \dots, N$, we fix a high threshold u_i and denote by $\mathcal{Z}_j^* = \{Z_{j,1}^*, \dots, Z_{j,T_j}^*\} \subset \mathcal{Z}_j$ the T_j values exceeding u_i at site $j \in \mathcal{N}_i$. Under the working assumption of independence across space and time, we then maximize the following likelihood function

$$L_i(\sigma_i, \xi_i \mid Z_{j,k}^*, j \in \mathcal{N}_i, k = 1, \dots, T_j) = \prod_{j \in \mathcal{N}_i} \prod_{k=1}^{T_j} \left\{ \frac{1}{\sigma_i} \left(1 + \xi_i \frac{Z_{j,k}^* - u_i}{\sigma_i} \right) \right\}_+^{-1/\xi_i - 1}, \quad (6)$$

and we optimize it numerically to obtain the maximum likelihood estimates (MLEs) $\hat{\sigma}_i$ and $\hat{\xi}_i$. Plugging the MLEs into $H_i(\cdot)$ yields the fitted GP distribution $\hat{H}_i(\cdot)$. The probability ζ_{u_i} can be estimated empirically by $\hat{\zeta}_{u_i} = T_i/T$. We can then approximate the upper tail of the distribution $F_i(\cdot)$ using (4) as $F_i(z) \approx G_i(z) := 1 - \hat{\zeta}_{u_i} \{1 - \hat{H}_i(z)\}$, for z large. For values below the threshold, the empirical distribution function $\hat{F}_i(\cdot)$ can be used instead. Therefore, the full distribution $F_i(z)$ can be estimated by $G_i(z)$, if $z > u_i$, and $\hat{F}_i(z)$, if $z \leq u_i$, and data can then be transformed to the unit Pareto scale via the probability integral transform, $Y_{i,t} = \{1 - F_i(Z_{i,t})\}^{-1}$, after replacing $F_i(\cdot)$ with its approximation. The whole marginal estimation procedure and standardization to a common unit Pareto scale are summarized in a flowchart displayed in Figure 3.

3.2 Extremal dependence modeling

3.2.1 Multivariate Pareto distributions

Our spatial dependence model is constructed from mixtures of tree-based multivariate Pareto distributions (MPDs) of Hüsler–Reiss type, so we give a brief background on MPDs, initially introduced by [Rootzén and Tajvidi \(2006\)](#) and then further studied by [Rootzén et al. \(2018\)](#), as well as the specific Hüsler–Reiss model, first introduced by [Hüsler and Reiss \(1989\)](#) and recently popularized by [Engelke and Hitz \(2020\)](#) in the context of extremal graphical models.

A random vector $\mathbf{Y} = (Y_1, \dots, Y_N)'$ with unit Pareto margins is said to be in the max-domain of attraction of the random vector $\mathbf{W} = (W_1, \dots, W_N)'$ if, for any $\mathbf{w} = (w_1, \dots, w_N)' \in \mathcal{F} = [0, \infty)^N \setminus \{\mathbf{0}\}$, we have

$$\lim_{T \rightarrow \infty} \mathbb{P} \left(\max_{t=1, \dots, T} Y_{1,t} \leq Tw_1, \dots, \max_{t=1, \dots, T} Y_{N,t} \leq Tw_N \right) = \mathbb{P}(\mathbf{W} \leq \mathbf{w}) = \exp\{-V(\mathbf{w})\}, \quad (7)$$

where $\mathbf{Y}_t = (Y_{1,t}, \dots, Y_{N,t})'$, $t = 1, 2, \dots$, are independent copies of \mathbf{Y} . If this convergence holds, \mathbf{W} must be max-stable ([Davison et al., 2019](#)) with standard Fréchet margins, i.e., $\mathbb{P}(W_i \leq w) = \exp(-1/w)$, $w \geq 0$, and $V(\mathbf{w}) := \Lambda([0, \infty]^N \setminus [\mathbf{0}, \mathbf{w}])$ is known as the exponent function and is constructed from a Radon measure Λ defined on the cone \mathcal{F} . If $\Lambda(\cdot)$ is absolutely continuous with respect to Lebesgue measure on \mathcal{F} , its Radon–Nikodym derivative is denoted by $\lambda(\cdot)$ and satisfies $\lambda(\mathbf{w}) = -\frac{\partial^N}{\partial w_1 \cdots \partial w_N} V(\mathbf{w})$. It can be shown that the intensity $\lambda(\cdot)$ has two main properties: (i) it is homogeneous of order $-(N+1)$, i.e., $\lambda(c\mathbf{w}) = c^{-(N+1)}\lambda(\mathbf{w})$ for any $c > 0$ and $\mathbf{w} \in \mathcal{F}$; and (ii) it has normalized margins in the sense that for any $i \in \{1, \dots, N\}$, $\int_{\{\mathbf{w} \in \mathcal{F}: w_i > 1\}} \lambda(\mathbf{w}) d\mathbf{w} = 1$.

Following [Resnick \(2008\)](#), the convergence in (7) is equivalent to $\lim_{u \rightarrow \infty} u\{1 - \mathbb{P}(\mathbf{Y} \leq u\mathbf{w})\} = V(\mathbf{w})$, $\mathbf{w} \in \mathcal{F}$; thus, for $\mathbf{w} \in \mathcal{F}$, the joint distribution of the multivariate threshold exceedances of \mathbf{Y} , defined through their infinity norm being large, satisfies

$$\mathbb{P}(\mathbf{X} \leq \mathbf{x}) = \lim_{u \rightarrow \infty} \mathbb{P}(\mathbf{Y}/u \leq \mathbf{x} \mid \|\mathbf{Y}\|_\infty > u) = [V(\min\{\mathbf{x}, \mathbf{1}\}) - V(\mathbf{x})]/V(\mathbf{1}), \quad (8)$$

where $\min\{\mathbf{x}, \mathbf{1}\}$ denotes the vector of component-wise minima between \mathbf{x} and $\mathbf{1}$. The limiting

random vector \mathbf{X} is said to follow an MPD (Rootzén and Tajvidi, 2006). The underlying exponent measure $\Lambda(\cdot)$ of the distribution of \mathbf{W} implicitly defines the MPD of \mathbf{X} , with support $\mathcal{L} = \{\mathbf{x} \in \mathcal{F} : \|\mathbf{x}\|_\infty > 1\}$. Further, we assume that the distribution of \mathbf{X} allows a positive and continuous density $f_{\mathbf{X}}(\cdot)$ on \mathcal{L} given by $f_{\mathbf{X}}(\mathbf{x}) = \frac{\partial^N}{\partial x_1 \cdots \partial x_N} \mathbf{P}(\mathbf{X} \leq \mathbf{x}) = \lambda(\mathbf{x})/V(\mathbf{1})$ for $\mathbf{x} \in \mathcal{L}$; the mixed derivative of $V(\min\{\mathbf{x}, \mathbf{1}\})$ in (8) cancels out because it is constant with respect to at least one x_i as $\|\mathbf{x}\|_\infty > 1$. The intensity $\lambda(\cdot)$ can thus be interpreted as the unnormalized MPD density, while the normalization constant $V(\mathbf{1}) = \Lambda([0, \infty]^N \setminus [\mathbf{0}, \mathbf{1}])$ ensures that it integrates to one over \mathcal{L} . Here $V(\mathbf{1})$ is called the N -variate extremal coefficient and serves as a summary of extremal dependence.

A particularly interesting parametric multivariate extreme-value model is the Hüsler–Reiss distribution, which arises as the only possible limit of renormalized maxima from suitably defined triangular arrays of multivariate Gaussian vectors (Hüsler and Reiss, 1989). Because of this connection with the multivariate Gaussian distribution, the Hüsler–Reiss model or its spatial extension known as the Brown–Resnick process (Kablichko et al., 2009) have been found to provide good fits in practical applications. Moreover, there are further fundamental links with the multivariate Gaussian distribution, especially in the context of graphical models (Engelke and Hitz, 2020), that make them appealing, and we shall explain in Section 3.2.2 how to construct tree-based Hüsler–Reiss MPDs. A N -dimensional Hüsler–Reiss distribution is parameterized by a variogram matrix $\mathbf{\Gamma}$ (i.e., symmetric, strictly conditionally negative-definite with non-negative entries Γ_{ij} and diagonal entries $\Gamma_{ii} = 0$, $i, j = 1, \dots, N$). For any $k \in \{1, \dots, N\}$, the corresponding intensity $\lambda(\cdot; \mathbf{\Gamma})$ is

$$\lambda(\mathbf{x}; \mathbf{\Gamma}) = x_k^2 \prod_{i \neq k} x_i^{-1} \phi_{N-1}(\tilde{\mathbf{x}}_{\setminus k}; \mathbf{\Sigma}^{(k)}), \quad \mathbf{x} \in \mathcal{L}, \quad (9)$$

where $\tilde{\mathbf{x}} = [\log(x_1/x_k) + \Gamma_{1k}/2, \dots, \log(x_N/x_k) + \Gamma_{Nk}/2]'$, $\tilde{\mathbf{x}}_{\setminus k}$ denotes the vector $\tilde{\mathbf{x}}$ with the k -th element removed, and $\phi_{N-1}(\cdot; \mathbf{\Sigma}^{(k)})$ is the $(N-1)$ -dimensional centered multivariate normal density with the strictly positive-definite covariance matrix

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

Note that, although the right-hand side of the density in (9) appears to depend on the choice of k , it is actually independent of k . The strength of dependence between the i -th and j -th component of a vector distributed according to the Hüsler–Reiss model is thus controlled by Γ_{ij} ; here $\Gamma_{ij} = 0$ leads to complete dependence between the components while the dependence weakens as $\Gamma_{ij} \rightarrow \infty$.

3.2.2 Tree-based multivariate Pareto distributions

We now briefly recall the approach of Engelke and Hitz (2020) for creating graphical models for multivariate extremes, and tree-based Hüsler–Reiss MPD models.

An undirected graphical model induces a multivariate density function $f(\mathbf{x})$ that factorizes according to a graph \mathcal{G} , defined by a collection of nodes \mathcal{V} and edges \mathcal{E} . In a probabilistic graphical model, each node $i \in \mathcal{V}$ denotes a random variable X_i and an edge $(i, j) \in \mathcal{E}$ is absent if and only if X_i is independent of X_j conditional on the other variables on the graph, \mathbf{X}_{-ij} . A clique C is a subset of \mathcal{V} where each pair of nodes is connected by an edge and let the set of all cliques of \mathcal{G} be denoted by \mathcal{C} . The Hammersley–Clifford Theorem (Lauritzen, 1996) then relates such a graphical structure to a factorization of $f(\mathbf{x})$ over \mathcal{C} given by $f(\mathbf{x}) = \prod_{C \in \mathcal{C}} \psi_C(\mathbf{x}_C)$, where $\psi_C(\cdot)$ is called a compatibility function, defined on a clique C . A graph \mathcal{G} is called decomposable if it can be written as a union of edge-disjoint subgraphs, and then, the above factorization can be written in terms of marginal densities as $f(\mathbf{x}) = \prod_{C \in \mathcal{C}} f_C(\mathbf{x}_C) / \prod_{D \in \mathcal{D}} f_D(\mathbf{x}_D)$, where \mathcal{D} is a multiset containing intersections between the cliques called separator sets. More details can be found in Lauritzen (1996), and Section 2.3 and Appendix A of Engelke and Hitz (2020).

Following Engelke and Hitz (2020), we here first consider graphs $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ that are decomposable and connected (trees in particular), with clique set \mathcal{C} and separator set \mathcal{D} , where all separators in \mathcal{D} are single nodes; thus, for \mathbf{X} as in Section 3.2.1, $f_D(\cdot)$ denotes the density of the univariate unit-Pareto distribution, i.e., $f_D(x_D) = x_D^{-2}$ for each $D \in \mathcal{D}$. The setting of decomposable graphs is computationally attractive but overly restrictive for spatial data, and we generalize it using tree-mixtures in Section 3.2.3. For each clique $C \in \mathcal{C}$, we can define a parametric MPD with intensity

$\lambda_C(\cdot; \boldsymbol{\theta}_C)$. Overall, the density over the whole graph \mathcal{G} can be expressed as

$$f_{\mathbf{X}}(\mathbf{x}; \boldsymbol{\theta}) = \frac{1}{V(\mathbf{1}; \boldsymbol{\theta})} \prod_{C \in \mathcal{C}} \frac{\lambda_C(\mathbf{x}_C; \boldsymbol{\theta}_C)}{\prod_{j \in C} x_j^{-2}} \prod_{i \in \mathcal{V}} x_i^{-2}, \quad (11)$$

where \mathbf{x}_C is the sub-vector of \mathbf{X} corresponding to the clique C and $\boldsymbol{\theta} = [\boldsymbol{\theta}_C]_{C \in \mathcal{C}}$ its parameters.

While the density in (11) factorizes over a graph up to the normalizing factor $V(\mathbf{1}; \boldsymbol{\theta})$, the resulting MPD random vector $\mathbf{X} = (X_i)_{i \in \mathcal{V}}$, $\mathcal{V} = \{1, \dots, N\}$, with support space \mathcal{L} , does not comply with the classical notion of conditional independence because \mathcal{L} is not a product space. However, \mathbf{X} still admits conditional independence relationships in the following modified sense. For any $k \in \mathcal{V}$, we can define the random vector $\mathbf{X}^{(k)}$ as \mathbf{X} conditioned on the event $\{X_k > 1\}$. Then, $\mathbf{X}^{(k)}$ has support on the product space $\mathcal{L}^{(k)} := \{\mathbf{x} \in \mathcal{L} : x_k > 1\}$ and has density

$$f_{\mathbf{X}^{(k)}}(\mathbf{x}) = \frac{f_{\mathbf{X}}(\mathbf{x})}{\int_{\mathcal{L}^{(k)}} f_{\mathbf{X}}(\mathbf{x}) d\mathbf{x}}, \quad \mathbf{x} \in \mathcal{L}^{(k)},$$

where $f_{\mathbf{X}}(\cdot)$ is as in Section 3.2.1. Let $\mathcal{A}, \mathcal{B}, \mathcal{C} \subset \mathcal{V}$ be non-empty disjoint subsets whose union is \mathcal{V} , as shown in the right panel of Figure 4, and also let the corresponding components of \mathbf{X} be $\mathbf{X}_{\mathcal{A}}$, $\mathbf{X}_{\mathcal{B}}$, and $\mathbf{X}_{\mathcal{C}}$, respectively. Similarly, let also $\mathbf{X}_{\mathcal{A}}^{(k)}$, $\mathbf{X}_{\mathcal{B}}^{(k)}$, and $\mathbf{X}_{\mathcal{C}}^{(k)}$ be the respective components of $\mathbf{X}^{(k)}$. Then, Engelke and Hitz (2020) show that graphical structure of \mathbf{X} is equivalent to the (classical) conditional independence $\mathbf{X}_{\mathcal{A}}^{(k)} \perp\!\!\!\perp \mathbf{X}_{\mathcal{B}}^{(k)} \mid \mathbf{X}_{\mathcal{C}}^{(k)}$ for all $k \in \mathcal{V}$, and in this case, they write $\mathbf{X}_{\mathcal{A}} \perp_e \mathbf{X}_{\mathcal{B}} \mid \mathbf{X}_{\mathcal{C}}$ to emphasize that this new notion adapted for multivariate extremes.

A tree $\mathcal{T} = (\mathcal{V}, \mathcal{E})$ with nodes \mathcal{V} and edges \mathcal{E} is a connected undirected graph without cycles, and cliques of maximum dimension two, and thus, the relation between the cardinalities of \mathcal{V} and \mathcal{E} is $|\mathcal{E}| = |\mathcal{V}| - 1$. In Figure 4, all the graphs are trees as there is a unique path between any two nodes; the first graph (left) is a directed tree and the second (middle) graph is an undirected tree. If \mathbf{X} is an extremal graphical model satisfying the global Markov property with respect to a tree \mathcal{T} , we obtain a simple stochastic representation for $\mathbf{X}^{(k)}$. Its extremal variogram rooted at node $k \in \mathcal{V}$ is defined as the matrix $\boldsymbol{\Gamma}^{(k)}$ with its (i, j) -th entry is given by $\Gamma_{ij}^{(k)} = \text{Var}\{\log X_i^{(k)} - \log X_j^{(k)}\}$, $i, j \in \mathcal{V}$, whenever the right-hand side is finite. Here $\Gamma_{ij}^{(k)}$ turns out to be an additive tree metric, and thus

it is a natural quantity to infer conditional independence in extremal tree models. [Engelke and Volgushev \(2022\)](#) define the extremal variogram for arbitrary MPDs, and in the special case of the Hüsler–Reiss distribution with variogram matrix $\mathbf{\Gamma}$, they show that $\Gamma_{ij} = \Gamma_{ij}^{(k)}$ for all $k = 1, \dots, N$, such that for tree-based Hüsler–Reiss MPD models,

$$\Gamma_{ij} = \sum_{(i^*, j^*) \in \mathcal{P}(i, j; \mathcal{T})} \Gamma_{i^* j^*}, \quad k = 1, \dots, d, \quad (12)$$

where $\mathcal{P}(i, j; \mathcal{T})$ denotes the set of edges on the unique path from node i to node j on tree \mathcal{T} .

3.2.3 Final model based on tree-mixtures

Our focus is on extending the extremal graphical modeling methodology of [Engelke and Hitz \(2020\)](#) to (potentially nonstationary) spatial extremes. Graph decomposability, a crucial assumption in [Section 3.2.2](#), is overly restrictive for spatial data, and tree structures in non-trivial settings do not allow spatial stationarity or realistic nonstationarity on regular lattice data in \mathbb{R}^2 . The most natural spatially-structured Markov random field on a regular lattice does not satisfy decomposability, making it difficult to extend the methodology to this case. [Yu et al. \(2016\)](#) and [Vettori et al. \(2020\)](#) propose using finite mixtures of tree-based models to retain computational tractability, and we build upon their approaches in the context of spatial extremes using tree-based MPDs.

Here, we assume that the dependence structure can be explained through a mixture of M tree-based MPD models. For a general undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, let the set of all spanning trees (thus, subgraphs of \mathcal{G}) be denoted by $\tilde{\mathcal{T}} = \{\mathcal{T}_1, \dots, \mathcal{T}_M\}$. While the nodes for all such trees are the same as they represent the same data locations, the set of edges are different, and we thus denote the trees by $\mathcal{T}_m = (\mathcal{V}, \mathcal{E}_m)$, $m = 1, \dots, M$. Since trees have the property that all cliques are of maximum dimension two and separators contain only one node, the MPD density based on \mathcal{T}_m is

$$f_{\mathbf{X}}(\mathbf{x} \mid \mathcal{T}_m; \boldsymbol{\theta}) = \frac{1}{V(\mathbf{1}; \boldsymbol{\theta})} \prod_{\{i, j\} \in \mathcal{E}_m} \frac{\lambda_{ij}(x_i, x_j; \boldsymbol{\theta}_{ij})}{x_i^{-2} x_j^{-2}} \prod_{i \in \mathcal{V}} x_i^{-2}, \quad (13)$$

where $\lambda_{ij} \equiv \lambda_{(i, j)}$ are the bivariate marginals of the exponent measure density $\lambda(\cdot; \boldsymbol{\theta})$ and $\boldsymbol{\theta}_{ij}$ is the dependence parameter vector corresponding to the edge $(i, j) \in \mathcal{E}$. For the specific Hüsler–Reiss

model, λ_{ij} can be deduced from (9), while the edge parameter is simply equal to $\theta_{ij} = \Gamma_{ij}$, the (i, j) -th element of the matrix $\mathbf{\Gamma}$. We construct our model based on an ensemble-of-tree mixture as

$$\tilde{f}_{\mathbf{X}}(\mathbf{x}; \boldsymbol{\theta}) = \sum_{m=1}^M p(\mathcal{T}_m) f_{\mathbf{X}}(\mathbf{x} | \mathcal{T}_m; \boldsymbol{\theta}), \quad (14)$$

where $p(\mathcal{T}_m)$'s are tree probabilities assigned by the decomposable prior over spanning trees proposed by Meilä and Jaakkola (2006), which takes the specific form

$$p(\mathcal{T}_m) = \frac{1}{C_{\tilde{\mathcal{T}}}} \prod_{(i,j) \in \mathcal{E}_m} \beta_{ij}, \quad m = 1, \dots, M, \quad (15)$$

where $C_{\tilde{\mathcal{T}}} = \sum_{m=1}^M \prod_{(i,j) \in \mathcal{E}_m} \beta_{ij}$, and β_{ij} are edge weights to be computed. We denote the edge weight matrix by $\boldsymbol{\beta}$ with its (i, j) -th element being β_{ij} . Here, $\boldsymbol{\beta}$ is symmetric and its diagonal entries are zero. Following Kirchhoff's matrix theorem (Beineke et al., 2004), we have $C_{\mathcal{T}} = \det\{\mathbf{Q}(\boldsymbol{\beta})\}$, where $\mathbf{Q}(\boldsymbol{\beta})$ denotes the first $(N - 1)$ rows and columns of the $(N \times N)$ -dimensional Laplacian matrix $\mathbf{L}(\boldsymbol{\beta}) = \text{diag}(\boldsymbol{\beta}\mathbf{1}_N) - \boldsymbol{\beta}$ corresponding to a graph with edge weight matrix $\boldsymbol{\beta}$. Here $\mathbf{1}_N$ is a N -length column vector of all ones. By substituting (15) with $C_{\mathcal{T}} = \det\{\mathbf{Q}(\boldsymbol{\beta})\}$ into (14), we have

$$\tilde{f}_{\mathbf{X}}(\mathbf{x}; \boldsymbol{\theta}) = \frac{1}{\det\{\mathbf{Q}(\boldsymbol{\beta})\}V(\mathbf{1}; \boldsymbol{\theta})} \prod_{i \in \mathcal{V}} x_i^{-2} \sum_{m=1}^M \prod_{\{i,j\} \in \mathcal{E}_m} \beta_{ij} \frac{\lambda_{ij}(x_i, x_j; \boldsymbol{\theta}_{ij})}{x_i^{-2} x_j^{-2}}. \quad (16)$$

Further, by denoting $\lambda_{ij}^*(x_i, x_j; \boldsymbol{\theta}_{ij}) = \lambda_{ij}(x_i, x_j; \boldsymbol{\theta}_{ij}) / (x_i^{-2} x_j^{-2})$, we obtain

$$\tilde{f}_{\mathbf{X}}(\mathbf{x}; \boldsymbol{\theta}) = \frac{1}{V(\mathbf{1}; \boldsymbol{\theta})} \prod_{i \in \mathcal{V}} x_i^{-2} \times \frac{\det\{\mathbf{Q}(\boldsymbol{\beta} \odot \boldsymbol{\lambda}^*)\}}{\det\{\mathbf{Q}(\boldsymbol{\beta})\}}, \quad (17)$$

where $\boldsymbol{\lambda}^*$ is a $(N \times N)$ -dimensional matrix with its (i, j) -th element being $\lambda_{ij}^*(x_i, x_j; \boldsymbol{\theta}_{ij})$, \odot denotes element-wise multiplication of two matrices, and $\mathbf{Q}(\cdot)$ denotes the same matrix operation as above.

Once the weights $\boldsymbol{\beta}$ are obtained (see Section 4.2), we can then compute the prior probabilities of each spanning tree in (15) and combine these tree-based MPD models as in (14) to construct the full mixture model. However, there are a number of drawbacks in building this mixture using all possible spanning trees. First, depending on the initial graph \mathcal{G} , the number of spanning trees can be huge, especially if the number of nodes $|\mathcal{V}|$ (i.e., number of sites in spatial applications) is

large or if \mathcal{G} is densely connected, making it computationally impractical. Second, not all these spanning trees might be meaningful in a spatial context, like the one we are interested in. While the second drawback can be somewhat mitigated by choosing the initial graph appropriately (e.g., as a simple lattice), the first drawback remains problematic in high-dimensional spatial settings. Therefore, instead of considering all possible spanning trees $\tilde{\mathcal{T}}$ as in (14), we propose setting some prior tree probabilities to zero and constructing the tree-mixture from a smaller sub-selection of suitably sampled random trees. We set the number of such random trees to be quite large, but typically much smaller than $M = |\tilde{\mathcal{T}}|$, which makes it computationally feasible, and then compute prior weights as in (15) but with a different normalizing factor to account for the fact that a large number of (unsampled) trees have prior probability zero.

There are various ways to generate random trees. Here, we proceed as follows to sample a variety of random trees that can collectively represent the spatial dependence structure of the data appropriately. Assuming that the initial graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a regular lattice (which is the case in our application), we first draw some random positive values ω_{ij} assigned to all edges $(i, j) \in \mathcal{E}$; we need to make sure that $\omega_{ji} = \omega_{ij}$ and $\omega_{ii} = 0$, but these values do not need to be the same for the different edges (in particular different edge types). For example, in our case, we draw these random values differently depending on whether the edge between two sites is horizontally (W–E), vertically (N–S), or diagonally (NW–SE and SW–NE) oriented on the lattice \mathcal{G} . This allows favoring trees along certain directions, which allows capturing some form of anisotropy and nonstationarity. Then, we compute the minimum spanning tree, which minimizes a weighted sum of the sampled random edge values ω_{ij} , and the computed edge weights β_{ij} . The minimum spanning tree is unique once the edge values have been sampled. However, by repeated sampling, our procedure can therefore generate a wide range of different random trees whose edges follow certain desired patterns.

3.2.4 Model properties

Given a tree \mathcal{T}_m , the marginal distribution for each component X_i of the MPD vector \mathbf{X} in (13), given that $X_i > 1$, is unit-Pareto with density x_i^{-2} , $x_i \geq 1$, which does not rely on the dependence

parameter vector $\boldsymbol{\theta}$. Therefore, the same holds for the tree-mixture in (14). Hence, our proposed tree-mixture model is effectively a copula model (on a scale different from uniform), which concerns the modeling of the extremal dependence structure.

For multivariate extremes, several summary statistics have been developed to measure the strength of dependence between the extremes of two variables X_i and X_j (Coles *et al.*, 1999). The most popular one is the χ -measure (recall its definition in Section 2), expressed as $\chi_{ij} = \lim_{u \rightarrow 1} \mathbb{P}\{F_i(X_i) > u \mid F_j(X_j) > u\}$ whenever the limits exists, where F_i and F_j denote the marginal distribution functions of X_i and X_j , respectively.

For a MPD random vector \mathbf{X} defined in (8), the χ -measure between the components X_i and X_j can be expressed as $\chi_{ij} = 2 - V_{ij}(1, 1)$, where V_{ij} denotes the underlying exponent function restricted to the subvector $(X_i, X_j)'$. In particular, for the Hüsler–Reiss distribution with variogram matrix $\boldsymbol{\Gamma}$, which is closed under marginalization, we get $\chi_{ij} = 2\bar{\Phi}(\sqrt{\Gamma_{ij}/2})$, where $\bar{\Phi}(\cdot)$ denotes the standard normal survival function and $\Gamma_{ij} = \text{Var}\{\log(X_i) - \log(X_j)\}$ is the (i, j) -th element of $\boldsymbol{\Gamma}$. In case of a tree-based Hüsler–Reiss distribution defined over a single tree $\mathcal{T}_m = (\mathcal{V}, \mathcal{E}_m)$, we thus get from (12) that for any $i, j \in \mathcal{V}$,

$$\chi_{ij} = 2\bar{\Phi}(\sqrt{\Gamma_{ij}}) = 2\bar{\Phi}\left(\sqrt{\sum_{(i^*, j^*) \in \mathcal{P}(i, j; \mathcal{T}_m)} \Gamma_{i^* j^*}}\right), \quad (18)$$

where $\mathcal{P}(i, j; \mathcal{T}_m)$ denotes the set of edges on the unique path from node i to node j on tree \mathcal{T}_m . Thus, for our final tree-mixture spatial model $\tilde{f}_{\mathbf{X}}(\cdot)$ in (17), we deduce from (18) and Equation (61) in Lemma 1 of Yu *et al.* (2016) that the χ -measure equals

$$\begin{aligned} \chi_{ij} &= \lim_{u \rightarrow 1} \sum_{m=1}^M p(\mathcal{T}_m) \mathbb{P}\{F_i(X_i) > u \mid F_j(X_j) > u, \mathcal{T}_m\} \\ &= 2 \sum_{m=1}^M p(\mathcal{T}_m) \bar{\Phi}\left(\sqrt{\sum_{(i^*, j^*) \in \mathcal{P}(i, j; \mathcal{T}_m)} \Gamma_{i^* j^*}}\right), \end{aligned} \quad (19)$$

where $p(\mathcal{T}_m)$ is defined as in (15). Note that the prior probabilities $p(\mathcal{T}_m)$ are properly modified when fewer trees are used, as explained in Section 3.2.3. Unlike Yu *et al.* (2016), who only derive

lower bounds for the χ -measure, χ_{ij} can in our case be obtained in closed form. Note that in the case of spatial data on a 2D lattice $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, the dependence structure from a single-tree Hüsler–Reiss MPD model in (18) is always (artificially) non-stationary, while that from the tree-mixture model in (19) is in general non-stationary, but it can also lead to stationarity depending on the model parameters and the prior tree probabilities. It is therefore a much richer class of models.

4 Inference

Marginal generalized Pareto distribution parameters are estimated by maximizing the likelihood function (6). Estimates are then used to transform the data to the unit-Pareto scale, as illustrated in Figure 3. In this section, we consider the tree-based MPD mixture (13) with trees supported over an initial graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and MPD mixture components of Hüsler–Reiss type parametrized by the variogram matrix $\mathbf{\Gamma}$; we first explain how to estimate $\mathbf{\Gamma}$ and then show how to compute the edge weights β in (15) used in the final model density (17). The Supplementary Material contains a simulation study that illustrates the good performance of our estimation approach.

4.1 Estimation of $\mathbf{\Gamma}$

From (13), it is clear that for a MPD constructed from a tree $\mathcal{T}_m = (\mathcal{V}, \mathcal{E}_m)$ spanning the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, the only relevant parameters in $\mathbf{\Gamma}$ that need to be estimated are those that correspond to all edges \mathcal{E}_m in the tree \mathcal{T}_m (which are the only cliques). Given that our final model is a mixture of these spanning tree-based MPDs, we need to estimate parameters Γ_{ij} over all edges of the initial graph \mathcal{G} . While it would be possible to impose some “spatial structure” to the parameters, e.g., by forcing all similar edge types (e.g., horizontal, etc.) to share the same parameter value across the graph, which would mimic stationarity, we have found that this solution is suboptimal in our case, and we thus estimate parameters Γ_{ij} edge-by-edge for all $(i, j) \in \mathcal{E}$. This approach leads to a less parsimonious but more flexible model that can capture complex forms of spatial nonstationarity.

Suppose $\mathbf{Y}_{i,j,t} = (Y_{i,t}, Y_{j,t})'$ denotes the bivariate vector of standardized Pareto observations corresponding to pixels i and j at time t , where the $Y_{i,t}$'s are obtained according to the marginal

transformation procedure outlined in Section 3.1. The corresponding limiting MPD that arise for high threshold exceedances according to (8) is here assumed to follow a bivariate Hüsler–Reiss distribution; recall (9). We thus fit the bivariate Hüsler–Reiss distribution to bivariate exceedances from $\mathbf{Y}_{i,j,t}$ such that $\|\mathbf{Y}_{i,j,t}\|_\infty > u$ for some high threshold $u > 1$ using a pairwise censored likelihood. From (8), the contribution of the t -th time point to the censored likelihood is

$$L_t(\Gamma_{ij}) = \begin{cases} \lambda(Y_{i,t}/u, Y_{j,t}/u; \Gamma_{ij})/V(1, 1; \Gamma_{ij}), & Y_{i,t} > u, Y_{j,t} > u, \\ -V_1(Y_{i,t}/u, 1; \Gamma_{ij})/V(1, 1; \Gamma_{ij}), & Y_{i,t} > u, Y_{j,t} \leq u, \\ -V_2(1, Y_{j,t}/u; \Gamma_{ij})/V(1, 1; \Gamma_{ij}), & Y_{i,t} \leq u, Y_{j,t} > u, \end{cases}$$

where $V_k(x_i, x_j; \Gamma_{ij}) = \partial V(x_i, x_j; \Gamma_{ij})/\partial x_k$, $k = i, j$, and $\lambda(x_i, x_j; \Gamma_{ij}) = -\partial^2 V(x_i, x_j; \Gamma_{ij})/(\partial x_i \partial x_j)$. Then, assuming independence over t , the full censored log-likelihood is $\ell(\Gamma_{ij}) = \sum_{t=1}^T \log\{L_t(\Gamma_{ij})\}$, and the corresponding parameter estimate is obtained as $\hat{\Gamma}_{ij} = \arg \max_{\Gamma_{ij} > 0} \ell(\Gamma_{ij})$. This optimization needs to be repeated for all edges $(i, j) \in \mathcal{E}$. Explicit forms of the censored likelihood for parametric models other than Hüsler–Reiss can be found, for example, in Coles *et al.* (2001).

4.2 Computation of β

Once we the parameters for each edge $(i, j) \in \mathcal{E}$ of the initial graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ are estimated, we build the mixture of tree-based MPDs in an optimal way by finding the edge weights β_{ij} in (14) that maximize the overall mixture likelihood; in other words, we adjust the prior tree probabilities $p(\mathcal{T}_m)$ in order to provide the best fit to the data. This is done by leveraging the expression (17) and thus computing the edge weight matrix β as

$$\hat{\beta} = \arg \max_{\beta} \left\{ \sum_{t=1}^T \log \det[Q(\beta \odot \hat{\mathbf{L}}_t)] - T \log \det [Q(\beta)] \right\}, \quad (20)$$

such that $\beta_{ij} \geq 0$ for all $(i, j) \in \mathcal{E}$ (with $\beta_{ij} = 0$ for all $(i, j) \notin \mathcal{E}$) and $\|U(\beta)\|_2 = 1$, where $\|U(\cdot)\|_2$ is the Euclidean norm for the upper triangular part of the input matrix, while $\hat{\mathbf{L}}_t$ is the corresponding matrix of estimated pairwise censored likelihood contributions for the t -th replicate (i.e., the values $L_t(\hat{\Gamma}_{ij})$ for each edge $(i, j) \in \mathcal{E}$, arranged in matrix form). For optimizing the likelihood in (20), we use a doubly-stochastic gradient-ascent (DSGA) algorithm, similar to Yu *et al.* (2016).

Algorithm 1 Stochastic Gradient Descent Algorithm

- 1: Initialize the parameter vector $\boldsymbol{\eta}$ as $\boldsymbol{\eta}^{\{0\}}$. Set iteration number to $k = 1$.
 - 2: Choose a mini-batch of the data, denoted by \mathcal{B} .
 - 3: Compute the gradient of the log-likelihood function (21) with respect to $\boldsymbol{\eta}$ using only the data in the mini-batch, denoted by $\nabla_{\boldsymbol{\eta}}g$.
 - 4: Update the parameter vector using the gradient as $\boldsymbol{\eta}^{\{k\}} = \boldsymbol{\eta}^{\{k-1\}} + \rho^{\{k-1\}}\tilde{\nabla}_{\boldsymbol{\eta}}g|_{\boldsymbol{\eta}=\boldsymbol{\eta}^{\{k-1\}}}$, where $\rho^{\{k-1\}}$ is the current step size.
 - 5: Project the updated parameter vector onto the feasible set, which consists of matrices with non-negative symmetric entries and unit upper-triangular Euclidean norm.
 - 6: Update iteration number to $k \mapsto k + 1$
 - 7: Repeat steps 2–6 until convergence.
-

The optimization problem in (20) is challenging due to the constraints on $\boldsymbol{\beta}$ and the complexity of the likelihood function. To ease the optimization, we relax the constraints in (20) and instead solve the regularized maximization problem

$$\hat{\boldsymbol{\eta}} = \arg \max_{\boldsymbol{\eta}=\log\{\boldsymbol{\beta}\}} \left\{ \sum_{t=1}^T \log \det[Q(\boldsymbol{\beta} \odot \hat{\mathbf{L}}_t)] - T \log \det [Q(\boldsymbol{\beta})] - a_1(\|U(\boldsymbol{\beta})\|_2^2 - 1)^2 \right\}, \quad (21)$$

for some fixed positive constant a_1 and where $\eta_{ij} = \log\{\beta_{ij}\}$ for all $(i, j) \in \mathcal{E}$. The DSGA algorithm is a type of stochastic gradient ascent algorithm that updates the parameter vector $\boldsymbol{\eta}$ in small steps based on the gradients of the log-likelihood function with respect to $\boldsymbol{\eta}$.

The stochastic gradient descent algorithm involves initializing the parameter vector $\boldsymbol{\eta}$ and choosing a mini-batch of data denoted by \mathcal{B} ; the algorithm then proceeds by iteratively updating the parameter vector along the gradient of the log-likelihood (see Algorithm 1 for details). The DSGA algorithm has two additional features that improve its convergence. First, it adds a row and column-normalization step after each update which ensures the resulting matrix to have unit norm. Second, it uses a step size that decreases over time to ensure convergence.

4.3 Bias-corrected extremal dependence estimate

Once the dependence parameters in $\boldsymbol{\Gamma}$ are estimated and the edge weights in $\boldsymbol{\beta}$ are computed, the estimated extremal dependence structure for our mixture model of tree-based MPDs can be easily obtained from (18). However, in spatial applications, trees are only an approximate representation

of reality and when the data are truly stationary and isotropic we have found that the tree-mixture model tends to underestimate the extremal dependence structure, especially at large distances; this is due to the fact that the dependence between two sites in the tree-mixture model must follow the paths between these two nodes (on the trees considered in the mixture), which are usually longer than their actual Euclidean spatial distance. This distortion of distances inevitably leads to an underestimation of the dependence strength. In order to mitigate this issue, we propose a simple post-processing step, where the estimated variogram matrix $\hat{\Gamma}$ is scaled by a factor $a \in (0, 1]$. The resulting χ -measure can be computed as in (18), but with the Γ_{ij} values replaced with $a\hat{\Gamma}_{ij}$, i.e.,

$$\hat{\chi}_{ij}(a) = 2 \sum_{m=1}^M \hat{p}(\mathcal{T}_m) \bar{\Phi} \left(\sqrt{\sum_{(i^*, j^*) \in \mathcal{P}(i, j; \mathcal{T}_m)} a \hat{\Gamma}_{i^* j^*}} \right). \quad (22)$$

To estimate the optimal value of a , we then minimize the L_2 -distance between the empirical χ_{ij} , say $\hat{\chi}_{ij}^{\text{emp}}$, and $\hat{\chi}_{ij}(a)$ in (22). For $\hat{\chi}^{\text{emp}} = \{\hat{\chi}_{ij}^{\text{emp}}\}_{(i, j) \in \mathcal{E}}$ and $\hat{\chi}(a) = \{\hat{\chi}_{ij}(a)\}_{(i, j) \in \mathcal{E}}$, the optimal a is thus obtained as $a_{\text{opt}} = \operatorname{argmin}_{a \in (0, 1]} \|\hat{\chi}^{\text{emp}} - \hat{\chi}(a)\|_2^2$. This simple bias-correction procedure works quite well as illustrated in a simulation study reported in the supplementary material of this paper.

5 Wildfire data application

Motivated by the exploratory analysis in Section 2, we analyze the extreme FFDI data using a local moving window approach to account for possible time trends. Specifically, we split the study period of 23 summer seasons (Nov., Dec., Jan., Feb.) into 14 decadal windows and perform window-specific analyses. Here, the first and the last windows are 1999–2009 and 2012–2022, respectively.

The first step of the analysis is to fit marginal distributions. We follow the procedure outlined in Section 3.1, whereby the GP distribution is fitted locally to high threshold exceedances (here, taken to be the empirical 95%-quantile at each site). For estimation, we pool data from neighboring sites, which helps to control spatial smoothness in estimated parameters and to keep marginal inference numerically stable. Figure 5 displays spatial maps of the estimated thresholds, GP scales, and GP shape parameters, for the first and the last windows. The estimated thresholds are spatially smooth

apart from a few “hotspots” near the southern coast. They show higher values near the middle and western regions of Australia and are lower near the eastern and northeastern coastal regions and a small southwestern region near the city of Perth. By contrast, the estimated GP parameters (and especially the shape) show more local variability, which is due to the relatively small sample size available to estimate these two tail parameters locally. The local data pooling approach is therefore crucial. From the maps, we can see that the GP scale parameter takes higher values near the southern coast and a small portion along the northern coast, while the GP shape parameter generally lies within the interval $[-0.25, 0.25]$ and is usually close to zero or slightly negative. While it is difficult to see significant differences between the estimates for the first and the last temporal windows, it can still be noticed that the GP scale for Window 14 is higher than for Window 1 across quite a large portion of the study domain, which is indicative of a temporal shift in FFDI data across the country. To investigate this further, we computed the marginal 99%-quantile based on our fitted model using (5) for both temporal windows; we found that estimated quantiles for Window 14 are higher than those for Window 1 at 1 949 out of 2 749 grid cells (i.e., about 71% of the domain, including major cities), thus confirming that wildfire risk has worsened over time.

The second step is transform the data to the unit Pareto scale, split the spatial domain into spatial clusters as explained in Section 2, and to fit our proposed MPD tree-mixture model from Section 3.2.3 in each spatio-temporal cluster-window combination. We consider 50 clusters for most of the analysis, but also consider lower and higher spatial resolutions (with 25 and 100 clusters, respectively); see Figure 8. The fitted χ_{ij} measure for the two different clusters (in the southwest and eastern part of the domain) considered in Figure 2 are displayed in Figure 6. The figure shows, for each cluster, the tree-based model estimates of χ_{ij} for each edge (i, j) of two different trees \mathcal{T}_m with the highest mixture probabilities $p(\mathcal{T}_m)$ (left and middle columns), as well as $\hat{\chi}_{ij}(a_{\text{opt}})$ for the graph \mathcal{G} based on the optimally bias-corrected tree-mixture MPD model (right column). The estimated dependence structures within these two clusters are clearly non-stationary given that different edges at the same physical distance can have very different χ_{ij} estimates. Figure 7 then

illustrates the behavior the estimated χ_{ij} -measure as a function of the geodesic distance (in km) between sites, for the first and last decadal windows and both spatial clusters considered previously. Interestingly, extremal dependence seems to have become stronger over time in the eastern cluster, where major cities are located. This suggests that extreme wildfires on the east coast have become wider in extent, perhaps due to climate change, which has implication for wildfire management and planning. By contrast, we do not see any significant changes for the cluster on the west coast.

Finally, Figure 8 displays cluster-wise aggregated wildfire risk for the first and last decadal window, computed by averaging samples from the fitted tree-mixture MPD model both over space and independent simulation replicates, after transforming them to the original FFDI scale using the fitted marginal model from Section 3.1. Figure 8 maps the results for Windows 1 and 14, as well as their difference, and does it for different cluster resolutions (low-resolution with 25 clusters, mid-resolution with 50 clusters, and high-resolution with 100 clusters). We observe generally lower wildfire aggregated risk along the eastern coast of Mainland Australia, but this is also the region where the risk has increased the most. This is concerning particularly because the eastern coast is the most densely populated area of Australia. Our results are helpful to understand the regional impact of climate change on the severity and spatial extent of extreme Australian wildfires, which is key for appropriate risk mitigation and planning.

6 Discussions and conclusions

In this paper, we introduce a novel approach for analyzing high-dimensional extremes by combining concepts from extreme-value theory and graphical models. Our proposed methodology extends the tree-based extremal graphical models introduced by Engelke and Hitz (2020) to tree-mixtures that may be used for multivariate or spatial extremes. More precisely, the method involves constructing a mixture of tree-based multivariate Hüsler–Reiss-type Pareto distribution models, with appropriately constructed mixture weights defined through a suitable prior over spanning trees, which enables the fast analysis of multivariate or spatial extremes while accounting for nonstationary tail

dependence structures. This flexible model effectively captures complex tail dependence structures. Our simulation results reported in the Supplementary Material showcase the high flexibility of our modeling approach, and the efficiency of our inference method, as demonstrated by the reasonably good match between the estimated and true extremal dependence profiles in a misspecified setting.

In our real data application, we utilize the newly proposed method to investigate the local spatial tail dependence characteristics of the FFDI fire index in Mainland Australia. To investigate climate change impacts on regional wildfire risk, we adopt a spatio-temporal local approach by analyzing the data separately within spatial clusters and overlapping decadal time windows. Through this approach, we can effectively capture spatio-temporal trends and dependence patterns in FFDI extremes. The results of our analysis provide valuable insights into the severity and spatial extent of extreme wildfire conditions, and how they have changed over time. Our findings highlight that the most densely populated areas of eastern Australia are likely to be most affected by climate change in terms of fire risk. This information is crucial for understanding the potential impacts of climate change on fire-prone regions and can aid in developing targeted mitigation strategies.

While we here focused on the modeling and risk assessment of wildfires, our methodology applies more generally to other types of natural hazards, such as heavy precipitation or extreme heatwaves.

Declarations

Conflicts of interest: The authors have no conflicts of interest to declare.

References

- Andrews, P. L., Loftsgaarden, D. O. and Bradshaw, L. S. (2003) Evaluation of fire danger rating indexes using logistic regression and percentile analysis. International Journal of Wildland Fire **12**(2), 213–226.
- Asenova, S., Mazo, G. and Segers, J. (2021) Inference on extremal dependence in the domain of attraction of a structured Hüsler–Reiss distribution motivated by a Markov tree with latent variables. Extremes **24**, 461–500.
- Balkema, A. A. and De Haan, L. (1974) Residual life time at great age. The Annals of Probability **2**(5), 792–804.
- Beineke, L. W., Wilson, R. J., Cameron, P. J. et al. (2004) Topics in algebraic graph theory. Volume 102. Cambridge University Press.
- Beyers, J. L., Neary, D. G., Ryan, K. C. and DeBano, L. F. (2005) Wildland fire in ecosystems: effects of fire on soil and water. United States Department of Agriculture, Forest Service, Rocky Mountain.
- Cisneros, D., Gong, Y., Yadav, R., Hazra, A. and Huser, R. (2023) A combined statistical and machine learning approach for spatial prediction of extreme wildfire frequencies and sizes. Extremes **26**(2), 301–330.
- Coles, S., Bawa, J., Trenner, L. and Dorazio, P. (2001) An Introduction to Statistical Modeling of Extreme Values. London: Springer.
- Coles, S., Heffernan, J. and Tawn, J. (1999) Dependence measures for extreme value analyses. Extremes **2**(4), 339–365.
- Davison, A. C. and Huser, R. (2015) Statistics of extremes. Annual Review of Statistics and its Application **2**, 203–235.
- Davison, A. C., Huser, R. and Thibaud, E. (2013) Geostatistics of dependent and asymptotically independent extremes. Mathematical Geosciences **45**(5), 511–529.
- Davison, A. C., Huser, R. and Thibaud, E. (2019) Spatial extremes. In Handbook of Environmental and Ecological Statistics, eds A. E. Gelfand, M. Fuentes, J. A. Hoeting and R. L. Smith. Boca Raton: CRC press.
- Davison, A. C. and Smith, R. L. (1990) Models for exceedances over high thresholds. Journal of the Royal Statistical Society: Series B (Methodological) **52**(3), 393–425.
- Dowdy, A. J., Mills, G. A., Finkele, K. and de Groot, W. (2010) Index sensitivity analysis applied to the Canadian forest fire weather index and the McArthur forest fire danger index. Meteorological Applications **17**(3), 298–312.
- Engelke, S. and Hitz, A. S. (2020) Graphical models for extremes. Journal of the Royal Statistical Society: Series B (Statistical Methodology) **82**(4), 871–932.
- Engelke, S. and Volgushev, S. (2022) Structure learning for extremal tree models. Journal of the Royal Statistical Society Series B: Statistical Methodology **84**(5), 2055–2087.

- Epskamp, S., Waldorp, L. J., Möttus, R. and Borsboom, D. (2018) The Gaussian graphical model in cross-sectional and time-series data. Multivariate Behavioral Research **53**(4), 453–480.
- Fiorucci, P., Gaetani, F. and Minciardi, R. (2008) Development and application of a system for dynamic wildfire risk assessment in Italy. Environmental Modelling & Software **23**(6), 690–702.
- Gissibl, N. and Klüppelberg, C. (2018) Max-linear models on directed acyclic graphs. Bernoulli **24**(4A), 2693–2720.
- Godfree, R. C., Knerr, N., Encinas-Viso, F., Albrecht, D., Bush, D., Christine Cargill, D., Clements, M., Gueidan, C., Guja, L. K., Harwood, T. et al. (2021) Implications of the 2019–2020 megafires for the biogeography and conservation of Australian vegetation. Nature Communications **12**(1), 1–13.
- Griffiths, D. (1999) Improved formula for the drought factor in McArthur’s forest fire danger meter. Australian Forestry **62**(2), 202–206.
- Hazra, A., Reich, B. J., Shaby, B. A. and Staicu, A.-M. (2018) A semiparametric spatiotemporal Bayesian model for the bulk and extremes of the Fosberg Fire Weather Index. arXiv preprint arXiv:1812.11699 .
- Huser, R. and Wadsworth, J. L. (2020) Advances in statistical modeling of spatial extremes. Wiley Interdisciplinary Reviews: Computational Statistics p. e1537.
- Hüsler, J. and Reiss, R.-D. (1989) Maxima of normal random vectors: between independence and complete dependence. Statistics & Probability Letters **7**(4), 283–286.
- Jhariya, M. K. and Raj, A. (2014) Effects of wildfires on flora, fauna and physico-chemical properties of soil-an overview. Journal of Applied and Natural Science **6**(2), 887–897.
- Kabluchko, Z., Schlather, M. and De Haan, L. (2009) Stationary max-stable fields associated to negative definite functions. The Annals of Probability **37**(5), 2042–2065.
- Klüppelberg, C. and Lauritzen, S. (2019) Bayesian networks for max-linear models. In Network Science, pp. 79–97. Springer.
- Koller, D. and Friedman, N. (2009) Probabilistic Graphical Models: Principles and Techniques. MIT press.
- Lauritzen, S. (1996) Graphical Models. Oxford Statistical Science: Series 17.
- Marlon, J. R., Bartlein, P. J., Gavin, D. G., Long, C. J., Anderson, R. S., Briles, C. E., Brown, K. J., Colombaroli, D., Hallett, D. J., Power, M. J. et al. (2012) Long-term perspective on wildfires in the western USA. Proceedings of the National Academy of Sciences **109**(9), E535–E543.
- Meilă, M. and Jaakkola, T. (2006) Tractable Bayesian learning of tree belief networks. Statistics and Computing **16**(1), 77–92.
- Noble, I., Gill, A. and Bary, G. (1980) McArthur’s fire-danger meters expressed as equations. Australian Journal of Ecology **5**(2), 201–203.
- Papastathopoulos, I. and Strokorb, K. (2016) Conditional independence among max-stable laws. Statistics & Probability Letters **108**, 9–15.

- Pickands III, J. (1975) Statistical inference using extreme order statistics. The Annals of Statistics **3**(1), 119 – 131.
- Ren, Z., Sun, T., Zhang, C.-H. and Zhou, H. H. (2015) Asymptotic normality and optimalities in estimation of large Gaussian graphical models. The Annals of Statistics **43**(3), 991–1026.
- Resnick, S. I. (2008) Extreme values, regular variation, and point processes. Volume 4. Springer Science & Business Media.
- Rootzén, H., Segers, J. and L Wadsworth, J. (2018) Multivariate peaks over thresholds models. Extremes **21**(1), 115–145.
- Rootzén, H. and Tajvidi, N. (2006) Multivariate generalized Pareto distributions. Bernoulli **12**(5), 917–930.
- Sanabria, L., Qin, X., Li, J., Cechet, R. and Lucas, C. (2013) Spatial interpolation of McArthur’s forest fire danger index across Australia: observational study. Environmental modelling & Software **50**, 37–50.
- Sandberg, D. V. (2003) Wildland fire in ecosystems: effects of fire on air. US Department of Agriculture, Forest Service, Rocky Mountain Research Station.
- Segers, J. (2020) One- versus multi-component regular variation and extremes of Markov trees. Advances in Applied Probability **52**(3), 855–878.
- Sirca, C., Spano, D., Pisanu, P., Duce, P., Delogu, G. and Cicalò, G. O. (2006) Ichnusa fire index: Development and preliminary evaluation at local scale. Forest Ecology and Management **234**(S), 250.
- Stephenson, C., Handmer, J. and Betts, R. (2013) Estimating the economic, social and environmental impacts of wildfires in Australia. Environmental Hazards **12**(2), 93–111.
- Thomas, D., Butry, D., Gilbert, S., Webb, D., Fung, J. et al. (2017) The costs and losses of wildfires. NIST Special Publication **1215**(11).
- Van Wagner, C., Forest, P. et al. (1987) Development and structure of the Canadian forest fire weather index system. In Canadian Forestry Service, Forestry Technical Report.
- Vettori, S., Huser, R., Segers, J. and Genton, M. G. (2020) Bayesian model averaging over tree-based dependence structures for multivariate extremes. Journal of Computational and Graphical Statistics **29**(1), 174–190.
- Westerling, A. L. (2016) Increasing western US forest wildfire activity: sensitivity to changes in the timing of spring. Philosophical Transactions of the Royal Society B: Biological Sciences **371**(1696), 20150178.
- Yu, H., Uy, W. I. T. and Dauwels, J. (2016) Modeling spatial extremes via ensemble-of-trees of pairwise copulas. IEEE Transactions on Signal Processing **65**(3), 571–586.

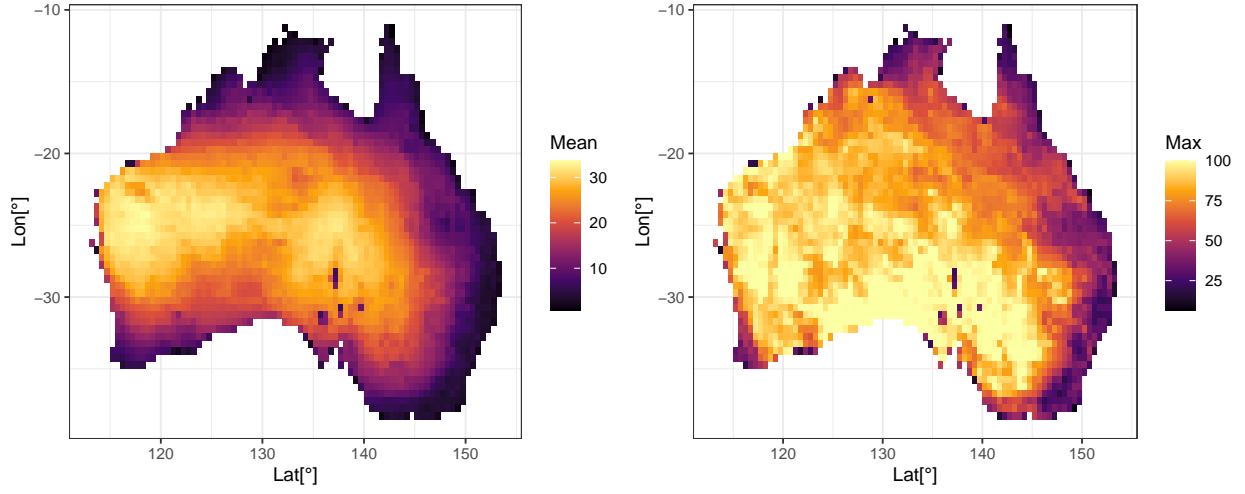


Figure 1: Grid cell-wise empirical mean (left) and maximum (right) of FFDI data.

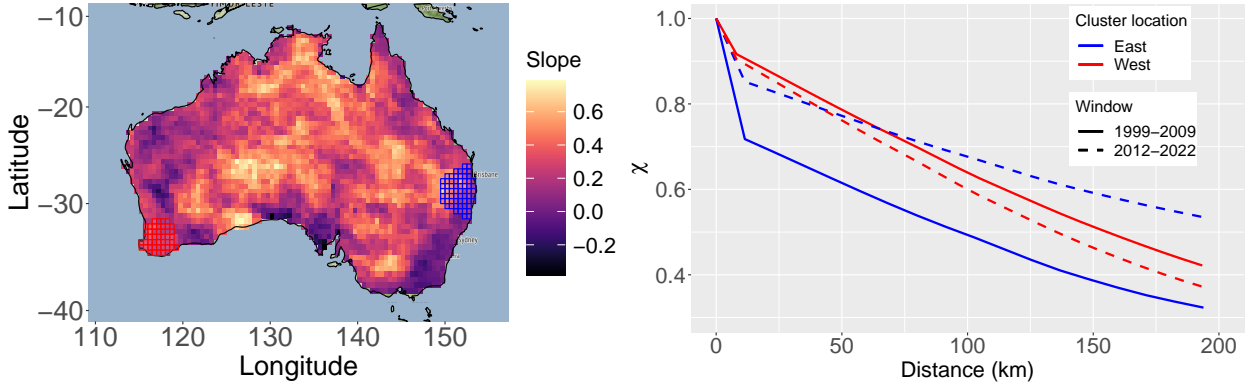


Figure 2: Grid cell-wise annual rate of change in the 98%-empirical quantile (left) and smoothed empirical extremal dependence profiles for two spatial clusters near the western and eastern coasts (highlighted in red and blue in the left panel), for the first and last decadal windows (right).

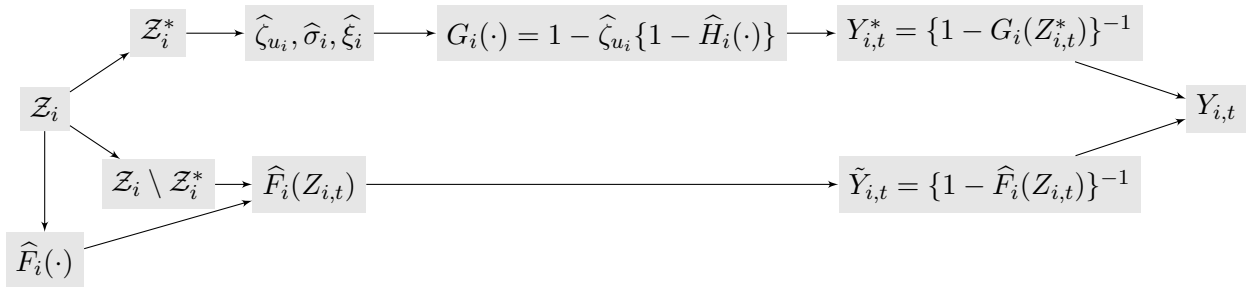


Figure 3: Flow chart of the marginal estimation procedure, and the transformation of data to the unit Pareto scale.

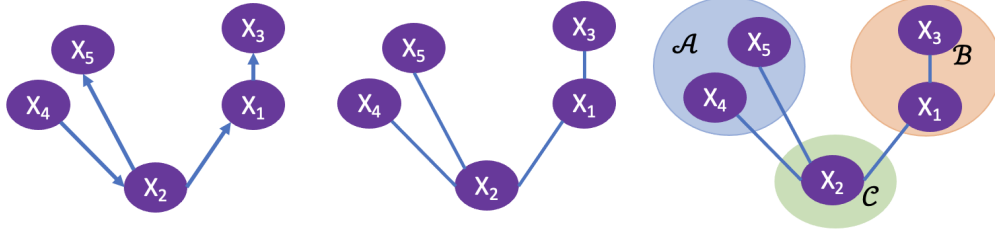


Figure 4: A directed graph (left) and the corresponding undirected graph (middle). In the right panel, for the graph in the middle, the separator set \mathbf{X}_C is highlighted in green, and \mathbf{X}_A and \mathbf{X}_B are highlighted in blue and orange, respectively.

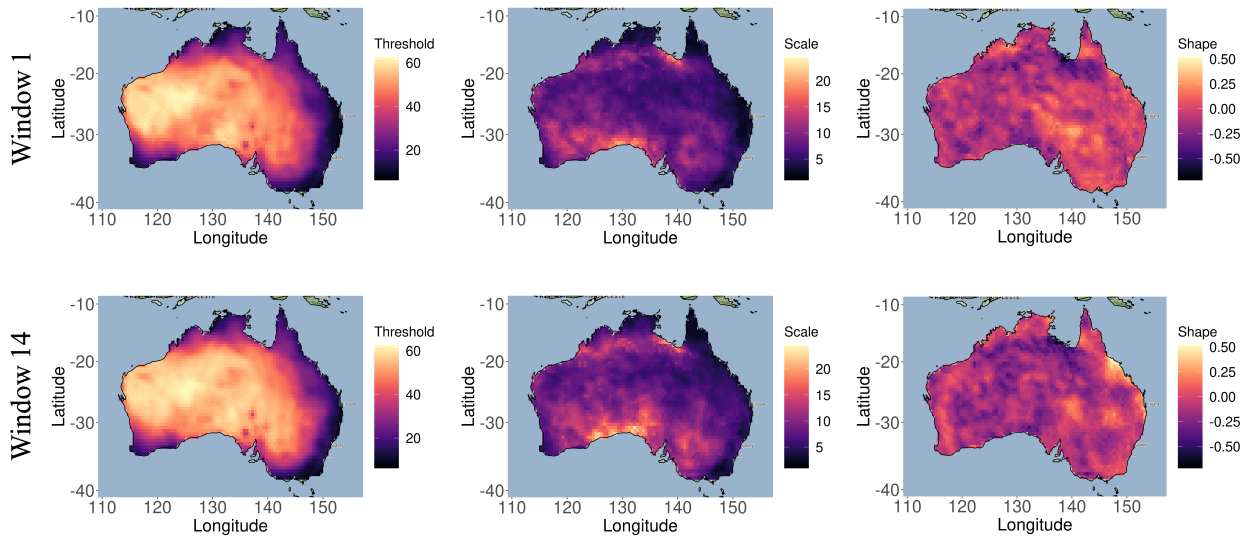


Figure 5: Spatial maps of the estimated thresholds (left), the generalized Pareto scale (middle), and shape (right) parameters for the first (top) and the last (bottom) temporal windows.

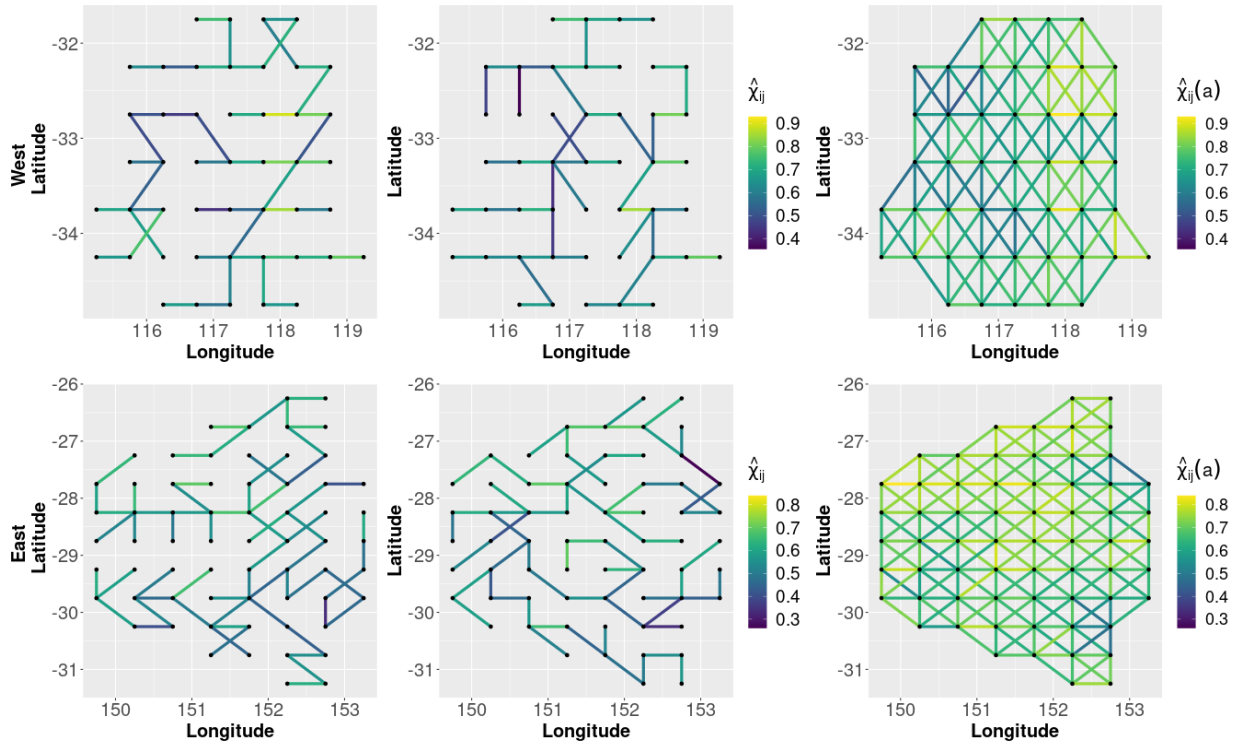


Figure 6: Model-based estimates of χ_{ij} for the tree-based MPD fitted within two clusters near the western (top) and the eastern (bottom) coastal regions, respectively, for the temporal window 2012–2022. The plots show the estimated tree-based χ_{ij} for the two trees with the highest (left) and the second highest (middle) estimated mixture probabilities, as well as the optimal $\hat{\chi}_{ij}(a_{\text{opt}})$ for the tree-mixture MPD model (right).

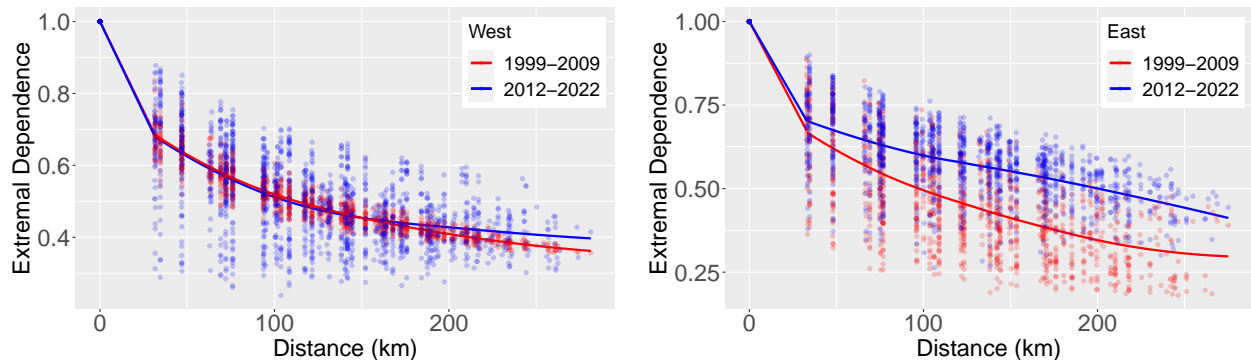


Figure 7: Smoothed bias-corrected $\chi_{ij}(a)$ estimates as a function of geodesic distance (in km) between sites, based on the proposed tree-mixture MPD model for two spatial clusters along the west (left) and east (right) coasts for the first (red) and last (blue) decade windows.

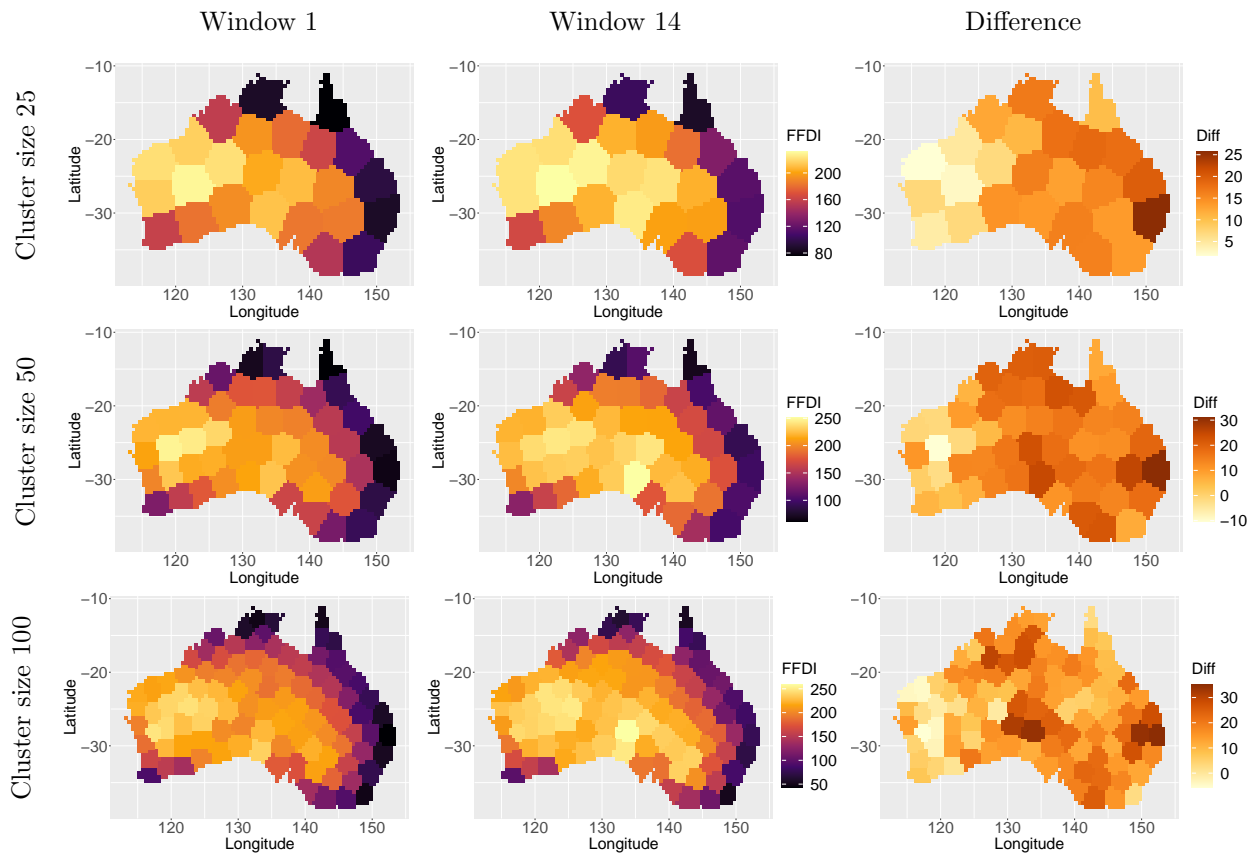


Figure 8: Spatial fire risk aggregates for two decade time Windows 1 and 14 (left and middle) and three cluster resolutions. The right panels show the difference in the spatial fire risk aggregates between the two windows.