

Exact Simulation of Max-Infinitely Divisible Processes

Peng Zhong^a, Raphaël Huser^{a,*}, Thomas Opitz^b

^a*Statistics Program, Computer, Electrical and Mathematical Sciences and Engineering (CEMSE) Division, King Abdullah University of Science and Technology (KAUST), Thuwal 23955-6900, Saudi Arabia*

^b*BioSP, INRAE, Avignon, 84914, France*

Abstract

Max-infinitely divisible (max-id) processes play a central role in extreme-value theory and include the subclass of all max-stable processes. They allow for a constructive representation based on the pointwise maximum of random functions drawn from a Poisson point process defined on a suitable function space. Simulating from a max-id process is often difficult due to its complex stochastic structure, while calculating its joint density in high dimensions is often numerically infeasible. Therefore, exact and efficient simulation techniques for max-id processes are useful tools for studying the characteristics of the process and for drawing statistical inferences. Inspired by the simulation algorithms for max-stable processes, theory and algorithms to generalize simulation approaches tailored for certain flexible (existing or new) classes of max-id processes are presented. Efficient simulation for a large class of models can be achieved by implementing an adaptive rejection sampling scheme to sidestep a numerical integration step in the algorithm. The results of a simulation study highlight that our simulation algorithm works as expected and is highly accurate and efficient, such that it clearly outperforms customary approximate sampling schemes. As a by-product, new max-id models, which can be represented as pointwise maxima of general location-scale mixtures and possess flexible tail dependence structures capturing a wide range of asymptotic dependence scenarios, are also developed.

Keywords: Adaptive rejection sampling, Exact simulation, Extremal function, Max-infinitely divisible process, Max-stable process

1. Introduction

Max-infinitely divisible processes (max-id) have gained much popularity in recent years for the flexible modeling of spatiotemporal extremes in phenomena such as precipitation, wind speeds and temperatures (Padoan, 2013; Huser et al., 2021; Bopp et al., 2021; Huser and Wadsworth, 2022; Zhong et al., 2022). Such models encompass max-stable processes (de Haan, 1984) as a subclass, which are currently the mainstream for modeling spatial and temporal extremes, as demonstrated by a large body of theoretical and applied literature (see, e.g., Davison et al., 2012; Opitz, 2013; Huser and Davison, 2014; Dey and Yan, 2016; Gissibl et al., 2018; Davison et al., 2019). Max-stability is a property that arises in asymptotic models for multivariate block maxima, but that is however often violated in finite samples (Huser and Wadsworth, 2022). Precisely, the dependence strength in max-stable processes does not depend on event magnitude, while most environmental and climatic processes suggest that it should become weaker at higher thresholds, with the most extreme events being typically more localized. To circumvent this limitation, more flexible max-id processes, which drop the restrictive max-stability assumption but retain natural properties of multivariate block maxima, have been considered. In particular, the max-id models proposed by Huser et al. (2021),

*Corresponding author

Email addresses: peng.zhong@kaust.edu.sa (Peng Zhong), raphael.huser@kaust.edu.sa (Raphaël Huser), thomas.opitz@inrae.fr (Thomas Opitz)

Bopp et al. (2021) and Zhong et al. (2022), extend certain popular classes of max-stable models (obtained as limiting cases on the boundary of the parameter space), in order to achieve a good compromise between the pragmatism of flexible max-id models and the strong theoretical foundations of the max-stable sub-class.

The theory behind max-id processes has been studied in depth (Brown and Resnick, 1977; Resnick, 1987; Giné et al., 1990). In particular, Giné et al. (1990) showed that max-id processes possess a functional Poisson point process (PPP) representation. More precisely, each max-id process $Z = \{Z(\mathbf{s})\}_{\mathbf{s} \in \mathcal{S}}$ defined over the region $\mathcal{S} \subset \mathbb{R}^d$, can be defined as the pointwise maximum of a potentially infinite number of random functions $\{\eta_i; i = 1, 2, \dots\}$ defined over \mathcal{S} , which are sampled according to a PPP with mean measure Λ , i.e.,

$$Z(\mathbf{s}) = \max_{i=1,2,\dots} \eta_i(\mathbf{s}), \quad \mathbf{s} \in \mathcal{S}, \quad (1)$$

where, by convention, $Z(\mathbf{s})$ takes value at the lower boundary of the support of $\eta_i(\mathbf{s})$ when the PPP $\{\eta_i\}$ contains no point (which can happen when Λ is finite). To avoid intricate theoretical issues dealing with the lower boundary, we here assume that Λ is an infinite measure, such that the number of Poisson points η_i in (1) is infinite. However, analytical forms of the multivariate density functions of such processes exist only in special cases, and by analogy with max-stable processes their numerical evaluation in high dimensions becomes impossible due to the combinatorial explosion of the number of terms to be calculated (Padoan et al., 2010; Castruccio et al., 2016; Huser et al., 2021). Consequently, simulation of max-stable processes is often crucial to study their dependence characteristics, or for simulation-based inference (Erhardt and Smith, 2012; Hainy et al., 2016; Lee et al., 2018), and the same is true for general classes of max-id processes. However, exact simulation of $Z(\mathbf{s})$ in (1) is not trivial, especially when the number of functions η_i is infinite and when these random functions are tricky to sample.

For the subclass of max-stable models, there are essentially two types of generic simulation algorithm that have been proposed based on variants of model representations stemming from (1) (Oesting and Strokorb, 2022). These algorithms exploit the specific structure of max-stable processes—the so-called *spectral representation* (de Haan, 1984), whereby on the unit Fréchet scale (i.e., $\Pr\{Z(\mathbf{s}) \leq z\} = \exp(-1/z)$, $z > 0$) the points of the Poisson process $\{\eta_i; i = 1, 2, \dots\}$ in (1) can be decomposed into two multiplicative and stochastically independent components as $\eta_i(\mathbf{s}) = R_i W_i(\mathbf{s})$. Precisely, $\{R_i; i = 1, 2, \dots\}$ here denotes a Poisson point process on $[0, \infty)$ with mean measure $r^{-2} dr$, independent of the random functions $\{W_i; i = 1, 2, \dots\}$, which are independent copies of a random process $W(\mathbf{s})$ defined on \mathcal{S} such that $E[\max\{0, W(\mathbf{s})\}] = 1$. This decomposition always exists for max-stable processes with unit Fréchet margins but is not unique. When the points $\{R_i; i = 1, 2, \dots\}$ have heavy-tailed intensity $\alpha r^{-(\alpha+1)} dr$ for some $\alpha > 0$, the resulting process Z remains max-stable, but with α -Fréchet margins. Moreover, the points $\{R_i^{-1}; i = 1, 2, \dots\}$ can be defined as the arrival times of a renewal process such that we can simulate them in decreasing order, i.e., $R_1 > R_2 > \dots$. When this is the case, the componentwise maximum of $\{R_j W_j; j = 1, \dots, i-1\}$ is less and less impacted by the i -th process $\eta_i = R_i W_i$ as $i > 1$ is iteratively incremented. If $W(\mathbf{s})$ is an almost surely bounded random process on \mathcal{S} , the contribution of $\eta_i = R_i W_i$ becomes even ultimately completely irrelevant, as i grows. Based on this idea, Schlather (2002) introduced the first exact simulation algorithm for max-stable processes by setting an appropriate stopping rule for the simulation of the Poisson points $R_i W_i$. If $W(\mathbf{s})$ is unbounded, we can use certain alternative representations to make it a bounded random process under a suitable change of measure (Dieker and Mikosch, 2015; Dombry et al., 2016; Oesting et al., 2018), but simulation from the bounded process is intricate and numerically challenging even in moderate dimensions for many of the customary models (Kablichko et al., 2009; Opitz, 2013). The second type of exact simulation algorithm for max-stable processes, introduced by Dombry et al. (2016), relies on the concept of extremal functions (Dombry and Eyi-Minko, 2013). With this method, it is possible to directly simulate the random functions $\{\eta_i\}$ on \mathcal{S} that contribute to the maximum at a finite number of locations. In practice, efficient conditional simulation of the random functions $\eta_i(\mathbf{s})$ given its value at a single location \mathbf{s}_0 is a prerequisite for efficient simulation with this approach.

However, none of the approaches described above tackled exact simulation for general max-id processes, and we intend to fill this gap. In this work, we demonstrate that the idea of the exact simulation algorithm based on extremal functions, developed in Dombry et al. (2016) for max-stable processes carries over to the general max-id case after suitable adjustments, and we illustrate it with max-id models that have found

interest in practical applications. The paper is organized as follows. In §2, we recall some theoretical background and formally define the notions of max-id processes and extremal functions. Then, in §3, we derive the specificities of the theory of max-id processes required for exact simulation based on extremal functions, and we present simulation algorithms focusing on two major classes of max-id processes. In particular, we start with some known max-id models of the form (1) where the random functions $\eta_i(\mathbf{s})$ can be represented as Gaussian scale mixtures, and we then also explore new max-id models based on Gaussian location mixtures that possess appealing tail dependence characteristics. In §4, we conduct a simulation study to demonstrate the performance of our proposed simulation algorithm. We conclude with a discussion in §5.

2. Max-id Processes and Extremal Functions

A process $\{Z(\mathbf{s})\}_{\mathbf{s} \in \mathcal{S}}$ is called max-infinitely divisible (max-id) if the joint cumulative distribution function (cdf) G of $\mathbf{Z}(K_N) = \{Z(\mathbf{s}_1), \dots, Z(\mathbf{s}_N)\}^\top$ at any finite collection of sites $K_N = \{\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_N\} \subset \mathcal{S}$ defines a valid cdf G^t (with the notation $G^t(\mathbf{z}) = \{G(\mathbf{z})\}^t$, $\mathbf{z} = (z_1, \dots, z_N)^\top \in \mathbb{R}^N$) for any $t > 0$. The distribution G^t does not necessarily stay within the same location-scale family; this property is only satisfied for max-stable distributions \tilde{G} , for which $\tilde{G}^m(\mathbf{a}_m \mathbf{z} + \mathbf{b}_m) = \tilde{G}(\mathbf{z})$, for any integer $m = 1, 2, \dots$, and appropriately defined normalizing vectors $\mathbf{a}_m \in (0, \infty)^N$ and $\mathbf{b}_m \in \mathbb{R}^N$. As shown in Equation (1), any max-id process can be constructed by taking pointwise maxima over a Poisson point process defined on a suitable function space. We assume that $\{Z(\mathbf{s})\}_{\mathbf{s} \in \mathcal{S}}$ is a sample-continuous max-id process in a function space \mathcal{C} , where \mathcal{C} denotes the space of continuous functions with compact support $\mathcal{S} \subset \mathbb{R}^d$ endowed with the uniform norm, i.e., $\|f\| = \sup_{\mathbf{s} \in \mathcal{S}} |f(\mathbf{s})|$, $f \in \mathcal{C}$, and $\{\eta_i(\mathbf{s}); i = 1, 2, \dots\}_{\mathbf{s} \in \mathcal{S}}$ in (1) is a Poisson point process with mean measure Λ on \mathcal{C} . Then, the extremal functions are defined as follows.

Definition 1. Let $K \subset \mathcal{S}$ be a nonempty subset of the domain \mathcal{S} , and Z a max-id process. A function $\phi \in \mathcal{C}$ is called K -extremal if there exists $\mathbf{s} \in K$ such that $Z(\mathbf{s}) = \phi(\mathbf{s})$, otherwise ϕ is called K -subextremal. We denote $\mathcal{C}_K^+(Z)$ the set of K -extremal functions and $\mathcal{C}_K^-(Z)$ the set of K -subextremal functions.

In other words, an extremal function is a Poisson point η_i in \mathcal{C} that contributes to the maximum process Z at one or more locations. Notice that both $\mathcal{C}_K^+(Z)$ and $\mathcal{C}_K^-(Z)$ depend on the max-id process Z , even though the realization of the max-id process Z on K is fully determined by $\mathcal{C}_K^+(Z)$ alone. Let ℓ be the vertex function (i.e., the lower boundary) of the max-id process Z , defined by

$$\ell(\mathbf{s}) = \sup\{z \in \mathbb{R} : \Pr(Z(\mathbf{s}) \geq z) = 1\} \in [-\infty, \infty), \quad \mathbf{s} \in \mathcal{S}. \quad (2)$$

Giné et al. (1990) showed that for sample continuous max-id processes in \mathcal{C} , if the vertex function ℓ is continuous, then it can be subtracted from Z . Therefore, we here assume $\ell(\mathbf{s}) = 0$ without loss of generality, and denote $\mathcal{C}_0 = \{f \in \mathcal{C}; f \neq 0, f \geq 0\}$. We further assume that for any fixed $\mathbf{s}_0 \in \mathcal{S}$,

$$\Lambda_{\mathbf{s}_0}(z) = \Lambda(\{\eta \in \mathcal{C}_0; \eta(\mathbf{s}_0) \in [z, \infty)\}) \text{ is continuous on } (0, \infty), \lim_{z \downarrow 0} \Lambda_{\mathbf{s}_0}(z) = \infty, \quad (3)$$

$$\Lambda_{\mathbf{s}_0}(\{\eta \in \mathcal{C}_0; \eta(\mathbf{s}_0) > \varepsilon\}) < \infty, \forall \varepsilon > 0. \quad (4)$$

Under these assumptions, Dombry and Eyi-Minko (2013) showed that for singletons $K = \{\mathbf{s}_0\}$ the set of extremal functions $\mathcal{C}_{\{\mathbf{s}_0\}}^+$ contains almost surely a single point (a function), denoted by $\phi_{\mathbf{s}_0}^+$, i.e., the resulting max-id “process” $Z(\mathbf{s}_0)$ is realized by one and only one Poisson point η_i at a given location \mathbf{s}_0 . This assumption implies that Z has continuous margins and the marginal distribution has no mass at the lower boundary $\ell(\mathbf{s}) = 0$.

Under the standing assumptions, Dombry and Eyi-Minko (2013, Lemma A.2) show that the mean measure Λ possesses a regular conditional probability measure $P_{\mathbf{s}_0}(z, \cdot)$ with respect to the location \mathbf{s}_0 ; i.e., given that $\eta_i(\mathbf{s}_0) = z > 0$ for a point η_i of the Poisson point process, the random function $\eta_i(\mathbf{s})$ for $\mathbf{s} \neq \mathbf{s}_0$ has (conditional) probability distribution given by $P_{\mathbf{s}_0}(z, \cdot)$. Dombry and Eyi-Minko (2013) further show in their Theorem 3.2 that $P_{\mathbf{s}_0}(z, \cdot)$ is also the conditional distribution for the extremal function $\phi_{\mathbf{s}_0}^+$ at \mathbf{s}_0

given $Z(\mathbf{s}_0) = z > 0$; moreover, conditionally on the values $Z(\mathbf{s}_1), \dots, Z(\mathbf{s}_n)$ of the max-id process at $K = \{\mathbf{s}_1, \dots, \mathbf{s}_n\}$, the subextremal functions in $\mathcal{C}_K^-(Z)$ define a Poisson point process that is independent of the extremal functions in $\mathcal{C}_K^+(Z)$. This restricted Poisson point process has again mean measure Λ but now with all the mass “removed” when one of the conditioning values $Z(\mathbf{s}_1), \dots, Z(\mathbf{s}_n)$ is exceeded, i.e., it has no mass on $\{\eta \in \mathcal{C}_0 : \max_{i=1}^n \eta(\mathbf{s}_i)/Z(\mathbf{s}_i) > 1\}$. This remarkable result is the key for the proposed simulation algorithm that consists in iteratively simulating extremal functions at a set of locations (see §3.1), and we summarize it in the following lemma.

Lemma 2. (*Dombry and Eyi-Minko, 2013, Theorem 3.2*) *Given the extremal functions $\mathcal{C}_K^+(Z)$ on K , the subextremal functions $\mathcal{C}_K^-(Z)$ form a Poisson point process on \mathcal{C}_0 with intensity $\mathbb{I}(f(K) < Z(K))\Lambda(df)$, where $\mathbb{I}(\cdot)$ is the indicator function.*

In particular, when the process $\{Z(\mathbf{s})\}_{\mathbf{s} \in \mathcal{S}}$ is max-stable with unit Fréchet margins, the Poisson point process $\{\eta_i; i = 1, 2, \dots\}$ can be decomposed as $\eta_i(\mathbf{s}) = R_i W_i(\mathbf{s})$, $i = 1, 2, \dots$, where $\{R_i; i = 1, 2, \dots\}$ is a Poisson point process on $(0, \infty)$ with mean measure $\kappa([r, \infty]) = r^{-1}$, $r > 0$, and $\{W_i(\mathbf{s})\}_{\mathbf{s} \in \mathcal{S}}$, $i = 1, 2, \dots$, are independent copies of a random process $\{W(\mathbf{s})\}_{\mathbf{s} \in \mathcal{S}}$ with $\mathbb{E}[\max\{W(\mathbf{s}), 0\}] = 1$, which are also independent of the points $\{R_i; i = 1, 2, \dots\}$ (de Haan, 1984; Schlather, 2002). A number of parametric max-stable models have been proposed for statistical applications, such as the Brown–Resnick model (Brown and Resnick, 1977; Kabluchko et al., 2009), and the extremal- t model (Opitz, 2013), where the intensity function associated to the mean measure Λ has a closed form expression. In these special cases, the multivariate conditional distributions required for exact simulation are relatively easy to simulate from, and Dombry et al. (2016) provide the exact forms of the distributions $P_{\mathbf{s}}(z, \cdot)$ required for simulating these max-stable models.

However, such simplifications are no longer available for general max-stable processes and for the classes of max-id (but not max-stable) process models proposed in the recent literature. Very often, the mean measure $\Lambda(\cdot)$ does not have a closed-form intensity function, and the corresponding conditional distributions are complicated; rather, calculations of relevant measures of the form $\Lambda(B)$ and intensities often involve computing a one-dimensional integral (Huser et al., 2021; Zhong et al., 2022). Therefore, simulating from $P_{\mathbf{s}}(z, \cdot)$ is more involved than in the max-stable case and requires further investigation into its form to propose efficient simulation techniques.

3. Exact Simulation

3.1. General simulation algorithm

We now extend the generic max-stable simulation algorithm based on extremal functions, as proposed by (Dombry et al., 2016), to the general max-id case. While our proposed algorithm in the max-id case is almost identical in its overall idea and general structure to Dombry et al.’s algorithm in the max-stable case (modulo certain technicalities and specificities in how the different steps are executed), it is still important to note that generalizing a specific simulation algorithm that works for max-stable processes to the much wider class of max-id models, while remaining an exact and feasible simulation procedure, is not necessarily trivial and requires careful investigation. This has never been explored before.

Our aim is, thus, to simulate the max-id process $\{Z(\mathbf{s})\}_{\mathbf{s} \in \mathcal{S}}$ at a finite number N of sites given as $K_N = \{\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_N\} \subset \mathcal{S}$. We define the subsets $K_n = \{\mathbf{s}_1, \dots, \mathbf{s}_n\}$ of the first n sites for $n = 1, \dots, N$, and we denote by $Z_n = \max_{\eta_i \in \mathcal{C}_{K_n}^+} \eta_i$ the process constructed from the extremal functions at the sites in K_n , such that $\mathbf{Z}(K_N) = \mathbf{Z}_N(K_N) := \{Z_N(\mathbf{s}_1), \dots, Z_N(\mathbf{s}_N)\}^\top$. To simulate from the joint distribution of the process $\{Z(\mathbf{s})\}$ at sites in K_N , we can iterate through the set of locations from \mathbf{s}_1 to \mathbf{s}_N in order to simulate the extremal functions at each location. More precisely, suppose we have already simulated the extremal functions on K_n . Recall that we denote by $\{\phi_{\mathbf{s}_n}^+\}_{1 \leq n \leq N}$ the sequence of extremal functions associated to the singletons $\{\mathbf{s}_n\}$, $n = 1, \dots, N$, where some of these functions may be identical to each other. The extremal function $\phi_{\mathbf{s}_{n+1}}^+$ at a new location $\mathbf{s}_{n+1} \notin K_n$ either coincides with one of the extremal functions on K_n already simulated, i.e., $\phi_{\mathbf{s}_{n+1}}^+ \in \mathcal{C}_{K_n}^+(Z)$, or it corresponds to a subextremal function on K_n , i.e., $\phi_{\mathbf{s}_{n+1}}^+ \in \mathcal{C}_{K_n}^-(Z)$, in which case it is a point of the restricted Poisson point process defined in Lemma 2.

Moreover, at \mathbf{s}_{n+1} the maximum of the values of the extremal functions on K_n , i.e., $Z_n(\mathbf{s}_{n+1})$, already provides a lower bound for the values of the new extremal function that could arise at \mathbf{s}_{n+1} . Therefore, we proceed as follows to find the extremal function $\phi_{\mathbf{s}_{n+1}}^+$, which directly provides $Z(\mathbf{s}_{n+1})$. We simulate from the unrestricted point process by assuming that its points are ordered in descending order of the values at \mathbf{s}_{n+1} , i.e., $\eta_1(\mathbf{s}_{n+1}) > \eta_2(\mathbf{s}_{n+1}) > \dots$. Simulation starts from the point η_1 with the largest value at \mathbf{s}_{n+1} and then continues in descending order. To keep only the points from the restricted point process of the subextremal functions of K_n , we reject the points violating the upper bound conditions at $\mathbf{s}_1, \dots, \mathbf{s}_n$ set by the restrictions given in Lemma 2. We simulate a point η_i by first simulating its value $\eta_i(\mathbf{s}_{n+1})$ at \mathbf{s}_{n+1} according to the marginal (unrestricted) Poisson point process with mean measure $\Lambda_{\mathbf{s}_{n+1}}$, and we then simulate the values at all the other sites in K_N using the regular conditional probability distribution $P_{\mathbf{s}_{n+1}}(\eta_i(\mathbf{s}_{n+1}), \cdot)$. Moreover, we can stop simulating as soon as a value $\eta_i(\mathbf{s}_{n+1})$ falls below the lower bound $Z_n(\mathbf{s}_{n+1})$ and we reject it, since this value (as well as all the following smaller ones) cannot correspond to the extremal function at \mathbf{s}_{n+1} . In any case, we stop as soon as we obtain a point η_i that is not rejected, and this point then corresponds to $\phi_{\mathbf{s}_{n+1}}^+$. If we have rejected all points until we stop, there is no new extremal function at \mathbf{s}_{n+1} , such that $\phi_{\mathbf{s}_{n+1}}^+ \in \mathcal{C}_{K_n}^+(Z)$.

The benefit of this approach is that it requires to simulate only a finite number of Poisson points to either obtain a new extremal function at \mathbf{s}_{n+1} , or to confirm that the extremal function at \mathbf{s}_{n+1} coincides with one of the already simulated extremal functions on K_n . Therefore, the approach exploits the fact that the already simulated extremal functions provide upper bounds at $\mathbf{s}_1, \dots, \mathbf{s}_n$ and a lower bound at \mathbf{s}_{n+1} . This idea then allows us to iteratively simulate all of the extremal functions at a finite set of locations $\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_N$, where we condition on the extremal functions at $\mathbf{s}_1, \dots, \mathbf{s}_n$ when identifying the extremal function at \mathbf{s}_{n+1} , $1 \leq n < N$. Since the extremal functions define the values of the max-id process at the locations $\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_N$, this is all we need. Dombry et al. (2016) have proposed this approach for max-stable processes, but their algorithm generalizes straightforwardly to max-id processes thanks to the general results of Dombry and Eyi-Minko (2013). The following theorem formalizes the simulation procedure detailed above by characterizing the conditional distributions of $\phi_{\mathbf{s}_n}^+$ given $\{\phi_{\mathbf{s}_k}^+\}_{1 \leq k \leq n-1}$ for extremal functions at position $n = 2, \dots, N-1$, which are required for iterative simulation.

Theorem 3 (Conditional distributions of extremal functions). *Consider a max-id process $\{Z(\mathbf{s})\}_{\mathbf{s} \in \mathcal{S}}$ as defined in §2. Given sites $K_N = \{\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_N\}$, the distribution of its extremal function $\phi_{\mathbf{s}_1}^+$ and of its conditional extremal functions $\phi_{\mathbf{s}_{n+1}}^+ \mid \{\phi_{\mathbf{s}_k}^+\}_{1 \leq k \leq n}$ for $n \geq 1$ can be characterized as follows:*

- (1) *Initial extremal function: given a realization z of $Z(\mathbf{s}_1)$, the extremal function $\phi_{\mathbf{s}_1}^+$ is distributed according to $P_{\mathbf{s}_1}(z, \cdot)$.*
- (2) *Conditional extremal functions for $n \geq 1$: given the extremal functions $\{\phi_{\mathbf{s}_i}^+\}_{1 \leq i \leq n}$ for sites in K_n , consider the Poisson measure of functions that do not exceed the already simulated max-id values of sites in K_n but that exceed the maximum of already simulated values at \mathbf{s}_{n+1} :*

$$\tilde{\Lambda}(\mathrm{d}f) = \mathbb{I}(f(\mathbf{s}_i) < Z_n(\mathbf{s}_i), 1 \leq i \leq n) \mathbb{I}(f(\mathbf{s}_{n+1}) > Z_n(\mathbf{s}_{n+1})) \Lambda(\mathrm{d}f). \quad (5)$$

Then, we denote the Poisson point process with mean measure $\tilde{\Lambda}$ by $\text{PPP}(\tilde{\Lambda})$, and the extremal function $\phi_{\mathbf{s}_{n+1}}^+$ conditional on $\{\phi_{\mathbf{s}_i}^+\}_{1 \leq i \leq n}$ is given as follows:

$$\tilde{\phi}_{\mathbf{s}_{n+1}}^+ = \begin{cases} \arg \max_{\phi \in \text{PPP}(\tilde{\Lambda})} \phi(\mathbf{s}_{n+1}), & \text{PPP}(\tilde{\Lambda}) \neq \emptyset, \\ \arg \max_{\phi \in \{\phi_{\mathbf{s}_1}^+, \dots, \phi_{\mathbf{s}_n}^+\}} \phi(\mathbf{s}_{n+1}), & \text{PPP}(\tilde{\Lambda}) = \emptyset. \end{cases}$$

Proof. The proof essentially follows the intuitive description of the algorithm above. It is almost identical (modulo some technicalities) to the proof of Theorem 2 in Dombry et al. (2016), which concerns max-stable processes, while we here treat max-id processes.

As mentioned, we proceed sequentially to simulate extremal functions. The distribution of the first extremal function $\phi_{\mathbf{s}_1}^+$ given $\phi_{\mathbf{s}_1}^+(\mathbf{s}_1) = z$ is exactly $P_{\mathbf{s}_1}(z, \cdot)$, as explained in §2 above, and $\phi_{\mathbf{s}_1}^+(\mathbf{s}_1)$ thus follows the marginal distribution of $Z(\mathbf{s}_1)$, which has now a more general form than for max-stable processes.

Notice also that given the continuity assumption in (3), there is almost surely exactly one extremal function at each location.

Then, to simulate the extremal function $\phi_{\mathbf{s}_{n+1}}^+$ at site \mathbf{s}_{n+1} , for $n \geq 1$, we recall that $\phi_{\mathbf{s}_{n+1}}^+$ is either a subextremal function in $\mathcal{C}_{K_n}^-$ or an extremal function in $\mathcal{C}_{K_n}^+$. According to Lemma 2, given $\mathcal{C}_{K_n}^+$, the conditional distribution of $\mathcal{C}_{K_n}^-$ follows the distribution of a Poisson point process with intensity

$$\mathbb{I}(f(\mathbf{s}_i) < Z_n(\mathbf{s}_i), 1 \leq i \leq n) \Lambda(df). \quad (6)$$

If $\phi_{\mathbf{s}_{n+1}}^+ \in \mathcal{C}_{K_n}^-$, we need $\phi_{\mathbf{s}_{n+1}}^+(\mathbf{s}_{n+1}) > Z_n(\mathbf{s}_{n+1})$, such that $\phi_{\mathbf{s}_{n+1}}^+$ is the extremal function at \mathbf{s}_{n+1} . Therefore, we can further restrict the Poisson process with intensity in (6) and consider the Poisson process $\text{PPP}(\tilde{\Lambda}) = \mathcal{C}_{K_n}^- \cap \{f \in \mathcal{C}_0 : f(\mathbf{s}_{n+1}) > Z_n(\mathbf{s}_{n+1})\}$, which has the mean measure $\tilde{\Lambda}$ given in Equation (5). If there is no point in $\text{PPP}(\tilde{\Lambda})$, i.e. $\phi_{\mathbf{s}_{n+1}}^+ \notin \mathcal{C}_{K_n}^- \cap \{f \in \mathcal{C}_0 : f(\mathbf{s}_{n+1}) > Z_n(\mathbf{s}_{n+1})\}$, it must be true that $\phi_{\mathbf{s}_{n+1}}^+ \in \mathcal{C}_{K_n}^+$ and $\phi_{\mathbf{s}_{n+1}}^+ = \arg \max_{\phi \in \{\phi_{\mathbf{s}_1}^+, \dots, \phi_{\mathbf{s}_n}^+\}} \phi(\mathbf{s}_{n+1})$. Otherwise, $\phi_{\mathbf{s}_{n+1}}^+(\mathbf{s}_{n+1})$ is a point of $\text{PPP}(\tilde{\Lambda}) \neq \emptyset$, and it can be identified according to the procedure outlined above: we simulate from the unrestricted Poisson process and reject the points violating the restriction in Equation (5) until either the restriction is satisfied or the simulated value falls below the lower bound $Z_n(\mathbf{s}_{n+1})$. \square

The pseudo code of an exact simulation algorithm based on Theorem 3 is given in Algorithm 1. Although the main structure of the algorithm resembles that proposed by Dombry et al. (2016) for max-stable processes, the key difficulty for general max-id processes resides in efficiently simulating from the distribution $P_{\mathbf{s}}(z, \cdot)$ given $Z(\mathbf{s}) = z$. For suitable constructions of the PPP $\{\eta_i; i = 1, 2, \dots\}$, we can investigate the specific structure of this conditional distribution. In this paper, we consider two different and very general structures for η_i , which are inspired from the construction of max-stable processes, yet they give rise to max-id models with much more flexibility in their joint tail decay rates. These two types of models, which encompass all of the commonly used max-stable models from the spatial extremes literature, are described in the following two subsections.

Algorithm 1 Exact simulation of a max-id process Z at locations $K_N = \{\mathbf{s}_1, \dots, \mathbf{s}_N\}$.

Require: Stationary marginal distribution G_0 of Z .

Require: Dependence characterization of Z allowing simulation of $P_{\mathbf{s}_n}$, $n = 1, \dots, N$.

- 1: Simulate $E_0 \sim \text{Exp}(1)$ and set $z = G_0^{-1}(\exp(-E_0))$.
 - 2: Simulate $Y \sim P_{\mathbf{s}_1}(z, \cdot)$ over K_N , which yields $\mathbf{Y}(K_N) = \{Y(\mathbf{s}_1), \dots, Y(\mathbf{s}_N)\}^\top$.
 - 3: Set $\mathbf{Z}(K_N) = \mathbf{Y}(K_N)$.
 - 4: **for all** $n = 2, \dots, N$ **do**
 - 5: Simulate $E_0 \sim \text{Exp}(1)$ and set $z = G_0^{-1}(\exp(-E_0))$.
 - 6: **while** $\{z > Z(\mathbf{s}_n)\}$ **do**
 - 7: Simulate $Y \sim P_{\mathbf{s}_n}(z, \cdot)$ over K_N , which yields $\mathbf{Y}(K_N) = \{Y(\mathbf{s}_1), \dots, Y(\mathbf{s}_N)\}^\top$.
 - 8: **if** $Y(\mathbf{s}_i) < Z(\mathbf{s}_i)$ for all $i = 1, \dots, n-1$ **then**
 - 9: update simulated values as $\mathbf{Z}(K_N) = \max(\mathbf{Z}(K_N), \mathbf{Y}(K_N))$ (componentwise).
 - 10: **end if**
 - 11: Simulate $E_1 \sim \text{Exp}(1)$ and update E_0 and z by setting $E_0 = E_0 + E_1$ and $z = G_0^{-1}(\exp(-E_0))$.
 - 12: **end while**
 - 13: **end for**
- Ensure:** Return $\mathbf{Z}(K_N)$
-

3.2. Representation of η_i as a Gaussian scale mixture

3.2.1. Model description

One possible structure of the general max-id process is achieved by representing η_i as a *Gaussian scale mixture*, i.e., $\eta_i = R_i W_i$, where $\{R_i\}$ are the points of a Poisson point process on \mathbb{R}^+ with mean measure κ such that $\kappa([0, \infty)) = \infty$, and $\{W_i\}$ are independent standard Gaussian random fields with mean zero and

unit variance on \mathcal{S} , which are not necessarily identically distributed and may depend on the overall random “magnitude” R_i . This construction mimics the max-stable spectral construction of extremal- t processes (Opitz, 2013) arising for a specific choice of Λ in (1) with $\{R_i\}$ and $\{W_i\}$ being independent. To extend max-stable processes, Huser et al. (2021) proposed using a more flexible mean measure for $\{R_i\}$ and Zhong et al. (2022) further relaxed the independence assumption between $\{R_i\}$ and $\{W_i\}$.

Specifically, they proposed a Weibull-tailed mean measure κ for $\{R_i\}$, where

$$\kappa(r) := \kappa([r, \infty)) = r^{-\beta} \exp\{-\alpha(r^\beta - 1)/\beta\}, \quad r > 0, \quad (\alpha, \beta)^\top \in (0, \infty)^2. \quad (7)$$

With this specification, $\kappa(r) \rightarrow r^{-\alpha}$ as $\beta \downarrow 0$. Therefore, the model remains in the “neighborhood” of max-stable processes for which $\kappa(r) = r^{-\alpha}$. Recall that the extremal- t max-stable process with α -Fréchet margins can indeed be obtained by taking $\eta_i = R_i W_i$ in (1) with $\{R_i\}$ a Poisson point process on $(0, \infty)$ with mean measure $\kappa(r) = r^{-\alpha}$, and $\{W_i\}$, independent zero mean Gaussian processes, that are rescaled such that $E[\max\{0, W_i(\mathbf{s})\}] = 1$, and that are also independent of $\{R_i\}$. In other words, if $\{R_i\}$ and $\{W_i\}$ are independent of each other and W_i is Gaussian, then the max-id model constructed from (7) can be arbitrarily close to the max-stable extremal- t dependence structure by varying the parameter β . Given $\{R_i\}$, a natural choice is, thus, to specify the processes $\{W_i\}$ to be standard Gaussian random fields with some correlation function $\rho(\mathbf{s}_1, \mathbf{s}_2; R_i)$, where R_i modulates the dependence range of W_i in such a way that large magnitudes R_i create processes W_i with weaker dependence. This gives extra flexibility for capturing the joint tail decay rate of the resulting max-id process Z in the asymptotic independence setting, i.e., when $\lim_{u \rightarrow 1} \Pr[G_0\{Z(\mathbf{s}_1)\} > u \mid G_0\{Z(\mathbf{s}_2)\} > u] = 0$ with G_0 being the marginal distribution of Z . For example, Zhong et al. (2022) chose a non-stationary correlation $\rho(\mathbf{s}_1, \mathbf{s}_2; R_i)$, which may be specified as

$$\rho(\mathbf{s}_1, \mathbf{s}_2; R_i) = (\lambda_{\mathbf{s}_1} \lambda_{\mathbf{s}_2}) \left(\frac{\lambda_{\mathbf{s}_1}^2 + \lambda_{\mathbf{s}_2}^2}{2} \right)^{-1} \exp \left\{ - \left(\frac{\lambda_{\mathbf{s}_1}^2 + \lambda_{\mathbf{s}_2}^2}{2} \right)^{-1/2} (1 + R_i)^\nu \|\mathbf{s}_1 - \mathbf{s}_2\| \right\}, \quad (8)$$

where $\lambda_{\mathbf{s}} > 0$ is a spatially-varying range parameter surface that can be associated with spatial covariates, and $\nu \geq 0$ is a parameter that controls how the scaling variable R_i influences the dependence range of W_i . In particular, when $\nu = 0$, R_i and W_i are independent. In the simulation examples below, we will illustrate some specific structures for $\lambda_{\mathbf{s}}$. When the value of R_i is large and $\nu > 0$, the dependence within the Gaussian random field W_i indeed decreases with increasing “magnitude” R_i . In this model, the mean measure of $\{\eta_i; i = 1, 2, \dots\}$ at the finite set of locations in $K \equiv K_N$ may be written as

$$\Lambda_K([\mathbf{0}, \mathbf{z}_K]^C) = \int_0^\infty \{1 - \Phi(\mathbf{z}_K/r; \boldsymbol{\rho}(K, K; r))\} \kappa(dr), \quad \mathbf{z}_K > \mathbf{0}, \quad (9)$$

where $\Phi(\cdot; \boldsymbol{\rho}(K, K; r))$ is the multivariate normal distribution function with zero mean and correlation matrix $\boldsymbol{\rho}(K, K; r)$ determined by the correlation function $\rho(\mathbf{s}_1, \mathbf{s}_2; r)$ for $\mathbf{s}_1, \mathbf{s}_2 \in K$, and $[\mathbf{0}, \mathbf{z}_K]^C \equiv ([0, z_1] \times \dots \times [0, z_N])^C$. Huser et al. (2021) showed that the marginal mean measure $\Lambda_{\mathbf{s}_0}$ is an infinite continuous measure and locally finite, thus it satisfies the assumptions in (3) stated in §2. The intensity function $g_{\Lambda_K}(\mathbf{z}_K)$ of Λ_K can be written as follows, where we use notation φ for the multivariate Gaussian density corresponding to Φ in (9):

$$g_{\Lambda_K}(\mathbf{z}_K) = \int_0^\infty r^{-N} \varphi(\mathbf{z}_K/r; \boldsymbol{\rho}(K, K; r)) \kappa(dr), \quad \mathbf{z}_K > \mathbf{0}. \quad (10)$$

For marginal intensities arising for singletons $K = \{\mathbf{s}\}$, we write $g_{\Lambda_{\mathbf{s}}}$. The density function of the probability measure $P_{\mathbf{s}_1}(z, \cdot)$ is then given by the conditional intensity

$$\mathbf{z}_{K \setminus \{\mathbf{s}_1\}} \mapsto g_{\Lambda_K}(z, \mathbf{z}_{K \setminus \{\mathbf{s}_1\}}) / g_{\mathbf{s}_1}(z) \propto g_{\Lambda_K}(z, \mathbf{z}_{K \setminus \{\mathbf{s}_1\}}). \quad (11)$$

To illustrate the tail dependence properties of this max-id process, we use the level-dependent extremal coefficient (Padoan, 2013; Huser et al., 2021). For a collection of sites $K_N \subset \mathcal{S}$ and $|K_N| = N$, the level-dependent extremal coefficient is defined as

$$\theta_N(z_0) = \frac{\log\{G(z_0)\}}{\log\{G_0(z_0)\}}, \quad (12)$$




Figure 1: Simulations (top row) and bivariate level-dependent extremal coefficients (bottom row) for models based on the Gaussian scale mixture construction. The parameters are chosen as $\nu = 0, 0.25, 0.5, 1$ (black, red, green, blue curves), where ν controls the dependence between R_i and W_i by using correlation function $\rho(\mathbf{s}_1, \mathbf{s}_2; R_i) = \exp\{-2(1 + R_i)^\nu \|\mathbf{s}_1 - \mathbf{s}_2\|\}$ conditional on R_i . The level z is shown on a logarithmic scale.

where $\mathbf{z}_0 = (z_0, \dots, z_0)^T \in \mathbb{R}^N$, G_0 is the marginal distribution of Z and G is the joint distribution of the random vector $\mathbf{Z}(K_N)$ with observations at the locations in the set $K \subset \mathcal{S}$. This coefficient can be expressed in terms of the measure Λ , and in the case where the margins are unit Fréchet it is specifically equal to $\theta_N(\mathbf{z}_0) = z_0 \Lambda_{K_N} \{(-\infty, \infty)^N \setminus (-\infty, z_0]^N\}$. From (12) it is easy to see that $\Pr\{\mathbf{Z}(K_N) \leq \mathbf{z}_0\} = G_0(z_0)^{\theta_N(\mathbf{z}_0)}$, such that this coefficient can be interpreted as the effective number of independent variables among the components of the vector $\mathbf{Z}(K_N)$ at level z_0 . Higher values of $\theta_N(\mathbf{z}_0)$ imply weaker dependence at quantile level z_0 . This dependence measure is widely used in extreme-value studies (mainly with max-stable models), where modeling the tail dependence structure is the main interest. As an example, we consider the correlation function of W_i conditional on the event magnitude R_i chosen as $\rho(\mathbf{s}_1, \mathbf{s}_2; R_i) = \exp\{-2(1 + R_i)^\nu \|\mathbf{s}_1 - \mathbf{s}_2\|\}$, which corresponds to the stationary and isotropic counterpart of (8) with spatially-constant $\lambda_{\mathbf{s}} = 1/2$. Figure 1 shows the realizations of $\{Z(s)\}_{s \in (0,1)}$ on a standard Gumbel scale generated using our proposed exact simulation algorithm, as well as the bivariate level-dependent extremal coefficients between two sites at distance $h = 0.5$ for the parameters $\alpha = 1$, $\beta = 0, 0.5, 1, 2$ and $\nu = 0, 0.25, 0.5, 1$. When $\beta = 0$ (interpreted here as $\beta \downarrow 0$), the extremal dependence persists with increasing level z . The curve of extremal dependence becomes a horizontal line when we ν and β are both zero. In this case, the max-id process is the max-stable extremal- t process. When β is different from zero, we get asymptotic independence (since $\theta_2(z_0) \rightarrow 2$ when $z \rightarrow \infty$), and the dependence level decreases as the values of β and ν increase. The realizations were generated using the same random seed, and we can discern stronger fluctuations of the realizations as the value of β increases, corresponding to a decreasing dependence strength.

3.2.2. Simulation algorithm

The generic Algorithm 1 based on extremal functions can be applied, and we now specifically describe how to sample from $P_{\mathbf{s}_1}(z, \cdot)$ at a finite set of locations K (need in lines 2 and 7 of Algorithm 1) for the max-id model constructed from Gaussian scale mixtures $\eta_i = R_i W_i$ as described in §3.2.1. Intuitively

speaking, we first sample a random variable \tilde{R}_z whose probability distribution corresponds to the intensity of R_i conditional on $\eta_i(\mathbf{s}_1) = R_i W_i(\mathbf{s}_1) = z$, and in the second step we sample the Gaussian field W_i at locations in K conditional on $\eta_i(\mathbf{s}_1) = z$ and \tilde{R}_z , i.e., we sample a Gaussian random vector conditional on its value z/\tilde{R}_z at \mathbf{s}_1 .

To obtain the distribution of \tilde{R}_z in the first step, we consider the joint intensity function of points $\{R_i\}$ and $\{W_i(\mathbf{s}_1)\}$ given by $(r, w) \mapsto \kappa(dr)\varphi(w)$. Next, we perform a change of variables from (r, w) to $(r, z) = (r, rw)$, which yields the intensity $g_{R,\eta}(r, z) = \kappa(dr)\varphi(z/r)/r$. Finally, we obtain the desired conditional density as $g_{\tilde{R}_z}(r; z) = g_{R,\eta}(r, z)/g_{\Lambda_{\mathbf{s}_1}}(z)$. Note that this defines indeed a valid probability density function since it corresponds to a conditional distribution of $P_{\mathbf{s}_1}(z, \cdot)$, which has been shown to be a proper probability measure.

In the second step, conditional on \tilde{R}_z , we write $\tilde{\eta}(K)$ for the conditional Gaussian random vector with mean vector $\boldsymbol{\rho}(\mathbf{s}_1, K; \tilde{R}_z)z$, common variance \tilde{R}_z^2 , and correlation matrix $\boldsymbol{\rho}(K, K; \tilde{R}_z) - \boldsymbol{\rho}(\mathbf{s}_1, K; \tilde{R}_z)\boldsymbol{\rho}(\mathbf{s}_1, K; \tilde{R}_z)^\top$, where $\boldsymbol{\rho}(\mathbf{s}_1, K; \tilde{R}_z)$ denotes the column vector of correlation values $\rho(\mathbf{s}_1, \mathbf{s}_j; \tilde{R}_z)$, $j = 1, \dots, N$, and $\boldsymbol{\rho}(K, K; \tilde{R}_z)$ is a matrix with entries $\rho(\mathbf{s}_i, \mathbf{s}_j; \tilde{R}_z)$, $i = 1, \dots, N$, $j = 1, \dots, N$. We denote this joint Gaussian density function as $g_{\tilde{\eta}(K)|\tilde{R}_z}$. Note that, after integrating out \tilde{R}_z , it is easy to formally show that the (unconditional) distribution of $\tilde{\eta}(K)$ is

$$\tilde{\eta}(K) \sim P_{\mathbf{s}_1}(z, \cdot).$$

Simulation from $P_{\mathbf{s}_1}(z, \cdot)$ at locations in K can thus be performed using a two-step procedure, where we need to simulate \tilde{R}_z first, and conditional on \tilde{R}_z , we then simulate $\tilde{\eta}(K)$ from a multivariate Gaussian distribution. We here explore the use of two simulation techniques where $g_{\tilde{R}_z}(r; z)$ only needs to be known up to a constant that depends on the value of z . Precisely, simulation from $P_{\mathbf{s}_1}(z, \cdot)$ proceeds according to the following successive steps:

- (1) Simulate a value \tilde{R}_z from the density function $g_{\tilde{R}_z}(r; z)$, $r > 0$, using either
 - the Metropolis–Hastings algorithm (e.g., [Robert and Casella, 2013](#))
 - the adaptive rejection sampling method proposed in [Gilks and Wild \(1992\)](#).
- (2) Given \tilde{R}_z , simulate a random vector $\tilde{\eta}(K)$ from the multivariate Gaussian density function $g_{\tilde{\eta}(K)|\tilde{R}_z}$.
- (3) Return $\tilde{\eta}(K)$.

Notice that the second step corresponds to simulating from a conditional multivariate normal distribution and can easily be performed exactly (e.g., [Varadhan, 2015](#)) in dimension N up to several thousands. Therefore, the key difficulty for an exact simulation procedure lies in the first step, where we have to simulate from the probability density $g_{\tilde{R}_z}$, known up to a constant only.

3.2.3. Options for simulating \tilde{R}_z

With the Metropolis–Hastings (MH) method in Step 1 of the above algorithm in §3.2.2, we use the following symmetric random walk proposal on a logarithmic scale:

$$\log(\tilde{R}_z^j) \mid \tilde{R}_z^{j-1} \sim \mathcal{N}(\log(\tilde{R}_z^{j-1}), \sigma^2), \quad j = 1, 2, \dots,$$

where σ^2 is the proposal variance and has to be tuned to achieve good mixing, and the initial value \tilde{R}_z^0 is either fixed by the user or drawn from some distribution that we can easily simulate from. MH is usually approximate (thus not exact) when the distribution of the initial value is different from the target distribution (which is always the case in practice), but its error becomes negligible when a sufficient number of iterations are performed. Here, the MH algorithm is performed on a single variable \tilde{R}_z at a time, so each Markov chain Monte Carlo (MCMC) chain involved in Algorithm 1 mixes well and converges very quickly to its stationary distribution, while iterations are of very small computational cost. Our experience shows that very few iterations are in fact needed in practice to get accurate results. We therefore regard Algorithm 1 combined with MH for simulating \tilde{R}_z as a general “quasi-exact” simulation method. [Propp and Wilson](#)




Figure 2: An illustration of the adaptive rejection sampling scheme, where the upper hull function $u(r)$ is tangent to the logarithmic target density $h(r)$ at locations $\mathbf{r} = \{r_1, r_2, r_3, r_4\}$.

(1996) further showed how to perform exact MCMC-based sampling by coupling Markov chains, but for simplicity, we shall here only consider the standard MCMC procedure outlined above, since it is easy to implement and already able to deliver very accurate simulation samples.

As an alternative to MH, we also consider adaptive rejection sampling (ARS), which is an exact simulation algorithm that is often used in combination with the Gibbs sampler. It can generate exact independent samples from the target distribution, here the density $g_{\tilde{R}_z}(r; z)$, $r > 0$ with fixed z . The only requirement for using this method is that $h(r) = \log g_{\tilde{R}_z}(r; z)$, $r > 0$ is concave in r with a connected support, and that it is continuous and differentiable everywhere within the support. In particular, for the previously presented parametric family of Gaussian scale mixture models based on (7), it is straightforward to show that $h(r)$ is indeed a continuous, differentiable and concave function over its support \mathbb{R} . The ARS algorithm adaptively constructs a piecewise upper hull function $u(r)$ and a piecewise lower bounding function $l(r)$ of $h(r)$, respectively, and then uses the distribution $s(r) = \exp\{u(r)\} / \int \exp\{u(r')\} dr'$ to generate proposals in the exact rejection sampling scheme. Figure 2 illustrates the adaptive rejection sampling scheme, where the upper hull function $u(r)$ is tangent to the target $h(r)$ at the points $\mathbf{r} = \{r_1, r_2, r_3, r_4\}$. Therefore, $u(r)$ and $l(r)$ are uniquely determined by the points in \mathbf{r} and the target function $g_{\tilde{R}_z}(r; z)$. As we increase the number of points in \mathbf{r} , the functions $u(r)$ and $l(r)$ will form an increasingly narrow envelope around the function $h(r)$, and the acceptance rate of the algorithm will increase. The steps of the ARS algorithm are as follows:

- (0) Initialization: choose a set of anchor points \mathbf{r} and construct the upper and lower bounds $u(r)$ and $l(r)$, respectively, as illustrated in Figure 2.
- (1) Sampling: sample $\tilde{R}_z^* \sim s(r)$, and sample $U \sim \text{Unif}(0, 1)$ independently of \tilde{R}_z^* .
- (2) First acceptance test: if $U \leq \exp\{l(\tilde{R}_z^*) - u(\tilde{R}_z^*)\}$ then accept \tilde{R}_z^* and go to 5.
- (3) Second acceptance test: If $U \leq \exp\{h(\tilde{R}_z^*) - u(\tilde{R}_z^*)\}$ then accept \tilde{R}_z^* and go to 5.

- (4) If the first and second acceptance tests both fail, reject \tilde{R}_z^* ; add \tilde{R}_z^* to list of anchor points and update the bounds $u(r)$ and $l(r)$. Go to 1.
- (5) Return \tilde{R}_z^* .

The distribution of the simulated value given acceptance is thus

$$\begin{aligned} \Pr(\tilde{R}_z^* \in A \mid \text{acceptance of } \tilde{R}_z^*) &= \frac{\int_A s(r) \left\{ \int_0^{\exp\{l(r)-u(r)\}} dw + \int_{\exp\{l(r)-u(r)\}}^{\exp\{h(r)-u(r)\}} dw \right\} dr}{\int_{\mathbb{R}} s(r) \left\{ \int_0^{\exp\{l(r)-u(r)\}} dw + \int_{\exp\{l(r)-u(r)\}}^{\exp\{h(r)-u(r)\}} dw \right\} dr} \\ &= \frac{\int_A \exp\{h(r)\} dr}{\int_{\mathbb{R}} \exp\{h(r)\} dr}, \end{aligned}$$

which corresponds exactly to the target distribution with density $g_{\tilde{R}_z}(r; z)$, $r > 0$. By exploiting the concavity property of h , ARS usually only needs a relatively small number of iterations before acceptance, and it may therefore be faster and more accurate than MH, as we will illustrate through the simulation experiments in §4.

3.3. Representation of η_i as a Gaussian location mixture

3.3.1. Model description

Another flexible and general max-id model family is the class of *Gaussian location mixtures*, where η_i has the form $\eta_i = R_i + W_i$. Here, the location variables $\{R_i\}$ are the points of a Poisson point process over \mathbb{R} with mean measure $\bar{\kappa}$, and $\{W_i\}$ are independent standard Gaussian random fields with mean zero and unit variance. Again, the random variables $\{R_i\}$ can be interpreted as overall event magnitudes, while the random processes $\{W_i\}$ vary over \mathcal{S} and can be viewed as spatial “profiles”. For a general and flexible parametric form of the distribution of event magnitudes $\{R_i\}$, we propose to define $\bar{\kappa}$ as

$$\bar{\kappa}(r) = \exp[-\alpha \{\mathbb{I}(r > 0)|r|^{\beta_1} - \mathbb{I}(r < 0)|r|^{\beta_2}\}], \quad r \in \mathbb{R}, \quad (\alpha, \beta_1, \beta_2)^\top \in (0, \infty)^2 \times (0, 2), \quad (13)$$

where $\bar{\kappa}(r)$ is the short-hand notation for $\bar{\kappa}([r, \infty))$. The parameter restrictions for α , β_1 and β_2 ensure that the mean measure Λ of the points $\{\eta_i\}$ is indeed Radon (i.e., locally finite) as required to defined a valid Poisson point process. Max-id models based on Gaussian location mixtures have never been explored in the literature, and the new model (13) is especially interesting in view of the flexible dependence structure that it yields and its close connection to the max-stable Brown–Resnick process (Kabluchko et al., 2009). The parameter β_1 controls the dependence strength in the upper tail, while β_2 controls the dependence strength in the lower tail. As in §3.2.1, we may further allow for dependence between the point R_i and the Gaussian field W_i for each i . For a general and flexible correlation function of W_i , we may again use the correlation function in (8) to drive the range of the Gaussian fields W_i through the location variable R_i , though other options are of course possible. In this case, the mean measure of $\{\eta_i; i = 1, 2, \dots\}$ at the finite set of locations $K \equiv K_N$ is given by

$$\Lambda_K([0, \mathbf{z}_K]^C) = \int_{\mathbb{R}} \{1 - \Phi(\mathbf{z}_K - r; \boldsymbol{\rho}(K, K; r))\} \bar{\kappa}(dr),$$

where $\Phi(\cdot; \boldsymbol{\rho}(K, K; r))$ and $\boldsymbol{\rho}(K, K; r)$ are defined as before. The expressions for unconditional and conditional intensity functions can be derived by analogy with the scale mixture case. When $\alpha = \beta_1 = \beta_2 = 1$ and $\{R_i\}$ and $\{W_i\}$ are independent of each other, the proposed model reduces to the popular Brown–Resnick max-stable model (Kabluchko et al., 2009; Wadsworth and Tawn, 2014), although this model is often defined through intrinsically stationary Gaussian processes $\{W_i\}$ with some variogram function $\gamma(h)$ rather than through a correlation function $\rho(\mathbf{s}_1, \mathbf{s}_2)$.

3.3.2. Simulation algorithm

By analogy with the approach developed for Gaussian scale mixtures, we can here again use Algorithm 1 for the max-id location mixture process. Moreover, to sample from $P_{\mathbf{s}_1}(z, \cdot)$ at a finite set of locations, we propose a similar simulation algorithm, where the first step consists in sampling a random variable \tilde{R}_z from the probability density $g_{\tilde{R}_z}(r; z) \propto \varphi(z - r)d\bar{\kappa}(dr)$, $r > 0$, and the second step is to sample a Gaussian random vector $\tilde{\eta}(K)$, conditional on \tilde{R}_z , with mean $\tilde{R}_z + \boldsymbol{\rho}(\mathbf{s}_1, K; \tilde{R}_z)(z - \tilde{R}_z)$, common unit variance, and correlation matrix $\boldsymbol{\rho}(K, K; \tilde{R}_z) - \boldsymbol{\rho}(\mathbf{s}_1, K; \tilde{R}_z)\boldsymbol{\rho}(\mathbf{s}_1, K; \tilde{R}_z)^\top$, where $\boldsymbol{\rho}(\mathbf{s}_1, K; \tilde{R}_z)$ denotes the column vector of correlation values $\rho(\mathbf{s}_1, \mathbf{s}_j; \tilde{R}_z)$, $j = 1, \dots, N$, and $\boldsymbol{\rho}(K, K; \tilde{R}_z)$ is the corresponding correlation matrix. As above, we denote the density of $\tilde{\eta}(K)$ conditional on \tilde{R}_z as $g_{\tilde{\eta}(K)|\tilde{R}_z}$. Again, the resulting (unconditional) distribution of $\tilde{\eta}(K)$ can be shown to be $P_{\mathbf{s}_1}(z, \cdot)$ at locations in K . However, here we observe that the density $g_{\tilde{R}_z}(r; z)$ is not logarithmically concave, and therefore we cannot use the adaptive resampling scheme to generate extremal functions at a finite set of locations in K from $P_{\mathbf{s}_1}(z, \cdot)$, but we can still use the quasi-exact Metropolis–Hastings algorithm, which provides high accuracy at low computational cost. Thus, our proposed procedure for simulating max-id processes of Gaussian location mixture type at a collection of sites K relies on Algorithm 1 combined with the following successive steps to sample from $P_{\mathbf{s}_1}(z, \cdot)$:

- (1) Simulate a value \tilde{R}_z according to the density $g_{\tilde{R}_z}(r; z)$ by using the Metropolis–Hastings method with random walk proposal distribution $\tilde{R}_z^j | \tilde{R}_z^{j-1} \sim \mathcal{N}(\tilde{R}_z^{j-1}, \sigma^2)$, $j = 1, 2, \dots$, and initial point \tilde{R}_z^0 .
- (2) Given \tilde{R}_z , simulate a vector $\tilde{\eta}(K)$ according to the multivariate Gaussian density function $g_{\tilde{\eta}(K)|\tilde{R}_z}$.
- (3) Return $\tilde{\eta}(K)$.

As in the Gaussian scale mixture case, the second step consists in simulating from a conditional multivariate normal distribution, which is fast and straightforward. Figure 3 shows examples of realizations of $\{Z(\mathbf{s})\}_{\mathbf{s} \in (0,1)}$ with standard Gumbel marginal distributions (generated using the proposed quasi-exact algorithm), as well as the bivariate level-dependent extremal coefficients at distance $h = 0.5$ for the parameters $\beta_1 = 0.5, 1, 1.5$, $\beta_2 = 0.5, 1, 1.5$ and $\alpha = 1$. The correlation function for W_i is here chosen to be exponential as $\rho(h) = \exp(-2h)$. We see that when both β_1 and β_2 are equal 1, the curve of the extremal coefficients is a horizontal line. In this case, the max-id process is the Brown–Resnick max-stable process. As the values for β_1 and β_2 increase, the dependence levels in the lower tail and the upper tail, which are controlled by β_2 and β_1 , respectively, decrease. When $\beta_1 > 1$, the extremal coefficient increases as the level $z \rightarrow \infty$. By contrast, when $\beta_1 < 1$, the extremal coefficient decreases as the level $z \rightarrow \infty$. Similar results are observed for the dependence level in the lower tail with β_2 . Therefore, this model shows great flexibility through its ability to capture asymptotic dependence and asymptotic independence in the upper tail and different extremal dependence structures in the lower tail. In these realizations generated with the same random seed, we discern that the fluctuations are increasing as the dependence level decreases, as is expected.

4. Simulation Study

In this section, we use the Gaussian scale mixture representation described in §3.2 as an example to illustrate Algorithm 1 and to assess the performance of the two algorithmic choices detailed in §3.2.2 for sampling from $P_{\mathbf{s}_0}(z, \cdot)$, i.e., of the Metropolis–Hastings scheme and the adaptive rejection sampling scheme. We also provide a comparison with respect to a naive simulation approach inspired from Schlather (2002), which consists in simulating the magnitudes R_i in descending order and truncating the maximum in (1) as $Z(\mathbf{s}) \approx \max_{i=1, \dots, n} \eta_i(\mathbf{s})$ with $\eta_i = R_i W_i(\mathbf{s})$ for a sufficiently large integer n .

We perform the simulation on a 7×7 regular grid in $(0, 1)^2$, corresponding to 49 locations in K . We fix $\alpha = 5, \beta = 2$ in (7) and $\nu = 3, \lambda_{\mathbf{s}} = \exp[2 - 0.5\Phi\{(s_x - 0.5)/0.25\}]$ in (8), where $\mathbf{s} = (s_x, s_y)^\top$ and $\Phi(\cdot)$ is the standard normal distribution function. We implemented three algorithms using the R software: the approximate naive sampling method (APPROX), Algorithm 1 combined with the quasi-exact random walk Metropolis–Hastings (MH), and Algorithm 1 combined with the exact adaptive rejection sampling (ARS). Our aim is to compare the simulation accuracy with respect to univariate distributions and to dependence




Figure 3: Simulations (top row) and bivariate level-dependent extremal coefficients (bottom row) for models based on the Gaussian location mixture construction. Parameters are chosen as $\alpha = 1$ and $\beta_1 = 0.5, 1, 1.5$ (blue, black, red). The correlation function for W_i is exponential and given by $\rho(h) = \exp(-2h)$. The level z is shown in logarithmic scale.

properties. We tuned the settings of these three algorithms such that they require approximately the same amount of running time to generate $N = 10^5$ independent samples of the max-id process on the grid. These settings are given as follows:

- (1) Naive method: Simulate $\eta_i = R_i W_i$, $i = 1, 2, \dots, n$ on the grid K , where $n = 100$.
- (2) Metropolis–Hastings scheme: In the MCMC, run 100 iterations to generate one sample \tilde{R}_z from the density function $g_{\tilde{R}_z}(r; z)$.
- (3) Adaptive rejection sampling scheme: Use 6 points to initialize the upper hull function and lower hull function.

Based on the generated samples, we compute the empirical Kullback–Leibler (KL) divergence $D_{s_0}(p||q)$ (Kullback and Leibler, 1951) in order to compare the accuracy of simulated samples at each site $s_0 \in K \subset \mathcal{S}$. Precisely, we compute

$$D_{s_0}(p||q) = \sum_{i=1}^N \hat{p}\{Z_i(s_0)\} \log \left[\frac{\hat{p}\{Z_i(s_0)\}}{q\{Z_i(s_0)\}} \right], \quad \hat{p}\{Z_i(s_0)\} = \frac{p\{Z_i(s_0)\}}{\sum_{i=1}^N p\{(Z_i(s_0))\}}, \quad (14)$$

where q places mass $1/N$ on each simulated sample as the empirical cdf, while p is the exact marginal density of Z . In practice, the samples are always treated as discrete random variables, and \hat{p} and q need to be probability mass functions instead of density functions. Therefore, we have 49 empirical KL divergences




Figure 4: Boxplots of empirical KL divergences for the 49 grid locations based on the simulated data sets generated by the naive method (APPROX), the Metropolis–Hastings scheme (MH) and the adaptive rejection sampling scheme (ARS).

in total, one for each point on the grid. Figure 4 shows the boxplot of the 49 empirical KL divergences using the simulated datasets generated by the three algorithms. Both the adaptive rejection sampling scheme and the Metropolis–Hastings scheme outperform the naive method by a significant margin. Moreover, the median of the empirical KL divergence for the adaptive rejection sampling scheme is notably lower than with the Metropolis–Hastings scheme. Generating a single max-id replication on the grid took approximately 1.2 seconds using a single-core machine.

To further evaluate the performance of each sampling scheme with respect to the dependence structure, we calculated the bivariate (level-dependent) empirical extremal coefficients $\hat{\theta}_2$. For each pair of sites, we compared the empirical extremal coefficient computed from the $N = 10^5$ simulated replicates with the theoretical counterpart computed by evaluating the mean measure Λ with a very accurate numerical unidimensional integral. Figure 5 shows the results for the three simulation techniques with the computed standard error, sd , of the approximation error shown in the title of each panel. The adaptive rejection sampling scheme appears to perform generally better than the naive scheme at all levels, especially at low quantile levels when the naive scheme is significantly biased. At high quantile levels ($q = 0.95$), all methods have a comparable performance. The pseudo-exact Metropolis–Hastings scheme appears to be as accurate as the adaptive rejection sampling method. Our results imply that if we draw statistical inferences based on simulated samples from the naive method, the strength of the dependence at relatively low quantile levels of extremes would be significantly overestimated in the current setting. Better results might be obtained for the naive method by increasing the number n of simulated fields $\{\eta_i\}$, but this comes with a computational cost. By contrast, our proposed algorithm always provides exact max-id simulations at a low computational cost, which may prove essential for simulation-based statistical inference.

As in [Dombry et al. \(2016\)](#), it can be shown that the expected number of random fields W_i to be simulated in the exact simulation algorithm equals to the number of simulation sites, i.e., the dimension $|K|$. To further illustrate the efficiency of our proposed algorithm, we simulated 10 replicas for different numbers of randomly sampled locations K on the unit square $(0, 1)^2$ using the ARS scheme and calculated the time and number of random fields W_i that needed to be simulated. The results are shown in Figure 6. The parameters are set to be $\alpha = 1, \beta = 1$ and the correlation function for W_i is chosen to be $\rho(h) = \exp(-h)$. Since we use the Cholesky decomposition for simulating the process W_i on K , whose complexity is $\mathcal{O}(|K|^3)$, the expected

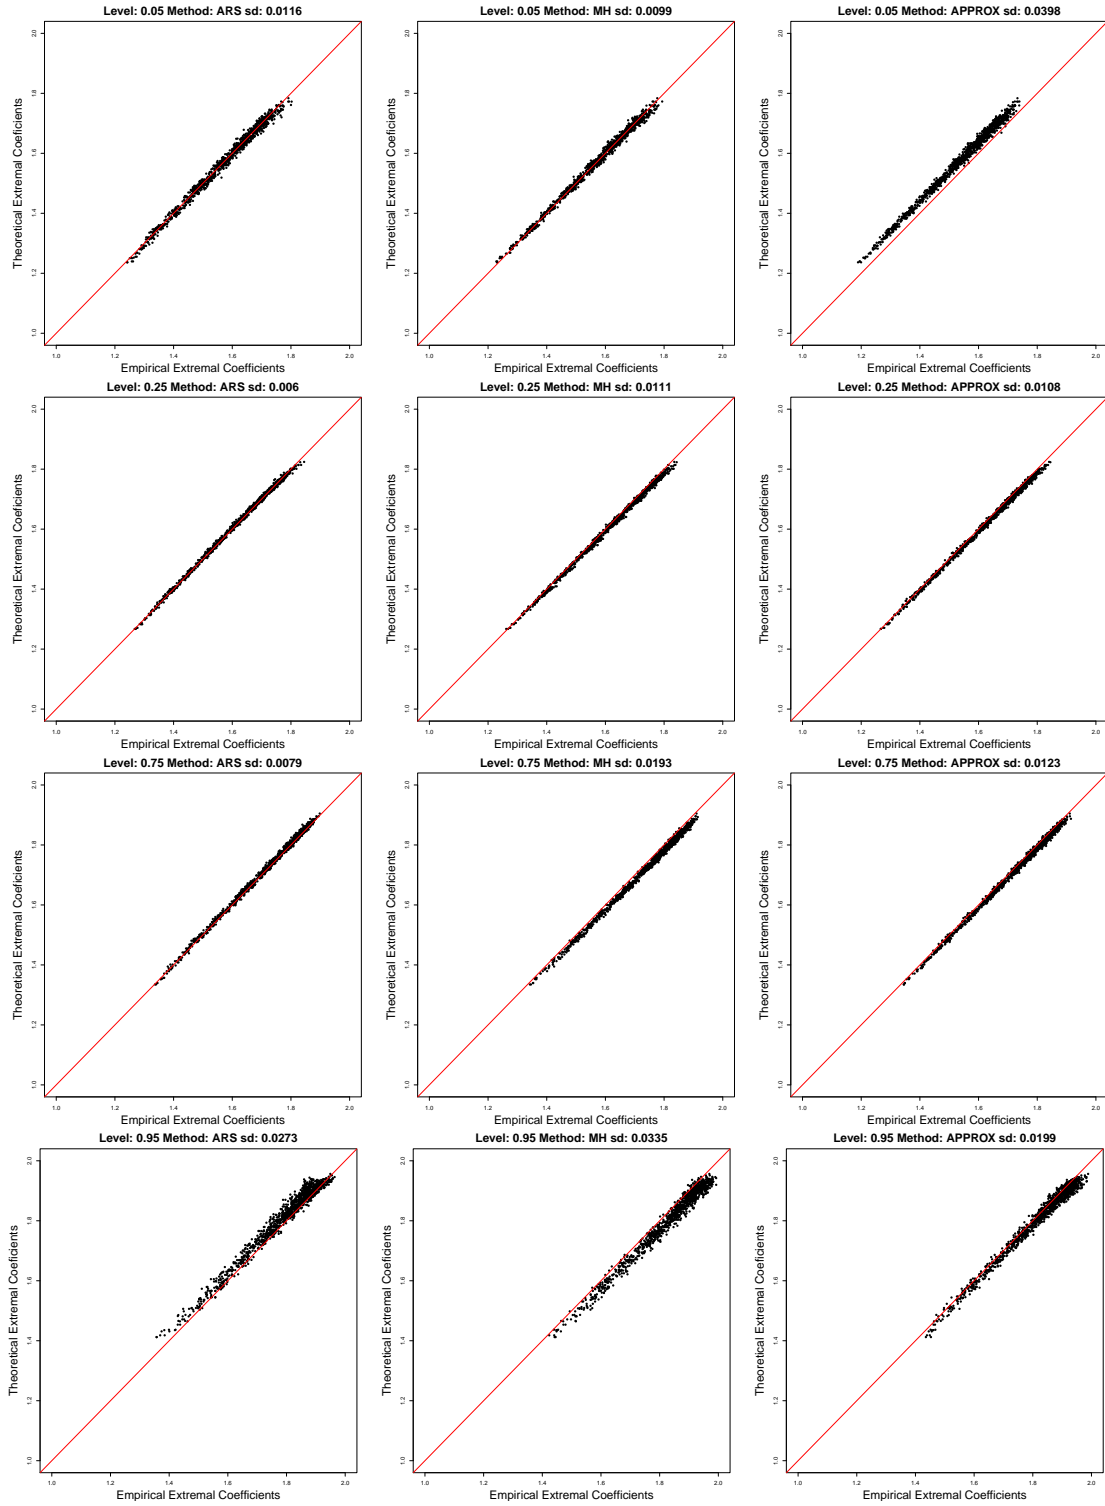


Figure 5: Bivariate scatterplots of theoretical against empirical bivariate level-dependent extremal coefficients at quantile levels 0.05, 0.25, 0.5, 0.75, 0.95 (from top to bottom) for all pairs of locations, based on the three sampling schemes (left: Algorithm 1 with adaptive rejection sampling (ARS), middle: Algorithm 1 with Metropolis–Hastings (MH), right: naive approximate method).




Figure 6: Left: Time in second needed to simulate one max-id replicate for different numbers of locations $|K|$. The red curve is the fit from a 4-th order polynomial. Right: Number of random fields W_i that were simulated for one max-id replicate for different numbers of locations $|K|$. The red line is the diagonal. The parameters are here set to be $\alpha = 1, \beta = 1$ and the correlation function for W_i is chosen to be $\rho(h) = \exp(-h)$.

overall complexity for simulating a max-id replicate is $\mathcal{O}(|K|^4)$. The red curve on the left panel of Figure 6 illustrates this. Here, the computational time is about two minutes for one replicate in dimension $|K| = 900$, which is already quite fast. If further speeds-up are required, approximate algorithms for simulating the Gaussian processes W_i may be employed (e.g., based on Gaussian Markov random fields). As expected, the right panel of Figure 6 shows that the expected number of simulations of W_i that are needed is indeed equal to $|K|$. We also investigated how the ordering of locations $\mathbf{s}_1, \dots, \mathbf{s}_n$ chosen in Algorithm 1 to iteratively sample from $P_{\mathbf{s}_0}(z, \cdot)$ might affect the algorithm's efficiency. Our results (not shown) suggest that there was very little difference among the four orderings considered (coordinate-wise, random, middle-out, or maximum-minimum orderings), in terms of the mean and variance of the number of processes W_i to sample in each simulation. In all simulations presented here, we used a coordinate-wise ordering for simplicity.

5. Discussion

In this paper, we have generalized the exact simulation algorithm for max-stable processes based on extremal functions to the more general class of max-id processes. Whenever possible, we recommend using the adaptive rejection sampling scheme for the conditional simulation of \tilde{R}_z , in order to achieve fast and exact simulations in cases where variables to be simulated do not correspond to some standard distribution from which samples are easily drawn. Although the specific implementation of the algorithm still depends on the chosen representation for η_i , our approach is general and can be applied to very broad classes of max-id processes while preserving computational efficiency. Moreover, the specific examples showcased in this paper already have very flexible dependence structures. In a simulation study, we used three different sampling algorithms to compare their speed and accuracy. Our novel exact simulation algorithm using the adaptive rejection sampling scheme clearly outperforms the two alternatives. In cases where the target density in the adaptive rejection sampling scheme is not logarithmically concave, it is possible to use an appropriately defined quasi-exact Metropolis–Hastings scheme to achieve high accuracy. Exploring exact MCMC procedures based on coupling of Markov chains (Propp and Wilson, 1996) would also be an interesting future research direction.

Besides standard uses of such fast and exact simulation techniques in Monte–Carlo-based statistical inferences involving quantities that cannot be calculated analytically, another interesting perspective concerns Approximate Bayesian Computation (ABC) for general max-id processes, which could be implemented by

analogy with the approach developed for max-stable processes in [Erhardt and Smith \(2012\)](#). As described above, the computation of multivariate density and distribution functions is often numerically very expensive due to integrals that have to be evaluated numerically, such that likelihood-free inference techniques are promising alternatives. Fast simulation of large numbers of samples of the model is required in likelihood-free estimation methods such as ABC, and in other simulation-based inference methods such as maximum simulated likelihood (e.g., [Hajivassiliou, 2000](#)).

Acknowledgments

We thank the Editor, Associate Editor, and two anonymous reviewers for constructive comments that helped improve the paper. This publication is based upon work supported by the King Abdullah University of Science and Technology (KAUST) Office of Sponsored Research (OSR) under Awards No. OSR-CRG2017-3434 and No. OSR-CRG2020-4394.

References

- Bopp G, Shaby BA, Huser R. A hierarchical max-infinitely divisible spatial model for extreme precipitation. *Journal of American Statistical Association* 2021;116:93–106.
- Brown BM, Resnick SI. Extreme values of independent stochastic processes. *Journal of Applied Probability* 1977;14:732–9.
- Castruccio S, Huser R, Genton MG. High-order composite likelihood inference for max-stable distributions and processes. *Journal of Computational and Graphical Statistics* 2016;25:1212–129.
- Davison AC, Huser R, Thibaud E. Spatial extremes. In: Gelfand AE, Fuentes M, Hoeting JA, Smith RL, editors. *Handbook of Environmental and Ecological Statistics*. CRC Press; 2019. p. 711–44.
- Davison AC, Padoan S, Ribatet M. Statistical modelling of spatial extremes (with Discussion). *Statistical Science* 2012;27(2):161–86.
- Dey DK, Yan J. *Extreme Value Modeling and Risk Analysis: Methods and Applications*. CRC Press, 2016.
- Dieker AB, Mikosch T. Exact simulation of Brown-Resnick random fields at a finite number of locations. *Extremes* 2015;18:301–14.
- Dombry C, Engelke S, Oesting M. Exact simulation of max-stable processes. *Biometrika* 2016;103:303–17.
- Dombry C, Eyi-Minko F. Regular conditional distributions of continuous max-infinitely divisible random fields. *Electronic Journal of Probability* 2013;18:1–21.
- Erhardt RJ, Smith RL. Approximate Bayesian computing for spatial extremes. *Computational Statistics & Data Analysis* 2012;56:1468–81.
- Gilks WR, Wild P. Adaptive rejection sampling for Gibbs sampling. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 1992;41:337–48.
- Giné E, Hahn MG, Vatan P. Max-infinitely divisible and max-stable sample continuous processes. *Probability Theory and Related Fields* 1990;87:139–65.
- Gissibl N, Klüppelberg C, Otto M. Tail dependence of recursive max-linear models with regularly varying noise variables. *Econometrics and Statistics* 2018;6:149–67.
- de Haan L. A spectral representation for max-stable processes. *The Annals of Probability* 1984;12:1194–204.
- Hainy M, Müller WG, Wagner H. Likelihood-free simulation-based optimal design with an application to spatial extremes. *Stochastic Environmental Research and Risk Assessment* 2016;30:481–92.
- Hajivassiliou V. Some practical issues in maximum simulated likelihood. In: Mariano R, Weeks M, Schuermann T, editors. *Simulation-Based Inference in Econometrics: Methods and Applications*. Cambridge University Press; 2000. p. 71–99.
- Huser R, Davison AC. Space-time modelling of extreme events. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 2014;76:439–61.
- Huser R, Opitz T, Thibaud E. Max-infinitely divisible models and inference for spatial extremes. *Scandinavian Journal of Statistics* 2021;48:321–48.
- Huser R, Wadsworth JL. Advances in statistical modeling of spatial extremes. *WIREs Computational Statistics* 2022;14:e1537.
- Kabluchko Z, Schlather M, de Haan L. Stationary max-stable fields associated to negative definite functions. *Annals of Probability* 2009;37:2042–65.
- Kullback S, Leibler RA. On information and sufficiency. *Annals of Mathematical Statistics* 1951;22:79–86.
- Lee XJ, Hainy M, McKeone JP, Drovandi CC, Pettitt AN. ABC model selection for spatial extremes models applied to South Australian maximum temperature data. *Computational Statistics & Data Analysis* 2018;128:128–44.
- Oesting M, Schlather M, Zhou C. Exact and fast simulation of max-stable processes on a compact set using the normalized spectral representation. *Bernoulli* 2018;24:1497–530.
- Oesting M, Storkorb K. A comparative tour through the simulation algorithms for max-stable processes. *Statistical Science* 2022;37:42–63.
- Opitz T. Extremal t processes: Elliptical domain of attraction and a spectral representation. *Journal of Multivariate Analysis* 2013;122:409–13.
- Padoan SA. Extreme dependence models based on event magnitude. *Journal of Multivariate Analysis* 2013;122:1–19.

- Padoan SA, Ribatet M, Sisson SA. Likelihood-based inference for max-stable processes. *Journal of the American Statistical Association* 2010;105:263–77.
- Propp JG, Wilson DB. Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Structures & Algorithms* 1996;1–2:223–52.
- Resnick SI. *Extreme Values, Regular Variation and Point Processes*. Springer-Verlag New York, 1987.
- Robert C, Casella G. *Monte Carlo statistical methods*. Springer Science & Business Media, 2013.
- Schlather M. Models for stationary max-stable random fields. *Extremes* 2002;5:33–44.
- Varadhan R. condMVNorm: Conditional Multivariate Normal Distribution; 2015. URL: <http://CRAN.R-project.org/package=condMVNorm>; r package version 2015.2-1.
- Wadsworth JL, Tawn JA. Efficient inference for spatial extreme value processes associated to log-Gaussian random functions. *Biometrika* 2014;101:1–15.
- Zhong P, Huser R, Opitz T. Modeling non-stationary temperature maxima based on extremal dependence changing with event magnitude. *Annals of Applied Statistics* 2022;In press.