Journal of Hydrology 529 (2015) 1451–1463

Contents lists available at ScienceDirect

Journal of Hydrology

journal homepage: www.elsevier.com/locate/jhydrol

Bayesian estimation of rainfall intensity–duration–frequency relationships

H. Van de Vyver *

Royal Meteorological Institute of Belgium, Ringlaan 3 Avenue Circulaire, B-1180 Brussels, Belgium

ARTICLE INFO

Article history:
Received 1 April 2015
Received in revised form 4 August 2015
Accepted 17 August 2015
Available online 28 August 2015
This manuscript was handled by Andras Bardossy, Editor-in-Chief, with the assistance of Peter F. Rasmussen, Associate Editor

Keywords:
Rainfall
IDF curves
Extreme value distributions
Bayesian estimation
Composite likelihood

SUMMARY

Rainfall intensity–duration–frequency (IDF) curves are one of the most commonly used tools in water resources engineering. They give an idea of how return levels of extreme rainfall intensities vary with duration over a range of return periods. It is assumed that the annual maximum intensity follows the generalised extreme value (GEV) distribution. Conventional methods of estimating IDF relationships do not provide estimates of uncertainty. We propose a Bayesian framework for handling uncertainties in IDF models. Firstly, we collect annual maximum intensity data over a relevant range of rainfall durations. Secondly, we define an approximate likelihood, the "independence" likelihood, in which the correlations have been ignored between maximum intensity data of different durations. Finally, we apply Bayesian inference to obtain the adjusted posterior, which accounts for likelihood misspecification. A comparison with earlier methods, without any adjustment amongst others, shows that the adjusted posteriors are considerably wider.

1. Introduction

Intensity–duration–frequency (IDF) relationships of extreme rainfall intensities are one of the most commonly used tools in water resources engineering for planning, design and operation of water resources projects. The construction of IDF curves has a long history, going back at least to the paper of Bernard (1932). Numerous commonly used IDF models have been published in the hydrological literature; an historical overview and discussion has been included in García-Bartual and Schneider (2001). Koutsoyiannis et al. (1998) have provided a general formula for the IDF relationship, consistent with the theoretical probabilistic foundation of the analysis of rainfall maxima.

The modelling of extreme events is of increasing interest, and there currently exists powerful statistical theories of extremes (Beirlant et al., 2004; Coles, 2001). The theory has been widely used in hydrology (Katz et al., 2002; Schliep et al., 2010; Van de Vyver, 2015), and other environmental sciences and finance as well (Chavez-Demoulin et al., 2006; Van de Vyver, 2012). In the past decades, many authors working on IDF curves have modelled annual maximum rainfall events with the generalised extreme value (GEV) distribution or Gumbel distribution, see Demarée (1985), Mohymont et al. (2004), Muller et al. (2008), Overeem et al. (2008) and Lehmann et al. (2013), to mention a few. Traditionally, in IDF analysis one first fits a probability distribution to each series of maximum intensity for a specific duration. Next, an IDF relationship is fitted to a range of quantiles of these distributions.

A commonly overlooked feature when estimating IDF relationships is the assessment of uncertainty. In statistical analysis, uncertainty is generally as important as estimates themselves, especially in extreme value analysis since the observations are, by definition, scarce. However, the typical estimation procedure of IDF relationships suffers the drawback that uncertainty in parameter estimates cannot be identified. To the best of our knowledge, the only contributions to this subject are provided in Overeem et al. (2008), and in a Bayesian framework, in Muller et al. (2008) and Lehmann et al. (2013).

Bayesian estimation techniques have steadily gained ground, and are now recognised as a legitimate alternative to classical statistics. A Bayesian framework is natural for handling uncertainties because inferences on model parameters are presented explicitly as a posterior distribution. Prior to the Bayesian analysis of IDF relationships, Muller et al. (2008) and Lehmann et al. (2013) defined a likelihood for extreme rainfall intensity, which is based on the assumption of independence amongst annual maximum

* Tel.: +32 2 373 05 43; fax: +32 2 373 05 48.
E-mail address: hvijver@meteo.be

http://dx.doi.org/10.1016/j.jhydrol.2015.08.036
0022-1694/

intensities over different durations. The resulting likelihood can be thought of as the most simplistic estimation equation, and is a typical example of a composite likelihood (Varin et al., 2011).

In this paper, we adjust the independence likelihood for appropriate inference, an idea which has been based on recent advances in Bayesian inference from composite likelihoods (Pauli et al., 2011; Ribatet et al., 2012). Alternatively, Muller et al. (2008) introduced another likelihood for IDF relationships by explicitly modelling the dependence with a bivariate GEV-distribution. The inferences based on both IDF likelihoods will be compared.

The remainder of the paper is structured as follows. Section 2 introduces IDF relationships, and details the general IDF relationship. Then we give an illustration of commonly used classical estimators in Section 3. Next, Section 4 and 5 outline the statistical methodologies we employ and how inference is implemented. Technical details are described more fully in the Appendix. Section 6 summarises the results and the paper concludes with a few discussion points in Section 7.

2. Intensity–duration–frequency (IDF) relationships

Let $R_l(d)$ be the total amount of precipitation (mm) that falls in the time interval $[t_l - d, t_l]$ (time expressed in h). The average intensity is then $I_l(d) = R_l(d)/d$ (mm h^{-1}). Next, we form the series of annual maximum intensities

$$I(d) = \max\{I_1(d), I_2(d), \ldots, I_m(d)\}. \quad (1)$$

In practice, the construction of the series of maximum intensities is performed for a number of M durations $d_k, k = 1, \ldots, M$, starting from a minimum duration (e.g., from 5–10 min to 1 h), and ending with a maximum duration of interest. The key quantity in the IDF analysis is the return level of $I(d)$ with return period T, here denoted by $i_T(d)$, for a given duration $d \in [d_1, d_M]$. Earlier IDF relationships are empirical relationships amongst $i_T(d)$ and d, for a specific T-value (Chow et al., 1988; García-Bartual and Schneider, 2001). Our study is based on the general IDF relationship of Koutsoyiannis et al. (1998), which is a semi-empirical relationship expressing $i_T(d)$ as a function of both d and T. IDF curves are logarithmic plots of $i_T(d)$ against duration d, for different fixed T-values. Fig. 1 (solid lines) serves an example of IDF curves for Uccle (Belgium), for durations starting from 10 min, and ending with 72 h.

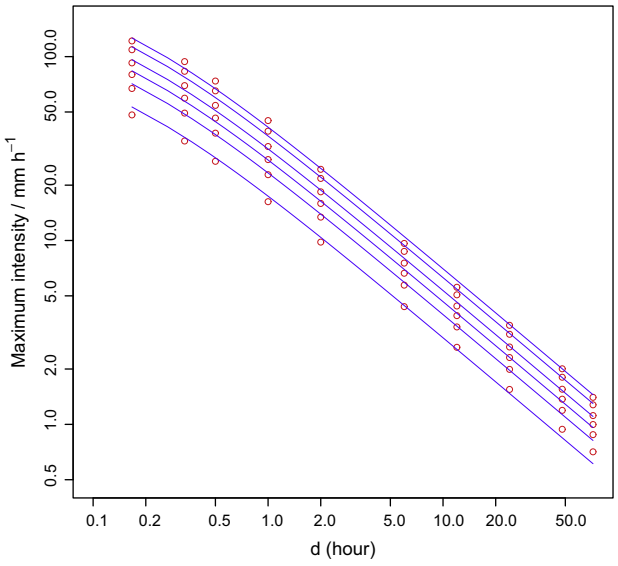

Fig. 1. IDF curves at Uccle (Belgium) for return periods $T = 2$, 5, 10, 20, 50 and 100 year.

2.1. General IDF relationship

For any rainfall duration d, we denote by $F_{I(d)}(i; d) = P\{I(d) \leqslant i\}$ the probability distribution of $I(d)$. The *return period* T associated with the *return level* $i_T(d)$ is computed with

$$T = \frac{1}{1 - F_{I(d)}(i_T; d)}. \quad (2)$$

Conversely, the *T-year return level*, $i_T(d)$, is defined as a value which, on average, is exceeded once in T years, and can be obtained by solving Eq. (2) for i_T. The *general IDF relationship* of Koutsoyiannis et al. (1998), is of the form

$$i_T(d) = \frac{a(T)}{b(d)}, \quad (3)$$

and has the simplicity of a separable functional dependence on T and d. Here, the function $b(d)$ is

$$b(d) = (d + \theta)^{\eta}, \quad \text{with} \quad \theta > 0,\ 0 < \eta < 1. \quad (4)$$

As the nominator $a(T)$ is independent of the rainfall duration d, the family of curves are parallel for different fixed T-values. The denominator $b(d)$ determines the shape of the curves: η is indicative of the slope of the straight part of the IDF curves, and θ is connected with the curvature change point. For the Uccle series, the general IDF relationship is, to a good approximation, valid for rainfall durations ranging between 10 min and 72 h. In Fig. 1, it can be seen that the curvature change point is around $d = 1$ h, and the IDF curves appear to be nearly straight lines in the region where $d \in [1, 72]$. We can thus suggest the following simplified IDF model when $d \in [1, 72]$:

$$b(d) = d^{\eta}, \quad \text{with } 0 < \eta < 1. \quad (5)$$

It will be confirmed later, in Section 6.2, that we may reject model Eq. (4) in favour of model Eq. (5) when modelling maxima intensity for these durations.

From Eq. (3), we can see that $a(T)$ is the T-year return level of the scaled maximum intensity $Y = I(d)\, b(d)$. We get

$$a(T) = F_Y^{-1}(1 - 1/T), \quad (6)$$

where $F_Y(y)$ is the probability distribution of Y. Obviously, $I(d)$ and Y follow the same family of distributions. The difference lies in the fact that the distribution parameters of $F_{I(d)}(i; d)$ are d-dependent, whereas those of $F_Y(y)$ are constant. Several distribution functions for $I(d)$ (for example Gumbel-, GEV-, Gamma-, log Pearson III distribution,...) have been included in the analysis of Koutsoyiannis et al. (1998).

2.2. General IDF relationship using the GEV distribution

Throughout this work, we assume that the annual maximum intensity, $I(d)$, follows to a good approximation the *generalised extreme value (GEV)* distribution. This is concerned with the statistical behaviour of block maxima (Beirlant et al., 2004; Coles, 2001),

$$M_m = \max\{X_1, \ldots, X_m\}, \quad (7)$$

where $X_1, \ldots, X_m$ is a sequence of independent and identically distributed (iid) random variables.

A key result is that, under regularity conditions, $\Pr\{M_m \leqslant z\}$ can be approximated by the GEV distribution for large m-values. This is a three-parameter family of functions

$$G(z; \mu, \sigma, \gamma) = \exp\left[-\left(1 + \gamma\,\frac{z - \mu}{\sigma}\right)^{-1/\gamma}\right], \quad (8)$$

where the location-parameter (μ) specifies the centre of the distribution; the scale-parameter (σ) determines the size of deviations around the location parameter; and the shape-parameter (γ) governs the tail behaviour of the distribution. We shortly denote $M_m \sim \text{GEV}[\mu, \sigma, \gamma]$. For sufficiently long time series, it is customary and convenient to extract the block maxima, generating a *block maxima-series*. Pragmatic considerations often lead to the adoption of one year blocks.

In the above notation, we have

$$I(d) \sim \text{GEV}[\mu(d), \sigma(d), \gamma]. \tag{9}$$

According to the analysis in Nadarajah et al. (1998), the shape parameter γ is modelled as a constant over d. More specifically, they proved that the next inequalities (which are trivially satisfied for rainfall depth $R(d)$):

$$\max_i\{R_i(d)\} \leqslant \max_i\{R_i(d')\} \leqslant (d'/d)\max_i\{R_i(d)\}, \quad \text{with } d' \text{ a multiple of } d,$$

imply that $\gamma_d = \gamma_{d'}$.

It follows that the scaled intensity $Y = I(d)\,b(d)$ obeys a GEV-distribution, i.e. $F_Y(y) = G(y; \mu, \sigma, \gamma)$, whereby $\mu = \mu(d)\,b(d)$, and $\sigma = \sigma(d)\,b(d)$. Putting everything together, the general IDF relationship Eqs. (3)–(6) becomes

- For $d \in [1/6, 72]$:

$$i_T(d) = \frac{\mu - \frac{\sigma}{\gamma}\left\{1 - \left[-\log\left(1 - \frac{1}{T}\right)\right]^{-\gamma}\right\}}{(d+\theta)^{\eta}}, \quad \text{with } \theta > 0,\ 0 < \eta < 1. \tag{10}$$

- For $d \in [1, 72]$:

$$i_T(d) = \frac{\mu - \frac{\sigma}{\gamma}\left\{1 - \left[-\log\left(1 - \frac{1}{T}\right)\right]^{-\gamma}\right\}}{d^{\eta}}, \quad \text{with } 0 < \eta < 1. \tag{11}$$

Eqs. (10) and (11) are denoted by $i_T(d) \propto (d+\theta)^{-\eta}$ and $i_T(d) \propto d^{-\eta}$, respectively.

An equivalent formulation to Eqs. (10) and (11), which will form a keypoint in our Bayesian methodology (see Sections 4 and 5), is to explicitly express $I(d) \sim \text{GEV}[\mu(d), \sigma(d), \gamma]$:

- IDF model valid for $d \in [1/6, 72]$:

$$\mu(d) = \frac{\mu}{(d+\theta)^{\eta}}, \quad \sigma(d) = \frac{\sigma}{(d+\theta)^{\eta}}, \text{ with } \theta > 0,\ 0 < \eta < 1. \tag{12}$$

- IDF model valid for $d \in [1, 72]$:

$$\mu(d) = \frac{\mu}{d^{\eta}}, \quad \sigma(d) = \frac{\sigma}{d^{\eta}}, \text{ with } 0 < \eta < 1. \tag{13}$$

3. Point estimators for IDF relationships

To highlight the need of a Bayesian approach for IDF models, we give a real world application of a classical point estimator. This involves the use of data to calculate a single value for the model parameters, which is to serve as a "best guess" or "best estimate". We introduce the 10 min precipitation series of Uccle, which is used for illustration throughout the paper. The series was recorded by a Hellmann–Fuess pluviograph in the period 1898–2007, and since 2008 by an automatic instrument. The Uccle series is worldwide unique, because not only of its length and high temporal resolution, but also because of the completeness and homogeneity (Demarée, 2003). Previous studies of IDF curves based on the Uccle rainfall series have been included in Demarée (1985), Mohymont et al. (2004) and Willems (2000).

Table 1
Classical point estimations for IDF model $i_T(d) \propto (d+\theta)^{-\eta}$, see Koutsoyiannis et al. (1998) (Section 3). Units are d/h, $i_T(d)$/mm h^{-1}.

Method	$\hat{\mu}$	$\hat{\sigma}$	$\hat{\gamma}$	$\hat{\eta}$	$\hat{\theta}$
Typical procedure	16.95	5.36	0.0535	0.80	0.11
Robust estimation	15.23	5.00	0.0400	0.75	0.07
One-step procedure	15.45	4.99	0.0470	0.75	0.08

Table 2
Classical point estimations for IDF model $i_T(d) \propto d^{-\eta}$, see Koutsoyiannis et al. (1998) (Section 3). Units are d/h, $i_T(d)$/mm h^{-1}.

Method	$\hat{\mu}$ (mm)	$\hat{\sigma}$ (mm)	$\hat{\gamma}$	$\hat{\eta}$
Typical procedure	16.30	4.62	0.0591	0.78
Robust estimation	14.53	4.24	0.0482	0.73
One-step procedure	14.62	4.22	0.0571	0.73

For a given variety of rainfall durations, the matrix of annual maximum intensities is given by $\mathbf{i} = (i_j(d_k)) = (i_{jk}) \in \mathbb{R}^{N\times M}$, where $j = 1, \ldots, N$ refers to the year, and $d_k, k = 1, \ldots, M$ to the rainfall duration. We denote by $\mathbf{i}(d_k) \in \mathbb{R}^N$ the vector of annual maximum intensities with duration d_k. The GEV-distribution of $I(d_k)$ has location $\mu(d_k) = \mu_k$, and scale $\sigma(d_k) = \sigma_k$. The typical estimation procedure consists of three steps (Koutsoyiannis et al., 1998).

- Determine the matrix $\mathbf{i}$ for durations 1/6, 1/3, 1/2, 1, 2, 6, 12, 24, 48 and 72 h. For each duration d_k, we fit the GEV$[\mu_k, \sigma_k, \gamma]$-distribution to the annual maximum series $\mathbf{i}(d_k)$ with maximum likelihood, i.e. maximise the logarithm of the *likelihood function*

$$L(\mathbf{i}(d_k); \psi_k) = \prod_{j=1}^{N} g(i_{jk}; \psi_k), \quad (g : \text{GEV-density}), \tag{14}$$

  with respect to $\psi_k = (\mu_k, \sigma_k, \gamma)$, resulting in the maximum likelihood estimator $\hat{\psi}_k$ (Coles, 2001). Beside maximum likelihood estimation, moment-based estimation methods are also popular (Hosking and Wallis, 1997).
- Calculate for each duration d_k the return level

$$\hat{i}_T(d_k) = \hat{\mu}_k - \frac{\hat{\sigma}_k}{\hat{\gamma}}\left\{1 - \left[-\log\left(1 - \frac{1}{T}\right)\right]^{-\hat{\gamma}}\right\}, \tag{15}$$

  for return periods $T = 2,\ 5,\ 10,\ 20,\ 50$ and 100 year. The return levels have been plotted in Fig. 1 (open circles).
- Estimate finally the parameters $(\mu, \sigma, \gamma, \eta, \theta)$ of model Eq. (10) by minimising the total squared difference between the IDF model and the return levels $\hat{i}_T(d_k)$.

The estimation results have been listed in Table 1. In addition, we have implemented two alternative estimation techniques of Koutsoyiannis et al. (1998) for the general IDF relationship; the results have also been listed in Table 1. Likewise, the estimation results for model $i_T(d) \propto d^{-\eta}$ have been listed in Table 2.

Because the traditional methodology is a multiple-step algorithm, it is difficult to assess the error propagation through the steps. On the other hand, Bayesian inference includes uncertainty in the probability model, yielding more complete predictions.

4. Likelihoods for IDF relationships

In all model-based statistical inference, the likelihood function is of central importance, since it expresses the probability of the

observed data under a particular statistical model. In this section we attempt to formulate approximate likelihoods for intensity data $\mathbf{i}$. Consider the M-dimensional random vector $I = (I(d_1), \ldots, I(d_M))$, with joint probability density function $g(i; \psi, \alpha)$. Here, ψ parameterises the marginal structure of each $I(d_k)$, and α parameterises the interdependence between $I(d_k)$ and $I(d_{k'})$. The form of the marginal distribution $I(d) \sim \text{GEV}[\mu(d), \sigma(d), \gamma]$ has been already expressed in Eqs. (12) and (13). We thus have $\psi = (\mu, \sigma, \gamma, \eta, \theta)$ or $\psi = (\mu, \sigma, \gamma, \eta)$, according to the IDF model. Under the assumption that the matrix of annual maximum intensities $\mathbf{i} = (i_j(d_k)) = (i_{jk})$ is based on N independent realisations of I, the *full likelihood* is

$$L(\mathbf{i}; \psi, \alpha) = \prod_{j=1}^{N} g(i_{j1}, \ldots, i_{jM}; \psi, \alpha). \tag{16}$$

The joint density function $g(i; \psi, \alpha)$ is, however, unknown. Furthermore, because the goal of the analysis is inference about IDF parameters ψ, and α is of no interest, the use of Eq. (16) may be considered impractical. In such a situation, one commonly assumes independence amongst the components of I, so that ψ can be inferred from the *independence likelihood*

$$L_{ind}(\mathbf{i}; \psi) = \prod_{j=1}^{N} \prod_{k=1}^{M} g(i_{jk}; \psi), \tag{17}$$

where $g(\bullet)$ is the density function of $\text{GEV}[\mu(d), \sigma(d), \gamma]$. Regarding to the estimation of IDF relationships, the likelihood Eq. (17) was already formulated in Muller et al. (2008), and in a spatial framework, in Lehmann et al. (2013). The term independence (log-) likelihood was taken from Chandler and Bates (2007), where it was used for dependent data structures in general. It is, in turn, a special case of a *composite likelihood* which is used in applications where the full likelihood is analytical unknown, or computationally prohibitive. For background on composite likelihoods, see Varin et al. (2011).

An alternative to the independence likelihood Eq. (17) has been formulated by Muller et al. (2008), who modelled the dependence explicitly by including only three rainfall durations: 1, 24 and 72 h. In fact, they considered yet a fourth duration, namely the daily maxima from a neighbouring station, but this is out of context for our illustrative purposes. Since the 1 h extreme rainfall occurs generally during a thunderstorm, whereas the 72 h extreme rainfall occurs generally during a frontal rainfall event, both events may be considered as independent. The dependence between 24 and 72 h annual maximum intensities is modelled by a bivariate extreme value distribution. A popular standard class, defined for marginal unit-Fréchet distributions, is the logistic family (Coles, 2001)

$$P\{X \leqslant x, Y \leqslant y\} = \exp\left[-\left(x^{-1/\alpha} + y^{-1/\alpha}\right)^{\alpha}\right], \quad x > 0,\ y > 0, \tag{18}$$

where the parameter α is a measure of dependence. It is easy to verify that the transformed intensity

$$U(I(d)) = \left(1 + \gamma(d)\,\frac{I(d) - \mu(d)}{\sigma(d)}\right)^{1/\gamma(d)}, \tag{19}$$

is unit-Fréchet distributed, i.e. $P\{U(I(d)) < z\} = \exp(-z^{-1})$. We arrive at the likelihood (Muller et al., 2008), which is termed by us the *trivariate likelihood*:

$$L_3(\mathbf{i}; \psi, \alpha) = \prod_{j=1}^{N} g\left(i_j(1); \psi\right) g\left(U\left(i_j(24)\right), U\left(i_j(72)\right); \psi, \alpha\right) \times U'\left(i_j(24)\right) U'\left(i_j(72)\right), \tag{20}$$

where $g(\bullet, \bullet)$ is the density function of the bivariate logistic distribution Eq. (18), and $U'(\bullet)$ presents the derivative of $U(I(d))$ with respect to the intensity $I(d)$. The vector of marginal parameters is $\psi = (\mu, \sigma, \gamma, \eta)$, and the matrix of annual maximum intensities is $\mathbf{i} = (\mathbf{i}(1), \mathbf{i}(24), \mathbf{i}(72))$. It should be noted that the trivariate likelihood is not a composite likelihood because it is not an approximation of the full likelihood Eq. (16). It is rather a good representation of the likelihood for maximum intensity data with durations 1, 24 and 72 h.

5. Bayesian inference for IDF relationships

There is now a large body of literature focusing on Bayesian statistics for extremes. For some general background, see Coles (2001) who devotes a section to this topic in Chapter 9. There is also the `R`-package `evdbayes` (Stephenson and Ribatet, 2012) which provides functions for the Bayesian analysis of univariate extreme value models. Recent advances in Bayesian modelling of spatial extremes are reviewed in Davison et al. (2012).

5.1. Bayesian framework

The set-up for a Bayesian analysis is the following. We assume annual maxima data $\mathbf{i}$ to be realisations of a random variable whose density falls within a parametric family $\{g(i; \psi) : \psi \in \Psi\}$. We can express the likelihood for ψ as $L(\mathbf{i}; \psi) = L(\mathbf{i}|\psi)$. For notational brevity we have omitted α, the vector of dependence parameters. The posterior distribution is given by

$$\pi(\psi|\mathbf{i}) \propto L(\mathbf{i}|\psi)\,\pi(\psi), \tag{21}$$

where $\pi(\psi)$ denotes a prior distribution. If a single estimate is required from the Bayesian inference, the mean, median or mode of the posterior distribution can be used. A Bayesian *credible interval (CI)* consists of posterior values that cannot be ruled out at some probability level.

Since the independence likelihood, $L_{ind}(\mathbf{i}; \psi)$, is not the probability of the observed data, the independence posterior distribution

$$\pi_{ind}(\psi|\mathbf{i}) \propto L_{ind}(\mathbf{i}; \psi)\,\pi(\psi), \tag{22}$$

can be misleading. Recently, two adjustments to the composite likelihood were proposed which retrieve some of the desirable properties of the full likelihood (Pauli et al., 2011; Ribatet et al., 2012). They are termed by Ribatet et al. (2012) the *magnitude-* and *curvature-adjustment*; technical details can be found in Appendix A. Inference from the trivariate likelihood Eq. (20) needs no adjustment because it is, as already noted, no genuine composite likelihood.

5.2. Implementation of the Bayesian inference procedures

Here we implement the Bayesian inference procedure developed in the foregoing. In Table 3 we specify the different prior distributions for the IDF parameters. The choices for μ and σ are similar to those in Coles (2001). About the shape prior, it is reasonable to restrict the choice to a plausible range, consistent with rainfall depths observed worldwide. For example, Papalexiou and Koutsoyiannis (2013) have analysed the annual maximum daily rainfall of more than 15,000 records from all over the world, and came to the conclusion that γ is expected to belong in a narrow range, approximately from 0 to 0.23. Taken in this context, a beta prior for γ, with support on $[-0.5, 0.5]$ and mean 0.1 seems realistic and not overly informative, a choice which was recommended by Martins and Stedinger (2000). Further, the prior distributions of η and θ were chosen in agreement with conditions $\theta > 0$ and $0 < \eta < 1$, see Eq. (4). For the additional dependence parameter in the trivariate likelihood, α, we have chosen a uniform prior with support $[0, 1]$.

We obtain approximate draws from the posterior distribution via an MCMC algorithm. For background on Bayesian inference

Table 3
Prior distribution of the IDF parameters.

μ	σ	γ	η	θ
$\mathcal{N}(0, 10^4)$	$\ln\mathcal{N}(0, 10^4)$	Beta$(\alpha = 6, \beta = 9, -0.5, 0.5)$	Unif$(0, 1)$	$\ln\mathcal{N}(0, 10^2)$

via MCMC, see Gelman et al. (2013). Two commonly used random walk MCMC methods are the *Metropolis–Hastings sampler* and the *Gibbs sampler*, see Appendices B.1 and B.2. In short, the Metropolis–Hastings method generates a random walk using a proposal density $q(.|\psi)$, and a method for rejecting some of the proposed moves. In this work, we use a Gaussian proposal distribution $q(.|\psi) \sim \mathcal{N}(\psi, \sigma^2_\psi)$. For many multi-dimensional problems, the Gibbs sampler has been found useful, and is defined in terms of subvectors of ψ. Suppose the parameter vector has been partitioned into G components or subvectors, $\psi = (\psi_1^T, \ldots, \psi_G^T)^T$. Each iteration of the Gibbs sampler cycles through the subvectors of ψ, drawing each subset conditional on the ψ-value of all the others with a Metropolis–Hastings-step. There are thus G Metropolis–Hastings-steps per Gibbs iteration. Different possibilities for Gibbs partitions are explored in Section 6.1.

Ribatet et al. (2012) presented two ways of incorporating an adjusted composite likelihood into the Gibbs sampler. The simplest one is termed by them the *overall Gibbs sampler*. The other one, the *adaptive Gibbs sampler*, improves the accuracy of the overall Gibbs sampler, but it is a much more computationally intensive method. Technical details are described more fully in Appendix B.3. With respect to this work, implementation issues have been summarised in the form of a pseudo-code at the end of Appendix B.4.

5.3. MCMC convergence diagnostics

To determine when it is safe to stop sampling, and use the samples to estimate characteristics of the posterior distribution of interest, we have used R-package CODA (Plummer et al., 2006) which contains a set of functions for convergence diagnostics for MCMC. We rely on the procedure of Heidelberger and Welch (1983), which includes the Cramer-von-Mises statistic to test the null hypothesis that the sampled values come from a stationary distribution. The test is successively applied, firstly to the whole chain, then after discarding the first 10%, 20%, ...of the chain until either the null hypothesis is accepted at significance level $p = 0.05$, or 50% of the chain has been discarded. The latter outcome constitutes "failure" of the stationarity test and indicates that a longer MCMC run is needed. The half-width test calculates a 95% confidence interval for the mean, using the portion of the chain which passed the stationarity test. Half the width of this interval is compared with the estimate of the mean. If the ratio between the half-width and the mean is lower than 0.1, the half-width test is passed. Otherwise, the length of the sample is deemed not long enough to estimate the mean with sufficient accuracy.

6. Application

6.1. Gibbs sampling: convergence, sensitivity, comparisons

Each statistical model was run for 10^5 MCMC iterations. In order to reduce the burn-in period, the MCMC algorithms were initialised with the values of a classical point estimator, obtained by a preliminary analysis. Tables 4–6 summarise the convergence diagnostics. For the simplest IDF model, $i_T(d) \propto d^{-\eta}$, convergence was reached by every MCMC scheme. For model $i_T(d) \propto (d+\theta)^{-\eta}$, only the overall Gibbs sampler including the curvature-adjusted independence likelihood appeared to have converged; in case of the magnitude-adjustment, the half-width test did not pass for the simulated θ-values, even not by 10^6 simulations. The adaptive Gibbs simulations may converge too slowly to be of practical value, or never arrive at the proper solution.

All the Gibbs samplers used blocks containing one model parameter, but other block dimensions are possible as well. Because the parameters of the general IDF relationship fall into two categories (those of function $a(T)$, and those of function $b(d)$), we might suggest to run two-block versions of the Gibbs sampler with $\psi_1 = (\mu, \sigma, \gamma)$, and $\psi_2 = (\eta, \theta)$ for $i_T(d) \propto (d+\theta)^{-\eta}$, or $\psi_2 = \eta$ for $i_T(d) \propto d^{-\eta}$. Concerning the independence likelihood, the adjusted posterior distributions provided by the one-element and two-block versions (if passed the Heidelberger and Welch (1983) test) are hardly different (not shown). In contrast, the three-block version $\psi_1 = (\mu, \sigma, \gamma), \psi_2 = \eta$ and $\psi_3 = \alpha$ for the trivariate likelihood did not converge.

To assess the sensitivity to the priors, we ran several additional MCMC simulations with different priors. We saw no substantive change in the posterior distributions for the μ-, σ- and η-parameter estimates. Our particular interest in rainfall extremes are the parameters γ and θ, which are known to be difficult to estimate. We proposed the following additional priors:

- $\pi(\gamma)$: (i) a normal distribution, $\mathcal{N}(0, \sigma^2)$, with variance $\sigma^2 = 0.16$, (ii) a uniform distribution on $[-0.5, 0.5]$, and (iii) an improper prior, Unif$(-\infty, \infty)$.

Table 4
Convergence diagnostics with the Heidelberger and Welch (1983) test of simulated model parameters, obtained with the overall Gibbs sampler with the adjusted independence likelihood. "–" indicates that the half-width test failed.

Parameters	Magnitude-adjusted			Curvature-adjusted		
	p-value	Mean	Half-width	p-value	Mean	Half-width
$i_T(d) \propto d^{-\eta}$						
μ	0.97	14.44	0.17	0.92	14.31	0.06
σ	0.35	4.25	0.04	0.47	4.18	0.02
γ	0.20	0.06	0.002	0.63	0.06	0.003
η	0.88	0.73	0.003	0.94	0.72	0.002
$i_T(d) \propto (d+\theta)^{-\eta}$						
μ	0.20	14.99	0.17	0.98	15.29	0.08
σ	0.09	4.96	0.038	0.83	5.04	0.03
γ	0.09	0.06	0.001	0.49	0.05	0.002
η	0.16	0.74	0.003	0.97	0.74	0.001
θ	0.12	–	–	0.94	0.08	0.001

Table 5
Convergence diagnostics with the Heidelberger and Welch (1983) test of simulated model parameters, obtained with the adaptive Gibbs sampler with the adjusted independence likelihood. "–" indicates that either the Cramer–von-Mises failed ($p < 0.05$), or the half-width test failed.

Parameters	Magnitude-adjusted			Curvature-adjusted		
	p-value	Mean	Half-width	p-value	Mean	Half-width
$i_T(d) \propto d^{-\eta}$						
μ	0.31	14.26	0.07	0.43	14.23	0.07
σ	0.28	4.18	0.02	0.23	4.16	0.02
γ	0.10	0.06	0.002	0.26	0.06	0.002
η	0.35	0.72	0.001	0.38	0.72	0.002
$i_T(d) \propto (d+\theta)^{-\eta}$						
μ	$9.39 \cdot 10^{-5}$	–	–	0.52	14.57	0.36
σ	$3.01 \cdot 10^{-3}$	–	–	$6.7 \cdot 10^{-4}$	–	–
γ	0.36	0.05	0.002	$2.1 \cdot 10^{-3}$	–	–
η	$1.35 \cdot 10^{-4}$	–	–	0.32	0.73	0.006
θ	$5.69 \cdot 10^{-7}$	–	–	0.08	–	–

Table 6
Convergence diagnostics with the Heidelberger and Welch (1983) test of simulated model parameters, obtained from the trivariate likelihood.

Parameters	p-value	Mean	Half-width
μ	0.56	14.22	0.03
σ	0.41	4.34	0.02
γ	0.53	0.09	0.003
η	0.53	0.72	0.0006
α	0.21	0.35	0.001

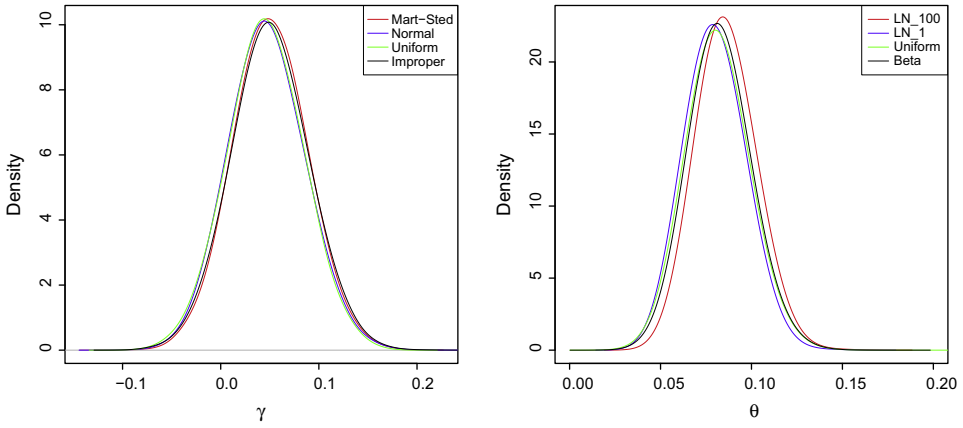

Fig. 2. Sensitivity analysis to priors. Posterior distributions of γ (left), and θ (right), corresponding to the priors of Table 3 and Section 6.1. Overall Gibbs sampler with the curvature-adjusted independence likelihood.

- $\pi(\theta)$: (i) a lognormal distribution, $\ln \mathcal{N}(0, \sigma^2)$, with variance $\sigma^2 = 1$, (ii) a uniform distribution on $[0, 1]$, and (iii) a Beta distribution, $\beta(2, 5)$.

Fig. 2 shows that the resulting posterior distributions are very similar for the overall Gibbs sampler including the curvature-adjusted independence likelihood. Similar results were obtained for the other MCMC algorithms (not shown).

The adaptive Gibbs sampler is expected to improve the estimation provided by the overall Gibbs sampler, so we compared the approximations of both algorithms. From Fig. 3 it is clearly seen that the posteriors obtained by both samplers differ slightly for the simplified model $i_T(d) \propto d^{-\eta}$. The question arises if it is really advantageous to implement the adaptive Gibbs sampler. A chain of 10^6 iterations provided by the overall Gibbs sampler takes a bit less than 1 h to run, whereas the adaptive Gibbs sampler takes 47 h to run.

6.2. Bayesian IDF model selection

We examined which of the two models, $i_T(d) \propto (d+\theta)^{-\eta}$ and $i_T(d) \propto d^{-\eta}$, is the most appropriate when considering $d \in [1, 72]$. Particular attention was paid to the prior choice of θ: because of $d \geqslant 1$, we have condition $\theta > -1$, instead of condition $\theta > 0$ in Eq. (4). The θ-prior in Table 3 was therefore being shifted leftward with one unit. Lower and upper bound estimates of the 95% CI were calculated by taking the pointwise 0.025 and 0.0975 empirical quantiles from the θ-draws. The overall Gibbs sampler with the adjusted independence likelihood gives a posterior mean $\bar{\theta} = 0.011$, and a 95% CI of $(-0.258, 0.354)$, which suggests that the specific choice $\theta = 0$ might be appropriate. We investigated this point further, and used the *deviance information criterion (DIC)* (Spiegelhalter et al., 2002) as a guide for model selection. The DIC produces a measure of the model fit, and a measure of model complexity p_D (i.e. the effective number of parameters), and sums

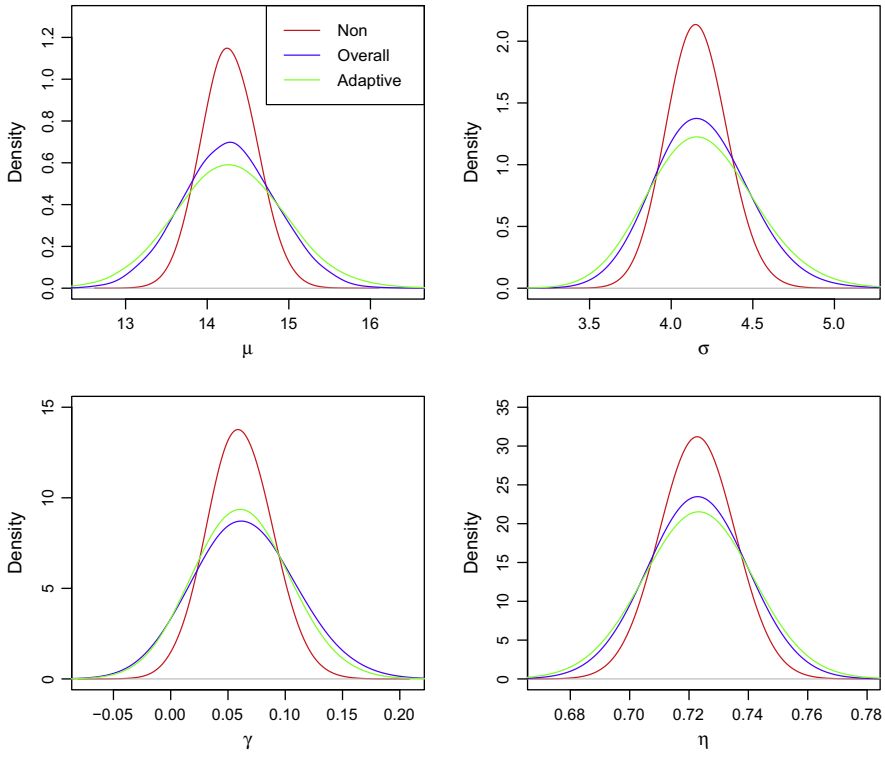

Fig. 3. Overall versus adaptive Gibbs sampler. Posterior distributions for the IDF parameters of model $i_T(d) \propto d^{-\eta}$. Inference from the magnitude-adjusted independence likelihood.

Table 7
Model comparison with DIC, in the case that $d \in [1, 72]$.

Likelihood	$i_T(d) \propto d^{-\eta}$		$i_T(d) \propto (d+\theta)^{-\eta}$	
	p_D	DIC	p_D	DIC
Trivariate	4.8	2456	4.8	2487
Independence	3.9	5930	4.8	5933

them to get an overall score (lower is better). The model comparison has been examined with both trivariate- and independence likelihood (Table 7). When using the trivariate likelihood, model $i_T(d) \propto d^{-\eta}$ is superior to model $i_T(d) \propto (d+\theta)^{-\eta}$, for rainfall durations $d \in [1, 72]$ (in terms of DIC performance). For the independence likelihood, a comparison of DIC suggests that both models are virtually indistinguishable. Pragmatically, we might prefer model $i_T(d) \propto d^{-\eta}$ because of its simplicity.

6.3. Comparison of independence- and trivariate likelihood

The trivariate likelihood only considers rainfall durations of 1-, 24- and 72 h. To make a fair comparison with the independence likelihood, we considered the same durations in the latter likelihood as well. The inference results in Table 8, and the posterior distributions of Fig. 4, indicate that estimations based on both likelihoods differ, in particular for μ and σ. A major observation is that for these parameters, the posterior distributions from the trivariate likelihood are more narrow. An extension of the trivariate likelihood definition could be considering the pairwise composite likelihood where each component is a bivariate likelihood based on maxima intensity data for durations d_k and $d_{k'}$. This should, however, substantially increases the complexity of the model since every pair needs to be modelled with a different dependence parameter α. For instance, modelling $i_T(d) \propto (d+\theta)^{-\eta}$ may require 10 rainfall durations, leading to at most $(10 \times 9)/2$ different α's.

Table 8
Trivariate-versus independence likelihood (including 1-, 24- and 72 h data). Overall Gibbs sampler with adjusted independence likelihood. Sample mean and standard deviation (in parentheses) of simulated values of IDF parameters of model $i_T(d) \propto d^{-\eta}$. Units are d/h, $i_T(d)$/mm h^{-1}.

	μ	σ	γ	η
Trivariate likelihood				
	14.25 (0.44)	4.35 (0.26)	0.091 (0.044)	0.72 (0.010)
Independence likelihood				
Non	14.13 (0.47)	4.09 (0.24)	0.080 (0.040)	0.72 (0.010)
Magn	14.14 (0.54)	4.11 (0.27)	0.081 (0.046)	0.72 (0.012)
Curv	14.13 (0.61)	4.10 (0.29)	0.080 (0.041)	0.72 (0.011)

The question whether or not this is technically feasible is reserved for future work. In the meantime, we proceed further with the independence likelihood.

6.4. Return level inference

Tables 9 and 10 show the posterior means and their approximate standard errors. The sample means provided by the new methods fairly match the classical point estimations (Tables 1 and 2). It is clearly seen that the variability of the unadjusted posterior is small with respect to the adjusted posterior. This is a direct reflection of the fact that we have lost information because the full likelihood was unavailable.

For operational purposes, the posterior distribution of the IDF parameters is difficult to use. Extreme behaviour is summarised by the T-year return level, $i_T(d)$. As we have draws from the posterior distribution, $\psi^{(1)}, \ldots, \psi^{(S)}$, we can naturally obtain draws from the posterior distribution of $i_T(d)$, viz. $i_T^{(1)}(d), \ldots, i_T^{(S)}(d)$ is computed by means of Eqs. (10) or (11). The posterior distributions of the return levels have been shown in Figs. 5 and 6. Again, the variability of the adjusted posteriors is higher. The probability level

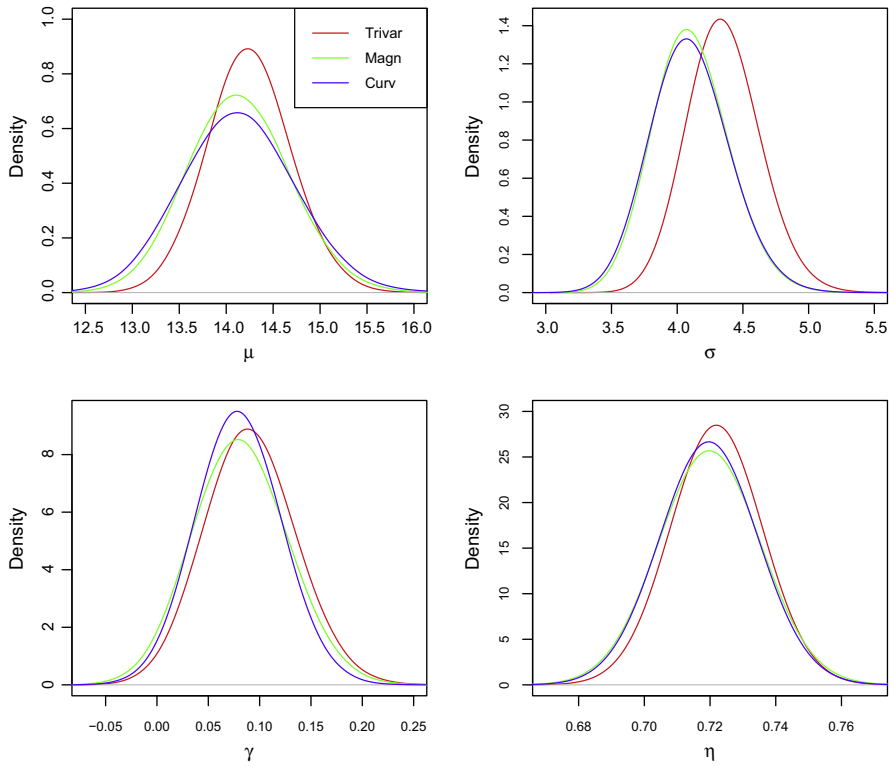

Fig. 4. Posterior marginal distributions for the IDF parameters of model $i_T(d) \propto d^{-\eta}$. Trivariate- and adjusted independence likelihood (overall Gibbs sampler) including 1-, 24- and 72 h data.

Table 9
Posterior means and their approximate standard errors (in parentheses) of simulated values of IDF parameters of model $i_T(d) \propto d^{-\eta}$. Units are d/h, $i_T(d)$/mm h^{-1}. Overall and adaptive Gibbs sampler with an adjusted independence likelihood.

	μ	σ	γ	η
Non	14.27 (0.33)	4.16 (0.16)	0.060 (0.027)	0.72 (0.008)
Overall Gibbs sampler				
Magn	14.27 (0.57)	4.19 (0.28)	0.065 (0.045)	0.72 (0.014)
Curv	14.24 (0.60)	4.16 (0.29)	0.062 (0.038)	0.72 (0.011)
Adaptive Gibbs sampler				
Magn	14.28 (0.68)	4.19 (0.31)	0.062 (0.041)	0.72 (0.016)
Curv	14.27 (0.67)	4.18 (0.31)	0.061 (0.042)	0.72 (0.016)

determines the credible interval CI=(CI$^-$,CI$^+$) for $i_T(d)$, and is conventionally set at 95%. The direct computation of the 95% CI with MCMC algorithms is not really product oriented because this can take a long time to run. For operational purposes, the logarithm of the bounds CI$^\pm$ can be fairly good approximated by a bivariate linear regression model

$$\begin{aligned} &\text{Model } i_T(d) \propto d^{-\eta}: && \log \mathrm{CI}^{\pm} \approx a_0 + a_1 \log d + a_2 \log T, \\ &\text{Model } i_T(d) \propto (d+\theta)^{-\eta}: && \log \mathrm{CI}^{\pm} \approx a_0 + a_1 \log(d+\theta) + a_2 \log T, \end{aligned} \tag{23}$$

for T ranging from 10 up to 1000 years. The estimated regression models, expressed in a multiplicative form, have been summarised in Table 11. In providing higher return level estimates, notice we extrapolate beyond the range of the data. As we can see from Table 11, extrapolation increases the amount of uncertainty of these estimates.

6.5. Model diagnostics

We evaluated how good the estimated models describe the available data. In the present context, classical QQ-plots are not useful as the maximum intensity data are not identically distributed. A possible extension of classical QQ-plots consists in transforming the data to identically distributed variables. Under the assumption $I(d) \sim \mathrm{GEV}[\mu(d), \sigma(d), \gamma]$, the transformation

$$\tilde{I}(d) = \frac{1}{\gamma} \log\left(1 + \gamma \frac{I(d) - \mu(d)}{\sigma(d)}\right), \tag{24}$$

results in a Gumbel distributed random variable (Beirlant et al., 2004; Coles, 2001), i.e. $P\{\tilde{I}(d) \leqslant z\} = \exp(-\exp(-z))$. The Gumbel distribution obtained does not any longer depend on d, and hence the random variable $\tilde{I}(d) =: \tilde{I}$ is identically distributed. The quantile function associated with the Gumbel distribution is given by

$$Q(p) = -\log(-\log p), \quad 0 < p < 1, \tag{25}$$

yielding the Gumbel QQ-plot coordinates

$$\left(-\log\left(-\log\frac{i}{K+1}\right), \tilde{I}_{(i)}\right), \quad i = 1, \ldots, K, \tag{26}$$

Table 10
Posterior means and their approximate standard errors (in parentheses) of simulated values of IDF parameters of model $i_T(d) \propto (d+\theta)^{-\eta}$. Units are d/h, $i_T(d)$/mm h^{-1}. Overall Gibbs sampler with the adjusted independence likelihood.

	μ	σ	γ	η	θ
Non	15.23 (0.46)	5.01 (0.19)	0.048 (0.023)	0.74 (0.010)	0.078 (0.011)
Curv	15.24 (0.71)	5.02 (0.35)	0.051 (0.037)	0.74 (0.013)	0.079 (0.015)

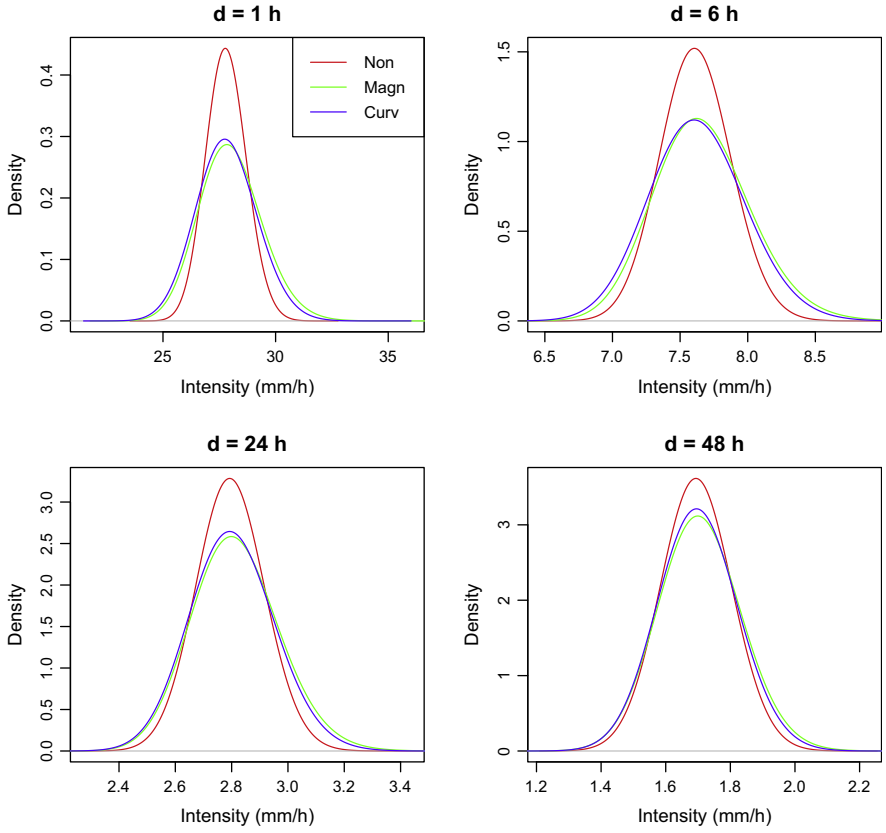

Fig. 5. Posterior distributions of 20-year return levels of the maximum intensity. Overall Gibbs sampler with the (adjusted) independence likelihood. IDF model: $i_T(d) \propto d^{-\eta}$.

where $\tilde{I}_{(1)} \leqslant \cdots \leqslant \tilde{I}_{(K)}$ are the corresponding order statistics of $\tilde{I}$, and $K = MN$ the total number of data. If the Gumbel model provides accurate description of the data, one expects the points on the Gumbel QQ-plot to be close to the leading diagonal. Observing Fig. 7, one can conclude that the IDF models reasonably describe the maximum intensities.

7. Conclusion

The statistical contribution of this work lies in developing and applying a Bayesian analysis for extreme rainfall IDF relationships. We reconsidered the independence likelihood of Muller et al. (2008) and Lehmann et al. (2013), assuming independence between annual maximum intensities over different durations. The main issue of this paper is about demonstrating the Ribatet et al. (2012) composite likelihood adjustments for IDF modelling.

We have examined two implementations of the Gibbs sampler (overall and adaptive version) for obtaining draws from the adjusted independence posterior. The overall Gibbs sampler is computationally efficient, and easy to implement. In contrast, the adapted Gibbs sampler requires an enormously computational effort. The modelling of $i_T(d) \propto d^{-\eta}$ did not give rise to particular problems. Both the overall and adaptive Gibbs sampler converged properly, and the posterior distributions are remarkably robust to the prior choice. There was virtually no difference in the estimates of the magnitude- and curvature-adjusted posteriors obtained by both Gibbs samplers, suggesting that both adjustments adequately capture the information contained in the independence likelihood. In addition, for the one-block version and a two-block version of the Gibbs sampler, there was likewise no difference in the estimates of the posteriors. The things change, however, when considering the model $i_T(d) \propto (d+\theta)^{-\eta}$ which covers rainfall durations less than one hour. Only the overall Gibbs sampler with the curvature-adjusted likelihood was able to draw from the independence posterior.

An alternative to the independence likelihood was proposed by Muller et al. (2008) who proceeded with a bivariate extreme value model, assuming, in particular, that a logistic model is appropriate. Their likelihood formulation, which is termed by us the trivariate likelihood, includes only annual maximum intensities over durations 1, 24 and 72 h. Although, the resulting posterior is more narrow than the adjusted independence posterior (based on 1, 24 and 72 h data as well), we recommend to use the independence likelihood because it features the ability to incorporate more rainfall durations. In particular, short-duration (10 min to 1 h) rainfall extremes, which are equally important for a number of purposes, cannot be modelled when using the trivariate likelihood.

Likelihood based methods allow for applying model selection with information criteria. The model comparison with DIC indicates that the simplification $\theta = 0$ in the model $i_T(d) \sim (d+\theta)^{-\eta}$ should be strongly preferred when considering durations $d \in [1, 72]$.

Part of the reason for adopting a Bayesian methodology was to obtain natural uncertainty estimates, and contrasts with classical analysis in which it is usual to calculate a point estimate. The sample mean of the adjusted posterior distribution agrees very well with the classical point estimations. The variance of the adjusted posterior is larger than that of the unadjusted posterior, illustrating the effect of not having the full likelihood available. Because MCMC simulations can be time-consuming, a more product oriented approach is to fit a bivariate regression model (as a function of d and T) to the estimated credible bounds of $i_T(d)$.

This study can be extended in various ways. First, the analysis was focused on the general IDF relationship (Koutsoyiannis et al., 1998). We could explore different IDF models, as given in Muller et al. (2008), Overeem et al. (2008) and Willems (2000) for

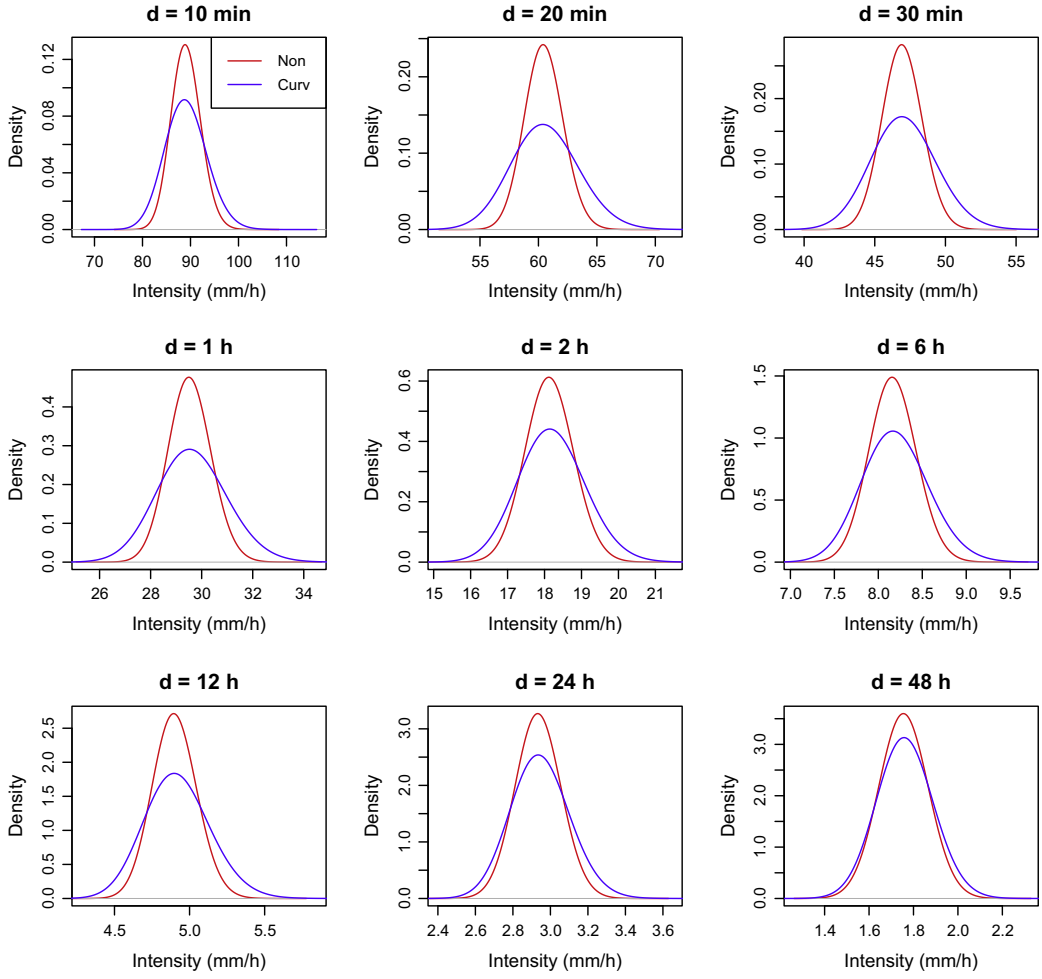

Fig. 6. Posterior distributions of 20-year return levels of the maximum intensity. Overall Gibbs sampler with the (adjusted) independence likelihood. IDF model: $i_T(d) \propto (d+\theta)^{-\eta}$.

Table 11
Regression relationships for the bounds $CI^{\pm}$ of the 95%-credible intervals of the maximum intensity return levels, $i_T(d)$, for return period $T = 10, \ldots, 1000$ year. Posterior distributions of $i_T(d)$ were obtained from an adjusted independence likelihood.

	95%-credible interval	Rel. Err. (%)
$i_T(d) \propto d^{-\eta}$, for $d \in [1,72]$		
Overall Gibbs sampler		
Magn	$CI^- \approx 17.674 \times d^{-0.727} \times T^{0.129}$	1.9
	$CI^+ \approx 16.652 \times d^{-0.718} \times T^{0.200}$	1.0
Curv	$CI^- \approx 17.336 \times d^{-0.726} \times T^{0.133}$	1.9
	$CI^+ \approx 17.001 \times d^{-0.717} \times T^{0.188}$	1.1
Adaptive Gibbs sampler		
Magn	$CI^- \approx 17.547 \times d^{-0.728} \times T^{0.132}$	2.1
	$CI^+ \approx 16.923 \times d^{-0.716} \times T^{0.191}$	1.1
Curv	$CI^- \approx 17.531 \times d^{-0.728} \times T^{0.131}$	2.0
	$CI^+ \approx 16.881 \times d^{-0.716} \times T^{0.191}$	1.1
$i_T(d) \propto (d+\theta)^{-\eta}$, for $d \in [1/6,72]$		
Overall Gibbs sampler		
Curv	$CI^- \approx 19.140 \times (d+0.079)^{-0.746} \times T^{0.137}$	2.2
	$CI^+ \approx 19.353 \times (d+0.079)^{-0.739} \times T^{0.188}$	1.3

example, and make a model comparison with DIC. Second, the trivariate likelihood can in principle be extended to the composite pairwise likelihood including all durations, but the computational cost will be appreciably higher. Third, despite the fact that the block maxima methodology is very popular in IDF analysis, it is a wasteful approach if other data on extremes are available. As an alternative, the peaks-over-threshold method could be employed in the IDF analysis, which incorporates more data, resulting in smaller uncertainty into return level prediction. Finally, another area for focus of future research is to add a spatial component in the independence likelihood for characterising rainfall extremes over a region of interest, as was already done by Lehmann et al.

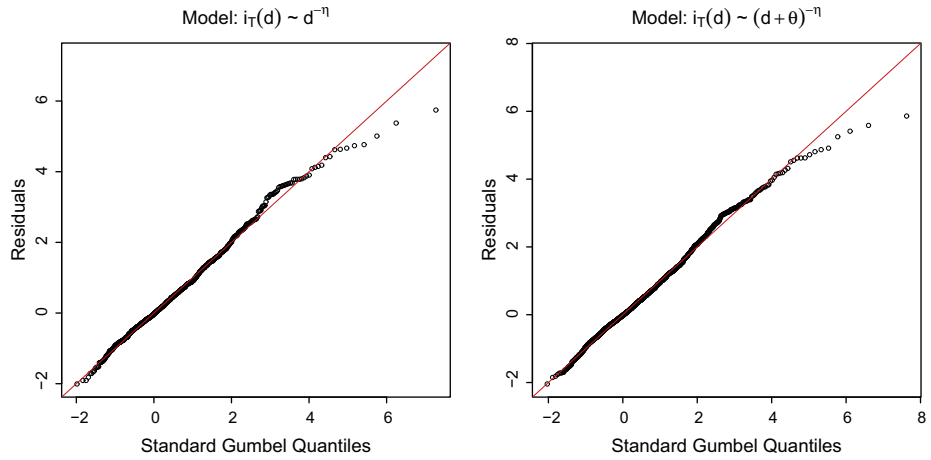

Fig. 7. Gumbel QQ-plot of maximum intensities from the Uccle series (1898-present). Modelled values are based on the overall Gibbs sampler with the curvature-adjusted independence likelihood. *Dots* Eq. (26). *Solid line* leading diagonal.

(2013) with unadjusted inference. When combining precipitation data of several sites, we hope to make a better inference, in particular for the difficult-to-estimate parameters γ and θ. The investigation of these problems is now in progress.

Acknowledgements

This work is supported, in part, by the Belgian Science Policy Office (BELSPO) under contract No. SD/RI/03A. The author wishes to thank two anonymous referees whose invaluable comments strongly aided me to develop the manuscript.

Appendix A. Bayesian estimation for composite likelihoods

Consider an M-dimensional random vector X, with probability density $g(x;\psi,\alpha)$. For N independent replicates of X, the full likelihood $L(x;\psi)$ is already defined in Eq. (16), but with $X = I$. In many research areas, one often uses complex models for which obtaining the full likelihood poses not only theoretical, but also computational challenges. In these situations, it is desirable to search for an approximation in the form of a composite likelihood. To some extent, a composite likelihood, $L_c(x;\psi)$, is an inference function derived by multiplying a collection of marginal or conditional densities. The simplest composite likelihood is obtained under independence assumptions, and consists of a product of marginal densities, see, for example, Eq. (17). The most general definition of composite likelihoods falls outside the scope of this paper, but the reader is referred to Varin et al. (2011). We briefly report recent developments on Bayesian inference from composite likelihoods (Pauli et al., 2011; Ribatet et al., 2012). The idea is to adjust the composite likelihood in such a way that the usual asymptotic χ^2-distribution for the likelihood ratio statistic is preserved. To be more precise, suppose that the vector of parameters $\psi \in \mathbb{R}^p$ is equal to ψ_0, and $\hat{\psi}$ maximises the full likelihood. Denote by $l_c(x;\psi) = \log L_c(x;\psi)$, the composite log-likelihood. Also, denote by $\hat{\psi}_c$ the composite likelihood estimator of ψ_0, obtained by maximising $l_c(x;\psi)$. Let $H(\psi_0) = -\mathrm{E}[\nabla^2 l_c(x;\psi_0)]$ be the expected Hessian of $l_c(x;\psi)$, and $J(\psi_0) = \mathrm{Var}[\nabla l_c(x;\psi_0)]$. Then, as $N \to \infty$,

$$\Lambda(\psi_0) = 2\left[l(x;\hat{\psi}) - l(x;\psi_0)\right] \xrightarrow{d} \chi^2_p. \tag{A.1}$$

In contrast, for the composite log-likelihood we have (Rotnitzky and Jewell, 1990):

$$\Lambda_c(\psi_0) = 2\left[l_c(x;\hat{\psi}_c) - l_c(x;\psi_0)\right] \xrightarrow{d} \sum_{i=1}^{p} \lambda_i X_i, \tag{A.2}$$

where $X_1,\ldots,X_p$ are independent χ^2_1 random variables, and $\lambda_1,\ldots,\lambda_p$ are the eigenvalues of $H(\psi_0)^{-1}J(\psi_0)$. Since ψ_0 is unknown, $H(\psi_0)$ is approximated by $H(\hat{\psi}_c)$, which has been included in the standard output of optimisation routines. The estimation of $V(\psi_0)$ poses more difficulties, since the naive estimator $V(\hat{\psi}_c)$ disappears when evaluated at the maximum likelihood estimator. A more detailed discussion on the estimation of $V(\psi_0)$ may be found in Varin (2008).

A.1. Magnitude adjustment (Pauli et al., 2011; Ribatet et al., 2012)

The idea is to determine k:

$$l_c^{magn}(x;\psi) = k\, l_c(x;\psi), \tag{A.3}$$

in such a way that the first and second order moments of

$$\Lambda_{magn}(\psi_0) = 2\left[l_c^{magn}(x;\hat{\psi}_c) - l_c^{magn}(x;\psi_0)\right], \tag{A.4}$$

converge to those of the χ^2_p-distribution. One gets

$$k = p\left(\sum_{i=1}^{p}\lambda_i\right)^{-1}. \tag{A.5}$$

A.2. Curvature adjustment (Ribatet et al., 2012)

Determine $C \in \mathbb{R}^{p\times p}$ in

$$l_c^{curv}(x;\psi) = l_c(x;\hat{\psi}), \quad \hat{\psi} = \hat{\psi}_c + C(\psi - \hat{\psi}_c), \tag{A.6}$$

such that the usual property Eq. (A.1) is recovered for $l_c^{curv}(x;\psi)$. One gets

$$C = M^{-1}M_A, \tag{A.7}$$

where $M^T M = H(\psi_0)$, and $M_A^T M_A = H(\psi_0)J(\psi_0)^{-1}H(\psi_0)$. If $p > 1$ the matrix square roots M and M_A are not unique; a possible solution was suggested by Chandler and Bates (2007).

Appendix B. Markov Chain Monte Carlo samplers

B.1. Metropolis–Hastings sampler

The idea behind one such method, the Metropolis–Hastings sampler, is to produce simulated values from the posterior distribution, $\pi(\psi|x) \propto L(x;\psi)\,\pi(\psi)$, in the following way: set an initial value $\psi^{(1)}$ and introduce a probability rule $q(\cdot|\psi^{(t)})$ to generate a proposal value ψ^* for $\psi^{(t+1)}$. Specifically, letting

$$\alpha_t = \min\left\{1, \frac{\pi(\psi^*)L(\psi^*|x)q(\psi^{(t)}|\psi^*)}{\pi(\psi^{(t)})L(\psi^{(t)}|x)q(\psi^*|\psi^{(t)})}\right\}, \tag{B.1}$$

we set

$$\psi^{(t+1)} = \begin{cases} \psi^* & \text{with probability} \quad \alpha_t, \\ \psi^{(t)} & \text{with probability} \quad 1-\alpha_t. \end{cases} \tag{B.2}$$

B.2. Gibbs sampler

Gibbs sampling can be regarded as a special case of the Metropolis–Hastings algorithm. Consider the partition of the parametric vector $\psi = (\psi_1^T, \ldots, \psi_G^T)^T \in \mathbb{R}^p$, where $\psi_j \in \mathbb{R}^{p_j}$ and $\sum_{j=1}^G p_j = p$. A Gibbs sampler successively draws from

$$\pi(\psi_j|\psi_{-j}, x) \propto L(\psi_j|\psi_{-j}, x)\pi(\psi_j), \quad j = 1, \ldots, G, \tag{B.3}$$

where ψ_{-j} is the parameter vector ψ with the elements of ψ_j removed. An iteration of the Gibbs sampler consists of cycling through the G parametric subvectors by drawing a Metropolis–Hastings sample of ψ_j, conditional on the value of ψ_{-j}. For every iteration step, say t, and corresponding set of values $\psi^{(t)} = (\psi_1^{(t)}, \ldots, \psi_G^{(t)})$, the algorithm proceeds as follows:

$$\begin{aligned} &\text{Draw}\, \psi_1^{(t+1)} \sim \pi(\psi_1|\psi_2^{(t)}, \ldots, \psi_G^{(t)}, x), \\ &\text{Draw}\, \psi_2^{(t+1)} \sim \pi(\psi_2|\psi_1^{(t+1)}, \psi_3^{(t)}, \ldots, \psi_G^{(t)}, x), \\ &\qquad \vdots \\ &\text{Draw}\, \psi_G^{(t+1)} \sim \pi(\psi_G|\psi_1^{(t+1)}, \ldots, \psi_{G-1}^{(t+1)}, x). \end{aligned} \tag{B.4}$$

B.3. Gibbs sampling from adjusted composite posteriors

Draws from the adjusted composite posterior can be obtained by replacing the full likelihood $L(x;\psi)$ by the adjusted composite likelihood $L_c^{adj}(x;\psi)$ in the Metropolis–Hastings step, Eqs. (B.1) and (B.2). The *overall Gibbs sampler*, calculates $L_c^{adj}(x;\psi)$ in advance (Ribatet et al., 2012). This requires only preliminary estimates of ψ_0, $H(\psi_0)$ and $V(\psi_0)$, which can be obtained by maximisation of the composite likelihood $L_c(x;\psi)$, see Appendices A.1 and A.2. However, at each Gibbs step, ψ_j is drawn using the conditional composite likelihood $L_c(\psi_j|\psi_{-j}, x)$. In the overall Gibbs sampler, an adjustment of the total $L_c(x;\psi)$ is made, not of the conditional $L_c(\psi_j|\psi_{-j}, x)$. Speaking in terms of accuracy, the sampler can thus be improved by adjusting $L_c(\psi_j|\psi_{-j}, x)$ based on current values of $\psi_{-j} = \psi_{-j}^{(t)}$. This requires maximum composite likelihood estimation with ψ_{-j} held fixed at $\psi_{-j}^{(t)}$, resulting in: $\hat{\psi}_{j,c}, \widehat{H}_{jj}(\hat{\psi}_{j,c}) = \nabla^2 l_c(\hat{\psi}_{j,c}|\psi_{-j}^{(t)}, x)$ and $\widehat{J}_{jj}(\hat{\psi}_{j,c})$, the sample covariance matrix of $\nabla l_c(\hat{\psi}_{j,c}|\psi_{-j}^{(t)}, x)$. It is, however, much more expensive than its overall counterpart, since every iteration step needs extra maximum likelihood estimates. This version is called the *adaptive Gibbs sampler* (Ribatet et al., 2012).

B.4. Pseudo-codes of Gibbs samplers using the independence IDF likelihood

Overall adjusted Gibbs sampler

Input: (i) Matrix $\mathbf{i}$ of maxima intensities, (ii) IDF model, i.e. $i_T(d) \propto (d+\theta)^{-\eta}$ or $i_T(d) \propto d^{-\eta}$, with parameters ψ, (iii) prior distribution $\pi(\psi)$ from Table 3, (iv) initial value $\psi^{(1)}$, provided by a classical point estimator, (v) partition of ψ in G blocks, $\psi = (\psi_1, \ldots, \psi_G)$ (see Section 6.1), (vi) for the ith parameter of the jth block, ψ_{ji}, define $q(\cdot|\psi_{ji}) \sim \mathcal{N}(\psi_{ji}, \sigma_{ji}^2)$, a Gaussian random walk process, (vii) adjusted independence likelihood, $L_{ind}^{adj}(\mathbf{i};\psi)$, see Appendices A.1 and A.2. (Unadjusted $L_{ind}(\mathbf{i};\psi)$ given in Eq. (17))
Output: a Markov chain realisation of length S.
for $t = 2$ **to** S **do**
 for $j = 1$ **to** G **do**
  for $i = 1$ **to** $|\psi_j|$ **do**
   $\psi_{ji}^* = \psi_{ji}^{(t)} + \epsilon_{ji}$ with $\epsilon_{ji} \sim \mathcal{N}(0, \sigma_{ji}^2)$;
  end
  $\alpha = \min\left\{1, \frac{\pi(\psi^*)L_{ind}^{adj}(\psi_j^*|\psi_{-j}^{(t)}, \mathbf{i})}{\pi(\psi^{(t)})L_{ind}^{adj}(\psi_j^{(t)}|\psi_{-j}^{(t)}, \mathbf{i})}\right\}$; $\quad (\psi_{-j}^{(t)}$: current value of $\psi_{-j})$
  $U \sim \text{Unif}(0,1)$;
  if $\alpha \leqslant U$ **then**
   $\psi_j^{(t+1)} = \psi_j^*$;
  else
   $\psi_j^{(t+1)} = \psi_j^{(t)}$;
  end
 end
end
return $\{\psi^{(t)}\}_{t=1,\ldots,S}$;

Adaptive adjusted Gibbs sampler

Input: (i) Matrix $\mathbf{i}$ of maxima intensities, (ii) IDF model, i.e. $i_T(d) \propto (d+\theta)^{-\eta}$ or $i_T(d) \propto d^{-\eta}$, with parameters ψ, (iii) prior distribution $\pi(\psi)$ from Table 3, (iv) initial value $\psi^{(1)}$, provided by a classical point estimator, (v) partition of ψ in G blocks, $\psi = (\psi_1, \ldots, \psi_G)$ (see Section 6.1), (vi) for the ith parameter of the jth block, ψ_{ji}, define $q(.|\psi_{ji}) \sim \mathcal{N}(\psi_{ji}, \sigma_{ji}^2)$, a Gaussian random walk process, (vii) independence likelihood $L_{ind}(\mathbf{i};\psi)$, as given in Eq. (17).
Output: a Markov chain realisation of length S.
for $t = 2$ **to** S **do**
 for $j = 1$ **to** G **do**
  Maximum independence likelihood estimation with ψ_{-j} held fixed at $\psi_{-j}^{(t)}$, resulting in: $\hat{\psi}_{j,ind}, \widehat{H}_{jj}(\hat{\psi}_{j,ind})$ and $\widehat{J}_{jj}(\hat{\psi}_{j,ind})$;
  Define adjusted $L_{ind}^{adj}(\psi_j|\psi_{-j}^{(t)}, \mathbf{i})$; (see Appendices A.1 and A.2)
  for $i = 1$ **to** $|\psi_j|$ **do**
   $\psi_{ji}^* = \psi_{ji}^{(t)} + \epsilon_{ji}$ with $\epsilon_{ji} \sim \mathcal{N}(0, \sigma_{ji}^2)$;
  end
  $\alpha = \min\left\{1, \frac{\pi(\psi^*)L_{ind}^{adj}(\psi_j^*|\psi_{-j}^{(t)}, \mathbf{i})}{\pi(\psi^{(t)})L_{ind}^{adj}(\psi_j^{(t)}|\psi_{-j}^{(t)}, \mathbf{i})}\right\}$; $\quad \left(\psi_{-j}^{(t)}\right.$: current value of $\left.\psi_{-j}\right)$
  $U \sim \text{Unif}(0,1)$;
  if $\alpha \leqslant U$ **then**
   $\psi_j^{(t+1)} = \psi_j^*$;
  else
   $\psi_j^{(t+1)} = \psi_j^{(t)}$;
  end
 end
end
return $\{\psi^{(t)}\}_{t=1,\ldots,S}$;

References

Beirlant, J., Goegebeur, Y., Segers, J., Teugels, J., 2004. Statistics of Extremes. Wiley Ltd., Sussex.

Bernard, M.M., 1932. Formulas for rainfall intensities of long duration. Trans. Am. Soc. Civ. Eng. 96, 592–624.

Chandler, R.E., Bates, S., 2007. Inference for clustered data using the independence loglikelihood. Biometrika 94, 167–183.

Chavez-Demoulin, V., Embrechts, P., Neslehova, J., 2006. Quantitative models for operational risk: extremes, dependence and aggregation. J. Bank. Financ. 30, 2635–2658.

Chow, V.T., Maidment, D.R., Mays, L.W., 1988. Applied Hydrology. McGraw-Hill, New York.

Coles, S., 2001. An Introduction to Statistical Modeling of Extreme Values. Springer-Verlag, Heidelberg.

Davison, A.C., Padoan, S.A., Ribatet, M., 2012. Statistical modeling of spatial extremes. Statist. Sci. 27, 161–186.

Demarée, G.R., 1985. Intensity–duration–frequency relationship of point precipitation at Uccle. Reference Period 1934-1983. L' Institut Royal Météorologique de Belgique, Publ. Ser. A 116, 52.

Demarée, G.R., 2003. Le pluviographe centenaire du plateau d'Uccle: son histoire, ses données et ses applications. La Houille Blanche 4, 1–8.

García-Bartual, R., Schneider, M., 2001. Estimating maximum expected short-duration rainfall intensities from extreme convective storms. Phys. Chem. Earth Part B 26, 675–681.

Gelman, A., Carlin, J.B., Stern, H.S., Dunson, D.B., Vehtari, A., Rubin, D.B., 2013. Bayesian Data Analysis, third ed. Chapman & Hall/CRC, Texts in Statistical Science.

Heidelberger, P., Welch, P.D., 1983. Simulation run length control in the presence of an initial transient. Oper. Res. 31, 1109–1144.

Hosking, J.R.M., Wallis, J.R., 1997. Regional Frequency Analysis: An Approach Based on L-moments. Cambridge University Press.

Katz, R.W., Parlange, M.B., Naveau, P., 2002. Statistics of extremes in hydrology. Adv. Water Resour. 25, 1287–1304.

Koutsoyiannis, D., Kozonis, D., Manetas, A., 1998. A mathematical framework for studying rainfall intensity-duration-frequency relationships. J. Hydrol. 206, 118–135.

Lehmann, E.A., Phatak, A., Soltyk, S., Chia, J., Lau, R., Palmer, M., 2013. Bayesian hierarchical modelling of rainfall extremes. In: Piantadosi, J., Anderssen, R.S., Boland, J. (Eds.), 20th International Congress on Modelling and Simulation (MODSIM), 2506–2512. Modelling and Simulation Society of Australia and New Zealand, Adelaide, Australia.

Martins, E.S., Stedinger, J.R., 2000. Generalized maximum-likelihood generalized extreme-value quantile estimators for hydrologic data. Water Resour. Res. 36, 737–744.

Mohymont, B., Demarée, G.R., Faka, D.N., 2004. Establishment of IDF-curves for precipitation in the tropical area of central Africa-comparison of techniques and results. Nat. Hazards Earth Syst. Sci. 4, 375–387.

Muller, A., Bacro, J.N., Lang, M., 2008. Bayesian comparison of different rainfall depth-duration-frequency relationships. Stoch. Environ. Res. Risk. Asses. 22, 33–46.

Nadarajah, S., Anderson, C.W., Tawn, J.A., 1998. Ordered multivariate extremes. J. Roy. Statist. Soc. B 60, 473–496.

Overeem, A., Buishand, T.A., Holleman, I., 2008. Rainfall depth-duration-frequency curves and their uncertainties. J. Hydrol. 348, 124–134.

Papalexiou, S.M., Koutsoyiannis, D., 2013. Battle of extreme value distributions: a global survey on extreme daily rainfall. Water Resour. Res. 49, 187–201.

Pauli, F., Racugno, W., Ventura, L., 2011. Bayesian composite marginal likelihoods. Statist. Sinica 21, 149–164.

Plummer, M., Best, N., Cowles, K., Vines, K., 2006. CODA: Convergence diagnosis and output analysis for MCMC. R News 6, 7–11.

Ribatet, M., Cooley, D., Davison, A.C., 2012. Bayesian inference from composite likelihoods, with an application to spatial extremes. Statist. Sinica 22, 813–845.

Rotnitzky, A., Jewell, N., 1990. Hypothesis testing of regression parameters in semiparametric generalized linear models for cluster correlated data. Biometrika 77, 485–497.

Schliep, E.M., Cooley, D., Sain, S.R., Hoeting, J.A., 2010. A comparison study of extreme precipitation from six different regional climate models via spatial hierarchical modeling. Extremes 13, 219–239.

Spiegelhalter, D.J., Best, N.G., Carlin, B.P., van der Linde, A., 2002. Bayesian measures of model complexity and fit. J. Roy. Statist. Soc. B 64, 583–639.

Stephenson, A.G., Ribatet, M., 2012. evdbayes: Bayesian analysis in extreme value theory, R package version 1.1-0.

Van de Vyver, H., 2012. Evolution of extreme temperatures in Belgium since the 1950s. Theor. Appl. Climatol. 107, 113–129.

Van de Vyver, H., 2015. On the estimation of continuous 24-h precipitation maxima. Stoch. Environ. Res. Risk Assess. 29, 653–663.

Varin, C., 2008. On composite marginal likelihoods. AStA-Adv. Stat. Anal. 92, 1–28.

Varin, C., Reid, N., Firth, D., 2011. An overview of composite likelihoods. Statist. Sinica 21, 5–42.

Willems, P., 2000. Compound intensity/duration/frequency-relationships of extreme precipitation for two seasons and two storm types. J. Hydrol. 233, 189–205.