

Bayesian hierarchical modelling of rainfall extremes

E.A. Lehmann^a, A. Phatak^a, S. Soltyk^b, J. Chia^a, R. Lau^a and M. Palmer^c

^a CSIRO Computational Informatics, Perth, WA, AUSTRALIA

^b Curtin University of Technology, Perth, WA, AUSTRALIA

^c 121 Lagoon Dr., Yallingup, WA, AUSTRALIA

E-mail: eric.lehmann@csiro.au

Abstract: Understanding weather and climate extremes is important for assessing, and adapting to, the potential impacts of climate change. The design of hydraulic structures such as dams, drainage and sewers, for instance, relies in part on accurate information regarding patterns of extreme rainfall occurring at various locations and at different durations. Deriving this kind of information is challenging from a statistical viewpoint because a lot of information must be extracted from very little data.

In this paper, we describe the use of a spatial Bayesian hierarchical model (BHM) for characterising rainfall extremes over a region of interest, using historical records of precipitation data from a network of rainfall stations. The rainfall extremes are assumed to have a generalised extreme value (GEV) distribution, with the shape, scale and location parameters representing the underlying variables of the BHM's process layer. These parameters are modelled as a linear regression over spatial covariates (latitude and longitude) with additive spatially-correlated random process. This spatial process leads to more precise estimates of rainfall extremes at gauged locations, and also allows the inference of parameters at ungauged locations. Furthermore, it also mitigates the limitations imposed by short rainfall records in that it allows the model to "borrow strength" from neighbouring sites, thereby reducing the uncertainty at both gauged and ungauged locations. Making use of r -largest order statistics in the data layer further allows the integration of multiple yearly rainfall amounts instead of the annual maximum only.

The proposed BHM uses a parametric representation that links the GEV scale parameter obtained for different accumulation durations. This approach leads to two additional process parameters, and allows the use of pluviometer data accumulated over a range of durations, thereby also increasing the amount of data available for inference. A main advantage of the Bayesian approach is that measures of variability arise naturally from the framework. These uncertainty measures represent information of crucial importance for a subsequent use of the estimated quantities.

We demonstrate this Bayesian approach using a dataset of pluviometer measurements recorded at 252 meteorological stations located on the Central Coast of New South Wales, Australia. For each station, the rainfall data is accumulated over 12 different durations ranging from 5 minutes to 72 hours, from which the two largest annual maxima are selected. Exploratory analyses of this rainfall dataset are carried out for various purposes, including: (i) basic quality control (removal of erroneous data values), and (ii) to provide insight into the relevance of the model structure and associated assumptions.

The proposed model is fitted using Markov chain Monte Carlo (MCMC) simulation, with several types of diagnostics plots used to assess the convergence properties of the resulting chains. We present numerical examples of estimated parameters resulting from the fitted model (regression coefficients, sill and range of the spatial correlation function) together with confidence intervals. Further results from this study are provided by calculating intensity–duration–frequency (IDF) curves for a few sites of interest (both gauged and ungauged) with associated estimates of uncertainty. These results are shown to be in good agreement with station-based maximum likelihood estimates, while achieving smoother curves with tighter uncertainty bands.

Keywords: *Pluviometer data, spatial statistical modelling, generalised extreme value distribution, Markov chain Monte Carlo, intensity–duration–frequency curve.*

1. INTRODUCTION

There is a general consensus that the occurrence, extent and intensity of extreme weather events (e.g., extreme rainfall or temperature, storms, etc.) are increasing and will continue to do so in the foreseeable future because of anthropogenic climate change (IPCC, 2007). This trend is likely to lead to an increase in natural hazards such as heat waves, floods and wildfires, which will have a significant impact on population health, food production, insurance costs, and damage to infrastructure and ecosystems. The need for an accurate and comprehensive understanding, analysis and forecasting of weather extremes and their consequences is therefore unquestionable, and scientific research in this field has recently become a priority for major research institutions and governments. In Australia, for instance, the Australian Rainfall and Runoff guideline document (Pilgrim, 1997), which provides widely used national guidance for the assessment of flood characteristics across the continent, is being updated in light of climate change (www.arr.org.au). This type of information is essential for risk-informed policy decisions and projects involving town planning, mining developments, flood warning and emergency management, operation of regulated river systems, and building of infrastructure such as roads, rail, airports, dams and stormwater systems.

Regional frequency analysis (RFA) is often used to estimate rainfall intensity–duration–frequency (IDF) relationships (Hosking and Wallis, 2005), which represent important inputs into models for impacts assessment. However, it is not always straightforward in RFA to model spatial and temporal variability of rainfall extremes, nor is it straightforward to obtain uncertainty estimates for IDF curves. Alternatives to RFA include spatial models, copulas and max-stable processes (see, e.g., Davison *et al.*, 2012).

Because extreme events are, by definition, rare, analysis of climatic extremes is based on very small amounts of data. Adding to this challenge is the fact that extrapolation of the analysis (forecasting) is typically required for locations with no direct observations. However, extreme weather events can be seen as the result of a spatial process (Banerjee *et al.*, 2004), with underlying climatic and topographic parameters varying smoothly over space and between neighbouring locations. Because it can lead to improved inference when using temporally and spatially sparse datasets, modelling this spatial correlation explicitly is desirable in any approach used to analyse such extreme events. Also, it is important for the inference framework to produce estimates of uncertainty, so as to provide meaningful inputs to subsequent processes relying on such analyses of weather extremes. These requirements call for the development of rigorous models that are both flexible and built upon strong statistical foundations.

Bayesian hierarchical modelling represents an approach that can be used for the implementation of such a flexible framework, and it allows for the integration of multiple sources of uncertainty. Several recent studies make use of BHM for the spatial modelling of extremes in various contexts (monitoring of hurricane winds, rainfall, wildfires, etc.) and various regions of the world (see, e.g., Davison *et al.*, 2012, and references therein). In this work, we focus on the analysis of a dataset of pluviometer observations recorded near Sydney, Australia, as described in Section 2. The basis for the BHM formulation is the Bayesian latent variable model (Schliep *et al.*, 2010; Davison *et al.*, 2012), which we review in Section 3 and subsequently extend to rainfall extremes at different durations. Section 4 provides an overview of some results obtained from the fitted model, such as parameter estimates and examples of IDF curves (at both gauged and ungauged locations), together with corresponding estimates of uncertainty. Finally, Section 5 concludes this paper with an overview of current limitations and future research directions.

2. RAINFALL DATA AND EXPLORATORY ANALYSES

The dataset of rainfall maxima used in this work was extracted from pluviometer records acquired at 252 stations located around the Sydney and Wollongong metropolitan areas in New South Wales, Australia, as shown in Figure 1. The extent of the study area is roughly 160 km by 340 km. Station records consist of rainfall depths (in mm) registered over 5 min intervals, with different record lengths ranging from 7 to 41 years of measurements during the 1959 – 2002 period. The pluviometer data at 5 min intervals were subsequently accumulated over 12 different durations, namely 5, 10, 15 and 30 min, and 1, 2, 3, 6, 12, 24, 48 and 72 hours. The 12 resulting time series were then used to determine the two largest annual rainfall amounts for the corresponding station, year, and accumulation duration. Overall, the dataset contains a total of 3683 years of precipitation maxima (data points) distributed across the 252 weather stations. In Section 4, results are shown for two stations of interest, represented by blue circles in Figure 1: the top circle was selected due to its long record length (41 years) and the bottom circle was chosen due to its geographical proximity and similarity to an (arbitrarily selected) ungauged location (red square in Figure 1).

A number of exploratory analyses were carried out prior to statistical modelling. We first calculated the maximum likelihood estimate (MLE) of generalised extreme value (GEV) distribution parameters (see Sec-

tion 3) at each station, and then used them to investigate: (i) the dependence of these GEV parameters on various covariates (e.g., latitude and longitude); (ii) the spatial correlation between GEV parameters of neighbouring stations; (iii) the dependence between the GEV parameters and accumulation duration; and (iv) the correlation between rainfall maxima at various durations.

The results from these analyses were used to: (i) gain insight into the potential relationships and correlation among the model variables (GEV parameters, covariates, etc.); (ii) identify and then remove erroneous data values (perhaps resulting from technical and/or human errors); (iii) provide MCMC starting values and prior information (see Section 3.3); and (iv) validate some of the assumptions made in formulating the model (e.g., use of covariates, definition of spatial processes and duration-dependent relationships; see Sections 3.2 and 3.3).

3. MODELLING RAINFALL EXTREMES

3.1. Generalised extreme value theory

The GEV distribution is often used to model rainfall extremes (Coles, 2001; Davison *et al.*, 2012). Under certain conditions, the annual maximum rainfall Y of fixed duration can be approximated by a GEV distribution with location $\mu \in (-\infty, +\infty)$, scale $\sigma > 0$, and shape $\xi \in (-\infty, +\infty)$, whose cumulative distribution function is given by:

$$P(Y \leq y) = F(y; \mu, \sigma, \xi) = \exp \left\{ - \left[1 + \xi \cdot \left(\frac{y - \mu}{\sigma} \right) \right]^{-1/\xi} \right\}, \quad (1)$$

where $1 + \xi(y - \mu)/\sigma > 0$. If the r -largest precipitation values are used instead of the annual maximum only, the corresponding density function can be expressed as (Coles, 2001):

$$\text{GEV}_r \left(y^{(1)}, \dots, y^{(r)}; \mu, \sigma, \xi \right) = \exp \left\{ - \left[1 + \xi \cdot \left(\frac{y^{(r)} - \mu}{\sigma} \right) \right]^{-1/\xi} \right\} \cdot \prod_{i=1}^r \frac{1}{\sigma} \left[1 + \xi \cdot \left(\frac{y^{(i)} - \mu}{\sigma} \right) \right]^{-1/\xi} \quad (2)$$

where the notation $y^{(i)}$ is used to denote the i -th largest rainfall amount: $y^{(1)} \geq y^{(2)} \geq \dots \geq y^{(r)}$ (associated observations of $Y^{(1)} \geq Y^{(2)} \geq \dots \geq Y^{(r)}$).

3.2. Re-parameterisation

Koutsoyiannis *et al.* (1998) re-parameterise the location parameter as $\tilde{\mu} = \mu/\sigma$ (sometimes called ‘dispersion’) and show that both $\tilde{\mu}$ and the shape parameter ξ can be reasonably assumed to remain constant for extreme rainfall data accumulated over various durations. Our exploratory analyses indicate that the GEV parameters of most stations satisfy this condition. However, the scale parameter σ depends on the accumulation duration d and Koutsoyiannis *et al.* (1998) suggest using the following relationship to model this dependence:

$$\sigma_d = \frac{\sigma \cdot d}{(d + \theta)^\eta}. \quad (3)$$

Again, exploratory investigations demonstrated that this relationship was adequate given the current dataset of precipitation extremes. This formulation links the scale parameter σ_d for pluviometer data accumulated over several durations. Consequently, in addition to the three standard GEV parameters $\tilde{\mu}$, σ and ξ , this

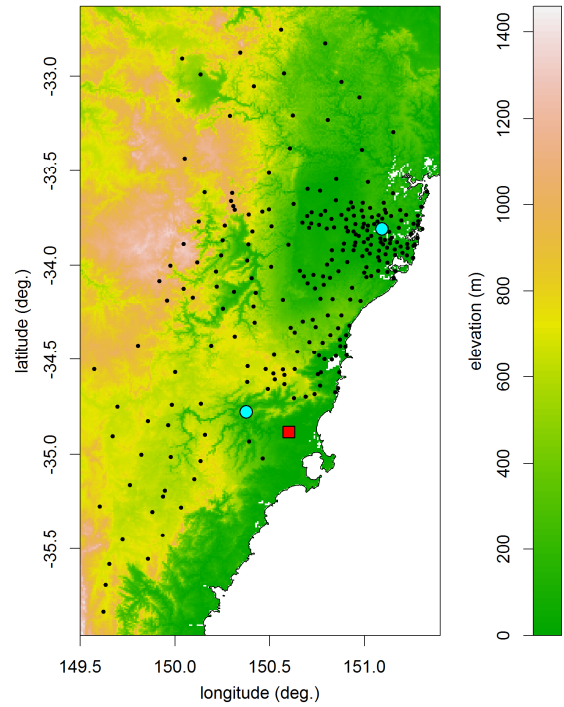


Figure 1. Study area showing the spatial locations of 252 pluviometer stations (black/blue circles). The blue circles and red square (ungauged location) correspond to three sites of interest considered in Section 4.

formulation requires two additional parameters, a “duration offset” term θ and a “duration exponent” term η , which will also need to be estimated within the BHM.

3.3. Spatial Bayesian hierarchical model

In this work, we use the spatial model presented in Davison *et al.* (2012) as a basis for our hierarchical framework. We extend this Bayesian formulation with the re-parameterisation given in Section 3.2 to account for the $D = 12$ available accumulation durations, and by making use of Eq. (2) to allow for $r = 2$ largest extreme rainfall values. The Bayesian hierarchical framework consists of three levels, as described below.

Data model. In this model, the precipitation maxima $Y_{s,t,d}$ for weather station $s = 1, \dots, S$ (with $S = 252$), year $t = 1, \dots, T_s$, and accumulation duration $d = 1, \dots, D$, are assumed independent conditional on the GEV parameters. They are modelled as $Y_{s,t,d} \sim \text{GEV}_r(\tilde{\mu}_s, \sigma_s, \xi_s, \theta_s, \eta_s)$ with $\text{GEV}_r(\cdot)$ denoting the r -largest GEV density given in Eq. (2), and using the relationship between the duration-dependent scale parameter $\sigma_{s,d}$ and the station-specific variables σ_s , θ_s and η_s given in Eq. (3). Note here that the rainfall record length T_s is allowed to vary across stations. As a result of the conditional independence assumption, the data likelihood can be written as:

$$p(\mathbf{Y} | \tilde{\boldsymbol{\mu}}, \boldsymbol{\sigma}, \boldsymbol{\xi}, \boldsymbol{\theta}, \boldsymbol{\eta}) = \prod_{s=1}^S \prod_{t=1}^{T_s} \prod_{d=1}^D \text{GEV}_r(\mathbf{y}_{s,t,d}; \tilde{\mu}_s, \sigma_s, \xi_s, \theta_s, \eta_s) \quad (4)$$

where \mathbf{Y} represents the dataset of all rainfall maxima, the vectors $\tilde{\boldsymbol{\mu}}, \boldsymbol{\sigma}, \boldsymbol{\xi}, \boldsymbol{\theta}$ and $\boldsymbol{\eta}$ contain the respective GEV parameters for each of the $S = 252$ stations, and with $\mathbf{y}_{s,t,d}$ denoting the vector of $r = 2$ largest precipitation maxima at station s for year t and accumulation duration d . For conciseness, the station subscript s will be omitted in the following when referring to the GEV parameters, and the symbol $\chi \in \{\tilde{\mu}, \sigma, \xi, \theta, \eta\}$ (respectively $\boldsymbol{\chi} \in \{\tilde{\boldsymbol{\mu}}, \boldsymbol{\sigma}, \boldsymbol{\xi}, \boldsymbol{\theta}, \boldsymbol{\eta}\}$) will be used to denote any of the GEV parameters.

Process model. Assuming that the GEV parameters vary smoothly over space, the model used in this work imposes a spatial process $\mathcal{P}(\cdot)$ on each of them, as follows:

$$h_{\chi}(\boldsymbol{\chi}) = \mathbf{X}_{\chi}^T \cdot \boldsymbol{\beta}_{\chi} + \mathcal{P}(\boldsymbol{\ell}, \boldsymbol{\alpha}_{\chi}, \lambda_{\chi}) \quad (5)$$

where $h_{\chi}(\cdot)$ is a nonlinear transformation for parameter χ , \mathbf{X}_{χ}^T is a matrix of N_{χ} covariates for each station, and $\boldsymbol{\beta}_{\chi}$ is a vector of corresponding N_{χ} regression coefficients. Since the right-hand side of Eq. (5) represents a real-valued vector, the function $h_{\chi}(\cdot)$ is used here to ensure that the GEV parameters remain within their respective range of values. For instance, using $\log(\boldsymbol{\sigma})$ on the left-hand side of Eq. (5) leads to strictly positive values for $\boldsymbol{\sigma}$, as per Section 3.1. Thus, we define $h_{\sigma}(\cdot) = h_{\theta}(\cdot) = \log(\cdot)$ and $h_{\eta}(\cdot) = \text{logit}(\cdot)$, while no transformation is necessary for the dispersion and shape parameters, i.e., $h_{\chi}(\boldsymbol{\chi}) = \boldsymbol{\chi}$ for $\boldsymbol{\chi} \in \{\tilde{\boldsymbol{\mu}}, \boldsymbol{\xi}\}$.

While there are indications that some GEV parameters might best be described as being constant over space (see, e.g., Koutsoyiannis *et al.*, 1998; Davison *et al.*, 2012, where ξ is modelled as independent of latitude/longitude), our implementation defines the term $\mathbf{X}_{\chi}^T \cdot \boldsymbol{\beta}_{\chi}$ for all GEV parameters as a regression over spatial coordinates with intercept (specifically, our implementation makes use of easting/northing coordinates in km). While this may not lead to the most parsimonious model definition, we essentially let the model itself converge to a state that ultimately indicates which of these regression coefficients are significant. Also, all covariates are first scaled to zero mean and unit variance prior to model simulation, so as to improve the inference (Gelman and Hill, 2007).

In Eq. (5), $\mathcal{P}(\cdot)$ represents a spatially correlated, zero-mean Gaussian random process (multivariate normal) with covariance matrix $\boldsymbol{\Sigma}_{\chi}(\cdot)$ defined on the set of stations’ locations contained in the variable $\boldsymbol{\ell}$:

$$\mathcal{P}(\boldsymbol{\ell}, \boldsymbol{\alpha}_{\chi}, \lambda_{\chi}) \sim \text{MVN}(\mathbf{0}, \boldsymbol{\Sigma}_{\chi}(\boldsymbol{\ell}, \boldsymbol{\alpha}_{\chi}, \lambda_{\chi})). \quad (6)$$

Here, we use the exponential family of correlation functions (Davison *et al.*, 2012) to model the smoothness of the GEV parameters, leading to the following definition of the covariance matrix entries:

$$[\boldsymbol{\Sigma}_{\chi}]_{ij} = \alpha_{\chi} \cdot \exp(-\|\boldsymbol{\ell}_i - \boldsymbol{\ell}_j\|/\lambda_{\chi}), \quad i, j = 1, 2, \dots, S, \quad (7)$$

with α_{χ} and λ_{χ} the sill and range, respectively, of the covariance function, and $\boldsymbol{\ell}_i$ denoting the spatial location of the i -th station. With this formulation, the range parameter λ_{χ} provides an indication of the distance beyond which the spatial correlation between values of a given GEV parameter χ at different spatial locations drops to a negligible level.

Prior distributions. Finally, the Bayesian hierarchical framework requires prior information for the parameters β_χ, α_χ and λ_χ . In accordance with Davison *et al.* (2012), and to reduce computational requirements, we select conjugate Gamma, inverse Gamma and multivariate-normal priors as follows:

$$\lambda_\chi \sim \text{Gamma}(\kappa_{\lambda_\chi}, \gamma_{\lambda_\chi}), \quad \alpha_\chi \sim \text{InvGamma}(\kappa_{\alpha_\chi}, \gamma_{\alpha_\chi}), \quad \beta_\chi \sim \text{MVN}(\boldsymbol{\mu}_{\beta_\chi}, \boldsymbol{\sigma}_{\beta_\chi}^2 \cdot \mathbf{I}_{N_\chi \times N_\chi}), \quad (8)$$

where κ and γ are the shape and scale hyper-parameters of the respective distributions, for which informative priors should be used according to Banerjee *et al.* (2004). In this work, we set these hyper-parameters to values such that the mean (or mode) of the respective prior distribution coincides with maximum likelihood estimates determined from the exploratory analyses, while setting the distributions' variance so as to cover some reasonable range of values around that mean or mode. As for the β_χ priors, we set their means to zero and variances to some relatively large value to obtain uninformative priors.

Bayesian inference. Using the data likelihood, process equations and prior densities defined in Eqs. (4), (5) and (8), the full conditional distributions for the model variables can be derived from the posterior density $p(\boldsymbol{\Omega} | \mathbf{Y})$, where $\boldsymbol{\Omega}$ is the set of all GEV parameters (for all 252 stations), their respective regression coefficients, as well as the sill and range parameters of their covariance functions (representing a total of 1285 parameters to estimate). The full conditionals are then used for inference of the model variables. This is carried out via Markov chain Monte Carlo (MCMC) simulation using standard Gibbs sampling for the α_χ and β_χ parameters (conjugate priors), while Metropolis–Hastings (MH) steps are necessary to sample the GEV parameters and λ_χ variables.

The results presented in Section 4 were obtained from MCMC chains simulated for a total of 150,000 iterations, with the first 20,000 iterations discarded (burn-in) and subsequently thinned by a factor of 35 so as to achieve mostly uncorrelated chains. The step sizes used in the MH sampling steps were set so as to ensure an average acceptance rate between 20 and 30%, as per the “golden acceptance rate” for high-dimensional models (Robert and Casella, 2010). Several types of diagnostics plots were used to assess the convergence properties of the chains, and they indicated that the chains usually converge within the first 5000 MCMC iterations. The model was found to be generally well-behaved, with no indication of identifiability issues, and it also remains essentially unaffected by a well-dispersed selection of the chains' starting values.

4. RESULTS

Figure 2 illustrates a typical result obtained from MCMC fitting of the hierarchical model. It shows the histogram of the MCMC chain (thinned and burn-in removed, resulting in a total of 3715 samples) for the range parameter λ_θ of the duration offset term θ . The blue line in this plot also shows the prior density selected for that parameter. This result is in contrast with previous literature on BHM for spatial extremes (Davison *et al.*, 2012; Sang and Gelfand, 2010) which suggests that it is not possible to learn from the data simultaneously about the sill and range parameters. As shown in Figure 2, the data has clearly informed our model about λ_θ , whose marginal posterior distribution is visibly different from its prior. Similar results are obtained for the range and sill parameters of all GEV variables. This may be because the dataset used in this work is spatially and temporally denser than in previous studies, and also because we combine rainfall extremes across many accumulation durations using the re-parameterisation of Section 3.2, thereby increasing the amount of available information.

Figure 2 further illustrates a major benefit from Bayesian modelling. Since samples are drawn from the entire posterior distribution, we can calculate not only point estimates such as the posterior mean, but also estimates of uncertainty. For example, the thick line at the bottom of the plot in Figure 2 shows the 95% credible interval (CI) for λ_θ , calculated as the 0.025 and 0.975 quantiles of the MCMC samples.

The posterior means of various model parameters and their 95% CIs are provided in Table 1. As discussed in Section 3.3, the regression coefficients associated with easting and northing (β_1 and β_2) do not appear to be significant for the duration-

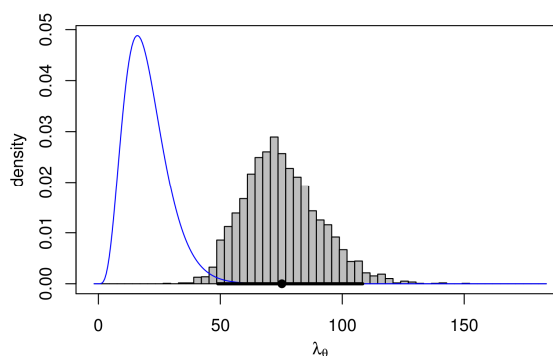


Figure 2. Histogram of MCMC samples for parameter λ_θ (posterior marginal distribution), with superimposed prior density (blue line). The black marker and thick line on the abscissa indicate the posterior sample mean and 95% CI extents, respectively.

Table 1. Bayesian estimates of model variables: regression coefficients β , sill α and range λ for each GEV parameter. Posterior means are in bold font, with 95% CI limits given on either side.

	β_0	β_1 (easting, km)	β_2 (northing, km)	α	λ (km)
$\tilde{\mu}$	2.71, 2.81 , 2.92	-0.22, -0.14 , -0.065	0.064, 0.11 , 0.17	0.028, 0.04 , 0.06	10.3, 16.22 , 24.11
σ	2.12, 2.19 , 2.25	0.13, 0.19 , 0.24	-0.12, -0.081 , -0.041	0.015, 0.022 , 0.031	9.52, 15.09 , 23.24
ξ	0.11, 0.12 , 0.13	0.017, 0.031 , 0.045	-0.027, -0.013 , 0.00084	0.0049, 0.0063 , 0.0082	0.29, 1.52 , 2.9
θ	-3.67, -2.89 , -2.09	-0.12, 0.25 , 0.62	-0.4, -0.058 , 0.27	0.36, 0.51 , 0.81	50.65, 75.61 , 105.46
η	0.42, 0.65 , 0.86	-0.14, -0.025 , 0.088	-0.041, 0.072 , 0.18	0.037, 0.063 , 0.11	34.51, 59.8 , 98.35

related GEV parameters η and θ . Values for the range parameter λ shown in Table 1 point to a moderate correlation distance for $\tilde{\mu}$ and σ , large correlation distances for θ and η , and minimal spatial correlation for the shape parameter ξ . These results are consistent with our exploratory analyses, especially for ξ , whose sample variogram was found to be essentially flat for any distance between about 6 and 120 km.

IDF curves are necessary inputs to impacts assessment, and they are computed as an extreme quantile of the (fitted) GEV distribution (Coles, 2001). The IDF curves in Figure 3 represent the precipitation intensity expected to be exceeded on average once in 100 years, for rainfall events recorded over various durations. In the left-hand plot, the black curve was obtained from the BHM-fitted GEV parameters at a station in Sydney (top blue circle in Figure 1). The red markers and associated error bars show the IDF curve and 95% confidence intervals based on MLEs of the GEV parameters (i.e., not from the BHM) using only that station’s 41 years of precipitation maxima. The BHM-based IDF curve lies within the MLE confidence interval, and it has a smaller uncertainty as the Bayesian inference draws on information from neighbouring stations. Moreover, it is a smoother curve due to the parametric relationship between scale and duration in Eq. (3).

In the plot on the right-hand side of Figure 3, the red line shows the IDF curve obtained for an ungauged location (red square in Figure 1), corresponding to the location of the town of Nowra on the South Coast of New South Wales. This was obtained by simulating from the spatial process to compute the posterior predictive distribution of the GEV parameters at that ungauged location. This illustrates the ability of the spatial model to make inferences at locations where no observations are available. For comparison, the black line shows the IDF curve at a nearby gauged location (blue circle at the bottom of Figure 1). Due to the spatial proximity, these two IDF curves are very similar, with the uncertainty being higher for the ungauged location, as would be expected at a location with no precipitation measurements.

5. DISCUSSION AND CONCLUSIONS

The results provided in this paper demonstrate that spatial Bayesian modelling is a powerful and flexible tool to investigate the characteristics of extreme weather events. Among others, the model described in this work

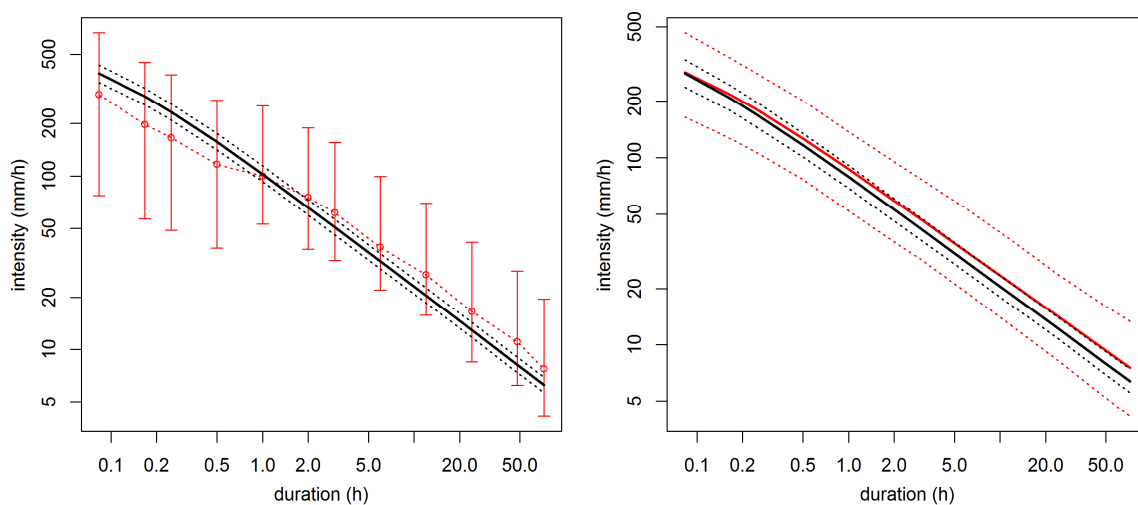


Figure 3. IDF curves for 100-year return period. *Left:* BHM-based IDF curve with 95% CI (black) and duration-by-duration MLE estimates with 95% confidence intervals (red), at a gauged location in Sydney (top blue circle in Figure 1). *Right:* BHM-based IDF curves with 95% CI for an ungauged location (red) and nearby gauged location (black), shown as the red square and bottom blue circle in Figure 1, respectively.

allows: (i) integration of r -largest pluviometer maxima accumulated over various durations; (ii) making the most of the limited amount of available observations, effectively pooling them together through the use of a spatial process (especially important for a stable estimate of the shape parameter ξ); (iii) drawing on the spatial model to make predictions at ungauged locations; and (iv) providing measures of uncertainty for all inferred quantities of interest. The model presented here has proved very robust and powerful in estimating a large number of variables (1285 parameters) from a relatively limited number of observations (3683 years of rainfall maxima). For instance, other authors (e.g., Cooley and Sain, 2010) make use of a penalisation term on ξ in the likelihood formulation to restrict the range of values for that parameter to within -0.5 and 0.5 , but this was not found to be necessary with our model.

The BHM that we describe here incorporates several simplifications. For instance, the current formulation assumes that all GEV parameters are stationary over the time span of pluviometer records (1959 – 2002). Future developments of the model will consider the introduction of time-varying parameters within the framework. Also, considering the spatial and topographical characteristics of the current dataset (see Figure 1), incorporating additional covariates such as elevation (height above sea level), distance to coastline and aspect (slope orientation) may provide additional information and improve the inference. As indicated by exploratory analyses, higher-order terms in the regressions could also improve the model fit of the existing covariates, and this will also be considered in future developments of this framework.

The approach used in this work is useful for investigating the characteristics of extremes derived from fitted marginal distributions, such as return levels and IDF curves. However, because of the assumption of conditional independence, it cannot be used to simulate realistic extreme rainfall surfaces (Sang and Gelfand, 2010) for the purposes of calculating depth-area curves; extending the method to do so is the subject of current research. Future research endeavours will also include: (i) combining precipitation maxima from both pluviometer and daily data (non-recording stations); (ii) integrating climate drivers that may have an impact on future rainfall extremes; and (iii) validating and comparing the BHM to other approaches.

ACKNOWLEDGEMENT

We gratefully acknowledge financial support from the Australian Government through Geoscience Australia, and the substantial in-kind support provided by the members of Engineers Australia.

REFERENCES

- Banerjee, S., Carlin, B.P. and Gelfand, A.E. (2004). *Hierarchical Modeling and Analysis for Spatial Data*. Chapman & Hall, CRC, New York.
- Coles, S. (2001). *An Introduction to Statistical Modeling of Extreme Values*. Springer-Verlag, London.
- Cooley, D. and Sain, S.R. (2010). Spatial hierarchical modeling of precipitation extremes from a regional climate model. *Journal of Agricultural, Biological, and Environmental Statistics*, 15(3):381–402.
- Davison, A.C., Padoan, S.A. and Ribatet, M. (2012). Statistical modeling of spatial extremes. *Statistical Science*, 27(2):161–186.
- Gelman, A. and Hill, J. (2007). *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Cambridge University Press, Cambridge.
- Hosking, J.R. and Wallis, J.R. (2005). *Regional Frequency Analysis: An Approach Based on L-Moments*. Cambridge University Press, Cambridge.
- Intergovernmental Panel on Climate Change (IPCC) (2007). *Climate Change 2007: Impacts, Adaptation and Vulnerability*. M.L. Parry, O.F. Canziani, J.P. Palutikof, P.J. van der Linden and C.E. Hanson, Eds., Cambridge University Press, Cambridge, UK, 976 pp.
- Koutsoyiannis, D., Kozonis, D. and Manetas, A. (1998). A mathematical framework for studying rainfall intensity–duration–frequency relationships. *Journal of Hydrology*, 206:118–135.
- Pilgrim, D.H. (1997). *Australian Rainfall & Runoff – A Guide to Flood Estimation*. Institution of Engineers, Australia, Barton, ACT, Australia.
- Robert, C.R. and Casella, G. (2010). *Introducing Monte Carlo Methods with R*. Springer, New York.
- Roberts, G., Gelman, A. and Gilks, W. (1997). Weak convergence and optimal scaling of random walk Metropolis algorithms. *Annals of Applied Probability*, 7:110–120.
- Sang, H. and Gelfand, A.E. (2010). Continuous spatial process models for spatial extreme values. *Journal of Agricultural, Biological, and Environmental Statistics*, 15(1):49–65.
- Schliep, E., Cooley, D., Sain, S. and Hoeting, J. (2010). A comparison study of extreme precipitation from six different regional climate models via spatial hierarchical modeling. *Extremes*, 13:219–239.