

A hierarchical Bayesian non-asymptotic extreme value model for spatial data

Federica Stolf | Antonio Canale

Department of Statistical Sciences,
University of Padova, Padova, Italy

Correspondence

Federica Stolf, Department of Statistical
Sciences, University of Padova, Padova,
Italy.

Email: federica.stolf@phd.unipd.it

Funding information

Fondazione Cassa di Risparmio di Padova
e Rovigo

Abstract

Spatial maps of extreme precipitation are crucial in flood prevention. With the aim of producing maps of precipitation return levels, we propose a novel approach to model a collection of spatially distributed time series where the asymptotic assumption, typical of the traditional extreme value theory, is relaxed. We introduce a Bayesian hierarchical model that accounts for the possible underlying variability in the distribution of event magnitudes and occurrences, which are described through latent temporal and spatial processes. Spatial dependence is characterized by geographical covariates and effects not fully described by the covariates are captured by spatial structure in the hierarchies. The performance of the approach is illustrated through simulation studies and an application to daily rainfall extremes across North Carolina (USA). The results show that we significantly reduce the estimation uncertainty with respect to state of the art techniques.

KEYWORDS

Bayesian hierarchical models, extreme value theory, rainfall extremes, return levels, spatial and temporal processes

1 | INTRODUCTION

Extreme value theory deals with stochastic behavior of extreme events, found in the tails of probability distributions, and it finds wide application in environmental sciences. Events such as extreme precipitation and storm wind speed are driven by complex spatio-temporal processes and are usually characterized by limited predictability. Understanding frequency and intensity of these phenomena is important in public-safety and long-term planning. Estimating the probability of extreme meteorological events is difficult because of limited temporal records and this issue is exacerbated in spatial settings, because forecasting entails the extrapolation in a large spatial domain.

A variety of statistical tools have been used for modeling extreme value and the book by Coles (2001) provides a gentle introduction to the topic. Traditional methods are based on the block maxima approach exploiting the generalized extreme value distribution (GEV) or on the peaks over threshold (POT) approach exploiting the generalized Pareto distribution (GPD). The GEV is a three-parameter family of distributions that describes the asymptotic behavior of suitably renormalized block maxima of a sequence of independent and identically distributed random variables. The shape parameter of this distribution plays the crucial role of determining the weight of the upper tail of the density. The GPD, used under a POT approach, is a two-parameter family of distributions that is used to model excesses over a suitably chosen

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2023 The Authors. *Environmetrics* published by John Wiley & Sons Ltd.

high threshold. The GEV and the GPD distributions are deeply connected through a point process characterization (Davison & Smith, 1990; Smith, 1989) and in most practical applications they are routinely employed regardless the suitability of their asymptotic arguments.

In the spatial context several approaches for extreme values have been proposed in the literature. Prediction at unobserved locations may be required, which can only be achieved with a proper spatial model. A natural generalization of GEV and GPD distributions to a multivariate context are respectively the max-stable processes (Padoan et al., 2010) and the r -Pareto processes (Fondeville de & Davison, 2018; Thibaud & Opitz, 2015). The max-stability or threshold-stability properties provide a strong theoretical justification for these methods, which have been used in a wide variety of environmental applications. However, these asymptotic models have a strong tail dependence that persists at increasingly high levels, while data usually suggest that it should weaken instead (Huser & Wadsworth, 2022). The max-stable and Pareto processes are always asymptotically dependent, i.e., extreme events have a positive probability of occurring simultaneously at distinct sites, no matter how extreme they are. This is a major limitation in practice, especially for environmental applications, where the spatial dependence tends to vanish for increasing quantile levels (Huser & Wadsworth, 2019). A good revision of the classical models of spatial extreme is provided by Davison et al. (2012). They highlight three main classes of models: models built on max-stable processes (Padoan et al., 2010; Reich & Shaby, 2012), copula models (Sang & Gelfand, 2010) and Bayesian hierarchical models (Cooley et al., 2007; Sang & Gelfand, 2008). For a discussion about more recent progress in the modeling and inference for spatial extremes see Huser and Wadsworth (2022).

To overcome the rigid asymptotic dependence structures of the max-stable and Pareto processes, some contributions have focused on developing asymptotic independence and more flexible hybrid models that can bridge the two asymptotic dependence regimes (Bopp et al., 2021; Huser & Wadsworth, 2019; Opitz, 2016). Alternatively, the conditional spatial extreme approach aims at describing the spatial behavior of a random process conditional on single points being large (Shooter et al., 2019; Wadsworth & Tawn, 2022).

Another limitation of GEV-based approaches is that the support of the GEV distribution depends on its parameter values. This feature can complicate inferential procedures, such as maximum likelihood estimation, and can be particularly problematic in a setting with multiple and interacting covariates. For example, Castro-Camilo et al. (2022) propose to substitute the GEV distribution with the blended generalized extreme value (bGEV), which has the right tail of a Fréchet distribution and the left tail of a Gumbel distribution, resulting in a heavy-tailed distribution with a parameter free support.

The wide popularity of the asymptotic methods led much of the extreme value literature to focus only on a small portion of data: the block maxima or few values above a threshold. The “ordinary” values from which the maxima are extracted are discarded, wasting much of the available information. In addition, in many environmental applications the number of yearly events is not sufficiently large for the asymptotic argument to hold as discussed, for example, by Koutsoyiannis (2004) and Marani and Ignaccolo (2015) in hydrology. Another limitation of traditional extreme value theory is the assumption of a constant distribution for the ordinary events over time, since many phenomena display changes in the event magnitude generation process that are well established (Marani & Ignaccolo, 2015). Zorzetto et al. (2020) introduced a Bayesian hierarchical model for extreme values building upon the so called metastatistical extreme value distribution (MEVD) (Marani & Ignaccolo, 2015; Marra et al., 2018; Marra et al., 2023). This is a non-asymptotic extreme value approach in which a compound parametric distribution describes the entire range of ordinary values, with parameters varying across blocks.

In this paper, we introduce a spatial hierarchical Bayesian model to analyze extreme values of rainfall intensity. Bayesian hierarchical models for extremes value modeling represent an active area of research (Bracken et al., 2018; Ghosh & Mallick, 2011; Zorzetto et al., 2020). One of the main advantages of employing a Bayesian approach in this context is the possibility of incorporating expert domain knowledge in the form of informative prior distributions. This is dramatically relevant in a field like extreme value analysis where observations are by nature scarce and particularly in environmental studies where reliable expert prior information about the geophysical processes at hand is often available. The proposed hierarchical Bayesian model focuses on modeling ordinary events and then estimating the cumulative distribution function (cdf) of maxima, making it intrinsically different from the models inspired by the GEV or max-stable processes. We refer to the proposed approach as a *non-asymptotic* model, in the sense that it does not assume a divergent number of events in a block or use the approximations of the cdf of block maxima, as the GEV. Our proposal does not characterize extremal dependence at data level, so in the sense of Huser and Wadsworth (2022) can be considered asymptotically independent, but it is capable of assessing marginal risks.

In Section 2, we introduce the general structure of the spatial hierarchical model and subsequently specialize it to the analysis of rainfall data. In Section 3, the proposed formulation is tested through an extensive simulation study. An

application to maximum rainfall data is described in Section 4. The paper ends with a final discussion. Code and data are available at <https://github.com/federicastolf/sHMEV>.

2 | A SPATIAL HIERARCHICAL BAYESIAN EXTREME VALUE MODEL

2.1 | Notation and general formulation

Let $x_{ij}(s)$ denote the magnitude of the i -th event within the j -th block for site s , where $j = 1, \dots, J$ with J the number of blocks in the observed sample, $i = 1, \dots, n_j(s)$ with $n_j(s)$ the number of events observed within the j -th block and $s \in \mathcal{A}$ with \mathcal{A} a spatial domain. Let S denote the total number of spatial points at which the data are observed within the spatial domain of interest \mathcal{A} , i.e., $S = |\mathcal{A}|$. Let $Y_j(s) = \max_i \{X_{ij}(s)\}$ denote the block maxima for each site s . As standard practice with environmental data, blocks have size of one year (Coles, 2001), thus by block maxima we refer to annual maxima. The main goal of extreme value analysis can be summarized in estimating the cumulative distribution function of $Y_j(s)$. In the classical extreme value framework, GEV distribution arises as the only possible limit model for block maxima. Specifically, if there exist sequences of constants $a_j(s) > 0$ and $b_j(s) \in \mathbb{R}$ such that, for $Y_j(s)^* = (Y_j(s) - b_j(s))/a_j(s)$, one has $\Pr(Y_j(s)^* \leq y) \rightarrow G(y)$ as $n_j(s) \rightarrow \infty$, where G is a non-degenerate distribution function, then G has the form

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

with $a_+ = \max(a, 0)$, $\mu \in \mathbb{R}$ (location), $\sigma \in \mathbb{R}^+$ (scale) and $\tau \in \mathbb{R}$ (shape), with support $\{y \in \mathbb{R} : 1 + \tau(y - \mu)/\sigma > 0\}$. This approach assumes that the limit GEV distribution is a good approximation for block maxima in finite samples.

Instead of relying on asymptotic arguments, we will exploit a fully generative hierarchical model. We focus on the ‘‘ordinary’’ values from which the maxima are extracted, in our application the daily records with non-zero rainfall, and we model the magnitude $x_{ij}(s)$ and the occurrence $n_j(s)$ of these events. Notably, the proposed approach exploits the information contained in all the sample of ordinary values and not only the one contained in the block maxima, or in the values above a pre-specified threshold, as done in the block maxima or peak over threshold approaches, respectively.

Specifically, we assume that $x_{ij}(s)$, conditionally on unobserved latent processes, are realizations of conditionally independent random variables $X_{ij}(s)$. We further assume that each $X_{ij}(s)$ has a location and time-specific distribution, but with all distributions belonging to the same parametric family with cdf $F(\cdot; \theta_j(s))$, with $\theta_j(s) \in \Theta$ unknown parameter vector of suitable dimension. We denote with $f(\cdot; \theta_j(s))$ the related probability density function. The main goal is to estimate the cdf of the block maxima:

$$\Pr(Y_j(s) \leq y) = F(y; \theta_j(s))^{n_j(s)}. \quad (2)$$

The following section specifies this general notation focusing on the case when $x_{ij}(s)$ are daily rainfall accumulations.

We assume that the response variables are influenced by both temporal and spatial underlying processes. The base layer models the magnitude of events at each spatial location for each block. The nested layer subsequently models the latent processes that drive the parameters of the external layer, both in time and space. Using a Bayesian approach, the last layer is associated with the prior distributions of the unknown parameters that govern the inner latent processes. The structure of the model is illustrated in Figure 1.

Let $n_j(s)$ be a realization of a stochastic process with conditional probability function $p\{\cdot; \lambda(s)\}$, where $\lambda(s)$ is an unknown parameter vector depending on the spatial location $s \in \mathcal{A}$. We further assume that $\theta_j(s)$ are realizations of a stochastic process with conditional probability density function $g\{\cdot; \eta(s)\}$, where $\eta(s)$ is an unknown vector of parameters. The proposed Bayesian spatial hierarchical model for extreme values (sHMEV) can be written in hierarchical form, for each $s \in \mathcal{A}$, as

$$\begin{aligned} x_{ij}(s) | n_j(s), \theta_j(s) &\sim f\{x_{ij}(s); \theta_j(s)\}, \quad \text{for } i = 1, \dots, n_j(s), \\ \theta_j(s) | \eta(s) &\sim g\{\theta_j(s); \eta(s)\}, \quad n_j(s) | \lambda(s) \sim p\{n_j(s); \lambda(s)\}. \end{aligned} \quad (3)$$

The latent spatial processes, $\eta(s)$ and $\lambda(s)$, can be driven by unknown parameters β_η and β_λ in a stochastic manner as

$$\lambda(s) | \beta_\lambda \sim k\{\lambda(s); \beta_\lambda\}, \quad \eta(s) | \beta_\eta \sim m\{\eta(s); \beta_\eta\}, \quad (4)$$




FIGURE 1 Graphical representation of the spatial hierarchical model described in (3) and (4).

where $k(\cdot; \beta_\lambda)$ and $m(\cdot; \beta_\eta)$ are suitable probability density functions. A simplified model assumes for $\eta(s)$ and $\lambda(s)$ simple deterministic functions of parameters β_η and β_λ and spatial covariates $z(s)$. The Bayesian representation of the model is completed by eliciting suitable prior distributions for the unknown parameters β_η and β_λ .

Having as target the cdf in (2) or one of its functionals, the estimation can be done in our setting marginalizing out the variables $n_j(s)$ and $\theta_j(s)$, that is,

$$\zeta\{y; \lambda(s), \eta(s)\} = \sum_{n \geq 0} \int_{\Theta} F\{y; \theta\}^n p\{n; \lambda(s)\} g\{\theta; \eta(s)\} d\theta. \tag{5}$$

2.2 | A specific formulation for modeling daily rainfall

Hereafter, we specify the approach described in the previous section to the case of annual maxima daily rainfall accumulations over an area of interest. Several parametric families have been employed to model rainfall accumulations and while most of the approaches are based on goodness of fit considerations, ours exploits physical knowledge of the phenomena. It is well established that precipitation is a heavy-tailed phenomenon (Katz et al., 2002) and Stechmann and Neelin (2014) suggest that the distribution of daily rainfall should follow a Gamma distribution, while Wilson and Toumi (2005) and Porporato et al. (2006) suggest that its right tail should decay as a stretched exponential (i.e., Weibull) distribution. Consistently with these arguments, we believe that the Weibull distribution is adequate for modeling the magnitudes of daily rainfall accumulations. In other applications, however, it is important to exercise prudence in selecting the parent distribution for event magnitudes, as this constitutes a critical component of our methodology. It is imperative to note that inaccuracies in this regard could potentially have an adverse impact on the final results.

Thus, we assume that $x_{ij}(s)$ (for $i = 1, \dots, n_j(s)$, $j = 1, \dots, J$ and $s = 1, \dots, S$) follow a Weibull distribution with parameter vector $\theta_j(s) = (\gamma_j(s), \delta_j(s))$, where $\delta_j(s) > 0$ denotes the scale parameter and $\gamma_j(s) > 0$ the shape parameter. To allow for the inter-block variability we assume that the latent variables $\gamma_j(s) \sim g_\gamma\{\gamma_j(s); \mu_\gamma(s), \sigma_\gamma\}$ and $\delta_j(s) \sim g_\delta\{\delta_j(s); \mu_\delta(s), \sigma_\delta\}$ are independent and follow a Gumbel distribution, a flexible yet parsimonious model allowing for possible asymmetry. We further assume that only the location parameters of the two Gumbel distributions are characterized by a spatial dependence and not their scale parameters, but more flexible alternatives are straightforward. For the location processes $\mu_\gamma(s)$ and $\mu_\delta(s)$ we assume a linear dependence from fixed spatial covariates $z(s)$, i.e.,

$$\mu_\gamma(s) = z(s)\beta_\gamma, \quad \mu_\delta(s) = z(s)\beta_\delta \tag{6}$$

where $\beta_\gamma = [\beta_{\gamma,0} \ \beta_{\gamma,1} \ \dots \ \beta_{\gamma,p}]^T$ and $\beta_\delta = [\beta_{\delta,0} \ \beta_{\delta,1} \ \dots \ \beta_{\delta,p}]^T$ and p is the number of available spatial covariates. The vectors β_γ and β_δ are independent and, in principle, we could also assume that they are of different size and related to different subsets of variables in $z(s)$. Although we focus on simple linear relations between the location processes and the spatial covariates, extensions to more flexible modeling structures are also straightforward, for example, adopting

basis expansions. The idea underlying this specification is that observations are characterized by a latent process defined by geographic covariates, while the residual spatial variability, not captured by covariates, is described by the random Gumbel variability.

For $n_j(s)$ we assume a binomial distribution, with a success probability $\lambda(s)$ and number of trials N_t equal to the block size (e.g., $N_t = 366$ in our application to annual maximum daily rainfall). The decision to employ a simple parametric model is supported by the consideration that the distribution of n_j mainly affects the distribution of extreme events only through its average value (Hosseini et al., 2020; Marra et al., 2019). We assume that the rainfall occurrence is also affected by the geographical characteristics of the sites, as previously done for the location parameters of the Gumbel distributions. Consistently with (6) we let

$$\text{logit}(\lambda(s)) = z(s)\beta_\lambda, \quad (7)$$

where the function $\text{logit}(x) = \log(x/(1-x))$ is employed to ensure $\lambda(s) \in (0, 1)$. Specific examples of (6) and (7) are discussed in Section 4.

We have assumed that the distributions of $n_j(s)$ and $x_{ij}(s)$ are independent, although one might expect a possible dependence between the number of the events that occur within a block and their distribution. Indeed, precipitation is composed of a mixture of high-intensity events that occur over short-time scales and low-intensity, long-time scale events, and there may be a negative relationship between the occurrence and the event distribution. It is important to note, however, that our data are daily and hence we are summing up all the short-term events occurring on the same day, thus partially losing the details of the most intense convective rainfall.

2.3 | Prior elicitation and posterior computation

The introduction of a hierarchical model that describes the entire distribution of daily rainfall accumulations allows to specify priors directly on the underlying distribution of the observed ordinary events, rather than on the distribution of block maxima. This is a great advantage, because it avoids the difficulty of prescribing a prior directly on the shape parameter ζ of the GEV distribution, to which it is difficult to attribute physical meaning.

In defining the prior distributions for the parameters σ_γ , σ_δ , β_δ , β_γ , and β_λ , we seek to harness information on the physical processes generating the data, avoiding, where possible, strongly uninformative priors. For example, Frisch and Sornette (1997) and Sornette (2006) provide geophysical motivations to expect the Weibull shape parameters $\gamma(s)$ to be centered around 2/3 for rain accumulation. Consistently with this, we fix the prior expectation of $\beta_{\gamma,0}$ to 2/3. In the present study, selecting a highly informative prior for $\beta_{\gamma,0}$ was predicated on a sound basis of expert domain knowledge. However, it is worth noting that such a justification may not be applicable in all contexts. In recognition of this potential limitation, we have undertaken a sensitivity analysis to examine the robustness of the posterior distribution of $\beta_{\gamma,0}$ in relation to the choice of hyperparameters for the prior. The analysis showed that results are robust to moderate and reasonable changes in the hyperparameter of the prior for $\beta_{\gamma,0}$. The remaining terms in the vector β_γ are assumed having null prior expectations, as the covariates have been standardized. As parametric distribution for each vector β_δ , β_γ and β_λ we chose independent Gaussian distributions. For the latent Gumbel scale parameters σ_γ and σ_δ , quantifying the between-block variability of the Weibull parameters, we opt for independent inverse Gamma distributions, with expectations equal to 25% and 5% of the respective location parameters (μ_δ and μ_γ). This choice reflects the expectation that the scale parameter varies across years more than the shape parameter, as its expected values should be more constrained by the behavior of the physical process of precipitation. For the choice of the hyperparameters of the distributions we adopt an empirical Bayesian approach (Casella, 1985) and a practical example is illustrated in Section 4.

Given the complex structure of the introduced model, the posterior distribution of the parameters is not available in closed form. Posterior approximation is then obtained using Markov Chain Monte Carlo (MCMC) methods. Specifically, we use the *Hamiltonian Monte Carlo* (HMC) approach (Betancourt, 2018) exploiting the flexibility of the Stan software (Carpenter et al., 2017).

In all the following examples, we run $n_c = 4$ parallel chains, with $n_g = 2000$ iterations in each chain, starting from different initial points. We discard the first half of each chain to account for the burn-in effect. Therefore, the final sample on which we perform inference is based on $B = n_c n_g / 2 = 4000$ draws, obtained by merging the draws from different chains, following standard practice (Gelman et al., 2013).

We can obtain an estimation of (5), the cumulative probability of block maxima approximating at each location $s \in \mathcal{A}$, with a two step procedure. We first compute the value of $\hat{\zeta}_s^{(b)}(y)$ at the generic b iteration for site s as

$$\hat{\zeta}_s^{(b)}(y) = \frac{1}{M_g} \sum_{j=1}^{M_g} F\{y; \theta_j^{(b)}(s)\}^{n_j^{(b)}(s)}, \quad (8)$$

where $\theta_j^{(b)}(s) = (\gamma_j^{(b)}(s), \delta_j^{(b)}(s))$ and $n_j^{(b)}(s)$, for $j = 1, \dots, M_g$, are drawn from the related posterior predictive distribution for site s in the j -th over M_g future blocks. Then, averaging over the B draws from the posterior distribution for each site s we have

$$\hat{\zeta}_s(y) = \frac{1}{B} \sum_{b=1}^B \hat{\zeta}_s^{(b)}(y). \quad (9)$$

3 | SIMULATION STUDY

3.1 | Description

To assess the empirical performance of the proposed model we perform a simulation study. A data set composed of $S = 27$ sites has been constructed, simulating uniformly inside a square of unitary side. Taking as reference point the bottom left edge of the square, the values on the x-axis and the y-axis represent the spatial coordinates, z_1 and z_2 .

We simulated the ordinary events and extract the annual maxima from them, since we believe that this would be more realistic than simulating the maxima directly from an asymptotic approximation, like the GEV distribution. Different synthetic data sets have been generated under three scenarios characterized by specific event magnitude distributions. In the first scenario (WEI) we assume a Weibull model in which the scale and shape parameters in each block follow Gumbel distributions. The location parameters of the Gumbel distributions are determined through a spatial trend, defined as:

$$t(s) = \beta_0 + \beta_1 z_1(s) + \beta_2 z_2(s), \quad (10)$$

with $\beta_0, \beta_1, \beta_2 \in \mathbb{R}$, and $s \in \mathcal{A}$. In the second scenario (WEI_{gp}) the model builds upon the previous one and in order to give some extra variability to the data, we add a Gaussian process with exponential correlation function to the regression trend defined in (10). Thus the location parameters of the Gumbel distributions are defined through the function

$$t(s) = \beta_0 + \beta_1 z_1(s) + \beta_2 z_2(s) + GP(s; \alpha, \nu), \quad (11)$$

where $GP(\cdot)$ is a zero mean Gaussian process with covariance function $\alpha \exp(-\|d\|/\nu)$ with parameters α and ν , while d denotes the pairwise Euclidean distance between locations. In the third specification (GM) we assume for the event magnitude a Gamma distribution, where the two parameters are generated by a spatial trend defined accordingly to Equation (10).

The number of events in each block is drawn from a binomial distribution with number of trials $N_t = 366$ and success probability determined through a logit transformation of the spatial trend defined in (10) with constant value of β_0, β_1 , and β_2 for all the three scenarios. Notably, only the first scenario reflects the structure of the proposed spatial hierarchical model. This will be used to assess the robustness of the proposed formulation under model misspecification.

Common to all scenarios, two independent data sets have been generated under the same specifications: the first one, which contains time series of length 20 years, will be used as the training set, while the second one, containing time series of length 100 years, will be used as the test set. Using limited temporal records for the estimation is representative of many real geophysical data sets.

The sHMEV model assumes, under correct specification of (10),

$$\begin{aligned} \mu_\gamma(s) &= \beta_{\gamma,0} + \beta_{\gamma,1} z_1(s) + \beta_{\gamma,2} z_2(s), & \mu_\delta(s) &= \beta_{\delta,0} + \beta_{\delta,1} z_1(s) + \beta_{\delta,2} z_2(s), \\ \text{logit}(\lambda(s)) &= \beta_{\lambda,0} + \beta_{\lambda,1} z_1(s) + \beta_{\lambda,2} z_2(s), \end{aligned}$$

where $z_1(s)$ and $z_2(s)$ are the spatial coordinates for site s , with $s \in \mathcal{A}$.

3.2 | Performance measures

To evaluate the performance of the proposed spatial hierarchical model we compare it with standard alternative methods. In particular, the methods used to benchmark sHMEV are Bayesian implementations of the classical generalized extreme value distribution (GEV) and the Bayesian hierarchical model (HMEV) described in Zorzetto et al. (2020). Between the competing models we did not include the peaks over threshold approach because the predictive results of the latter model were similar to those of GEV.

For both the models used to benchmark sHMEV we take informative priors. More precisely, for the HMEV model we follow the specification of Zorzetto et al. (2020), while for the GEV model the prior distribution for the shape parameter is centered around the value 0.114 and has a standard deviation of 0.125; these values have been determined from investigations of rainfall records at the global scale (Papalexiou & Koutsoyiannis, 2013).

To compare the different methods, we evaluate the predictive accuracy in estimating the distribution of block maxima on the test set. We introduce different criteria to measure the predictive performance of the methods for quantiles above a given non exceedance probability, that are computed marginally for each station.

Specifically, we consider the fractional squared error (FSE) defined by

$$\frac{1}{m_T} \sum_{j=1}^{M_x} \mathbb{1}_{(\tilde{T}, \infty)}\{T_{js}\} \sqrt{\frac{1}{B} \sum_{b=1}^B \left(\frac{\hat{\zeta}_s^{(b)^{-1}}(p_{js}) - y_{js}}{y_{js}} \right)^2}, \quad (12)$$

where $\mathbb{1}_A(x)$ is the indicator function that equals 1 if x belongs to A , $\hat{\zeta}_s^{(b)^{-1}}(\cdot)$ is the quantile function of the specific model at the b -th MCMC iteration for site s , T_{js} is the empirical return time of y_{js} defined as $T_{js} = (1 - p_{js})^{-1}$, with $p_{js} = \text{rank}(y_{js}) / (M_x + 1)$ and M_x is the number of blocks used to compute the FSE. In this section M_x corresponds to the number of blocks in the test set, that is, $M_x = 100$. The value m_T represents the number of observations in the test set with empirical return time equal to or larger than \tilde{T} , i.e., $m_T = \sum_{j=1}^{M_x} \mathbb{1}_{(\tilde{T}, \infty)}(T_{js})$. The FSE index represents an average measure of a standardized distance between model-estimated quantiles and empirical quantiles for return times larger than \tilde{T} . In the following analysis we define $\tilde{T} = 2$, consistent with the range of exceedance probability of interest in many practical applications.

To separately assess the precision and the variability of extreme value quantile estimates, we employ two additional measures, namely the average bias and the average width of the 90% posterior predictive credible intervals, defined respectively as

$$b_q = \frac{1}{m_T} \sum_{j=1}^{M_x} \mathbb{1}_{(\tilde{T}, \infty)}\{T_{js}\} \frac{1}{B} \sum_{b=1}^B \left(\frac{\hat{\zeta}_s^{(b)^{-1}}(p_{js}) - y_{js}}{y_{js}} \right), \quad (13)$$

$$\Delta_{q90} = \frac{1}{m_T} \sum_{j=1}^{M_x} \mathbb{1}_{(\tilde{T}, \infty)}(T_{js}) (\hat{q}_{95}(p_{js}) - \hat{q}_5(p_{js})), \quad (14)$$

where the quantities $\hat{q}_{95}(p_{js})$ and $\hat{q}_5(p_{js})$ are the upper and lower bounds of the posterior credibility interval for the quantile $\hat{\zeta}_s^{(b)^{-1}}(p_{js})$.

3.3 | Results

Figure 2 shows the empirical distribution of the FSE over the sites, computed on the test set. The proposed sHMEV outperforms the competitors in all scenarios, even if we can see that in the GM scenario the interquartile variability of the FSE index for different sites is higher with respect to the other two models. In the WEI_{gp} scenario the GEV model has some rather high FSE values, probably a consequence of a consistent inter-block variability for the distribution of $x_{ij}(s)$.

To gain a deeper understanding of this general behavior, Figure 3 reports the results of the two criteria introduced in (13) and (14). Generally, the best performance for the bias appears to be specification dependent. In the WEI and GM scenarios, sHMEV tends to slightly underestimate the posterior predictive quantiles and HMEV gets the best results. In the WEI_{gp} scenario, instead, sHMEV shows the lower bias. For the width of the credibility intervals, sHMEV is consistently the

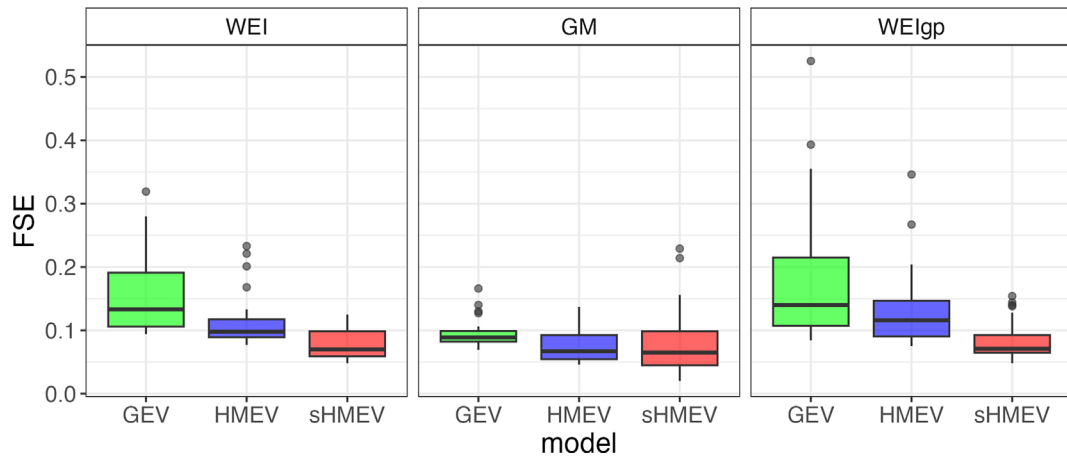


FIGURE 2 Fractional square error computed for the three different model specifications.

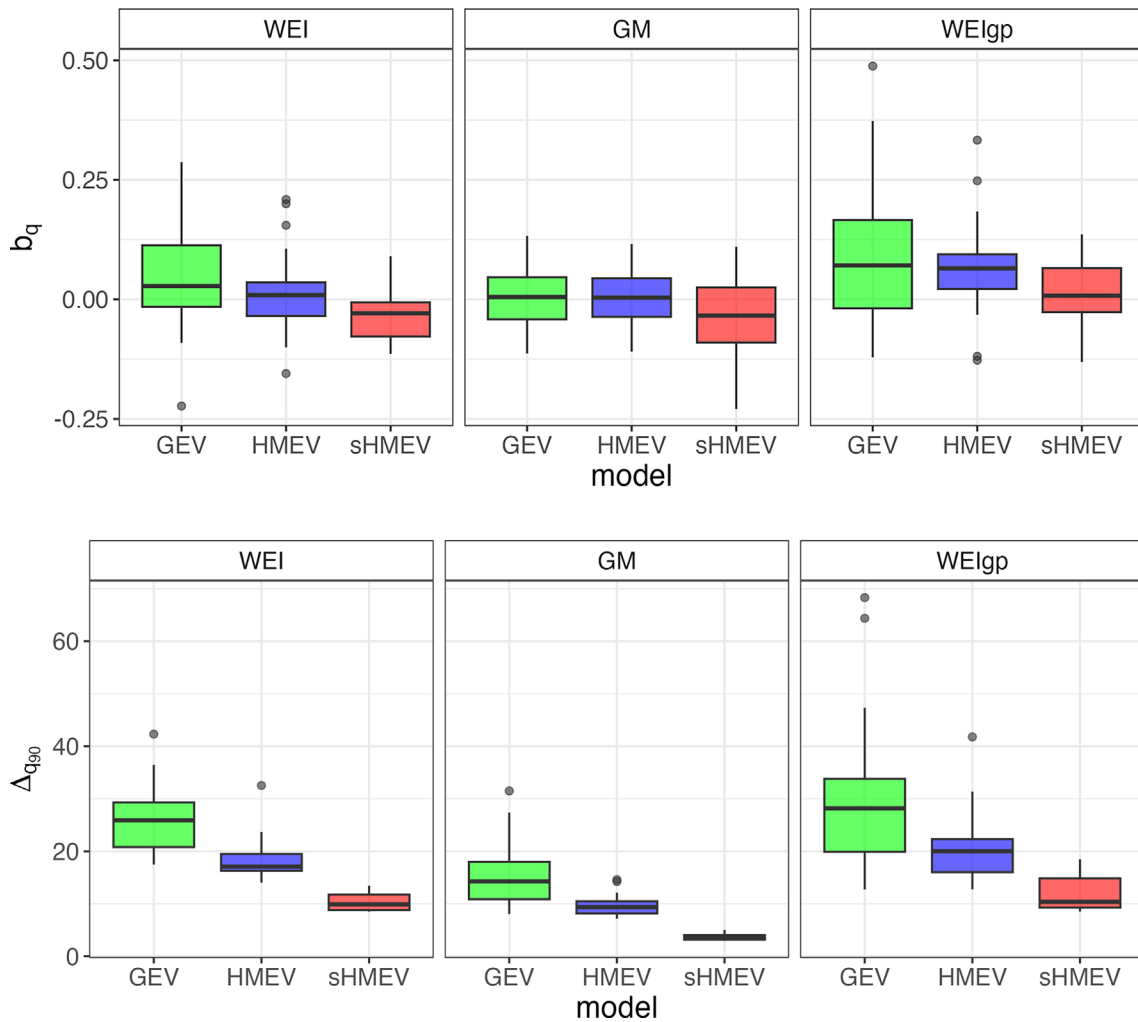


FIGURE 3 Mean bias (top row) and mean credibility interval width (bottom row) computed for the three different model specifications.

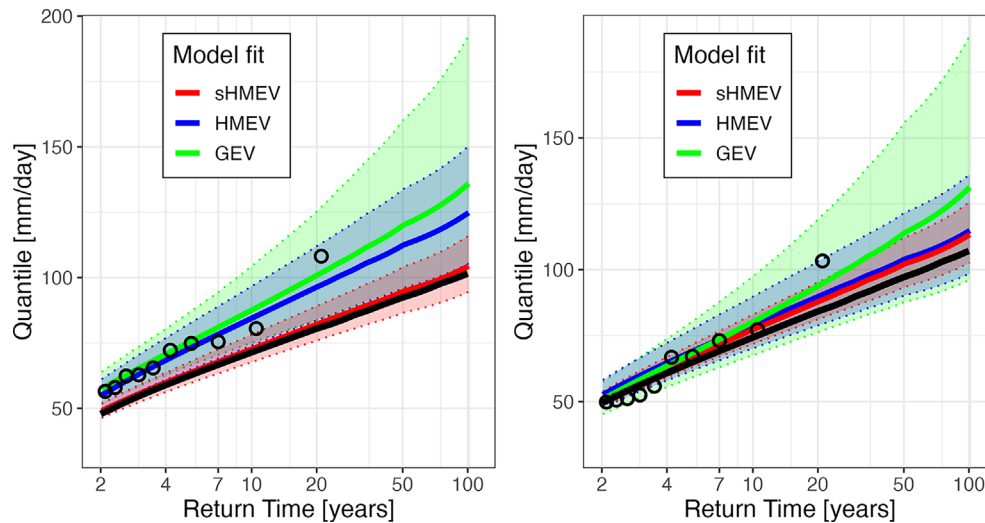


FIGURE 4 Quantiles predicted for two sites by the GEV (green), HMEV (blue), and sHMEV (red) models based on data simulated in the first scenario (WEI). Solid lines show the expected value of the quantile for a given return time, while dashed lines represent the bounds of 90% credibility intervals. Circles represent the observed block maxima on the training set, while the black lines report the quantiles computed from the true sHMEV model.

most efficient procedure, producing narrower credibility intervals. These considerations suggest that the variability of the estimates affects the calculation of the FSE index more than the bias, since the latter is reduced for all three competitors.

To visualize this global behavior, Figure 4 shows a representative example of the performance of the methods. Specifically, it reports the quantile versus return time plots obtained for the different methods applied to the data generated under hypothesis of correct model specification (WEI) for two randomly selected sites. The plots for all the remaining sites are shown in Figures S1 and S2 in the supplementary materials. The quantiles computed from the true model are overall satisfactorily captured by the proposed model. The sHMEV model yields quantile estimates with narrower credibility intervals, especially compared to the GEV model, and it is characterized by lighter tails than the other two methods. Note that the GEV model, despite the informative prior used, appears to be more sensitive to the largest observations in the training set and tends to overestimate the true quantile function, as noted for the site in the right panel of Figure 4. This behavior is expected, given the limited length of the training samples used here, and consistent with previous studies (Zorzetto et al., 2020). For the site considered in the left panel of Figure 4 the maxima observed on the training set are quite different from the expected values on the test set. This situation depicts a common challenge in extreme value analysis, where limited temporal records result in biased predictions due to insufficient information. In this case sHMEV, also exploiting information from the other sites (*borrowing strength*), obtains more accurate and less variable estimations than the competitors.

4 | APPLICATION: RAINFALL IN NORTH CAROLINA

4.1 | United States historical climatological network data

The data analyzed in this section are extracted from the United States historical climatological network (USHCN) data, that are freely available from the National Centers for Environmental Information (NCEI) of the National Oceanic and Atmospheric Administration (NOAA) (Menne et al., 2012). The data consist of daily precipitation records for all the available weather stations in North Carolina, for the time period 1870 through 2021, with a significant fraction of the available records being longer than 100 years. The region is characterized by heterogeneity in morphological and climatic features: it varies from the plain areas near the coastline, to the hilly and mountainous zones in the west of the region. This allows us to test the proposed model in different climates and precipitation regimes.

The records characterized by non-blank quality flags were removed, as well as the years characterized by more than 30 daily missing observations. Therefore, for the subsequent analysis we select only stations that contain more than 73

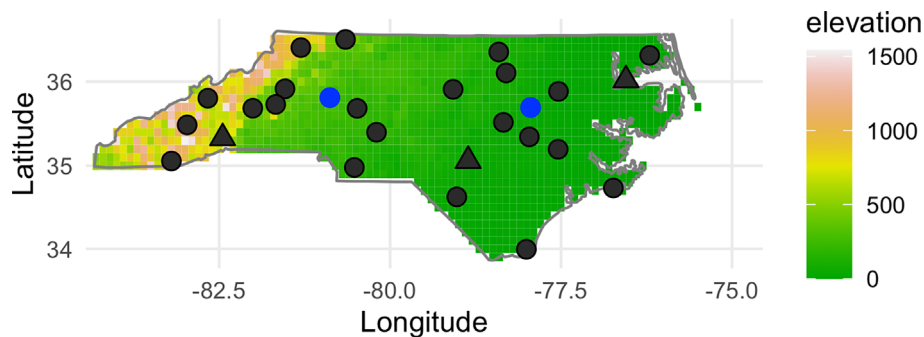


FIGURE 5 Map of North Carolina showing the sites and altitude in meters above sea level of the weather stations. The sites marked by black symbols were used to fit the model, the remaining to validate the model.

years of data, for a total of 27 stations. We randomly chose 25 stations to fit our model and for each one we take the first 20 years.

Figure 5 shows the station locations, with black points and triangles indicating stations included in the training set and blue points indicating stations included in the test set. Black triangles indicate three stations that will be used for illustrative purposes in the following. One of these three stations is on the coastline (Edenton), another one is in the middle of the region (Fayetteville) and the last one is in the mountainous area (Hendersonville).

To assess the temporal dependence, we plot the autocorrelations of the daily rainfall accumulations time series observed at each station. Figure S3 in the supplementary materials shows the ACF plots for the three stations taken as examples, but the trend is the same for all stations considered in the analysis. As the locations fail to show any significant temporal dependence, operations to enhance the events pseudo-independent (e.g., declustering) are not necessary.

We tested for spatial dependence in the positive daily precipitation residuals using an exponential variogram with a linear trend and found a low dependence between stations within 20 km. Since the two closest stations are 24 km apart, the choice to model the spatial dependence with a latent process driven by spatial covariates seems appropriate.

Figure S4, in the supplementary materials, depicts the annual maxima of precipitation for all stations, where the color and the size of the circles are respectively proportional to the mean and the standard error of the annual maxima. The greatest intensity and variability of the maxima occur for stations on the coast and for some in the mountainous area. Therefore, we decided to consider as geographical covariates to be included in the model, in addition to latitude and longitude, also the altitude and the distance from the coast in km. These, in fact, appear to have a marginal effect on both the two parameters of the Weibull distribution, scale and shape, estimated via the method of moments for positive accumulations and on the number of annual rainy days.

4.2 | Fit of the sHMEV model

We apply to the data the sHMEV model specified in Section 2.2, using for the estimation the first 20 years of 25 stations, which means $J = 20$ and $S = 25$. According to the considerations made in the previous section, we define $\mu_\gamma(s)$ as

$$\mu_\gamma(s) = \beta_{\gamma,0} + \beta_{\gamma,1}\text{lat}(s) + \beta_{\gamma,2}\text{lon}(s) + \beta_{\gamma,3}\text{alt}(s) + \beta_{\gamma,4}\text{dist}(s), \quad (15)$$

where $\text{lat}(s)$, $\text{lon}(s)$, $\text{alt}(s)$, and $\text{dist}(s)$ are respectively latitude, longitude, altitude, and distance from the coast of site s . All the covariates have been standardized. The functions for the parameters δ and λ are similarly defined. We focus on a simple specification, since more flexible formulations that include Gaussian processes to model additional unexplained heterogeneity had obtained worse predictive performance.

For prior elicitation, we follow what is reported in Section 2.3, avoiding the use of particularly uninformative distributions. We describe below the procedure followed for choosing the hyperparameters of the normal distributions for β_δ , but similar reasoning was used for the selection of the other hyperparameters. Recalling that the covariates were standardized, the value of the intercept $\beta_{\delta,0}$ refers to the average case, that is the case in which all predictors take on average value. For this parameter, we define a prior distribution centered on the mean for the 25 stations of the parameter δ estimated on the data by the method of moments. The variance of the distribution is chosen such that the probability that

$\beta_{\delta,0}$ is between two values considered reasonable is greater than 0.95. For the parameter $\beta_{\delta,1}$ we take a prior distribution centered on the least squares estimation of the simple regression of δ (always estimated for the 25 stations by the method of moments) on latitude. The variance was chosen equivalently to what was done for $\beta_{\delta,0}$. We follow a similar procedure for $\beta_{\delta,2}$, $\beta_{\delta,3}$, and $\beta_{\delta,4}$.

Convergence of the MCMC chains has been assessed using several diagnostic techniques (Brooks & Gelman, 1998). Traceplots are reported in Figure S5 of the supplementary materials. After establishing the convergence of the chains we assess whether the parametric assumptions of the proposed model provide a good fit to the observed data. We perform posterior predictive checks (Gelman et al., 2013), comparing relevant quantities, such that y_i or x_{ij} , with their corresponding posterior predictive densities. The posterior predictive distributions are not analytically available, but it is straightforward to simulate new data from them. Examination of the posterior predictive distributions for the annual maxima, number of events, and daily rainfall magnitudes for two of the example stations are given in Figure S6, in the supplementary materials. Overall these posterior predictive distributions are satisfactorily captured by sHMEV, even if there are some discrepancies for the distribution of $n_j(s)$. However, as discussed in Section 2.2, the distribution of $n_j(s)$ mainly affects the estimation of extreme events only through its average value, which appears to be adequately captured by the model. Note that a discrepancy appears for small values of daily rainfall magnitudes, due to the sensitivity of the measurement instruments. However, the right tail of the daily precipitation distribution, which plays an important role in determining extreme values, is well captured by sHMEV.

The average time for running the MCMC sampling for this dataset is below 18 min on a Apple M1 CPU laptop with 8 GB of RAM.

4.3 | Results

A summary of the posterior distributions for β_δ , β_γ , β_λ , σ_δ , and σ_γ is given in Table 1. Overall, the distributions show low dispersion and the effect of the covariates on precipitation intensity and occurrence agrees with what was observed in the exploratory analysis.

Our goal is to estimate the posterior distribution for the return level for every location in North Carolina. Since the cdf of annual maxima is a function of the parameter γ , δ , and λ , it is sufficient to estimate the posteriors of these processes. We divide the study region into a grid of points and we consider the values of latitude, longitude, altitude, and distance from the coast for each point. With the posterior distributions of β_δ , β_γ , and β_λ and the values of the covariates it is immediate to compute μ_γ , μ_δ , and μ_λ for each point of the map. Doing this for each iteration b , $b = 1, \dots, B = 4000$, provides draws from the posterior distribution of λ and simulating from a Gumbel distribution we get also draws from the posterior distribution of γ and δ .

Figure S7 in the supplementary materials shows the pointwise mean and the pointwise interquartile range of the posterior draws over the B iterations. The Weibull shape parameter, γ , is lower for the mountainous area and tends to increase as latitude increases. For the scale parameter δ , a positive effect of longitude and distance from the coast is observed. Finally, the map for λ shows that the number of annual rainy days, as might be expected, is higher in mountainous areas and decreases along the coast. The level of uncertainty for all three parameters is greatest at some mountain locations.

TABLE 1 Summary statistics for the posterior distributions of the latent process parameters referred to γ and δ , respectively the shape and scale parameters of the Weibull distribution and λ , the success probability of the binomial distribution for the number of events. The posterior means and the associated 95% credible intervals (parentheses) are displayed.

	γ	δ	λ
β_0	0.86 (0.85, 0.87)	10.5 (10.3, 10.7)	-0.93 (-0.94, -0.92)
$\beta_1(\text{lat})$	0.02 (0.01, 0.03)	0.01 (-0.25, 0.26)	0.14 (0.11, 0.17)
$\beta_2(\text{lon})$	0.2 (-0.04, 0.07)	0.49 (-0.21, 1.17)	-0.7 (-0.85, -0.55)
$\beta_3(\text{alt})$	-0.03 (-0.05, -0.01)	0.01 (-0.46, 0.45)	0.13 (0.11, 0.15)
$\beta_4(\text{coast})$	0.01 (-0.06, 0.07)	-0.28 (-1.02, 0.43)	-0.7 (-0.85, -0.55)
σ	0.09 (0.08, 0.10)	2.19 (2.01, 2.38)	

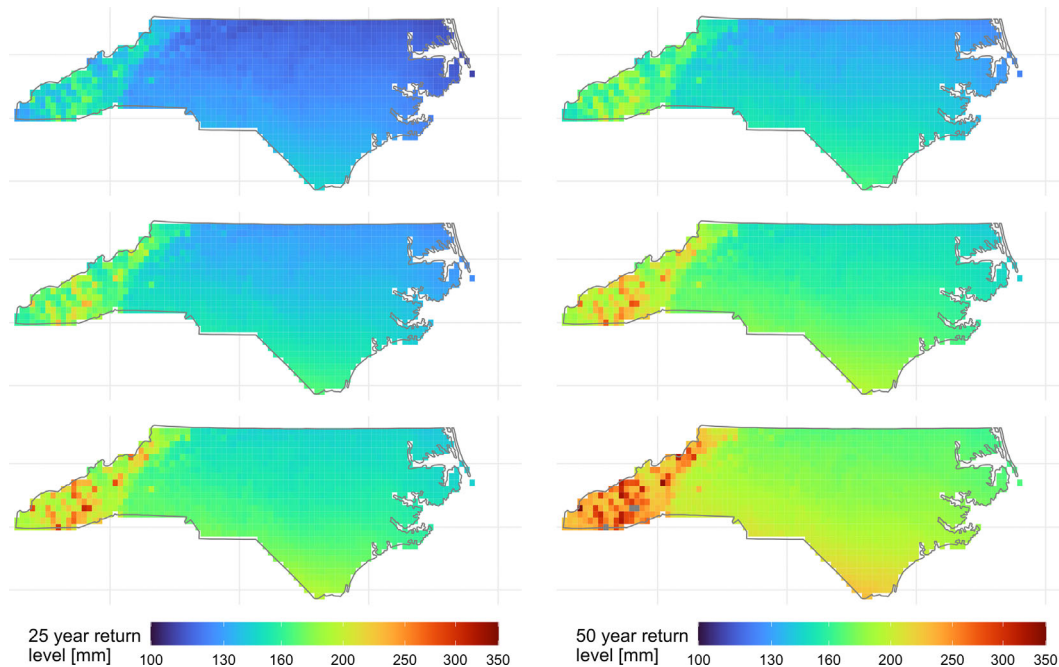


FIGURE 6 Maps of the predictive pointwise 25 and 50 year return level estimates for rainfall (mm) obtained from the sHMEV model. The top and bottom rows show the upper and lower bounds of the 90% pointwise credible intervals, the middle rows show the predictive pointwise posterior mean.

Figure 6 shows maps of the predictive pointwise posterior mean for the 25 and 50 year return levels, with pointwise 90% credible intervals. Adopting a Bayesian methodology, indeed, allows to obtain natural uncertainty estimates for the return levels, taking the pointwise 0.05 and 0.95 empirical quantiles from the return level draws. The return levels were calculated as described in Section 2.3 using (8) and (9). We observe higher values for the mountainous and southwest areas, where there are few stations and the model is forced to extrapolate. In general, there also appears to be a slightly decreasing trend with latitude.

We want to evaluate the performance of the proposed model in predicting extreme values for sites not in the data, that is, by extrapolating to space. For this purpose, we calculate for several return times the quantiles for the two stations in the test set in space, which were not used to estimate the model. Figure 7 reports the quantile versus return time plots for the stations of Statesville and Wilson. For Statesville the quantiles predicted by the sHMEV model match observed annual maxima very well, for Wilson, instead, the model predictions turn out to be less good. However, for this site there are three rather high observed values that are quite difficult to predict. These values were observed during storms and cyclones that caused significant economic and social damages for Wilson county: the tropical storm Brenda on July 29, 1960, the hurricane Floyd on September 16, 1999, and the hurricane Matthew on October 8, 2016. Thus we are not dealing only with heavy rainfall, but with widespread phenomena that involve several extreme meteorological events and require a larger physical investigation. We also recall that the model was estimated using only the first 20 years of the time series and exploiting all the available temporal information could lead to a more accurate spatial extrapolation.

In the following, we compare extreme value quantiles obtained from the HMEV, GEV, and sHMEV models, evaluating the predictive accuracy in estimating the empirical distribution of annual maxima in the test set over time. In fact, all three models were trained on just the first 20 years of record.

For the GEV model the shape parameter τ in Equation (1) plays a crucial role, since it determines the behavior of the tails of the distribution. Similar to other dataset on extreme rainfall, the estimates of the shape parameter are positive, corresponding to the heavy-tailed Fréchet case, but not strongly so (Davison et al., 2012; Papalexiou & Koutsoyiannis, 2013). Furthermore, these estimates agree with the ones obtained by estimating a GEV model on the data simulated from the posterior predictive distribution of the sHMEV model.

Figure S8 in the supplementary materials shows the empirical distribution of the indexes FSE, b_q and Δ_{q90} over the 27 stations. Regarding FSE, the sHMEV model provides the best results for a larger number of stations; however, it has rather high values for some sites. Looking specifically at these sites, they are mainly located on the coast. Probably for

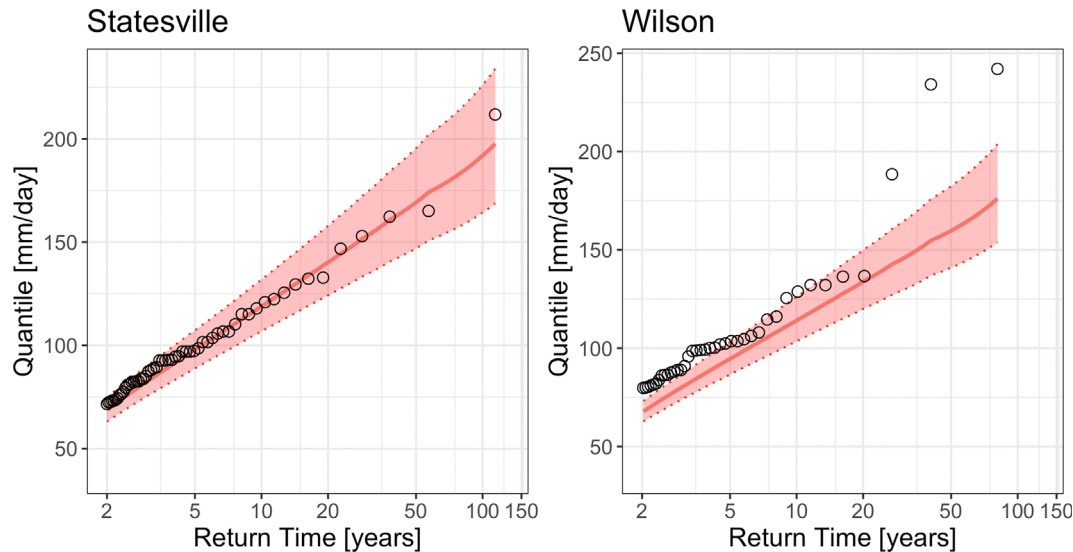


FIGURE 7 Quantiles predicted for the two stations (Statesville and Wilson) taken as the test set on the space. Solid lines show the expected value of the quantile for a given return time, while dashed lines represent the bounds of 90% credibility intervals. Circles represent the observed maxima.

these stations, which exhibit a specific behavior, the sharing of information among different sites results in an excessive shrinkage to the mean, leading to higher errors in the prediction of extreme values. Regarding estimation bias, both sHMEV and HMEV tend to underestimate quantiles. As regards variability, the sHMEV model is the most efficient and it provides considerably thinner credible intervals than the GEV model. This reduction in variability results in a small bias in point estimation, consistent with a well-known tradeoff in statistics known as the variance-bias tradeoff. However, when considering a global measure such as the fractional squared error the proposed approach still outperforms the alternatives. These results are consistent with those obtained in the simulation study.

Figure S9 in the supplementary materials reports the quantile versus return time plots of the different competing methods for two of the stations taken as an example. Models estimates differ, with sHMEV exhibiting, as previously observed from our simulation study, narrower credibility intervals with respect to HMEV and GEV models. For the Fayetteville station the sHMEV model presents an overall good agreement with the empirical frequencies associated with the annual maxima extracted from the entire record. The GEV model, instead, is more influenced by the specific training set used and tends to overestimate the values. For the Edenton station none of the three models adequately captures the distribution of the annual maxima. It should be noted that this station is on the coast, and is likely to experience very unusual and difficult to predict precipitation.

5 | DISCUSSION

We introduced a spatial hierarchical Bayesian model to analyze environmental extreme values. The proposed approach, which extends and generalizes the hierarchical model of Zorzetto et al. (2020), avoids the asymptotic arguments of classical extreme values models and exploits most of the information contained in the data through the ordinary events. The spatial dependence has been induced in the parameters determining the ordinary events and, specifically, modeled through a linear combination of spatial covariates with unknown regression parameters. The Bayesian approach allowed the inclusion of valuable prior information that is often present in environmental modeling. The performance of the method are competitive with the state of the art methods and with the original proposal of Zorzetto et al. (2020) not exploiting any spatial information.

While we focused on simple formulations for the parameters related to the ordinary events, more complex formulations are also possible including semiparametric specifications, for example exploiting spline regressions, or including random noise through suitable stochastic process modeling the residual spatial dependence, for example, Gaussian processes.

The inclusion of spatial covariates sheds light on possible extensions. For example, time dependence can be introduced through a suitable function of covariates or including trend and seasonality. One might also consider a regularized version of the method, by setting a suitable shrinkage prior on the regression coefficients, see Carvalho et al. (2022).

ACKNOWLEDGMENTS

The authors thank the Editor, the Associate Editor, and three anonymous Referees for thoughtful suggestions on an early version of the manuscript. Antonio Canale was supported by the CARIPARO Foundation through the Excellence Grant 2021 to the “Resilience” Project.

REFERENCES

- Betancourt, M. (2018). “A Conceptual Introduction to Hamiltonian Monte Carlo”. In: *arXiv preprint arXiv:1701.02434*.
- Bopp, G. P., Shaby, B. A., & Huser, R. (2021). A hierarchical max-infinitely divisible spatial model for extreme precipitation. *Journal of the American Statistical Association*, 116(533), 93–106.
- Bracken, C., Holman, K. D., Rajagopalan, B., & Moradkhani, H. (2018). A Bayesian hierarchical approach to multivariate nonstationary hydrologic frequency analysis. *Water Resources Research*, 54(1), 243–255.
- Brooks, S. P., & Gelman, A. (1998). General methods for monitoring convergence of iterative simulations. *Journal of Computational and Graphical Statistics*, 7(4), 434–455.
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, B., Brubaker, M., Guo, J., Li, P., & Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, 76(1), 1–32.
- Carvalho, M., Pereira, S., Pereira, P., & Zea Bermudez, P. (2022). An extreme value Bayesian lasso for the conditional left and right tails. *Journal of Agricultural, Biological and Environmental Statistics*, 27(2), 222–239.
- Casella, G. (1985). An introduction to empirical Bayes data analysis. *The American Statistician*, 39, 83–87.
- Castro-Camilo, D., Huser, R., & Rue, H. (2022). Practical strategies for generalized extreme value-based regression models for extremes. *Environmetrics*, 33(6), e2742.
- Coles, S. (2001). *An introduction to statistical modeling of extreme values* Springer Series in Statistics. Springer-Verlag.
- Cooley, D., Nychka, D., & Naveau, P. (2007). Bayesian spatial modeling of extreme precipitation return levels. *Journal of the American Statistical Association*, 102, 824–840.
- Davison, A. C., Padoan, S. A., & Ribatet, M. (2012). Statistical modeling of spatial extremes. *Statistical Science*, 27(2), 161–186.
- Davison, A. C., & Smith, R. L. (1990). Models for exceedances over high thresholds. *Journal of the Royal Statistical Society*, 52(3), 393–442.
- Fondeville, R., de, & Davison, A. C. (2018). High-dimensional peaks-over-threshold inference. *Biometrika*, 105(3), 575–592.
- Frisch, U., & Sornette, D. (1997). Extreme deviations and applications. *Journal de Physique I*, 7(9), 1155–1171.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian data analysis* (3rd ed.). Chapman and Hall/CRC.
- Ghosh, S., & Mallick, B. (2011). A hierarchical Bayesian spatio-temporal model for extreme precipitation events. *Environmetrics*, 22, 192–204.
- Hosseini, S. R., Scaioni, M., & Marani, M. (2020). Extreme Atlantic hurricane probability of occurrence through the metastatistical extreme value distribution. *Geophysical Research Letters*, 47, 2019GL086138.
- Huser, R., & Wadsworth, J. L. (2019). Modeling spatial processes with unknown extremal dependence class. *Journal of the American Statistical Association*, 114(525), 434–444.
- Huser, R., & Wadsworth, J. L. (2022). Advances in statistical modeling of spatial extremes. *WIREs Computational Statistics*, 14(1), e1537.
- Katz, R. W., Parlange, M. B., & Naveau, P. (2002). Statistics of extremes in hydrology. *Advances in Water Resources*, 25(8), 1287–1304.
- Koutsoyiannis, D. (2004). Statistics of extremes and estimation of extreme rainfall: I. Theoretical investigation. *Hydrological Sciences Journal*, 49(4), 575–590.
- Marani, M., & Ignaccolo, M. (2015). A metastatistical approach to rainfall extremes. *Advances in Water Resources*, 79, 121–126.
- Marra, F., Nikolopoulos, E., Anagnostou, E., & Morin, E. (2018). Metastatistical extreme value analysis of hourly rainfall from short records: Estimation of high quantiles and impact of measurement errors. *Advances in Water Resources*, 117, 27–39.
- Marra, F., Amponsah, W., & Papalexiou, S. M. (2023). Non-asymptotic Weibull tails explain the statistics of extreme daily precipitation. *Advances in Water Resources*, 173, 104388.
- Marra, F., Zoccatelli, D., Armon, M., & Morin, E. (2019). A simplified MEV formulation to model extremes emerging from multiple nonstationary underlying processes. *Advances in Water Resources*, 127, 280–290.
- Menne, M. J., Durre, I., Vose, S. R., Gleason, B. E., & Houston, T. (2012). An overview of the global historical climatology network-daily database. *Journal of Atmospheric and Oceanic Technology*, 29, 897–910.
- Opitz, T. (2016). Modeling asymptotically independent spatial extremes based on Laplace random fields. *Spatial Statistics*, 16, 1–18.
- Padoan, S. A., Ribatet, M., & Sisson, S. A. (2010). Likelihood-based inference for max-stable processes. *Journal of the American Statistical Association*, 105(489), 263–277.
- Papalexiou, S. M., & Koutsoyiannis, D. (2013). Battle of extreme value distributions: A global survey on extreme daily rainfall. *Water Resources Research*, 49(1), 187–201.
- Porporato, A., Vico, G., & Fay, P. A. (2006). Superstatistics of hydro-climatic fluctuations and interannual ecosystem productivity. *Geophysical Research Letters*, 33(15), L15402.
- Reich, B. J., & Shaby, B. A. (2012). A hierarchical max-stable spatial model for extreme precipitation. *The Annals of Applied Statistics*, 6(4), 1430–1451.

- Sang, H., & Gelfand, A. (2008). Hierarchical modeling for extreme values observed over space and time. *Environmental and Ecological Statistics*, *16*, 407–426.
- Sang, H., & Gelfand, A. (2010). Continuous spatial process models for spatial extreme values. *Journal of Agricultural, Biological, and Environmental Statistics*, *15*, 49–65.
- Shooter, R., Ross, E., Tawn, J., & Jonathan, P. (2019). On spatial conditional extremes for ocean storm severity. *Environmetrics*, *30*(6), e2562.
- Smith, R. L. (1989). Extreme value analysis of environmental time series: An application to trend detection in ground-level ozone. *Statistical Science*, *4*, 367–377.
- Sornette, D. (2006). *Critical phenomena in natural sciences: Chaos, fractals, selforganization and disorder: concepts and tools*. Springer Science & Business Media.
- Stechmann, S., & Neelin, J. (2014). First-passage-time prototypes for precipitation statistics. *Journal of the Atmospheric Sciences*, *71*, 3269–3291.
- Thibaud, E., & Opitz, T. (2015). Efficient inference and simulation for elliptical Pareto processes. *Biometrika*, *102*(4), 855–870.
- Wadsworth, J. L., & Tawn, J. A. (2022). Higher-dimensional spatial extremes via single-site conditioning. *Spatial Statistics*, *51*, 100677.
- Wilson, P. S., & Toumi, R. (2005). A fundamental probability distribution for heavy rainfall. *Geophysical Research Letters*, *32*(14), L14812.
- Zorzetto, E., Canale, A., & Marani, M. (2020). “Bayesian non-asymptotic extreme value models for environmental data”. In: *arXiv preprint arXiv:2005.12101*.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Stolf, F., & Canale, A. (2023). A hierarchical Bayesian non-asymptotic extreme value model for spatial data. *Environmetrics*, e2806. <https://doi.org/10.1002/env.2806>