

# INLA goes extreme: Bayesian tail regression for the estimation of high spatio-temporal quantiles

Thomas Opitz<sup>1</sup>, Raphaël Huser<sup>2\*</sup>, Haakon Bakka<sup>2</sup> and Håvard Rue<sup>2</sup>

\*Email: raphael.huser@kaust.edu.sa; Phone: +966 12 8080682

February 6, 2018

## Abstract

This work has been motivated by the challenge of the 2017 conference on Extreme-Value Analysis (EVA2017), with the goal of predicting daily precipitation quantiles at the 99.8% level for each month at observed and unobserved locations. To propose a general approach to this specific problem, we develop a Bayesian generalized additive modeling framework tailored to estimate complex trends in marginal extremes observed over space and time. Our approach is based on a set of regression equations linked to the exceedance probability above a high threshold and to the size of the excess, the latter being modeled using the asymptotic generalized Pareto (GP) distribution suggested by Extreme-Value Theory. Latent random effects are modeled additively and semi-parametrically using Gaussian process priors, which provides high flexibility and interpretability. Fast and accurate estimation of posterior distributions may be performed thanks to the Integrated Nested Laplace approximation (INLA), efficiently implemented in the R-INLA software, which we also use for determining a nonstationary threshold based on a model for the body of the distribution. We show that the GP distribution meets the theoretical requirements of INLA, and we then develop a penalized complexity prior specification for the tail index, which is a crucial parameter for extrapolating tail event probabilities. This prior concentrates mass close to a light exponential tail while still allowing heavier tails by penalizing the distance to the exponential distribution at a constant rate. We illustrate this new methodological framework through the modeling of spatial and seasonal trends in daily precipitation data provided by the EVA2017 challenge, comprising observations at 40 stations in the Netherlands during the period 1972–2016. Capitalizing on R-INLA's fast computation capacities and powerful distributed computing resources, we conduct an extensive cross-validation study to select model parameters governing the smoothness of trends, which are critical for accurate prediction. Our results clearly outperform simple benchmarks and are comparable to the best-scoring approach among the other teams.

**Keywords:** Bayesian hierarchical modeling; Extreme-Value Analysis conference challenge; Extreme-Value Theory; generalized Pareto distribution; high quantile estimation; Integrated Nested Laplace Approximation (INLA).

---

<sup>1</sup>INRA, UR546 Biostatistics and Spatial Processes, 228, Route de l'Aérodrome, CS 40509, 84914 Avignon, France.

<sup>2</sup>Computer, Electrical and Mathematical Sciences and Engineering (CEMSE) Division, King Abdullah University of Science and Technology (KAUST), Thuwal 23955-6900, Saudi Arabia.

# 1 Introduction

At its origins, Extreme-Value Theory focused on describing the asymptotic behavior of sample maxima (see, e.g., the monograph of [Gumbel, 1958](#)), and equivalent theoretical statements were later established for threshold exceedances ([Balkema and de Haan, 1974](#); [Pickands, 1975](#)). Statistical methodology for threshold exceedances based on a limiting point process representation or on the related generalized Pareto (GP) distribution was developed by [Hosking and Wallis \(1987\)](#) and [Davison and Smith \(1990\)](#) among others, and numerous applications of threshold-based approaches to environmental extremes observed over space and/or time have since been published; see, e.g., [Huser and Davison \(2014\)](#), [Thibaud and Opitz \(2015\)](#) and [Bacro \*et al.\* \(2017\)](#) for recent applications to precipitation data, [Chavez-Demoulin and Davison \(2005\)](#) for temperature data, and [Northrop and Jonathan \(2011\)](#) and [Jonathan \*et al.\* \(2014\)](#) for oceanographic data. In contrast with the block maximum approach, the use of threshold exceedances allows detailed modeling of trends, seasonality and extremal clustering characteristics due to short-term dependence, while also giving more flexibility for balancing bias and variance; however, choosing a good threshold remains challenging ([Frigessi \*et al.\*, 2002](#); [Scarrott and MacDonald, 2012](#)), which has lead some authors to implement Bayesian algorithms accounting for threshold uncertainty ([Tancredi \*et al.\*, 2006](#)).

Fully Bayesian modeling approaches for spatial and/or temporal extremes often rely on latent processes embedded into the GP parameters to capture trends and dependence ([Cooley \*et al.\*, 2007](#); [Bopp and Shaby, 2017](#)). In particular, [Cooley \*et al.\* \(2007\)](#) use Gaussian processes to capture spatial dependence and covariate-driven trends in precipitation data, taking advantage of simulation-based Markov chain Monte-Carlo (MCMC) methods for the estimation of posterior distributions. Here, we adopt a similar model structure for capturing spatio-temporal trends in precipitation data and for providing extreme quantile predictions. However, our statistical inference approach relies on the Integrated Nested Laplace Approx-

imation (INLA) of posterior distributions, which is much faster and sidesteps convergence issues in simulation-based techniques, while providing highly accurate results (Rue *et al.*, 2009; Martins *et al.*, 2013; Rue *et al.*, 2017). Furthermore, our choice of prior distributions allows us to appropriately smooth predicted quantiles over space and time, while borrowing strength across nearby space-time points, which is especially important when predicting rare events. We also derive the penalized complexity (PC) prior distribution for the GP tail index, which provides a principled prior choice (Simpson *et al.*, 2017) for this crucial parameter by penalizing the distance to a baseline model possessing a light, exponentially-decaying, tail.

Our strategy for modeling space-time trends in marginal precipitation extremes can be decomposed into three stages, each consisting of a suitable univariate response distribution combined with a regression equation capturing systematic temporal and spatial effects. The latter are described semi-parametrically in terms of appropriate Gaussian process priors. In the first modeling stage, we fit a Gamma regression to the positive precipitation intensities, in order to fix a suitable high spatio-temporal threshold. Next, we estimate the overall excess probability above this threshold through Bernoulli regression. Finally, we fit a GP distribution to the observed threshold exceedances, assuming a constant tail index and a scale parameter varying in space and time. We then predict extreme spatio-temporal quantiles by combining the posterior mean predictions from the Bernoulli and GP stages. For reasons of modeling and computational complexity, our Bayesian regression models are based on the working assumption that the data are conditionally independent with respect to the latent spatio-temporal trend components. By imposing this model simplification, the computational efficiency of INLA allows us to conduct an extensive cross-validation study for selecting certain crucial model parameters.

This new methodology, whose development was motivated by the challenge organized for the 10th Extreme Value Analysis conference (EVA2017) held in June 2017 in Delft, Nether-

lands, has been applied to Dutch precipitation data, with the ultimate goal of predicting spatio-temporal quantiles at the 99.8% level; see [Wintenberger \(2018\)](#) for more details on the dataset, the evaluation criterion, and the results of this competition.

In the remainder of the paper, [Section 2](#) presents the dataset and some preprocessing steps. The Bayesian space-time regression framework for tail modeling is developed in [Section 3](#), where the derivation of the penalized complexity prior for the tail index is presented in [Section 3.3](#). We explain the Bayesian estimation approach using INLA in [Section 4](#). The cross-validation study and our final results for the Dutch precipitation application are reported in [Section 5](#). Some concluding remarks with an outlook towards possible future developments are summarized in [Section 6](#).

## 2 Data

The complete dataset consists of daily precipitation accumulations (in inches) measured at 40 stations during the period 1972–2016. The data were divided into a training set (1972–1995) made available to the teams participating to the EVA2017 challenge, and a validation set (1996–2016) used to assess the high quantile predictions of the different teams, and therefore revealed only afterwards. Some of the stations were active during most of the training period, but many were not, resulting in a mixed dataset comprising a few time-rich stations and many time-poor ones; the available sample size varies from  $n = 0$  at Stations 6–10 and 37 to  $n = 8387$  at Station 39; see [Figure 1](#) and [Wintenberger \(2018\)](#) for more details on the data. In terms of modeling, this implies that it is crucial to build a spatial model that borrows strength across nearby stations to obtain robust and reliable predictions at unobserved, or scarcely observed, locations.

The exact coordinates of the stations were not disclosed before the end of the competition to prevent reverse engineering. However, shifted coordinates were provided and are displayed

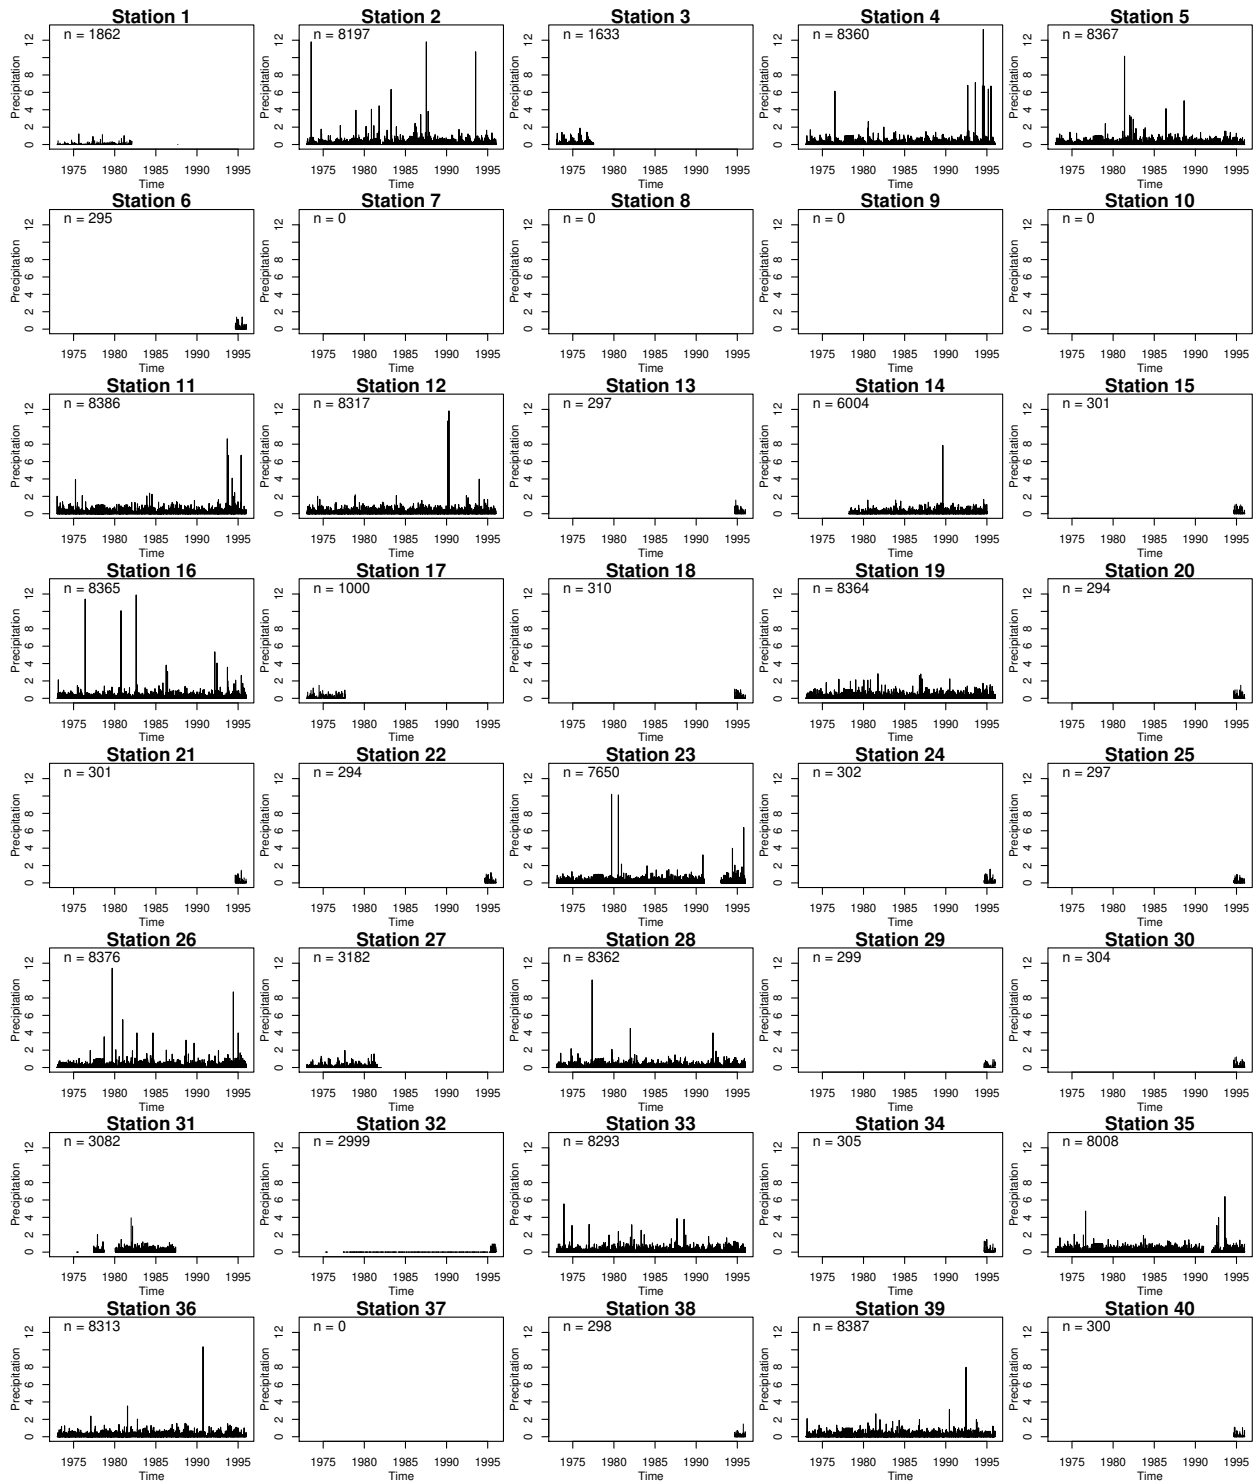


Figure 1: Original precipitation time series measured at 40 stations during the training period. All time series are displayed using a common x-axis (time) and y-axis (precipitation amounts in inches). For each station, the sample size, ranging from  $n = 0$  to  $n = 8387$  for Station 39, is indicated in the top left corner of the corresponding panel.

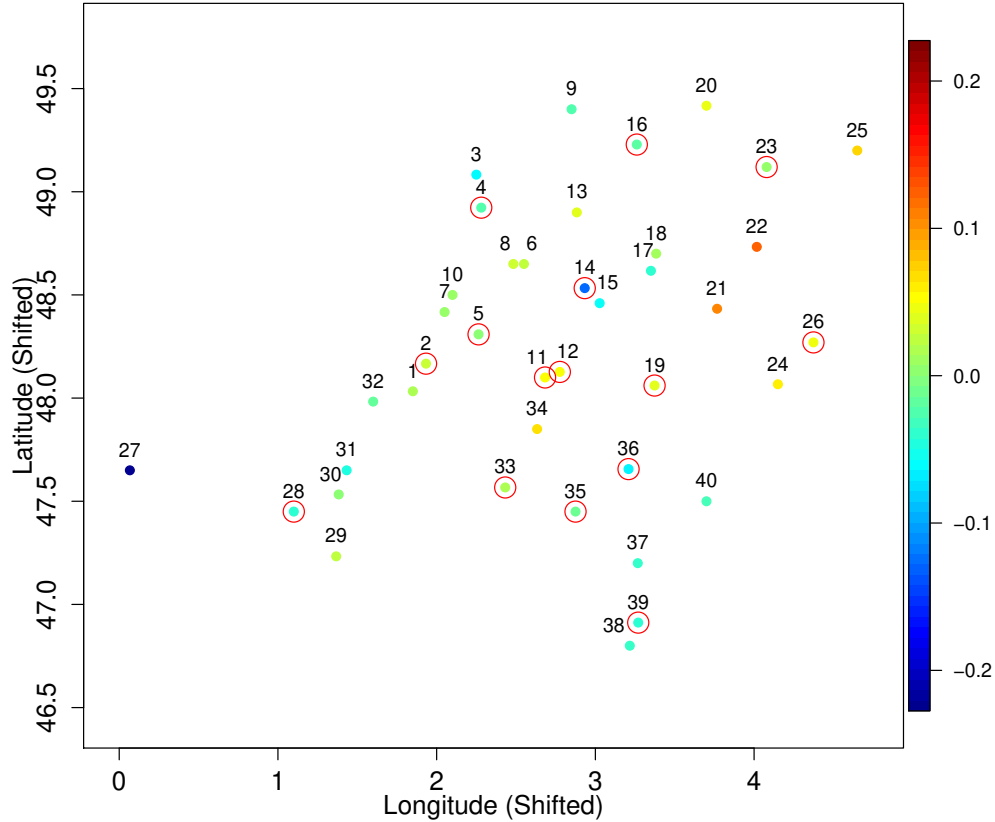


Figure 2: Map of monitoring stations, plotted using the shifted coordinates provided by the EVA2017 challenge, and colored according to the estimated spatial random effect for Stage 1 (Gamma distribution) in our “best model”; see §3.2 and 5.2 for further details. Red circles indicate time-rich stations with at least 3650 observations (i.e., about 10 years of data) over the training period. [Wintenberger \(2018\)](#) displays the true locations and geographical map.

in Figure 2. As shown in [Wintenberger \(2018\)](#), the region of study, revealed only afterwards, covers the Netherlands almost entirely. Although this country is quite small and mostly flat, which suggests that a stationary (or mildly non-stationary) process over space might be reasonable in this specific example, we opted to build a flexible Bayesian model for extreme events able to capture potentially complex trends in space and time; in this way, our general modeling and estimation approach may be applied to a wide range of scenarios. As Figure 2 shows, time-rich stations are scattered across the study region, which is essential for borrowing strength across locations and efficiently estimating spatial trends.

Further exploratory plots (not shown) reveal that the distribution of precipitation varies

slightly over time, with the strongest precipitation usually occurring during the summer. To capture this, our model described in §3 features a seasonal effect defined over weeks.

Figure 1 also reveals that some parts of the data still contain erroneous measurements. For example, Stations 1 and 2 are geographically very close to each other, but their time series are very dissimilar; the data recorded in Station 1 appear in fact to be extremely small compared to the rest of the stations. Furthermore, Station 32 has a very long series of zeros from 1977 to 1995. Finally, Stations 4, 5, 11, 19, 23, 26, 28, 35 and 39 have a series of constant measurements during the period 1977–1978. We therefore removed all these highly suspicious data before proceeding further with the model fitting.

## 3 Modeling

### 3.1 Asymptotic tail models

In classical Extreme-Value Theory, the Fisher–Tippett–Gnedenko theorem characterizes the convergence of renormalized maxima taken over a sequence of increasingly large blocks of independent and identically distributed (i.i.d.) random variables; see Coles (2001), Chapter 3. The class of possible limit distributions may be summarized in terms of a single parametric family known as the generalized extreme-value (GEV) distribution. Equivalently, the generalized Pareto (GP) distribution may be used to model high threshold exceedances; see Coles (2001), Chapter 4, and Davison and Smith, 1990. The block maximum and threshold exceedance approaches can be unified using point process theory (Coles, 2001, Chapter 7, Davison and Huser, 2015). One assumes in practice that the asymptotic distribution provides a reasonable model for sample maxima or threshold exceedances for a fixed and finite block size or threshold, respectively. Provided that the asymptotic GP distribution holds *exactly* to describe the tail of a random variable  $Y \sim F$  above a predefined high threshold  $u$ , the upper tail may be represented in terms of the exceedance probability  $p_u = \text{pr}(Y > u) = 1 - F(u)$

and the GP distribution of positive threshold exceedance  $Y_u^+ = Y - u \mid Y > u$ , i.e.,

$$\Pr(Y - u > y \mid Y > u) = \frac{1 - F(u + y)}{1 - F(u)} = 1 - \text{GP}(y; \sigma, \xi) = (1 + \xi y / \sigma)_+^{-1/\xi}, \quad y > 0, \quad (1)$$

where  $a_+ = \max(0, a)$ ,  $\xi \in \mathbb{R}$  is the tail index and  $\sigma > 0$  is a scale parameter. When  $\xi = 0$ , the expression in (1) is interpreted as the limit  $\xi \rightarrow 0$  and corresponds to the exponential survivor function  $\exp(-y/\sigma)$ ,  $y > 0$ . When  $\xi > 0$ , the GP distribution has a power-law decay with infinite upper endpoint, and when  $\xi < 0$ , it has finite and parameter-dependent upper endpoint  $-\sigma/\xi > 0$ , which leads to certain complications with likelihood-based inference. As our main goal is to model precipitation extremes, which are usually heavy-tailed, we restrict ourselves in this paper to  $\xi \geq 0$  and thus exclude the case  $\xi < 0$  below.

To avoid confounding problems due to the correlation between estimated GP parameters, it may be convenient to reparametrize the GP distribution using the tail index  $\xi$  and a specific  $q$ -quantile  $\kappa_q$  for some fixed probability of interest  $q \in (0, 1)$ , i.e.,

$$\text{GP}(y; \kappa_q, \xi) = \begin{cases} 1 - [1 + \{(1 - q)^{-\xi} - 1\}y/\kappa_q]_+^{-1/\xi}, & \xi \neq 0, \\ 1 - (1 - q)^{y/\kappa_q}, & \xi = 0, \end{cases} \quad y > 0. \quad (2)$$

In our Bayesian spatio-temporal analysis of extreme precipitation described in §5, we use the parametrization (2) in terms of the median.

In practice, the equality between the left- and right-hand sides of (1) should only be interpreted as an approximation for large  $u$ , and a careful bias-variance assessment must be performed to fix a suitable threshold  $u$  above which observations are deemed to be *extreme*; too low a threshold may cause sub-asymptotic bias, while too high a threshold implies large estimation variance due to the small number of threshold exceedances; see Davison and Smith (1990), Northrop and Jonathan (2011) and Scarrott and MacDonald (2012) for standard diagnostics. In our application, we select  $u$  by cross-validation; see §5.

We now present our general three-stage modeling strategy for estimating high quantiles, relaxing the above description to account for spatio-temporal trends in marginal extremes.



## 3.2 Bayesian tail regression framework

The modeling of univariate extremes using the GP distribution has been popularized by [Davison and Smith \(1990\)](#), who advocate to capture systematic variation in extreme events by including fixed covariate effects into the GP scale  $\sigma$  (and potentially, if strongly supported by the data, also the tail index  $\xi$ ). Similarly, using the equivalent point process representation of high threshold exceedances, [Coles and Tawn \(1996\)](#) propose to include spatial covariates into the extreme-value location and scale parameters, while treating the tail index as constant. With a different objective in mind, [Beirlant \*et al.\* \(1999\)](#) formulate an exponential regression model for the (positive) tail index, and suggest that it may be adapted to incorporate exogeneous covariate information. To flexibly handle non-stationarity at high levels, [Davison and Ramesh \(2000\)](#) advocate a local likelihood approach for time series extremes, while [Chavez-Demoulin and Davison \(2005\)](#) propose a generalized additive modeling (GAM) framework, whereby smoothing splines are incorporated into GP parameters. In the Bayesian framework, [Casson and Coles \(1999\)](#) and [Cooley \*et al.\* \(2007\)](#) build hierarchical models for high threshold exceedances, which include latent spatial random effects fitted using expensive simulation-based MCMC methods.

The spatio-temporal Bayesian hierarchical approach that we develop here is similar in spirit to the purely spatial model of [Cooley \*et al.\* \(2007\)](#); however, the fast and accurate inference method based on INLA presented in §4 makes it possible to consider very high-dimensional data and complex space-time models, comprising fixed and random effects that may potentially be organized hierarchically within the regression equations.

Let  $Y(s, t)$  denote daily precipitation observed at location  $s \in \mathcal{S}$  and time  $t \in \{1, \dots, n\}$ , where  $\mathcal{S} \subset \mathbb{R}^d$  is our study region. Our modeling and estimation strategy relies on the following three stages:

1. Positive precipitation intensities  $Y_0^+ = Y(s, t) \mid Y(s, t) > 0$  are first modeled assuming

a Gamma distribution with mean  $\mu(s, t) > 0$  varying in space and time and constant shape parameter  $k > 0$ , i.e.,

$$Y_0^+ \sim \text{Gam}\{y; \mu(s, t), k\} := \frac{k^k}{\mu(s, t)^k \Gamma(k)} y^{k-1} \exp\{-ky/\mu(s, t)\}, \quad y > 0, \quad (3)$$

where  $\Gamma(\cdot)$  is the Gamma function. For the purpose of using INLA, our parametrization of the Gamma distribution is slightly changed with respect to the classical one. A high space-time threshold  $u(s, t)$  is then chosen as the  $p_+$ -quantile of positive precipitation intensities, deduced after fitting model (3) to the data, where  $p_+ \in (0, 1)$  is a fixed probability. The parameter  $p_+$  defines the level of the threshold at each space-time location, and therefore controls the accuracy of the asymptotic GP approximation and the number of threshold exceedances in the following stages. As  $p_+$  involves a crucial bias-variance trade-off, we selected it by cross-validation in §5.

2. Using the space-time threshold  $u(s, t)$  obtained in Stage 1, exceedance indicators are modeled as Bernoulli random variates, i.e., setting  $Z_u(s, t) = \mathbb{I}\{Y(s, t) > u(s, t)\}$ ,

$$Z_u(s, t) \sim \text{Ber}\{z; p_u(s, t)\} := p_u(s, t)^z \{1 - p_u(s, t)\}^{1-z}, \quad z \in \{0, 1\}, \quad (4)$$

where  $p_u(s, t) \in (0, 1)$  denotes the *overall* probability of exceeding the threshold, i.e., taking into account both the positive and zero parts of the precipitation distribution.

3. Using the space-time threshold  $u(s, t)$  obtained in Stage 1, positive threshold exceedances  $Y_u^+(s, t) = Y(s, t) - u(s, t) \mid Y(s, t) > u(s, t)$  are assumed to follow the GP distribution parametrized in terms of its  $q$ -quantile  $\kappa_q(s, t)$  and tail index  $\xi \geq 0$ ; recall (1) and (2). As the tail index usually lacks information for being accurately estimated, we treat it as constant over space and time. Moreover, assuming that the non-stationary patterns in the bulk and the tail of the distribution might be quite similar to each other, we further specify  $\kappa_q(s, t) = \mu(s, t)r_q(s, t)$ , where  $\mu(s, t)$  is the

Gamma mean estimated in the first stage using a rich training dataset, and  $r_q(s, t)$  is a spatio-temporal correction. This allows to borrow strength from the bulk for more accurate tail estimation. Since we expect  $r_q(s, t)$  to be close to stationary in practice, estimation is simplified in a Bayesian framework, where informative priors can be used.

The Gamma distribution in Stage 1 has been used a lot in hydrology as a whole statistical model for (positive) precipitation intensities (Wilks, 2006). However, its tail decays exponentially fast, and is therefore often found to be too light and to underestimate probabilities associated with extreme events (Katz *et al.*, 2002); by contrast, the GP model used in Stage 3 is motivated by asymptotic theory, and can capture heavy tails. Therefore, Stage 1 is only used here to select a suitable threshold  $u(s, t)$  in a potentially highly non-stationary context and to get a reasonable prior for the size and shape of non-stationaries in the tail, while Stages 2 and 3 provide a complete characterization of the tail in terms of the threshold exceedance probability  $p_u(s, t)$ , the GP  $q$ -quantile  $\kappa_q(s, t)$  and the tail index  $\xi$ . Note that Stage 1 is “context-dependent”, in the sense that the Gamma distribution might need to be replaced by another distribution if the data were different, while Stages 2 and 3 are “context-independent”, as Extreme-Value Theory holds under general assumptions. The overall  $\alpha$ -quantile  $y_\alpha(s, t)$  of the precipitation distribution, for  $\alpha > 1 - p_u(s, t)$ , is then

$$\begin{aligned}
 y_\alpha(s, t) &= u(s, t) + \text{GP}^{-1}\{1 - (1 - \alpha)/p_u(s, t); \kappa_q(s, t), \xi\} \\
 &= \begin{cases} u(s, t) + \kappa_q(s, t) [\{(1 - \alpha)/p_u(s, t)\}^{-\xi} - 1] / \{(1 - q)^{-\xi} - 1\}, & \xi \neq 0, \\ u(s, t) + \kappa_q(s, t) \log\{(1 - \alpha)/p_u(s, t)\} / \log(1 - q), & \xi = 0, \end{cases} \quad (5)
 \end{aligned}$$

where  $\text{GP}^{-1}$  denotes the GP survivor function.

In order to capture systematic space-time variations in the data’s tail behavior, we use a regression formulation with suitable link functions and an additive structure in the predictor components. To keep the model fairly simple and robust while allowing for high flexibility, we assume that each predictor includes (additionally to the intercept) a spatial random effect

and a temporal random effect chosen to be separable, i.e.,

$$\begin{aligned}\log\{\mu(s, t)\} &= \beta_0^{\text{Gam}} + x^{\text{Gam}}(s) + x^{\text{Gam}}(t), \\ \text{logit}\{p_u(s, t)\} &= \beta_0^{\text{Ber}} + x^{\text{Ber}}(s) + x^{\text{Ber}}(t), \\ \log\{\kappa_q(s, t)\} &= \log\{\mu(s, t)\} + \beta_0^{\text{GP}} + x^{\text{GP}}(s) + x^{\text{GP}}(t),\end{aligned}$$

where  $\beta_0^{\text{Gam}}, \beta_0^{\text{Ber}}, \beta_0^{\text{GP}}$  denote fixed intercepts,  $x^{\text{Gam}}(s), x^{\text{Ber}}(s), x^{\text{GP}}(s)$  are spatial random effects defined at each station, and  $x^{\text{Gam}}(t), x^{\text{Ber}}(t), x^{\text{GP}}(t)$  are temporal random effects defined on a weekly basis and cyclic with a yearly period. The superscripts refer to each model stage. In our model, we include the logarithm of the Gamma mean  $\log\{\mu(s, t)\}$  (or, in practice, of its estimated posterior mean) as an additive offset into the GP  $q$ -quantile, such that  $\log\{r_q(s, t)\} = \beta_0^{\text{GP}} + x^{\text{GP}}(s) + x^{\text{GP}}(t)$  can be interpreted as a residual effect after having removed the scaling implied by the Gamma distribution for positive precipitation. If the residual space and time effects  $x^{\text{GP}}(s)$  and  $x^{\text{GP}}(t)$ , respectively, are not significantly nonzero, we can conclude that such trends in threshold exceedances are already well explained by the Gamma model. Since we adopt a Bayesian framework with Gaussian process priors centered at zero, including this offset also prevents an overly strong influence of the prior specification on posterior distributions when only few exceedances are observed. Routine calculations show that the GP density with the above link function is log-concave with respect to the predictor  $\log\{\kappa_q(s, t)\}$ , which is beneficial for INLA and optimization routines in general.

Each spatial effect, denoted by  $x(s)$  for simplicity, is assumed to be driven by a zero mean Gaussian process prior with precision  $\tau_s > 0$  and Matérn correlation function (Stein, 1999; Lindgren *et al.*, 2011; Lindgren and Rue, 2015), i.e.,

$$\text{cov}\{x(s_1), x(s_2)\} = \tau_s^{-1} \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\sqrt{2\nu}h/\psi\right)^\nu K_\nu\left(\sqrt{2\nu}h/\psi\right), \quad h = \|s_1 - s_2\|, \quad (6)$$

where  $K_\nu$  denotes the modified Bessel function of second kind of order  $\nu > 0$ ,  $\psi > 0$  is the spatial range parameter, and  $h$  is the distance between two locations  $s_1, s_2 \in \mathcal{S}$ . The

spatial effect allows us to borrow strength across locations, its main purpose being to make predictions at unobserved (or scarcely observed) stations. Its small-scale behavior related to sample path regularity is of relatively minor importance here, and so we fix  $\nu = 1$ . The range parameter  $\psi$ , however, controls the amount of spatial smoothing, and we select it by cross-validation in §5 to optimize spatial prediction. Unlike  $\psi$ , the precision  $\tau_s$  is estimated, rather than post-selected by cross-validation.

Each temporal effect, denoted by  $x(t)$  for simplicity, is assumed to be driven by a second-order Gaussian random walk prior  $x^*(\omega)$  defined over weeks with precision  $\tau_t > 0$  for its weekly innovations (Lindgren and Rue, 2008 and Rue and Held, 2005, Chapter 3). Let  $\omega_t = \{1, \dots, \lfloor n/7 \rfloor\}$  denote the week associated to time  $t \in \{1, \dots, n\}$ ; one has  $x(t) = x^*(\omega_t)$ , where, simultaneously for all weeks  $\omega$ ,

$$x^*(\omega + 1) - 2x^*(\omega) + x^*(\omega - 1) \sim \mathcal{N}(0, \tau_t^{-1}), \quad (7)$$

and the Gaussian process  $x^*(\omega)$  is restricted to sum to zero for identifiability and to be cyclic over a period of one full year (i.e., 52 “weeks” with approximate length of 7 days). Because the precision  $\tau_t$  controls the smoothness of the weekly effect, we select it by cross-validation in §5 to optimize temporal prediction.

Although the inference approach based on INLA may be exploited with additional fixed effects and more complicated random effects, potentially including space-time interactions, we restricted ourselves to the above structure, which was found to be flexible enough and to provide robust and interpretable results in our application. For the shape of the Gamma distribution, we fix a moderately informative Gamma prior distribution with shape 2 and mean 1, and due to the large sample size we can expect a negligible influence of the prior distribution on posterior estimates. The intercepts have noninformative Gaussian priors centered at 0. The prior choice for the tail index is discussed in §3.3. Moreover, the priors for the various hyperparameters and additive regression components are mutually independent.

### 3.3 Penalized complexity prior for the tail index

The tail index plays a crucial role in high quantile estimation. Light exponential tails arise with  $\xi = 0$ , while heavier power-law tails arise with  $\xi > 0$ . When  $\xi \geq 1$ , the mean is infinite, and when  $\xi \geq 1/2$ , the variance is infinite. Because of this, too large values of  $\xi$  are unrealistic for many data types and typically lead to uncertainty inflation for high quantiles. It is therefore important to choose a suitable prior distribution for  $\xi$ , giving priority to light and moderately heavy tails while properly downweighting unrealistically heavy tails.

An elegant approach to choosing priors in a principled way when little expert knowledge is available consists in using the penalized complexity (PC) priors introduced by [Simpson \*et al.\* \(2017\)](#), which penalize the “distance” from a base model at constant rate, independently of its parametrization. These priors take the geometry induced by the choice of model parametrization into account, therefore avoiding intricate interpretation problems that arise otherwise. Furthermore, these priors are designed to allow for shrinkage towards a simpler reference model, which helps to prevent over-fitting by penalizing complex models. They also provide an “objective” (i.e., “automatic”) way of choosing the prior distribution family, while keeping some degree of subjectivity in selecting the penalization rate parameter.

Following [Simpson \*et al.\* \(2017\)](#), a natural definition of the “distance”  $d(f_\xi, f_{\xi_0})$  of a model  $f_\xi$  with respect to a reference model  $f_{\xi_0}$  (here assumed to be in the same family for simplicity) may be based on the Kullback-Leibler divergence (KLD) by setting

$$d(f_\xi, f_{\xi_0}) = \sqrt{2 \text{KLD}(f_\xi || f_{\xi_0})}, \quad \text{with} \quad \text{KLD}(f_\xi || f_{\xi_0}) = \int f_\xi(y) \log \{f_\xi(y)/f_{\xi_0}(y)\} dy. \quad (8)$$

For the GP distribution (1) with  $0 \leq \xi < 1$ , a natural choice of base model is the exponential distribution which arises when  $\xi = 0$ . In this way, a prior assuming constant decay rate in terms of the Kullback-Leibler divergence (8) penalizes departure from an exponential tail, which usually makes sense for environmental data; the heavier the tail, the stronger the

penalty. Let  $f_\xi(y) := \sigma^{-1} (1 + \xi y/\sigma)^{-1/\xi-1}$ ,  $y > 0$ ,  $\sigma > 0$ ,  $\xi > 0$ , be the GP density and  $f_{\xi_0}(y) := \lim_{\xi \rightarrow 0} f_\xi(y) = \sigma^{-1} \exp(-y/\sigma)$ ,  $y > 0$ ,  $\sigma > 0$ ,  $\xi_0 = 0$ , be the exponential density. Knowing that the GP distribution (1) with parameters  $\sigma, \xi$  has mean  $\sigma/(1 - \xi)$  when  $\xi < 1$ , and using the change of variable  $t = \xi^{-1} \log(1 + \xi y/\sigma)$ , one obtains

$$\begin{aligned}
\text{KLD}(f_\xi||f_{\xi_0}) &= \int f_\xi(y) \log \{f_\xi(y)/f_{\xi_0}(y)\} dy \\
&= \int_0^\infty \sigma^{-1} (1 + \xi y/\sigma)^{-1/\xi-1} \log \left\{ \frac{\sigma^{-1} (1 + \xi y/\sigma)^{-1/\xi-1}}{\sigma^{-1} \exp(-y/\sigma)} \right\} dy \\
&= \sigma^{-1} \int_0^\infty y \sigma^{-1} (1 + \xi y/\sigma)^{-1/\xi-1} dy - \left(1 + \frac{1}{\xi}\right) \int_0^\infty \log(1 + \xi y/\sigma) \sigma^{-1} (1 + \xi y/\sigma)^{-1/\xi-1} dy \\
&= \sigma^{-1} \frac{\sigma}{1 - \xi} - \left(1 + \frac{1}{\xi}\right) \xi \int_0^\infty t \exp(-t) dt = \frac{1}{1 - \xi} - (1 + \xi) \\
&= \frac{\xi^2}{1 - \xi}, \quad 0 \leq \xi < 1,
\end{aligned} \tag{9}$$

which does not depend on the scale parameter  $\sigma$ , as expected. To define a PC prior  $\pi(\xi)$  for the tail index  $\xi$ , we assume that a given model for the data  $f_\xi$  is penalized at constant rate in terms of its “distance”  $d(f_\xi, f_{\xi_0}) = \sqrt{2\text{KLD}(f_\xi||f_{\xi_0})}$  to the reference model  $f_{\xi_0}$ , therefore involving the exponential distribution in the “metric” space defined through  $d(f_\xi, f_{\xi_0})$ . Because the KLD (9) converges to infinity as  $\xi \rightarrow 1$ , such a prior will put zero mass on  $\xi \geq 1$ , hence preventing infinite-mean models to occur. This can be seen as an advantage for applications in hydrology, where the tail index is usually quite small. We here propose two possible prior choices, which are based on (i) Equation (9) and (ii) an approximation of (9) as  $\xi \rightarrow 0$ :

- (i) The first possibility uses the *exact* formula (9), which yields  $d(f_\xi, f_{\xi_0}) = \sqrt{2}\xi/(1 - \xi)^{1/2}$  and implies that the corresponding PC prior for  $\xi$  with support  $[0, 1)$  is

$$\begin{aligned}
\pi(\xi) &= \lambda \exp(-\lambda d(f_\xi, f_{\xi_0})) \left| \frac{\partial d(f_\xi, f_{\xi_0})}{\partial \xi} \right| = \sqrt{2}\lambda \exp \left\{ -\sqrt{2}\lambda \frac{\xi}{(1 - \xi)^{1/2}} \right\} \left\{ \frac{1 - \xi/2}{(1 - \xi)^{3/2}} \right\} \\
&= \tilde{\lambda} \exp \left\{ -\tilde{\lambda} \frac{\xi}{(1 - \xi)^{1/2}} \right\} \left\{ \frac{1 - \xi/2}{(1 - \xi)^{3/2}} \right\}, \quad 0 \leq \xi < 1,
\end{aligned} \tag{10}$$

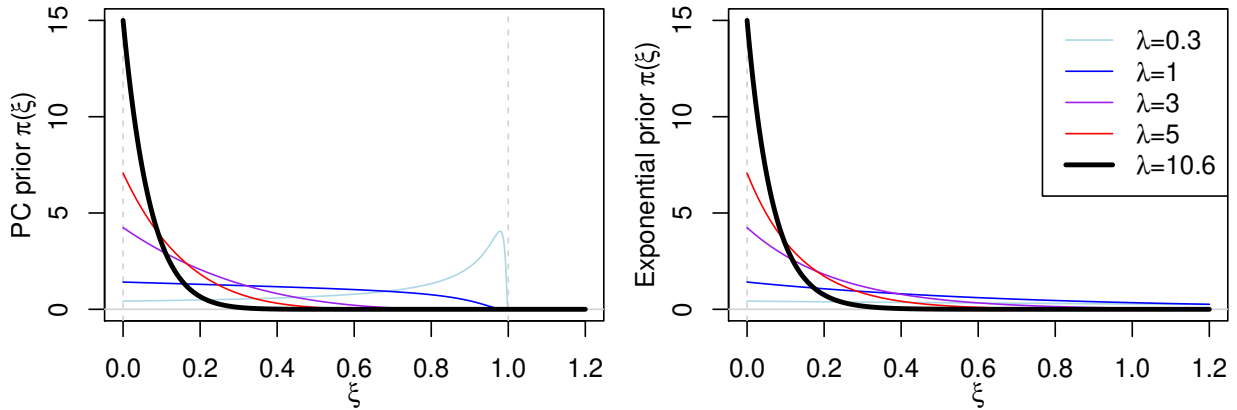


Figure 3: PC priors  $\pi(\xi)$  for the GP tail index  $\xi$  using the exact KLD formula in (10) (left) or the approximate KLD formula in (11) (right), i.e., using an exponential prior, setting the penalization rate parameter to  $\lambda = 0.3$  (light blue),  $\lambda = 1$  (dark blue),  $\lambda = 3$  (purple),  $\lambda = 5$  (red), and  $\lambda = 10.6$  (thick black), which was combined with (11) in our application.

where the penalization rate parameter is  $\lambda = \tilde{\lambda}/\sqrt{2} > 0$ .

- (ii) The second possibility is to approximate the KLD (9) by  $\xi^2$ , as  $\xi \rightarrow 0$ . This yields  $d(f_\xi, f_{\xi_0}) = \sqrt{2}\xi$  and implies that the corresponding approximate PC prior for  $\xi$ , with support  $[0, \infty]$ , is exponential with rate  $\tilde{\lambda} = \sqrt{2}\lambda > 0$ , i.e.,

$$\pi(\xi) = \sqrt{2}\lambda \exp(-\sqrt{2}\lambda\xi) = \tilde{\lambda} \exp(-\tilde{\lambda}\xi), \quad \xi \geq 0. \quad (11)$$

Our two proposed priors are illustrated in Figure 3. These priors are very similar to each other when the penalization rate parameter  $\lambda$  is large (giving strong preference to  $\xi \approx 0$ ) but they differ significantly when  $\lambda$  is small. The PC prior using the exact KLD formula in (10) gives zero probability to  $\xi \geq 1$ , while the prior (11) gives exponentially decreasing but positive probability to all  $\xi \geq 0$ . This bounded/unbounded support distinction between the two proposed priors is more apparent for small  $\lambda$ ; the PC prior (10) becomes concentrated around  $\xi = 1$  as  $\lambda \rightarrow 0$ , which is not very intuitive. In practice, the shrinkage rate  $\lambda$  may be elicited from prior information in specific applications. For example, when the tail is known to be relatively light,  $\lambda \geq 3$  seems to be a reasonable choice, and both priors should give



similar results. In this case, both densities resemble an exponential density, which shrinks the tail index  $\xi$  towards  $\xi = 0$ ; large values of  $\xi$  will only arise if the data are heavy-tailed and provide strong evidence for it. In our precipitation example in §5, we use a fairly informative exponential prior for  $\xi$  with rate  $\tilde{\lambda} = 15$  (i.e.,  $\lambda = 15/\sqrt{2} \approx 10.6$ ), which sets the prior probability for tail indices larger than  $\xi = 0.2$  to be only about 5%; see Figure 3.

## 4 Bayesian inference

### 4.1 Posterior distributions

The three regression models of our application are fitted separately, and the following exposition concentrates on this approach. It would be possible to have some of the fixed and random effects in common between the GP exceedance distribution and the Bernoulli distribution of exceedance indicators; in this case, we would fit those two regressions as a single model with two types of data/likelihood combinations. However, the Gamma regression must always be fitted separately in a first step, since it determines the data used for the other two regressions. By putting prior distributions on the fixed and random effects of the regression and on the various hyperparameters, we can apply Bayes' formula and infer the posterior distributions of model components of interest.

In what follows, the notation is independent of the data/likelihood combination, given in each of the three stages as (1) positive precipitation intensities/Gamma, (2) exceedance indicators/Bernoulli and (3) threshold exceedances/GP, respectively. For the observed data vector with components  $y(s_i, t_i)$ ,  $i = 1, \dots, m$ , we write  $y = (y_1, \dots, y_m)$  in short, and the vector of latent Gaussian predictors is denoted as  $\eta = (\eta_1, \dots, \eta_m)^T$  such that  $\eta_i = \beta_0 + x(s_i) + x(t_i)$ . We denote hyperparameters related to the likelihood by the vector  $\theta_y$ , in our case given by the Gamma shape parameter and GP tail index, while the Bernoulli has no hyperparameter. The vector  $x$  regroups all variables arising in the fixed and random effects, which includes

$\beta_0$ , weekly effects  $x(t)$ , which are based on a discretization of the year using 52 variables, and spatial effects  $x(s)$  for 40 observation and prediction sites  $s$ . Hyperparameters related to the Gaussian priors of predictor components are collected in a vector  $\theta_x$ ; they include the Matérn precision  $\tau_s$  and the precision  $\tau_t$  of random walk innovations. The predictor vector  $\eta$  is linearly related to the latent effects through an observation matrix  $A$  determined by the structure of observation points and latent components, such that  $\eta(x) = Ax$ . The prior distribution on  $x$  is Gaussian with zero mean and precision matrix (i.e., inverse covariance matrix)  $Q(\theta_x)$ . Owing to prior independence between the predictor components  $\beta_0$ ,  $x(s)$  and  $x(t)$ ,  $Q(\theta_x)$  is block-diagonal with three blocks. For parameters in  $\theta_x$  and  $\theta_y$  that are not fixed to a specific value, we define a prior distribution  $\pi(\theta)$  with  $\theta = (\theta_y, \theta_x)$ , which is assumed to factorize over the components of  $\theta$ . The Gaussian density of  $x$  is written  $\pi(x | \theta_x)$ .

Using  $\pi(y_i | \eta_i, \theta_y)$  to denote the likelihood of a data point  $y_i$  conditional on its predictor  $\eta_i$  and likelihood hyperparameters  $\theta_y$ , the joint density of  $y$ ,  $x$  and  $\theta$  may be expressed as

$$\pi(y, x, \theta) = \pi(\theta) \times \pi(x | \theta_x) \times \prod_{i=1}^m \pi(y_i | \eta_i(x), \theta_y). \quad (12)$$

For notational convenience, we now suppose that the predictor vector  $\eta$  is a subvector of  $x$  with  $\eta_i = x_i$ ,  $i = 1, \dots, m$ . We are mainly interested in the posterior distributions of the variables in  $x$  and  $\theta$ , whose joint posterior density conditional on the data is calculated through the Bayes formula as

$$\pi(x, \theta | y) \propto \exp \left\{ -\frac{1}{2} x^T Q(\theta_x) x + \sum_i \log \pi(y_i | x_i, \theta_y) + \log \pi(\theta) + \log K(\theta_x) \right\} \quad (13)$$

where  $K(\theta_x)$  is a normalizing constant related to the hyperparameters of the latent Gaussian field. Specifically, we are interested in the univariate posterior distributions of predictors  $\eta_i$  and components of  $\theta$ , which requires integrating the density (13) with respect to all components of  $x$  and  $\theta$  except the one of interest. Since we cannot provide an exact analytical solution of this high-dimensional integral, we could resort to simulation-based MCMC tech-

niques (see, e.g., [Banerjee et al., 2004](#), Chapter 5) for producing a large representative sample of  $x$  and  $\theta$ . Instead, we propose using INLA, which exploits suitably chosen variants of analytical Laplace approximations of integrand functions ([Tierney and Kadane, 1986](#)), such that integrals are over Gaussian densities and their evaluation becomes straightforward.

## 4.2 Integrated Nested Laplace Approximation (INLA)

INLA ([Rue et al., 2009, 2017](#); [Opitz, 2017](#)) has become a cutting-edge tool for fast and accurate inference in complex and hierarchically structured GAM-like models with latent Gaussian structure and sparse precision matrix, such as our three models described above. For a vector  $v$  whose  $k$ th component has been removed, we write  $v_{-k}$ . INLA provides accurate approximations of the following two types of univariate posterior densities:

$$\pi(\theta_k | y) = \int \pi(x, \theta | y) dx d\theta_{-k}, \quad (14)$$

$$\pi(x_i | y) = \int \int \pi(x, \theta | y) dx_{-i} d\theta = \int \pi(x_i | \theta, y) \pi(\theta | y) d\theta, \quad i = 1, \dots, \tilde{m} \quad (15)$$

where  $\tilde{m}$  is the length of  $x$ . If a hyperparameter has been fixed to a specific value, then its prior density can be interpreted as a Dirac mass of 1 at this value, and the related integral vanishes. In general, the number of estimated hyperparameters should be kept small, since integration with respect to components of  $\theta$  is carried out numerically through an astute choice of discretization points. The main obstacle remains integration with respect to the Gaussian components  $x$ , for which the Laplace approximation is applied in a nested way: first, to approximate the posterior of the hyperparameters  $\pi(\theta | y)$ , and second to integrate with respect to the Gaussian vector  $x_{-i}$ . The log-concavity of the likelihood  $\pi(y_i | \eta_i, \theta_y)$  with respect to  $\eta_i$  is crucial since it ensures that a useful Gaussian approximation to  $\pi(x | \theta, y)$  in [\(13\)](#) can be calculated by matching the mode and the curvature around it using an iterative Newton–Raphson scheme based on second-order Taylor expansions. Indeed, owing to the conditional independence of  $y_i$  with respect to  $\eta_i$ , the precision matrix in the approximation

is of the form  $Q(\theta_x) + \text{diag}(c)$ , where the components  $c_i$  in the added diagonal matrix must be positive and are closely related to the second derivative of  $\log(y_i | x_i, \eta_y)$  with respect to  $x_i$ ; hence the requirement of log-concavity. Notice that a simplified approximation scheme skips the second Laplace approximation by simply using a conditional Gaussian distribution for  $x_{-i} | x_i, \theta$  based on the first Laplace approximation, but slight accuracy problems may arise with non-Gaussian univariate likelihoods. For more details on the approximation mechanism and its implementation, see [Rue \*et al.\* \(2009, 2017\)](#) and [Opitz \(2017\)](#).

## 5 Results for the Dutch precipitation data application

### 5.1 Cross-validation study

The goal of the EVA2017 challenge was to predict the overall  $\alpha_{\text{target}}$ -quantile of daily precipitation, with  $\alpha_{\text{target}} = 0.998$ , for each month and station  $s \in \mathcal{C}_j \subset \mathcal{S}$ , where  $\mathcal{C}_j$ ,  $j = 1, 2$ , denote two different sets of stations (so-called Challenges 1 and 2, respectively, see [Wintemberger, 2018](#)). The target quantile is so high that fully non-parametric methods would likely perform quite badly and be outperformed by parametric alternatives; furthermore, we expect extreme-value models (such as the GP distribution (1)) to provide a good fit at this level. The evaluation criterion used to rank the predictions submitted by all teams was based on the quantile loss function defined as

$$\ell_\alpha(y, q) = \begin{cases} \alpha(y - q), & y > q, \\ -(1 - \alpha)(y - q), & y \leq q, \end{cases} \quad (16)$$

where  $y$  and  $q$  represent observations and predicted quantiles, respectively. As the  $\alpha$ -quantile of the random variable  $Y$  minimizes the expected risk function  $q \mapsto \text{E}\{\ell_\alpha(Y, q)\}$  ([Koenker, 2005](#)), this justifies using (16) for assessing quantile predictions, although alternative risk measures may also be used ([Daouia \*et al.\*, 2018](#)). To optimize our spatio-temporal prediction performance at stations with poor time coverage, we select the most crucial hyperparameters, i.e., the Matérn range parameter  $\psi$  in (6), the temporal precision parameter  $\tau_t$  in (7), and the

threshold probability  $p_+$ , using three different cross-validation (CV) criteria based on (16). An alternative would have been to estimate such parameters directly in the Bayesian model by specifying suitable prior distributions, but since we do not capture stochastic dependence in observations at nearby space-time points in our models, overfitting could have been an issue with such an approach. As these hyperparameters determine the effective amount of information that is used or borrowed from neighboring space-time locations, they may strongly affect our predictions and thus need to be selected with care; recall that  $p_+$  controls the bias-variance trade-off of the asymptotic GP approximation,  $\psi$  controls the amount of spatial smoothing, and  $\tau_t$  controls the amount of temporal smoothing of the corresponding random effects. Let  $y_\alpha(s, t)$  denote the true  $\alpha$ -quantile for station  $s$  and time  $t$ , and  $\widehat{y}_\alpha(s, t)$  denote our estimate of  $y_\alpha(s, t)$  based on Equation (5), obtained by fitting the three stage-model described in §3 to the entire dataset using INLA as explained in §4. Let also  $\widehat{y}_\alpha^{-s}(s, t)$  and  $\widehat{y}_\alpha^{-j}(s, t)$  denote these predicted quantiles obtained after removing station  $s$  and year  $j$ , respectively. We consider the following CV criteria designed to optimize spatial prediction, temporal prediction, or a combination of both:

$$\text{CV}_{\text{space}}(\alpha) = \sum_{s \in \mathcal{C}_2} \sum_{j=1972}^{1995} \sum_{t \in \mathcal{T}(j)} \ell_\alpha\{Y(s, t), \widehat{y}_\alpha^{-s}(s, t)\}, \quad (17)$$

$$\text{CV}_{\text{time}}(\alpha) = \sum_{s \in \mathcal{C}_2} \sum_{j=1972}^{1995} \sum_{t \in \mathcal{T}(j)} \ell_\alpha\{Y(s, t), \widehat{y}_\alpha^{-j}(s, t)\}, \quad (18)$$

$$\text{CV}_{\text{space-time}}(\alpha) = \text{CV}_{\text{space}}(\alpha) + \text{CV}_{\text{time}}(\alpha); \quad (19)$$

here,  $\mathcal{T}(j) \in \{1, \dots, n\}$  denotes the set of time points (days) within a specific year  $j$ . The lower the CV criteria the better the model. As the target probability  $\alpha_{\text{target}} = 0.998$  is extreme, it may be better in practice to evaluate  $\text{CV}_{\text{space}}(\alpha)$ ,  $\text{CV}_{\text{time}}(\alpha)$  and  $\text{CV}_{\text{space-time}}(\alpha)$  for some  $\alpha < \alpha_{\text{target}}$  in order to prevent overfitting. We tried to use  $\alpha = 0.995, 0.998$  but without any major difference in the results. The results presented below and in §5.2 are based on  $\alpha = \alpha_{\text{target}} = 0.998$ . The final monthly quantile predictions were then extracted

from our selected “best model” by averaging the daily predictions for that month.

After projecting the (shifted) longitude-latitude coordinates to a proper metric system, we considered  $\psi \in \{25, 50, 75, 100, 125, 150, 175, 200, 225, 250\}$ km,  $\tau_t = \sigma_t^{-2}$  with  $\sigma_t \in \{0.005, 0.010, 0.015, 0.020, 0.025\}$ , and  $p_+ \in \{0.90, 0.91, 0.92, 0.93, 0.94, 0.95, 0.96, 0.97, 0.98, 0.99\}$ , and fitted our three-stage hierarchical Bayesian tail regression model for all 500 combinations of these hyperparameters. Notice that  $p_+ = 0.99$  (for positive precipitation intensities) roughly corresponds to an overall probability of  $\alpha = 0.995$ , as there were about 50% of zeros observed at each station. Although model estimation based on INLA is much faster than classical simulation-based MCMC methods and it bypasses convergence assessment issues, the total number of fits to perform in order to compute the cross-validation criteria (17), (18) and (19) is equal to 29500 (500 models times 59 hold-out samples, corresponding to the sum of 35 stations and 24 years). In our case, a single fit (for all three stages) took about 8.5 hours overall on 2 cores, and we used the KAUST Supercomputer Shaheen II to compute them all in parallel. The total number of core-hours used was about half a million, which corresponds roughly to 57 years of computation on a single-core machine. This cross-validation study underlines the importance of fast and accurate inference methods coupled with large distributed computing resources.

Figure 4 displays the space-time cross-validation scores (19) for all models, while Table 1 presents a summary of the five best models. Overall, it appears that smaller probabilities  $p_+$  yield much better predictions, thanks to the reduced variability due to the larger effective sample size of threshold exceedances. As for the other hyperparameters  $\psi$  and  $\sigma_t$ , it is difficult to draw general conclusions, as the quite “noisy” diagnostics shown in Figure 4 make interpretation difficult. Nonetheless, the model that provides the best space-time predictions according to (19) has  $p_+ = 0.92$ ,  $\psi = 50$  and  $\sigma_t = 0.01$  (Model 122), and it yields reasonable results; hence, we decided to proceed with Model 122 as our “best model”. As the

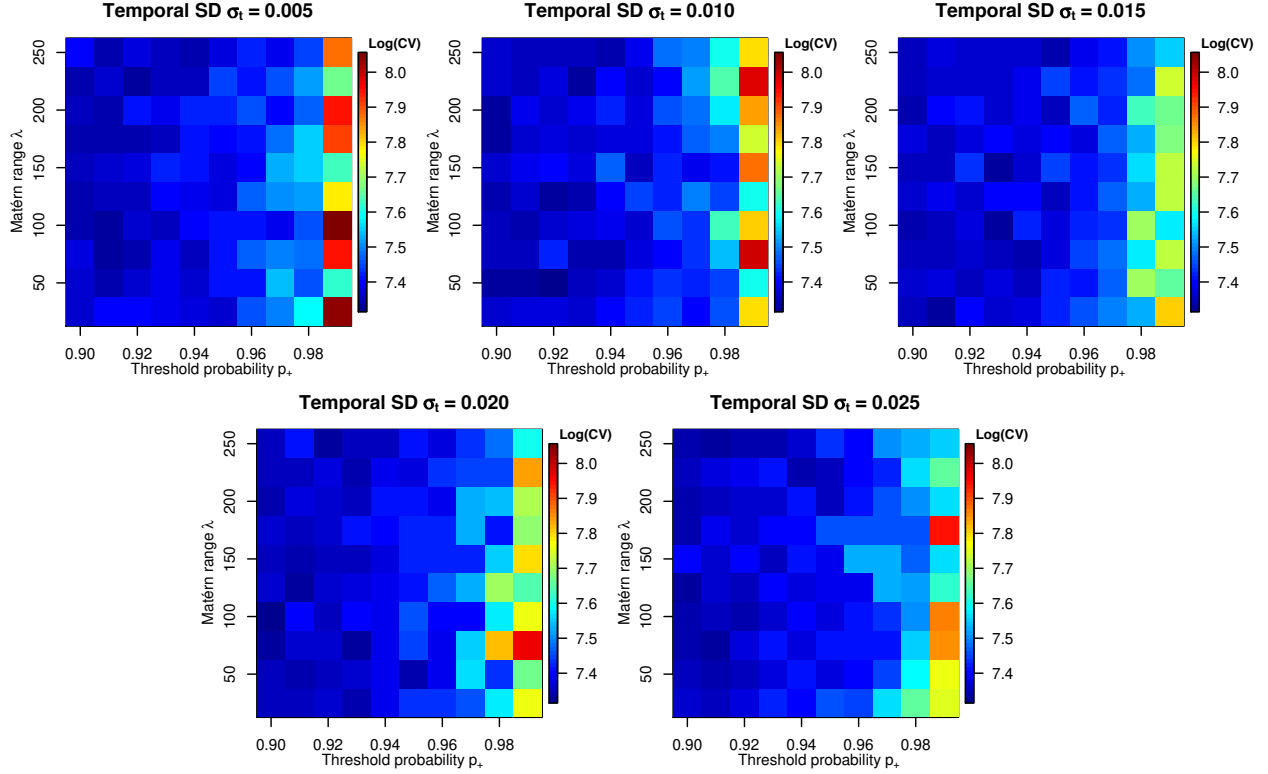


Figure 4: Space-time cross-validation scores (19) with  $\alpha = 0.998$ , displayed on a common logarithmic color scale. Each panel displays the CV scores in terms of  $\psi$  and  $p_+$  for specific values of  $\sigma_t$  (see panel titles). Lower values (i.e., darker blue cells) are better models.

Table 1: Five “best models” ranked according to the space-time CV criterion (19) with  $\alpha = 0.998$ . Columns report the model ID, the Matérn spatial range  $\psi$ , the temporal standard deviation (SD)  $\sigma_t = \tau_t^{-1/2}$ , the threshold probability  $p_+$  for positive precipitation intensities, and the CV scores (17), (18) and (19), respectively.

Model ID	Range $\psi$	Temp. SD $\sigma_t$	Proba. $p_+$	Cross-validation criteria		
				Space (17)	Time (18)	Space-time (19)
122	50	0.010	0.92	759.4	751.9	1511.3
304	100	0.020	0.90	760.9	753.4	1514.2
333	75	0.020	0.93	772.9	747.3	1520.2
211	25	0.015	0.91	767.3	759.4	1526.6
420	250	0.025	0.91	767.1	759.6	1526.8

models reported in Table 1 mostly yield a similar interpretation in terms of the estimated random effects, the next section showcases the results obtained for Model 122.

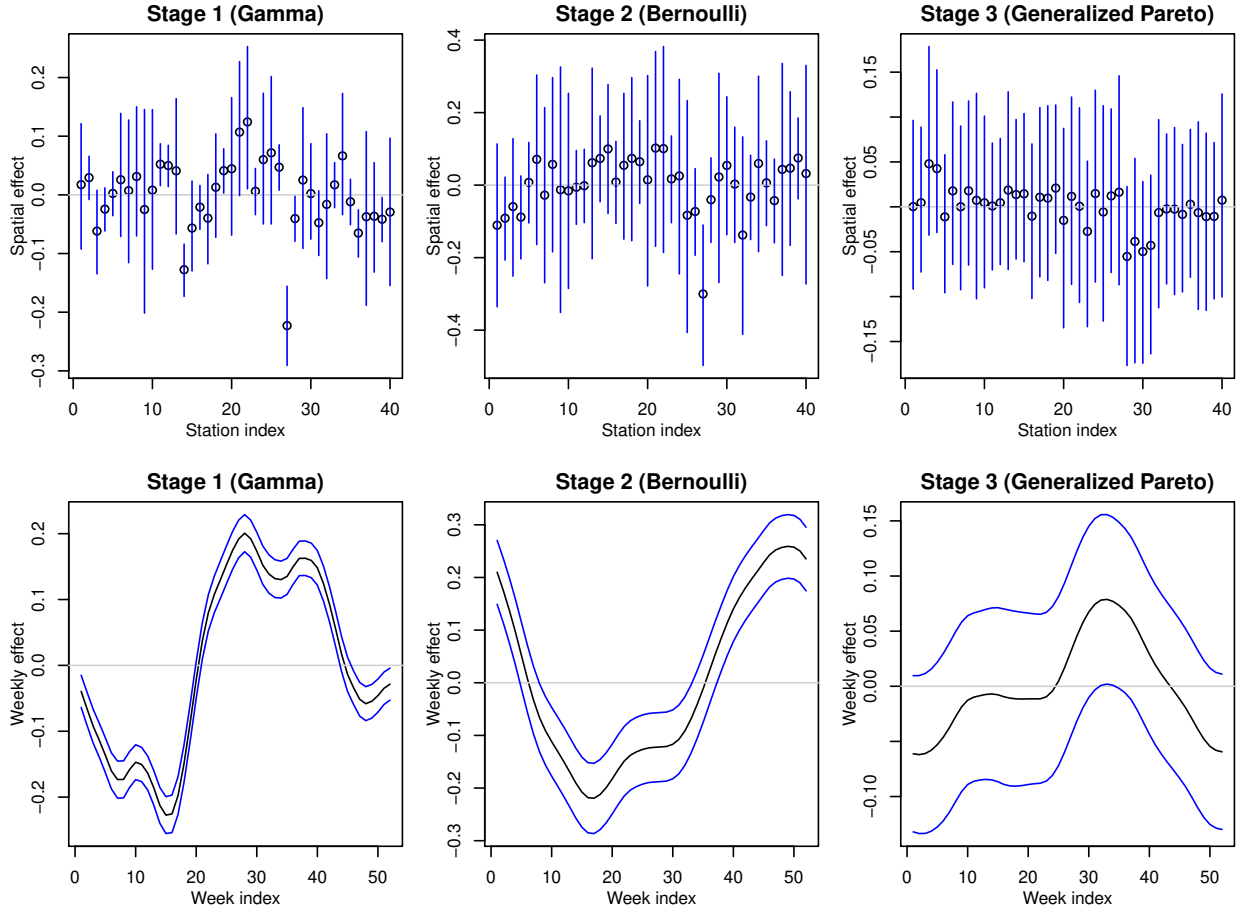


Figure 5: Spatial (top) and weekly (bottom) effects displayed for the three stages (left to right) of our selected “best model” (Model 122 in Table 1). Black dots and curves show the posterior means, while blue segments and curves are 95% pointwise credible intervals.

## 5.2 Final results and interpretation

Here, we provide further results and interpretation for our selected model. Figure 5 shows the posterior means and 95% pointwise credible intervals estimated for the spatial and weekly random effects in each of the three stages (i.e., Gamma, Bernoulli and GP models), while Figure 2 maps the spatial effect estimated in the first stage (Gamma). Overall, the uncertainty in the Gamma model is much lower than the Bernoulli and GP models, as it uses more information. Although these results neglect model selection uncertainty and model misspecification due to the conditional independence assumption of the data given the la-



tent parameters, the weekly effect appears overall highly significantly different from zero for the Gamma model, while the GP model fitted to threshold exceedances (i.e., to the top 8% of positive precipitation data) seems close to stationary in space and time when taking the uncertainty into account. As the GP scale parameter involves the Gamma mean as a multiplicative offset, this confirms that shape and size of non-stationary patterns are similar in the bulk and the tail of the distribution. The spatial patterns look quite stationary overall.

As expected, the credible intervals for the spatial effect in the Gamma model at time-rich stations are much narrower than those at time-poor stations; for example, Station 12, with  $n = 8317$  observations and surrounded by many other stations in the middle of the study region has the narrowest credible interval, while Station 9, with no observation and located on the Northern border of the study region has the widest credible interval. Interestingly, Stations 21 and 22, with very poor observational records, but surrounded by Stations 19, 23 and 26 with rich time series, have significant (or almost significant) positive effects, which confirms that our model succeeds in borrowing strength across nearby locations. Furthermore, the flexibility of our model is demonstrated by Station 27, located in the sea and far from the other stations, which stands out with a strong and highly significant negative effect; the corresponding time series shown in Figure 1 indeed reveals that this particular station is less exposed to intense precipitation than stations inland.

The weekly effect of the Gamma model is quite smooth and well estimated. It shows that precipitation intensity tends to be quite mild in April and stronger during summer, especially in July and September–October. The weekly effect of the Bernoulli model, however, has an opposite pattern and shows that the threshold  $u(s, t)$  is more frequently exceeded from September to February than during the other months, which suggests that there are more wet days during the fall and winter, as expected. The GP weekly effect finally reveals that the largest exceedances over the threshold (and hence, the most intense precipitation events)

tend to occur during the summer, though the associated uncertainty is high.

The GP tail index has posterior mean  $\xi = 0.34$ , with 95% credible interval  $(0.31, 0.38)$ , which shows a significant departure from an exponential density (arising with  $\xi = 0$ ) and reveals that precipitation data in the Netherlands are heavy-tailed with finite first and second moments. This also suggests that the posterior distribution for  $\xi$  is quite far from its prior displayed in Figure 3; although the prior for  $\xi$  was chosen to be quite informative with strong shrinkage towards light tails, the thousands of threshold exceedances available to fit the model provided strong evidence for heavier tails.

To conclude our analysis, Figure 6 displays monthly boxplots of the data together with our final  $\alpha_{\text{target}}$ -quantile predictions, for each month and station; the results are also compared to empirical quantiles computed for each month by pooling all stations together under the assumption of spatial stationarity. While our estimated weekly effect is quite weak but clearly visible, the difference between stations appears to be almost negligible, which suggests that we could have perhaps simplified the model by removing the spatial effect. Nevertheless, our Bayesian approach based on penalized complexity priors is designed to shrink the model towards a simple reference model, which helps in getting robust estimates and avoiding overfitting. Figure 6 also illustrates the ability of our spatio-temporal modeling approach based on Extreme-Value Theory to extrapolate predictions to extreme quantiles, even at locations where no data have been recorded.

## 6 Conclusion

Estimating extreme spatio-temporal quantiles from non-stationary data is not an easy task, and we think that the Extreme-Value Analysis conference 2017 challenge, which has motivated this work, has positively contributed to advancing the existing methods and the current literature.



Figure 6: Boxplots of the cubic root of precipitation intensities for each month and station. Red curves are our model-based predictions of the  $\alpha_{\text{target}}$ -quantile, with  $\alpha_{\text{target}} = 0.998$ , while blue curves are the empirical  $\alpha_{\text{target}}$ -quantiles pooling all stations together. Stations, where prediction was required, are indicated by a  $\star$  (Challenge 1) and  $\star\star$  (Challenges 1 & 2).

In this paper, we have combined the generalized Pareto distribution from Extreme-Value Theory with a flexible Bayesian latent Gaussian modeling approach. Our proposed three-stage model features spatial and temporal random effects that can capture systematic variations in the data. Our proposed generalized additive structures embedded in each stage provide interesting insight into the data’s bulk and tail behaviors and could be made more complex if required by the context. Although the dataset studied for this competition is fairly small and weakly non-stationary, our very fast and accurate inference approach based on INLA could be applied to extremely high-dimensional and highly non-stationary space-time data with additional hierarchical levels. Moreover, the combination of INLA and access to large distributed computing resources made it possible to run an extensive cross-validation analysis to select crucial hyperparameters and optimize our space-time quantile predictions.

Despite the model complexity, our approach based on penalized complexity (PC) priors guarantees stable and robust quantile predictions and contributes to avoid overfitting, which is especially important when predicting extremes. In this paper, we have derived the PC prior for the tail index, which controls the tail decay rate and therefore highly impacts extrapolations to high levels. Our proposed prior favors light exponential-like tail decay rates while penalizing unrealistically heavy tails.

Although the distribution in the first stage might need to be modified with other types of environmental data, our framework based on Extreme-Value Theory is very general and the methodology could be easily adapted to different contexts. Model estimation is executed using the R package `R-INLA`, which is convenient and easy to use. For reproducibility purposes, our code is illustrated on <http://www.r-inla.org/examples/case-studies>.

Although predicted quantiles from our model were quite good overall, certain aspects could have been improved. In particular, excluding months instead of years in our cross-validation study might have improved the estimation of the temporal effect, defined on a

weekly basis. Furthermore, it would have been possible to cross-validate the tail index, as well, which is a crucial parameter for predicting extreme quantiles. However, because of our limited (though large) computational resources, we opted to perform our cross-validation study on parameters that control bias-variance trade-offs relevant for optimizing spatio-temporal prediction. Finally, we could have considered a more flexible model by allowing, for example, the tail index to vary over seasons or months, instead of being constant over space and time. However, this would increase the number of hyperparameters in the model, whose estimation might be more tricky and computer-demanding to perform using INLA.

In future work, it would be interesting to replace the generalized Pareto distribution by more flexible sub-asymptotic tail models (see, e.g., [Papastathopoulos and Tawn, 2013](#); [Naveau \*et al.\*, 2016](#)), which could potentially be applied at much lower thresholds. Furthermore, as our model assumes that the data are conditionally independent given the latent parameters similarly to [Cooley \*et al.\* \(2007\)](#), it is not well-suited to properly capture space-time dependence, and it would be interesting to extend our methodological framework to more complex spatial ([Wadsworth and Tawn, 2012](#); [Thibaud \*et al.\*, 2013](#); [Thibaud and Opitz, 2015](#); [Huser \*et al.\*, 2017](#); [Castro Camilo and Huser, 2017](#); [Huser and Wadsworth, 2018](#)) and spatio-temporal ([Davis \*et al.\*, 2013](#); [Huser and Davison, 2014](#); [Bacro \*et al.\*, 2017](#)) models for extremes. However, such extremal models are much more difficult and expensive to fit, and when the primary goal is to estimate high *marginal* quantiles as in this work, the exact characterization of the dependence structure is a secondary issue, or even perhaps a nuisance.

## Acknowledgements

The research reported in this publication was supported by funding from King Abdullah University of Science and Technology (KAUST). Support from the KAUST Supercomputing Laboratory and access to Shaheen II is gratefully acknowledged.

## References

- Bacro, J.-N., Gaetan, C., Opitz, T. and Toulemonde, G. (2017) Hierarchical space-time modeling of exceedances with an application to rainfall data. arXiv preprint 1708.02447.
- Balkema, A. A. and de Haan, L. (1974) Residual life time at great age. *Annals of Probability* **2**(5), 792–804.
- Banerjee, S., Carlin, B. P. and Gelfand, A. E. (2004) *Hierarchical Modeling and Analysis for Spatial Data*. Boca Raton, FL: Chapman & Hall/CRC.
- Beirlant, J., Dierckx, G., Goegebeur, Y. and Matthys, G. (1999) Tail index estimation and an exponential regression model. *Extremes* **2**(2), 177–200.
- Bopp, G. P. and Shaby, B. A. (2017) An exponential–gamma mixture model for extreme Santa Ana winds. *Environmetrics* **28**(8).
- Casson, E. and Coles, S. (1999) Spatial regression models for extremes. *Extremes* **1**(4), 449–468.
- Castro Camilo, D. and Huser, R. (2017) Local likelihood estimation of complex tail dependence structures in high dimensions, applied to US precipitation extremes. arXiv preprint 1710.00875.
- Chavez-Demoulin, V. and Davison, A. C. (2005) Generalized additive modelling of sample extremes. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **54**(1), 207–222.
- Coles, S. (2001) *An Introduction to Statistical Modeling of Extreme Values*. London: Springer.
- Coles, S. G. and Tawn, J. A. (1996) Modelling Extremes of the Areal Rainfall Process. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **58**(2), 329–347.
- Cooley, D., Nychka, D. and Naveau, P. (2007) Bayesian spatial modeling of extreme precipitation return levels. *Journal of the American Statistical Association* **102**(479), 824–840.

- Daouia, A., Girard, S. and Stupfler, G. (2018) Estimation of tail risk based on extreme expectiles. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **80**(2), 263–292.
- Davis, R. A., Klüppelberg, C. and Steinkohl, C. (2013) Max-stable processes for modeling extremes observed in space and time. *Journal of the Korean Statistical Society* **42**(3), 399–414.
- Davison, A. C. and Huser, R. (2015) Statistics of extremes. *Annual Review of Statistics and its Application* **2**, 203–235.
- Davison, A. C. and Ramesh, N. (2000) Local likelihood smoothing of sample extremes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **62**(1), 191–208.
- Davison, A. C. and Smith, R. L. (1990) Models for exceedances over high thresholds (with discussion). *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **52**(3), 393–442.
- Frigessi, A., Haug, O. and Rue, H. (2002) A dynamic mixture model for unsupervised tail estimation without threshold selection. *Extremes* **5**(3), 219–235.
- Gumbel, E. J. (1958) *Statistics of Extremes*. New York: Columbia University Press.
- Hosking, J. R. and Wallis, J. R. (1987) Parameter and quantile estimation for the generalized Pareto distribution. *Technometrics* **29**(3), 339–349.
- Huser, R. and Davison, A. C. (2014) Space-time modelling of extreme events. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **76**(2), 439–461.
- Huser, R., Opitz, T. and Thibaud, E. (2017) Bridging asymptotic independence and dependence in spatial extremes using gaussian scale mixtures. *Spatial Statistics* **21**(A), 166–186.
- Huser, R. and Wadsworth, J. L. (2018) Modeling spatial processes with unknown extremal dependence. *Journal of the American Statistical Association* To appear.
- Jonathan, P., Randell, D., Wu, Y. and Ewans, K. (2014) Return level estimation from non-stationary spatial data exhibiting multidimensional covariate effects. *Ocean Engineering* **88**, 520–532.

- Katz, R. W., Parlange, M. and Naveau, P. (2002) Extremes in hydrology. *Advances in Water Resources* **25**(8–12), 1287–1304.
- Koenker, R. (2005) *Quantile Regression*. Cambridge: Cambridge University Press.
- Lindgren, F. K. and Rue, H. (2008) A note on the second order random walk model for irregular locations. *Scandinavian Journal of Statistics* **35**(4), 691–700.
- Lindgren, F. K. and Rue, H. (2015) Bayesian spatial modelling with R-INLA. *Journal of Statistical Software* **63**(19), 1–25.
- Lindgren, F. K., Rue, H. and Lindström, J. (2011) An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach (with discussion). *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **73**(4), 423–498.
- Martins, T. G., Simpson, D. P., Lindgren, F. K. and Rue, H. (2013) Bayesian computing with INLA: New features. *Computational Statistics & Data Analysis* **67**, 68–83.
- Naveau, P., Huser, R., Ribereau, P. and Hannart, A. (2016) Modeling jointly low, moderate, and heavy rainfall intensities without a threshold selection. *Water Resources Research* **52**(4), 2753–2769.
- Northrop, P. J. and Jonathan, P. (2011) Threshold modelling of spatially-dependent non-stationary extremes with application to hurricane-induced wave heights (with discussion). *Environmetrics* **22**(7), 799–809.
- Opitz, T. (2017) Latent Gaussian modeling and INLA: A review with focus on space-time applications. *Journal of the French Statistical Society* **158**(3), 62–85.
- Papastathopoulos, I. and Tawn, J. A. (2013) Extended generalised Pareto models for tail estimation. *Journal of Statistical Planning and Inference* **143**(1), 131–143.
- Pickands, J. (1975) Statistical inference using extreme order statistics. *Annals of Statistics* **3**(1), 119–131.
- Rue, H. and Held, L. (2005) *Gaussian Markov Random Fields: Theory and Applications*. London: Chapman & Hall/CRC.



- Rue, H., Martino, S. and Chopin, N. (2009) Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of the Royal Statistical Society (Series B)* **71**(2), 319–392.
- Rue, H., Riebler, A., Sørbye, S. H., Illian, J. B., Simpson, D. P. and Lindgren, F. K. (2017) Bayesian computing with INLA: A review. *Annual Review of Statistics and Its Application* **4**, 395–421.
- Scarrott, C. and MacDonald, A. (2012) A Review of Extreme Value Threshold Estimation And Uncertainty Quantification. *REVSTAT* **10**(1), 33–60.
- Simpson, D. P., Rue, H., Riebler, A., Martins, T. G. and Sørbye, S. H. (2017) Penalising model component complexity: A principled, practical approach to constructing priors. *Statistical Science* **32**(1), 1–28.
- Stein, M. L. (1999) *Interpolation of Spatial Data: Some Theory for Kriging*. New York: Springer.
- Tancredi, A., Anderson, C. and O’Hagan, A. (2006) Accounting for threshold uncertainty in extreme value estimation. *Extremes* **9**(2), 87–106.
- Thibaud, E., Mutzner, R. and Davison, A. C. (2013) Threshold modeling of extreme spatial rainfall. *Water Resources Research* **49**(8), 4633–4644.
- Thibaud, E. and Opitz, T. (2015) Efficient inference and simulation for elliptical Pareto processes. *Biometrika* **102**(4), 855–870.
- Tierney, L. and Kadane, J. B. (1986) Accurate approximations for posterior moments and marginal densities. *Journal of the American Statistical Association* **81**(393), 82–86.
- Wadsworth, J. L. and Tawn, J. A. (2012) Dependence modelling for spatial extremes. *Biometrika* **99**(2), 253–272.
- Wilks, D. (2006) *Statistical Methods in the Atmospheric Sciences*. Oxford: Elsevier.
- Wintenberger, O. (2018) Editorial: special issue on the Extreme Value Analysis conference challenge “Prediction of extremal precipitation”. *Extremes* To appear.