

# MODELLING EXTREMES OF SPATIAL AGGREGATES OF PRECIPITATION USING CONDITIONAL METHODS

BY JORDAN RICHARDS<sup>1,2</sup>, JONATHAN A. TAWN<sup>1</sup> AND SIMON BROWN<sup>3</sup>

<sup>1</sup>*STOR-i Centre for Doctoral Training, Department of Mathematics and Statistics, Lancaster University, LA1 4YR, UK*

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

<sup>3</sup>*Hadley Centre for Climate Prediction and Research, Met Office, FitzRoy Road, Exeter, EX1 3PB, UK*

Inference on the extremal behaviour of spatial aggregates of precipitation is important for quantifying river flood risk. There are two classes of previous approach, with one failing to ensure self-consistency in inference across different regions of aggregation and the other imposing highly restrictive assumptions. To overcome these issues, we propose a model for high-resolution precipitation data, from which we can simulate realistic fields and explore the behaviour of spatial aggregates. Recent developments have seen spatial extensions of the [Heffernan and Tawn \(2004\)](#) model for conditional multivariate extremes, which can handle a wide range of dependence structures. Our contribution is twofold: extensions and improvements of this approach and its model inference for high-dimensional data; and a novel framework for deriving aggregates addressing edge effects and sub-regions without rain. We apply our modelling approach to gridded East-Anglia, UK precipitation data. Return-level curves for spatial aggregates over different regions of various sizes are estimated and shown to fit very well to the data.

**1. Introduction.** Fluvial flooding is typically not caused by high intensity extreme rainfall at single locations, but by the extremes of precipitation events which are aggregated over spatial catchment areas. Accurate modelling of such events for both the present and the future can help to mitigate the financial impacts associated with floods, for example if river defences are built to withstand a  $T$ -year event of this kind. Approaches to quantifying the tail behaviour of spatial aggregates exist in the literature; however, these techniques are often simplistic or make unrealistic assumptions about the behaviour of the process for which they are trying to model. We present a novel methodology for making inference on the tail behaviour of spatial aggregates, which we apply in the context of extreme precipitation aggregates.

We define a spatial process  $\{Y(s) : s \in \mathcal{S}\}$  for some spatial domain  $\mathcal{S}$ . Our interest lies in the upper tail behaviour of the spatial aggregate  $R_{\mathcal{A}}$  on regions  $\mathcal{A} \subset \mathcal{S} \subset \mathbb{R}^2$ ,

$$(1) \quad R_{\mathcal{A}} = \int_{\mathcal{A}} Y(s) ds,$$

for different, possibly overlapping,  $\mathcal{A}$ , and the joint behaviour of  $(R_{\mathcal{A}}, R_{\mathcal{B}})$  for  $\mathcal{A}, \mathcal{B} \subset \mathcal{S}$ . Typically, the data we would have available for inference are realisations of  $\mathbf{Y}_t = (Y_t(s_1), \dots, Y_t(s_d))$  for  $t = 1, \dots, n$ , which are observations of said process  $\{Y(s)\}$  at  $d$  sampling locations  $\mathbf{s} = (s_1, \dots, s_d) \subset \mathcal{S}$  at  $n$  sampling times. Note that  $\mathbf{s}$  need not be point locations; they can instead be non-overlapping grid-boxes. Data produced by climate models are often available in this form and observations of  $Y(s_i)$ ,  $i = 1, \dots, d$ , correspond to spatial

---

*Keywords and phrases:* extremal dependence, extreme precipitation, spatial aggregates, spatial conditional extremes.

aggregates themselves, as they are typically presented as an average over the grid box  $s_i$ . In these circumstances, the integral in (1) can be replaced with the equivalent summation, but our methodology is still applicable; see Section 4.

Modelling extreme rainfall events over different spatial aggregation regions is difficult due to spatially-localised and high intensity convection rainfalls. For summer precipitation, convection events are the primary cause of extreme spatial aggregates of rainfall, even over regions much larger than the scale of convection events. This poses significant challenges when using data from any observation network of rain gauges, with such localised events being either recorded at a single rain gauge or missed entirely, with the latter more likely if the network is sparse or the observation window is short. In contrast, radar networks provide spatially complete coverage but are comparatively short in duration and have their own issues, such as calibration and attenuation (Harrison, Driscoll and Kitchen, 2000). The use of climate/numerical models to provide inference on extremes of real-world processes is now a standard approach, see e.g., Cooley and Sain (2010); Northrop and Jonathan (2011); Brown (2020) and Huser and Wadsworth (2020), which we adopt. However, early generation climate models failed to model convective rainfall due to insufficient spatial resolution. Currently, the only viable approach to capture the physics driving extreme convection events, and any future changes, is to use a new generation of high-resolution climate models that explicitly resolve convection (Kendon et al., 2014). In addition to enabling the study of extreme convective rainfall events, these models reproduce well the empirical properties of observational data, benefit from long temporal runs, have no missing data, and provide the opportunity to explore future climatic change scenarios.

We are interested in modelling extreme aggregated precipitation data for regions in East-Anglia, UK. We use data from the UKCP18 local model projections (Kendon et al., 2019; Met Office Hadley Centre, 2019). Analysing such high-resolution data poses significant new computational challenges given the large number of sites in our chosen domain, with  $d = 934$  exceeding the vast majority of previous spatial extremes studies. The lack of missing data enables us to perform simple and fair comparisons of our approach with existing methods.

We assume that both the full marginal behaviour, and dependence, of  $\{Y_t(s)\}$  is stationary with respect to time. Marginally, the upper tail behaviour of  $Y(s)$  is assumed to be characterised by a generalised Pareto distribution (GPD) with scale and shape parameters,  $v(s) > 0$  and  $\xi(s)$ , respectively, that vary smoothly over  $s \in \mathcal{S}$  (Davison and Smith, 1990). Dependence in  $\{Y(s)\}$  is characterised through a marginal transformation to the process  $\{X(s) : s \in \mathcal{S}\}$ , which has standardised margins; see Section 2.2. Letting  $F_{Y(s)}(\cdot)$  and  $F_X(\cdot)$  be the marginal CDFs of  $\{Y(s)\}$  and  $\{X(s)\}$ , respectively, we rewrite (1) as

$$R_{\mathcal{A}} = \int_{\mathcal{A}} Y(s) ds = \int_{\mathcal{A}} F_{Y(s)}^{-1} \{F_X[X(s)]\} ds.$$

We focus on the situation where  $\{X(s)\}$  is a stationary process; an assumption that we find holds well for our application (see Section 4.4). If this assumption did not hold, both non-parametric, and parametric, methods exist that account for non-stationarity in extremal dependence (Huser and Genton, 2016; Castro-Camilo and Huser, 2020; Richards and Wadsworth, 2021), and these methods can easily be incorporated into our methodology.

As  $R_{\mathcal{A}}$  is essentially a sum of variables, a natural starting point is to consider extreme value theory for sums, i.e., when  $\mathcal{A}$  is a finite collection of indices. Consider the case where  $\xi(s)$  varies over  $\mathcal{A}$  and let  $\xi_{\mathcal{A}} = \max_{s \in \mathcal{A}} \xi(s)$ . If  $\xi_{\mathcal{A}} > 0$ , then Breiman's lemma (Breiman, 1965) gives that the upper-tail of  $R_{\mathcal{A}}$  is determined solely by the behaviour at sites  $\mathcal{A}^* = \{s \in \mathcal{A}, \xi(s) = \xi_{\mathcal{A}}\}$ . Connected results are given by Richards and Tawn (2021) for  $\xi_{\mathcal{A}} \leq 0$ . When the value of  $\xi_{\mathcal{A}} - \max_{s \in \mathcal{A} \setminus \mathcal{A}^*} \xi(s)$  is small, convergence to the limiting extremal behaviour of  $R_{\mathcal{A}}$  will be slow as the behaviour at the sites not in  $\mathcal{A}^*$  is completely ignored. Another

perspective is to view  $R_{\mathcal{A}}$  as an infinite sum, and so the central limit theorem may apply when  $\xi_{\mathcal{A}} < \frac{1}{2}$  (which ensures finite marginal variances); which if correct would imply that the upper-tail distribution of  $R_{\mathcal{A}}$  should be increasingly more likely to be exponential as  $\|\mathcal{A}\|$  increases, i.e.,  $\xi = 0$  (due to characteristics of maxima of Gaussian random variables), irrespective of whether or not  $\xi(s)$  varies over  $\mathcal{A}$ . Generally this is a false argument, due to Breiman's lemma, when  $\xi_{\mathcal{A}} > 0$ ; however, it is supported in certain circumstances by results in [Richards and Tawn \(2021\)](#) when  $\xi_{\mathcal{A}} < 0$ . We consider regions  $\mathcal{A}$  which are small relative to lags at which spatial independence occurs, so we lack sufficient independence for central limit theorem arguments to apply.

There are three main existing modelling approaches for inference on the upper tails of  $R_{\mathcal{A}}$ : univariate methods, spatial approaches that focus on modelling all of the data, and spatial approaches that focus on modelling only the extremes; our approach falls in the latter class, making less restrictive assumptions than previous methods of this type.

We first consider the univariate case. Within an extreme value analysis framework, univariate methods for estimating the size of  $T$ -year events are well studied and cemented in asymptotic theory ([Coles, 2001](#)). If we can create a sample of observations of  $R_{\mathcal{A}}$ , we can use univariate methods to make inference on its upper tail, i.e., fit a GPD to exceedances of a sample of  $R_{\mathcal{A}}$  above some fixed threshold and then extrapolate to high quantiles. However, creating this sample can be challenging if  $s$  are irregularly spaced point locations. Further complications arise if we have partially missing observations of  $\{Y(s)\}$ . Even if these issues are overcome, when using univariate methods we lose the information present in the margins of  $\{Y(s)\}$  and dependence of  $\{X(s)\}$ . If the process we are considering is precipitation, this can lead to inference that is not self-consistent and may be physically unrealistic; a trait that can be undesirable to practitioners. To explain this further, observe that, for precipitation,  $\{Y(s)\}$  is non-negative everywhere, i.e.,  $Y(s) \geq 0$  for all  $s \in \mathcal{S}$ . Trivially it follows that  $R_{\mathcal{A}} \geq R_{\mathcal{B}}$  for all  $\mathcal{B} \subseteq \mathcal{A} \subset \mathcal{S}$  and hence return levels should be similarly ordered. This natural ordering may not follow if we take a simple univariate approach to modelling the upper tail behaviour of the  $R_{\mathcal{A}}$  and  $R_{\mathcal{B}}$  aggregates separately ([Nadarajah, Anderson and Tawn, 1998](#)).

In the context of precipitation aggregates, one richly studied approach has developed a class of stationary stochastic processes to model the whole precipitation intensity process, continuous in both time and space. These models typically describe the intensity as the accumulation, at each point in time and space, over a random number of simple shaped individual stochastic rain cells, which cluster in time and space, and move on stochastic trajectories. These models were first developed for a single site by [Rodriguez-Iturbe, Cox and Isham \(1987\)](#), then developed spatially and for fine-scale processes, see [Benoit, Allard and Mariethoz \(2018\)](#). These models are typically estimated by optimising the fit against a range of characteristics of observed fields. As a result, these models can capture well the broad-scale features of typical precipitation fields. However, for deriving the upper-tail of  $R_{\mathcal{A}}$ , the models and their inference have limitations as there is no guarantee that models for the body of a process fit well to the extremes. Yet it is precipitation fields that are extreme somewhere in  $\mathcal{A}$  that yield extremes of  $R_{\mathcal{A}}$  unless  $\mathcal{A}$  is very large relative to the range of spatial dependence.

From the perspective of spatial extremal methods,  $R_{\mathcal{A}}$  is simply a functional of a spatial field which is extreme in some way. A typical approach to modelling extreme fields is the use of max-stable models, see [Padoan, Ribatet and Sisson \(2010\)](#); [Westra and Sisson \(2011\)](#); [Reich and Shaby \(2012\)](#). These models are predominately fit to component-wise block maxima, typically annual maxima, at sampling locations. Typically annual maxima do not occur concurrently for all sampling locations, do not give zeros, and so aggregating over realisations from a max-stable process is not appropriate for inference on aggregates of extreme precipitation events. [Coles \(1993\)](#) rectified some of these issues by using a point-process representation of a max-stable field to derive the spatial profile of concurrent events.

Coles and Tawn (1996) used this formulation to derive closed form results for the tail behaviour of  $R_{\mathcal{A}}$  where the tail parameters are determined by the marginal GPD parameters of  $\{Y(s)\}$  and its dependence structure; Ferreira, de Haan and Zhou (2012) formalise these results and provide some non-parametric extensions. Further extensions of this framework by Engelke, De Fondeville and Oesting (2018) relate not only the extremal behaviour of  $\{Y(s)\}$  and the aggregates  $R_{\mathcal{A}}$ , but also the joint behaviour of aggregates over different regions,  $\mathcal{A}$ . de Fondeville and Davison (2021) use functional Pareto processes to model the dependence in  $\{Y(s)\}$  and Palacios-Rodríguez et al. (2020) illustrate non-parametric Pareto process modelling to simulate extreme precipitation fields, re-sampling event profiles from observed, gridded data. All of these modelling approaches rely on the marginal shape parameters of  $\{Y(s)\}$  to be spatially homogeneous i.e.,  $\xi(s) = \xi$  for all  $s \in \mathcal{A}$ , for each  $\mathcal{A}$  of interest. Furthermore they all have major limitations for applications due to their restrictive choice of dependence structure, as described below.

The co-occurrence of extremal events is often quantified through the upper tail index  $\chi(s_A, s_B)$  (Joe, 1997) for all  $s_A, s_B \in \mathcal{S}$ , which can be defined for  $\{Y(s)\}$  as  $\chi(s_A, s_B) = \lim_{q \uparrow 1} \chi_q(s_A, s_B)$ , where

$$(2) \quad \chi_q(s_A, s_B) = \Pr\{Y(s_B) > F_{Y(s_B)}^{-1}(q) | Y(s_A) > F_{Y(s_A)}^{-1}(q)\}.$$

Positively associated max-stable, or Pareto, processes are asymptotically dependent (Coles, 2001) at all spatial distances, i.e.,  $\chi(s_A, s_B) > 0$  for all  $s_A, s_B \in \mathcal{S}$ . These models are then unable to account for cases where we have positive association, but  $\chi(s_A, s_B) = 0$  for some  $s_A, s_B \in \mathcal{S}$  which holds for all Gaussian processes when  $s_A \neq s_B$ ; we refer to this scenario as asymptotic independence, i.e., the tendency for extreme events to occur increasingly independently as the magnitude of the events gets larger.

Wadsworth and Tawn (2019) have developed a conditional approach to spatial extremes which allows for both asymptotic dependence and asymptotic independence at different spatial distances. They provide a spatial extension of the multivariate Heffernan and Tawn (2004) model, which enables the modelling of processes given that the process is extreme at least at one location. Dependence parameters within the Heffernan and Tawn (2004) model are represented as smooth functions, parametric or splines, of distance between variables at the site of interest and the conditioning site, and the residual process is driven by a latent Gaussian process, see Section 2.2. Extensions of this approach are given by Shooter et al. (2020); Simpson, Opitz and Wadsworth (2021); Huser and Wadsworth (2020); Simpson and Wadsworth (2021). These papers provide alternative parametric norming functions, implement Bayesian inference, spatio-temporal extensions, and involve small numbers of sampling locations ( $d < 300$ ), so that full inference is computationally feasible. One exception is Simpson, Opitz and Wadsworth (2021) who detail an inference approach using INLA for much larger  $d$ ; however, this imposes restrictions on the dependence structure parameters that are not appropriate in our application. Across the papers, they perform extremal modelling of air and sea temperature fields and spatial wave heights, each of which is spatially smoother than precipitation and are without the issue of zeros. In particular, their primary descriptor of the extremal dependence comes through a location function, see (6).

Compared to the use of limit models for  $R_{\mathcal{A}}$ , a much better sub-asymptotic approximation is likely to come from modelling the margins and copula of  $\{Y(s)\}$  and deriving the high quantile behaviour of  $R_{\mathcal{A}}$  through computational methods, such as simulation. This is the approach we take but for  $\{Y(s)\}$  as a continuous spatial process. We build on the spatial conditional extremes methods and adapt this approach for modelling extreme precipitation fields, i.e., addressing the issues caused by the zeros in the field. We present a novel and flexible solution to the problem of modelling the extremes of spatial aggregated variables and provide improvements to the modelling and inference of spatial extreme fields generally.

Using this spatial model we construct a sub-asymptotic tail model for  $R_{\mathcal{A}}$ . Our approach overcomes the problems of the existing univariate method with ordering  $R_{\mathcal{A}}$  over different  $\mathcal{A}$  and dealing with partially missing data, and is more flexible than approaches that restrict the spatial dependence in  $\{Y(s)\}$  to be either asymptotic dependence or asymptotic independence everywhere or impose that  $\xi(s)$  is constant over  $s \in \mathcal{A}$ . We find that when using Monte-Carlo methods to approximate  $\Pr\{R_{\mathcal{A}} > r\}$ , that the position and size of  $\mathcal{A}$  within  $\mathcal{S}$  are important considerations in the inference for  $R_{\mathcal{A}}$ , which require new methods to detect and account for edge effects. Due to the high value of  $d$  in our application, there are serious problems with evaluation of likelihood and confidence interval construction for quantiles of  $R_{\mathcal{A}}$ , so we adapt recently proposed stratified sampling schemes for model fitting. We also propose new parametric forms for the spatial dependence functions. We observe high variability in the fitted model for  $\{Y(s)\}$  as we move further away from the centre of an event, which corresponds to the observed roughness in events that generate extreme precipitation, which unlike applications for other physical processes requires specific care with modelling the scale function, see (7). We find strong indication that our model for  $\{Y(s)\}$  fits well and that we are able to handle comfortably the structure of observed extremes fields, including their zeroes, and estimate very well the distribution of  $R_{\mathcal{A}}$  for our data.

The layout of this paper is as follows: Section 2 describes our model for the process  $\{Y(s)\}$ . We describe methods for model inference and simulation of events in Section 3, which includes our censoring technique for handling zero values. In Section 4 we discuss the marginal and dependence model fits for the precipitation data, and inference on the tail behaviour of spatial aggregates of these data. We compare the results from our approach with those using GPD fitted to the sample aggregates and using a spatial asymptotically dependent model in Section 4.4. We end with further discussion and model extensions in Section 5.

## 2. Modelling the extremes of the spatial process.

2.1. *Marginal Model.* The site-wise marginals of  $\{Y(s)\}$  can be modelled using a fitted GPD above some high threshold and the empirical distribution below (Coles, 2001). We extend this approach by incorporating a third component, which we denote  $p(s)$ , that describes the probability that there is no rain at site  $s$ . The marginal distribution function of  $Y(s)$  for each  $s \in \mathcal{S}$  is

$$(3) \quad F_{Y(s)}(y) = \begin{cases} p(s) & \text{if } y = 0 \\ \frac{1-\lambda(s)-p(s)}{F_{Y_+(s)}(q(s))} F_{Y_+(s)}(y) + p(s) & \text{if } 0 < y \leq q(s) \\ 1 - \lambda(s) \left[ 1 + \frac{\xi(s)(y-q(s))}{v(s)} \right]_+^{-1/\xi(s)} & \text{if } y > q(s), \end{cases}$$

where  $v(s) > 0$  and  $F_{Y_+(s)}(y)$  denotes the distribution function of strictly positive values of  $Y(s)$  and  $p(s) \geq 0$ ,  $\lambda(s) > 0$  and  $p(s) + \lambda(s) < 1$ ; this ensures that the marginal distribution is continuous across components. We expect spatial smoothness over  $F_{Y(s)}$  and so define full spatial models for the three components of (3) which also enable us to make inference about  $F_{Y(s)}$  for all  $s \in \mathcal{S}$ , i.e., including where  $Y(s)$  is not observed.

We first consider the distribution of  $Y(s)$  above  $q(s)$ . Following the approach of Youngman (2019), we fit a generalised additive GPD model (GAM) to exceedances  $Y(s) - q(s)$ . This allows us to represent the GPD scale and shape parameters,  $v(s)$  and  $\xi(s)$ , respectively, through a basis of smooth splines. We set  $\lambda(s) = \lambda$  for all  $s \in \mathcal{S}$ , allowing us to estimate  $q(s)$  for  $s \in \mathcal{S}$ . We use a non-parametric approach, and simply fit a thin-plate spline to point-wise estimates of  $q(s)$  for the associated  $\lambda$ , as the parametric method of Youngman (2019) fails as we have point masses below  $q(s)$ , caused by rounding of data.

We estimate  $p(s) = \Pr\{Y(s) = 0\}$  for  $s \in \mathcal{S}$  as a spatially smooth surface by using a logistic GAM (Wood, 2006); that is, we fit  $\text{logit}\{\mathbb{E}[p(s)]\} = g(s)$  where  $g(\cdot)$  is a smoothing spline. The degree of smoothness in the surface  $p(s)$  is determined by the choice of spline used for  $g$ . We estimate  $F_{Y_+(s)}$  using the empirical distribution of strictly positive values of  $Y(s)$ , which we denote  $\hat{F}_{Y_+(s)}(\cdot)$ . We use the site-wise empirical distribution for fitting and recognise that, should we require simulation of  $Y(s)$  for  $s \in \mathcal{S} \setminus \mathbf{s}$ , we can use additive quantile regression (Fasiolo et al., 2021) to compute the empirical distributions at unobserved locations. This can be performed, if necessary, using only a local neighbourhood of sampling locations, as this method is too computationally expensive for large  $d$  and so is not viable without dimension reduction.

We use (3) to perform a site-wise transformation of the margins of the data to a standardised distribution, albeit with discrete mass in each lower-tail. For modelling dependence within a process  $\{X(s) : s \in \mathcal{S}\}$  using the Wadsworth and Tawn (2019) conditional extremes framework, we require its margins to have standard exponential upper tails, i.e.,  $\Pr\{X(s) > x\} \sim C \exp(-x)$  for some  $C > 0$ , as  $x \rightarrow \infty$  and for all  $s \in \mathcal{S}$ . We follow Wadsworth and Tawn (2019) and use Laplace margins. We handle discrete components in the margins of  $\{X(s)\}$  using censoring techniques, see Section 2.2.3.

**2.2. Dependence Modelling.** Wadsworth and Tawn (2019) model the underlying extremal dependence in our standardised process  $\{X(s)\}$ , given that it is extreme for some  $s \in \mathcal{S}$ , by first conditioning on the process being above some high threshold  $u$  at a specified site  $s_O \in \mathcal{S}$ . We introduce the function  $h(s_A, s_B) = \|s_A - s_B\|$  for  $s_A, s_B \in \mathcal{S}$ , where  $\|\cdot\|$  is some distance metric (we use the anisotropic measure (12)). Under the assumption that there exists normalising functions  $\{a : (\mathbb{R}, \mathbb{R}_+) \rightarrow \mathbb{R}\}$ , with  $a(x, 0) = x$ , and  $\{b : (\mathbb{R}, \mathbb{R}_+) \rightarrow (0, \infty)\}$ , such that as  $u \rightarrow \infty$ , they assume that for each  $s_O \in \mathcal{S}$

$$(4) \quad \left( \left\{ \frac{X(s) - a\{X(s_O), h(s, s_O)\}}{b\{X(s_O), h(s, s_O)\}} : s \in \mathcal{S} \right\}, X(s_O) - u \right) \Big| \left( X(s_O) > u \right) \xrightarrow{d} \left( \left\{ Z(s|s_O) : s \in \mathcal{S} \right\}, E \right),$$

where  $E$  is a standard exponential variable and process  $\{Z(s|s_O)\}$  which is non-degenerate for all  $s \in \mathcal{S}$  where  $s \neq s_O$ . That is, there is convergence in distribution of the normalised process to  $\{Z(s|s_O) : s \in \mathcal{S}\}$ , termed the residual process, which is independent of  $E$ , and  $Z(s_O|s_O) = 0$  almost surely. Characterisations of the normalising functions,  $a$  and  $b$  and the residual process  $Z(s|s_O)$  are given in Sections 2.2.1 and 2.2.2, respectively.

To make inference on the upper tail of  $R_{\mathcal{A}}$  for any  $\mathcal{A} \subset \mathcal{S}$  we require the process  $\{X(s)\}$  given an extreme value somewhere in the domain  $\mathcal{S}$ , i.e.,

$$(5) \quad \left\{ X(s) : s \in \mathcal{S} \right\} \Big| \left( \max_{s \in \mathcal{S}} X(s) > u \right)$$

for large  $u$ . Limit (4) conditions only on observing an exceedance at a specific site  $s_O \in \mathcal{S}$ , so cannot be immediately used. However, this limit provides a core building block for what is required when combined with a limiting model for  $\{X(s_O) > u : s_O \in \mathcal{S}\} | (\max_{s \in \mathcal{S}} X(s) > u)$  as  $u \rightarrow \infty$ . For a stationary process, this limiting model will be invariant to  $s_O$  for all  $s_O \in \mathcal{S}$  which are sufficiently far from the boundaries of  $\mathcal{S}$ . Wadsworth and Tawn (2019) show how simulation from process (5) can be achieved through an importance sampling method; the outline of this is given in Section 3.3.

For the process  $\{X(s)\}$  to exhibit long-range independence over  $\mathbb{R}^2$ , we need conditions on  $a$  and  $b$  and  $Z(s|s_O)$  so that independence is achieved as  $h = h(s, s_O) \rightarrow \infty$  for any  $s_O \in \mathcal{S}$  and suitably distanced  $s \in \mathcal{S}$ . This requires for any fixed  $x > 0$ , that  $a(x, h) \rightarrow 0$  and  $b(x, h) \rightarrow 1$  as  $h \rightarrow \infty$ . Further, the residual process  $Z(s|s_O)$  must have identical margins to  $X(s)$  as  $h \rightarrow \infty$ ; in particular, we require standard Laplace margins for  $Z(s|s_O)$  as  $h \rightarrow \infty$ .

2.2.1. *Normalising functions.* For inference, we assume parametric forms for the location and scaling functions  $a$  and  $b$ , respectively, in limit (4). [Wadsworth and Tawn \(2019\)](#) provide a discussion of these normalising functions and provide some suggestions for their possible parametric forms. We considered the range of parametric forms discussed by [Wadsworth and Tawn \(2019\)](#); [Shooter et al. \(2020\)](#), but for brevity we report only the models that provided the best fit. We let

$$(6) \quad a(x, h) = x\alpha(h), \quad \text{with } \alpha(h) = \begin{cases} 1, & h \leq \Delta, \\ \exp(-\{(h - \Delta)/\kappa_{\alpha_1}\}^{\kappa_{\alpha_2}}), & h > \Delta, \end{cases}$$

where  $\Delta \geq 0$  and  $\kappa_{\alpha_1}, \kappa_{\alpha_2} > 0$  which allows  $\{X(s)\}$  to be asymptotically dependent up to distance  $\Delta$  from  $s_O$ , and asymptotically independent thereafter. We also take

$$(7) \quad b(x, h) = x^{\beta(h)}, \quad \text{with } \beta(h) = \kappa_{\beta_3} \exp(-\{h/\kappa_{\beta_1}\}^{\kappa_{\beta_2}})$$

for  $\kappa_{\beta_1}, \kappa_{\beta_2} > 0$  and  $\kappa_{\beta_3} \in [0, 1]$ , and so  $b(0, x) = x^{\kappa_{\beta_3}}$ . Long-range independence holds for  $\{X(s)\}$ , as  $a(x, h) \rightarrow 0$  and  $b(x, h) \rightarrow 1$  as  $h \rightarrow \infty$  for fixed  $x > 0$ , whatever the parameters.

2.2.2. *Residual process  $\{Z(s|s_O)\}$ .* We follow [Shooter et al. \(2020\)](#) by imposing that the residual process  $\{Z(s|s_O)\}$  has delta-Laplace margins; a random variable follows a delta-Laplace distribution, i.e.,  $DL(\mu, \sigma, \delta)$ , with location, scale and shape parameters  $\mu \in \mathbb{R}$ ,  $\sigma > 0$  and  $\delta > 0$ , respectively, if its density is

$$(8) \quad f(z) = \frac{\delta}{2k\sigma\Gamma(\frac{1}{\delta})} \exp\left\{-\left|\frac{z - \mu}{k\sigma}\right|^\delta\right\}, \quad (z \in \mathbb{R})$$

with  $\Gamma(\cdot)$  as the standard gamma function and  $k^2 = \Gamma(1/\delta)/\Gamma(3/\delta)$ . The scaling by  $k$  is used to improve identifiability between  $\sigma$  and  $\delta$ , as the random variable has expectation  $\mu$  and variance  $\sigma^2$  regardless of the value of  $\delta$ . Use of the delta-Laplace distribution introduces flexibility in the marginal choice for  $Z(s|s_O)$ , as for  $\delta = 1$  or  $2$ , we have the Laplace or Gaussian distributions respectively. As with the normalising functions, we parametrise the delta-Laplace parameters as smooth functions of distance from the conditioning site  $s_O$ . That is,  $Z(s|s_O) \sim DL(\mu\{h(s, s_O)\}, \sigma\{h(s, s_O)\}, \delta\{h(s, s_O)\})$ , with

$$(9) \quad \begin{aligned} \mu(h) &= \kappa_{\mu_1} h^{\kappa_{\mu_2}} \exp\{-h/\kappa_{\mu_3}\}, & (\kappa_{\mu_2} > 0, \kappa_{\mu_3} > 0), \\ \sigma(h) &= \sqrt{2}(1 - \exp\{-(h/\kappa_{\sigma_1})^{\kappa_{\sigma_2}}\}), & (\kappa_{\sigma_1} > 0, \kappa_{\sigma_2} > 0), \\ \delta(h) &= \max\{1, 1 + (\kappa_{\delta_1} h^{\kappa_{\delta_2}} - \kappa_{\delta_4}) \exp\{-h/\kappa_{\delta_3}\}\}, & (\kappa_{\delta_1} \geq 0, \kappa_{\delta_2} > 0, \kappa_{\delta_3} > 0, \kappa_{\delta_4} < 1), \end{aligned}$$

for  $h \geq 0$ . These functions satisfy the constraint that  $\mu(0) = \sigma(0) = 0$ , which ensures that  $Z(s_O|s_O) = 0$  holds and provides a flexible modelling choice for  $\delta$ . Use of the *max* operator in  $\delta$  ensures that the tails of  $Z(s|s_O)$  are not heavier than  $X(s)$ , i.e.,  $\delta\{h(s, s_O)\} \geq 1$  for all  $s \in \mathcal{S}$ . Furthermore, long-range independence in  $X(s)$  is achieved as  $\mu(h) \rightarrow 0$ ,  $\sigma^2(h) \rightarrow 2$ , and  $\delta(h) \rightarrow 1$  as  $h \rightarrow \infty$ , where the variance of a standard Laplace random variable is 2.

Following the approach of [Shooter et al. \(2020\)](#), dependence in  $\{Z(s|s_O)\}$  is induced by first considering the process  $\{W(s|s_O)\} = \{W(s)|(W(s_O) = 0)\}$  for all  $s \in \mathcal{S}$ , where  $\{W(s)\}$  is a standard stationary Gaussian process with correlation function  $\rho(h)$ . We set  $\{Z(s|s_O)\} = \{F_{Z(s|s_O)}^{-1}\{\Phi[W(s|s_O)]\}\}$  for all  $s \in \mathcal{S}$ , where  $\Phi(\cdot)$  and  $F_{Z(s|s_O)}$  are the CDFs of a standard Gaussian distribution and  $Z(s|s_O)$ , respectively. The corresponding density function to  $F_{Z(s|s_O)}$  is  $f_{Z(s|s_O)}$ , defined by (8) and (9).

To illustrate the dependence in  $\{Z(s|s_O) : s \in \mathcal{S}\}$ , we consider the joint distribution of  $Z(s|s_O)$  in a finite-dimensional setting, which we achieve by using a Gaussian copula model.

Consider any  $s_O \in (s_1, \dots, s_d)$ , and without loss of generality, rewrite the sampling locations as  $s_O, (s_1, \dots, s_{d-1})$ , i.e., here we illustrate with  $s_O = s_d$ . The joint distribution of  $\{Z(s_1|s_O), \dots, Z(s_{d-1}|s_O)\}$  for  $s_O$  is, for  $\mathbf{z} = (z_1, \dots, z_{d-1})$ ,

$$(10) \quad F_{s_O}(\mathbf{z}) = \Phi_{d-1} \left\{ \Phi^{-1}(F_{Z(s_1|s_O)}(z_1)), \dots, \Phi^{-1}(F_{Z(s_{d-1}|s_O)}(z_{d-1})) \right\}; \mathbf{0}, \Sigma \},$$

where  $\Phi_{d-1}(\cdot; \mathbf{0}, \Sigma)$  is the CDF of a  $(d-1)$ -dimensional Gaussian distribution with mean  $\mathbf{0}$ . The correlation matrix  $\Sigma$  must account for the conditioning  $W(s)|(W(s_O) = 0)$ . To create  $\Sigma$ , we initialise a stationary correlation matrix  $\Sigma^*$  using correlation function  $\rho(\cdot)$  evaluated for all pairwise distances, and we condition on observing  $W(s_O) = 0$ . That is, the correlation matrix  $\Sigma$  has  $(i, j)$ -th element

$$(11) \quad \Sigma_{ij} = \frac{\Sigma_{ij}^* - \Sigma_{i0}^* \Sigma_{j0}^*}{(1 - \Sigma_{i0}^{*2})^{1/2} (1 - \Sigma_{j0}^{*2})^{1/2}}.$$

Note the elements of  $\Sigma$  are normalised such that the diagonal elements are equal to one. In our application,  $\rho(\cdot)$  is taken to be the Matérn correlation function

$$\rho(h) = \frac{2^{1-\kappa_{\rho_2}}}{\Gamma(\kappa_{\rho_2})} \left( \frac{2h\sqrt{\kappa_{\rho_2}}}{\kappa_{\rho_1}} \right)^{\kappa_{\rho_2}} K_{\kappa_{\rho_2}} \left( \frac{2h\sqrt{\kappa_{\rho_2}}}{\kappa_{\rho_1}} \right), \quad (\kappa_{\rho_1} > 0, \kappa_{\rho_2} > 0),$$

where  $K_{\kappa_{\rho_2}}(\cdot)$  is the modified Bessel function of the second kind of order  $\kappa_{\rho_2}$ .

To account for spatial anisotropy in the extremal dependence structure of  $\{X(s)\}$  we use the transformation of coordinates

$$(12) \quad s^* = \begin{pmatrix} 1 & 0 \\ 0 & 1/L \end{pmatrix} \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} s,$$

where  $\theta \in [-\pi/2, 0]$  controls rotation and  $L > 0$  controls the coordinate stretching effect; with  $L = 1$  recovering the isotropic model. We define our distance metric  $\|s_A - s_B\| = \|s_A^* - s_B^*\|_*$ , where  $\|\cdot\|_*$  denotes great-circle, or spherical, distance.

**2.2.3. Extensions for censored precipitation data.** The non-zero probability of zeroes for precipitation  $Y(s)$  causes the  $X(s)$ , for all  $s \in \mathcal{S}$ , to have non-zero mass at a finite lower endpoint. Consequently, the Gaussian copula and delta-Laplace marginal model described in (10) are not appropriate for the transformed precipitation data in these lower tail regions. To circumvent this issue, we apply censoring at all points where  $Y_t(s) = 0$  for all  $s \in \mathcal{S}$  and  $t = 1, \dots, n$ ; an approach adopted by [Benoit, Allard and Mariethoz \(2018\)](#), albeit for small  $d$ . A spatially-varying censoring threshold  $c(s)$  is attained by transforming  $p(s)$  in Section 2.1 to the Laplace scale, using  $c(s) = F_L^{-1}\{p(s)\}$ , where  $F_L(\cdot)$  is the standard Laplace CDF. We then assert that  $Y_t(s) = 0 \Leftrightarrow X_t(s) \leq c(s)$ , which in turn implies that  $Z_t(s|s_O) \leq c_t^{(s_O)}(s)$  where  $c_t^{(s_O)}(s)$  is dependent on the value observed at the conditioning site and given in (13).

For inference, the number of censored components varies at each time point; the number, locations and censoring values can all vary, with a maximum value of  $d-1$  locations with censoring. If the number of censored components is large, it is clear from (10) that evaluation of a censored distribution function will be computationally expensive. We take a pseudo-likelihood approach to inference, which we detail in Section 3.1; this requires only a bivariate density. We also detail its multivariate analogue; although this is not used for inference, it has a variety of uses, i.e., infilling of extreme events with missing observations or inference at sites  $s \in \mathcal{S} \setminus s$ . These features are detailed in Appendix A of the Supplementary Materials ([Richards, Tawn and Brown, 2021](#)).

### 3. Inference and Simulation.



3.1. *Model Fitting.* Our censored triplewise likelihood approach for model fitting is based on the pseudo-likelihood approach of [Padoan, Ribatet and Sisson \(2010\)](#); their pairwise approach provides unbiased estimation of model dependence parameters. Recall that some observations are right-censored at different sampling locations with varying rate of occurrence over time. To define a single likelihood contribution at time  $t$ , we begin by considering a single conditioning site amongst the observed sites  $s_i \in (s_1, \dots, s_d)$  such that  $y_t(s_i) > F_{Y(s_i)}^{-1}\{F_L(u)\}$  and hence  $x_t(s_i) > u$ . We then define the set of all such times by  $\mathcal{T}^{(s_i)} = \{t = 1, \dots, n : x_t(s_i) \geq u\}$ . For the observed sites, we define  $h_{i,j} = h(s_i, s_j)$  for  $i, j = 1, \dots, d$  with  $i \neq j$ . Then for each site  $s_j, j = 1, \dots, d, j \neq i$ , we define the residual for time  $t \in \mathcal{T}^{(s_i)}$  for conditioning site  $s_i$  as

$$(13) \quad z_t^{(s_i)}(s_j) = \begin{cases} [x_t(s_j) - a\{x_t(s_i), h_{i,j}\}]/b\{x_t(s_i), h_{i,j}\} & \text{if } x_t(s_j) > c(s_j), \\ c_t^{(s_i)}(s_j), & \text{otherwise,} \end{cases}$$

where  $a$  and  $b$  are described in Section 2.2.1, and  $c_t^{(s_i)}(s_j) = \frac{c(s_j) - a\{x_t(s_i), h_{i,j}\}}{b\{x_t(s_i), h_{i,j}\}}$  is the censored residual for site  $s_j$  with conditioning site  $s_i$  and  $t \in \mathcal{T}^{(s_i)}$ . The full pseudo-likelihood is given, for residuals  $\mathbf{z}_t^{(s_i)}$  with conditioning site  $s_i$  and parameter vector  $\boldsymbol{\psi}$ , by

$$(14) \quad L_{CL}(\boldsymbol{\psi}) = \prod_{i=1}^d \prod_{t \in \mathcal{T}^{(s_i)}} L_{CL}^{s_i}(\boldsymbol{\psi}; \mathbf{z}_t^{(s_i)}) = \prod_{i=1}^d \prod_{t \in \mathcal{T}^{(s_i)}} \prod_{j < k; j \& k \neq i} \frac{g_{s_i}(z_t^{(s_i)}(s_j), z_t^{(s_i)}(s_k), c_t^{(s_i)}(s_j), c_t^{(s_i)}(s_k))}{J(z_t^{(s_i)}(s_j), z_t^{(s_i)}(s_k), c_t^{(s_i)}(s_j), c_t^{(s_i)}(s_k))},$$

where  $L_{CL}^{s_i}$  is the censored likelihood contribution for  $s_i$  and the bivariate density  $g$  is defined in Appendix A of the Supplementary Materials ([Richards, Tawn and Brown, 2021](#)). The Jacobian term is

$$J\left(z_t^{(s_i)}(s_j), z_t^{(s_i)}(s_k), c_t^{(s_i)}(s_j), c_t^{(s_i)}(s_k)\right) = b\{x_t(s_i), h_{i,j}\}^{1\{z_t(s_j) > c_t^{(s_i)}(s_j)\}} b\{x_t(s_i), h_{i,k}\}^{1\{z_t(s_k) > c_t^{(s_i)}(s_k)\}},$$

where  $1\{\cdot\}$  is the indicator function. The parameter vector  $\boldsymbol{\psi}$  contains all parameters of the normalising functions  $a$  and  $b$ , the residual parameter functions  $\mu, \sigma$  and  $\delta$ , the correlation  $\rho$  and the anisotropy. Estimation of  $\boldsymbol{\psi}$  can be achieved by maximising (14).

3.2. *Stratified sampling regime.* Clearly, maximising (14) is computationally infeasible if  $d$  is large, as evaluation of (14) requires  $(d-1)(d-2) \sum_{i=1}^d |\mathcal{T}^{(s_i)}|/2$  evaluations of  $g_{(s_i)}$ , which can require double integrals and grows as  $O(d^3 n)$ . We detail a stratified sampling regime to create a pseudo-likelihood that circumvents the computational issue. To construct a sub-sample of data that can be used to estimate the parameters  $\boldsymbol{\psi}$  via pseudo-likelihood estimation, we first need to consider what these parameters represent. These parameters control characteristics of the dependence functions described in Sections 2.2.1 and 2.2.2, which are functions of either distance to the conditioning site  $h_{i,j}$  or pairwise distances  $h_{j,k}$  for  $s_i, s_j, s_k \in \mathbf{s}$ . Thus, we construct our sub-sample by drawing triples of sites  $(s_i, s_j, s_k) \in \mathbf{s}$ ; there are  $d(d-1)(d-2)/2$  possible triples of sites, and our sub-sample must adequately represent the distribution of the distances in the full data. However, not all distances can be represented in the sub-sample. If we pick triples randomly, then we are more likely to pick sites with larger pairwise distances. This has two disadvantages: pairs with larger distances are not informative about the dependence parameters, as at these distances the process may exhibit near-independence, and we are also unlikely to learn about the dependence for small distances, as less pairs with smaller pairwise distances are sampled. To ensure this is not the case, each triple  $(s_i, s_j, s_k)$  is chosen so that the distances  $h_{i,j}, h_{i,k}$  and  $h_{j,k}$  do not exceed a specified threshold. This is a natural extension of the approach of [Huser and Davison \(2013\)](#)

who suggest using only pairs of locations that are within some low distance  $h_{max} > 0$  of each other - we instead impose this constraint on triples. We investigate the reliability of our approach in Appendix B of the Supplementary Materials (Richards, Tawn and Brown, 2021).

To sub-sample  $d_s \ll d(d-1)(d-2)/2$  triples of locations for inference, we begin by uniformly sampling a conditioning site  $s_i \in (s_1, \dots, s_d)$ . A pair of sites are then drawn randomly from the set  $\{s_j, s_k \in \mathcal{S} \setminus s_i, j < k : \max\{h_{i,j}, h_{i,k}\} < h_{max}\}$  without replacement, and the process is repeated. In sampling in this way, only sites within distance  $h_{max}$  of the conditioning site are used for inference. That is, we estimate the spatial functions of the dependence parameters for  $h_{i,j} < h_{max}$  and  $h_{i,k} < h_{max}$ , and so  $h_{j,k} < 2h_{max}$  only, and then extrapolate to larger distances. There is a trade-off involved in choosing the value of  $h_{max}$ : if too low, then extrapolations to larger distances are likely to be poor; if too high, fit at small distances is compromised. In Section 4.3, we describe a heuristic technique for choosing  $h_{max}$  and we find that this works well in practice. The number of triples  $d_s$  is chosen to be as large as possible, whilst pseudo-likelihood estimation remains computationally feasible.

**3.3. Simulation of an event.** We now detail a technique that will allow us to draw realisations of  $\{Y(s) : s \in \mathcal{S}\}$ . First, we note that the model in Section 2.2 does not describe the dependence in all of  $\{Y(s)\}$ ; instead, it describes

$$(15) \quad \left\{ Y(s) : s \in \mathcal{S} \right\} \left| \left( \max_{s \in \mathcal{S}} \{F_L^{-1}(F_{Y(s)}\{Y(s)\})\} > v \right) \equiv \left\{ F_{Y(s)}^{-1}(F_L\{X(s)\}) : s \in \mathcal{S} \right\} \left| \left( \max_{s \in \mathcal{S}} \{X(s)\} > v \right),$$

for  $v \geq u$  with  $u$  used for fitting in Section 3.1. Thus, to create a realisation of  $\{Y(s) : s \in \mathcal{S}\}$ , we draw realisations of (15) with probability

$$(16) \quad \Pr \left\{ \max_{s \in \mathcal{S}} \{F_L^{-1}(F_{Y(s)}\{Y(s)\})\} > v \right\},$$

and otherwise draw realisations of

$$(17) \quad \left\{ Y(s) : s \in \mathcal{S} \right\} \left| \left( \max_{s \in \mathcal{S}} \{F_L^{-1}(F_{Y(s)}\{Y(s)\})\} < v \right).$$

As we do not expect realisations of (17) to contribute to the tail behaviour of  $R_A$ , we simply draw realisations of (17) from the observed data. We estimate (16) empirically; although this could be inferred using the parametric model of Section 2. If  $\mathcal{S}$  does not correspond to the set of sampling locations, then we would have to approximate (17) though some form of infilling, i.e., using the quantile regression technique (Fasiolo et al., 2021) discussed in Section 2.1.

We now describe a simulation technique that will allow us to draw realisations of (15). That is, the field  $\{Y(s) : s \in \mathcal{S}\}$  given that an extreme value above a threshold is observed anywhere in the domain. This threshold varies with  $s$  and corresponds to the relative quantile  $v$  on the Laplace scale. Wadsworth and Tawn (2019) detail the procedure for achieving this. There are three steps: drawing conditioning sites  $s_O \in \mathcal{S}$ , simulating the fields  $\{Y(s) : s \in \mathcal{S}\} | (F_L^{-1}(F_{Y(s_O)}\{Y(s_O)\}) > v)$  using the fitted model described in Section 2, and then using importance sampling to approximate (15), see Algorithm 1, Step 2. The first step requires random sampling of conditioning sites  $s_O$  for some  $s_O \in \mathcal{S}$ ; we do this uniformly, which provides a good first approximation of the occurrence of these sites in  $\mathcal{S}$  and then improve on this via the importance sampling regime described below.

To simulate  $N$  realisations from process (15), we follow Wadsworth and Tawn (2019) and draw an initial  $N' > N$  realisations of the process  $\{X(s) : s \in \mathcal{S}\}$  on the Laplace scale. Then, using importance sampling, we sub-sample  $N$  realisations from  $\{X(s) : s \in \mathcal{S}\} | \max_{s \in \mathcal{S}} X(s) > v$ , and transform the margins of the sample to the original scale,  $\{Y(s)\}$ . The sub-sampling regime adds extra weight to realisations for which the conditioning site is near the boundary

of the domain. This is to alleviate the edge effect caused by not using conditioning sites outside of the boundaries of  $\mathcal{S}$ . A discussion of a related issue is given in Section 3.4. We found that setting  $N' \approx 5N$  was sufficient for our application, although this may be dependent on the size of  $\mathcal{S}$  and value of  $N$ .

---

**Algorithm 1** Simulating (15)
 

---

1. For  $i = 1, \dots, N'$  with  $N' > N$ :
  - a) Draw a conditioning location  $s_O^{(i)}$  from  $\mathcal{S}$  with uniform probability density  $1/|\mathcal{S}|$ .
  - b) Simulate  $E^{(i)} \sim Exp(1)$  and set  $x_i(s_O^{(i)}) = v + E^{(i)}$ .
  - c) Simulate a field  $\{z_i(s|s_O^{(i)}) : s \in \mathcal{S}\}$  from the residual process model defined in Section 2.2.2.
  - d) Set  $\{x_i(s) : s \in \mathcal{S}\} = a\{x_i(s_O^{(i)}), h(s, s_O^{(i)})\} + b\{x_i(s_O^{(i)}), h(s, s_O^{(i)})\} \times \{z_i(s|s_O^{(i)}) : s \in \mathcal{S}\}$ .
2. Assign each simulated field  $\{x_i(s) : s \in \mathcal{S}\}$  an importance weight of

$$\left\{ \int_{\mathcal{S}} 1\{x_i(s) > v\} ds \right\}^{-1},$$

for  $i = 1, \dots, N'$ , and sub-sample  $N$  realisations from the collection with probabilities proportional to these weights.

3. Transform each  $\{x_i(s) : s \in \mathcal{S}\}$  to  $\{y_i(s) : s \in \mathcal{S}\}$  using the marginal transformation (3). If  $x_i(s) \leq c(s)$ , set  $y_i(s) = 0$ , where for some  $s' \in \mathcal{S}$ ,  $y_i(s')$  is above its  $F_L(v)$ -th quantile.
- 

3.4. *Inference on Spatial Aggregates.* Using the sample of realisations of  $\{Y(s) : s \in \mathcal{S}\}$  generated in Section 3.3, we make inference about the tail behaviour of  $R_{\mathcal{A}}$  in (1) or the corresponding sum; here we focus on the latter, but a discussion of the integral is given in Section 5. The possible size of the aggregation region  $\mathcal{A}$  in relation to the region  $\mathcal{S}$  is of particular interest. Trivially, we require  $\mathcal{A} \subseteq \mathcal{S}$ . However, we cannot have  $\mathcal{A} = \mathcal{S}$ , as if we did, the simulation algorithm will never generate an event for which the conditioning site lies outside of the boundaries of  $\mathcal{S}$ , but we still observe an extreme event somewhere inside  $\mathcal{A}$ . To avoid such edge-effects, we require the boundaries of  $\mathcal{A}$  to be far enough inside the interior of  $\mathcal{S}$ , such that the distribution (15) does not change if the size of  $\mathcal{S}$  increases. Informally, we require a buffer zone between the boundaries of  $\mathcal{A}$  and  $\mathcal{S}$  which is large enough, such that any event with conditioning site outside of  $\mathcal{S}$  has negligible effect on the distribution within  $\mathcal{A}$ . We select the width  $\tau$  of this buffer zone by using the measure  $\chi_q(s_A, s_B)$  given in (2) and stationarity. We choose  $\tau$  such that for any  $s \in \mathcal{A}$  and  $s_O \in \mathbb{R}^2 \setminus \mathcal{S}$ , such that for  $h(s, s_O) > \tau$ , we have  $\chi_q(s_O, s) < \gamma$  for small  $\gamma > 0$  and for all large  $q$ ; that is, we have a small probability, less than  $\gamma$ , of observing a large event at  $s$  given that there is an extreme event at any site outside of  $\mathcal{S}$ . This measure can be evaluated empirically or by simulating from the fitted model; we take the latter approach in Section 4.4.

## 4. Application.

4.1. *Data.* We consider data consisting of average hourly precipitation rate (mm/hour) taken from the UK climate projections 2018 (UKCP18) (Lowe et al., 2018). Data are from a convection permitting model which produces values over hourly intervals between the years 1980 and 2000, using the observed atmospheric conditions. The sampling locations are  $(5km)^2$  grid boxes corresponding to the British National Grid from Ordnance Survey (OSGB). The spatial domain  $\mathcal{S}$  of interest is East-Anglia, UK (see Figure 1) and only data sampled over land have been included, leaving 934 sampling locations. Each observation corresponds to the average over the assigned spatio-temporal grid-box. The data represent the

average in each grid-box, and so a natural quantity of interest is  $\bar{R}_A := R_A/|A|$ , rather than  $R_A$ , but we present results on  $R_A$  as this variable must satisfy the ordering constraints discussed in Section 1. To remove any seasonal effect observed in the data, we use summer, i.e., July-August, observations only, leaving 43200 fields<sup>1</sup>. We chose to take summer precipitation events as these typically exhibit higher intensity than winter events (Sharkey and Winter, 2019). We treat the centre of each grid box as a sampling location, and as the grid-boxes are non-overlapping and contiguous, we can approximate the integral  $R_A$  in (1) using a sum. We use the great-circle distance as our distance metric described in Section 2.2.2.

4.2. *Marginal Analysis.* Initial analysis shows that the data consists of 8.7% hours with zero precipitation, but much of the data with non-zero values exhibits noise around zero produced by the climate model. Thus, the data less than  $1 \times 10^{-5}$  mm/hour were set to zero<sup>2</sup>, increasing the average number of dry hours to 83.7%. Figure 1 gives a spatial map of the estimated probability of zero precipitation  $p(s)$  within a given hour; this is estimated using the logistic regression GAM detailed in Section 2.1. We observe some spatial variation in  $p(s)$ , with slightly lower estimates being found along the north-east coast.

We fit the spatial marginal model detailed in Section 2.1. We take  $\lambda(s) = 0.005$  for all  $s \in \mathcal{S}$  in (3) and the corresponding GPD threshold  $q(s)$ , estimated using a thin-plate spline, is illustrated in Figure 1 with  $q(s)$  varying roughly over  $\mathcal{S}$ ; larger values are found along the east coast. The GPD GAM model with spatially smooth estimate parameters is then fit to site-wise exceedances above  $\hat{q}(s)$  at each site; a spatial map of the shape parameters is given in Figure 1. We take the approach of Youngman (2019) and use as many knots as is computationally feasible in the thin-plate splines, which is 300. This creates a potentially overly rough spline which may overfit the data and not capture true physical smoothness; however, our primary interest is in the dependence structure when studying aggregates as this is the novel element of our model, so we chose this approach to ensure that the empirical marginal distributions are as well modelled as possible. We observe  $\hat{\xi}(s) > 0$  for all  $s \in \mathcal{S}$ , and so the marginal upper tails are unbounded at each site. Q-Q plots of the marginal fits at five randomly sampled locations are presented in Figure S3 of the Supplementary Materials (Richards, Tawn and Brown, 2021), all showing good fits. To evaluate the fit over all locations, we use a pooled Q-Q plot, transforming all data onto standard exponential margins using the fitted model, see Figure S3. Again the fit is remarkably good, although confidence intervals are not provided due to the spatial dependence in the pooled data.

4.3. *Dependence Model.* All dependence models are fitted by taking the exceedance threshold  $u$  in (4) to be the standard Laplace 98% quantile. This leaves 864 fields for fitting the extremal dependence model given an observed extreme at a single conditioning site. A lower threshold  $u$  was considered; however, we found that this leads to poorer model fits as the data exhibits a partial mixing of dependence structures. We believe this is due to the presence of multiple data generating processes in the climate model. Precipitation is typically generated by either high intensity events with localised spatial profiles, i.e., convective cells, or low intensity events with much large spatial profiles, i.e., frontal storms (Thomassen et al., 2020). In the absence of covariates to distinguish between these events in the data, we use a higher exceedance threshold to remove any frontal events; this is discussed further in Section 5. The empirical estimate (and 95% confidence interval) for the probability in (16) is 0.273 (0.257, 0.290); this corresponds to the proportion of all observed fields used for fitting

<sup>1</sup>Note that the UKCP18 data uses a 360 day calendar, and so each month is composed of 30 days.

<sup>2</sup>This level of precipitation would be recorded as zero by a rain gauge.

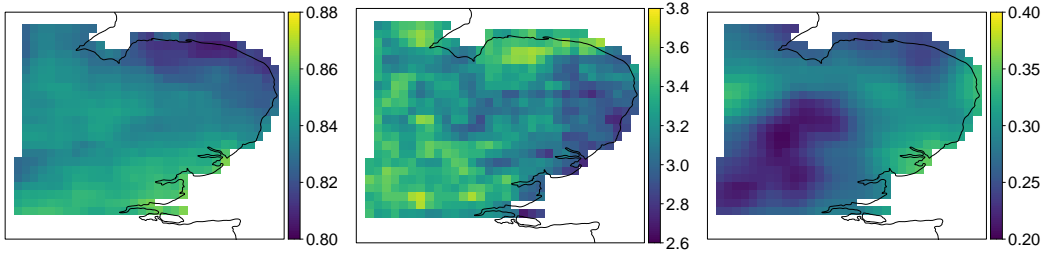


FIG 1. Spatially smoothed marginal distribution parameter estimates for East Anglia. Left:  $\hat{p}(s)$ , centre:  $\hat{q}(s)$ , right:  $\hat{\xi}(s)$ .  $\hat{v}(s)$  is illustrated in Figure S2 of the Supplementary Materials (Richards, Tawn and Brown, 2021).

when we pool over all 934 conditioning sites. Confidence intervals for (16) were created using a stationary bootstrap (Politis and Romano, 1994) with 1000 samples and expected block size  $m_K$ , here taken to be 48 hours. To compute a single bootstrap sample, we repeat the following until a sample of length greater than or equal to  $n$  is obtained; draw a starting time  $t^* \in \{1, \dots, n\}$  uniformly at random and a block size  $K$  from a geometric distribution with expectation  $m_K$ , then add the block of observations  $\{y_t(s) : s \in \mathcal{S}, t \in \{t^*, \dots, t^* + K - 1\}\}$  to the bootstrap sample. However, in cases where  $t^*$  is generated with  $t^* + K - 1 > n$ , we instead add  $\{y_t(s) : s \in \mathcal{S}, t \in \{1, \dots, t^* + K - n - 1\} \cup \{t^*, \dots, n\}\}$ . We then truncate the sample to have length  $n$ .

We proceed with an initial analysis by fitting a simple version of the model of Section 2.2 to these data. We fit the model with two caveats: we make the temporary assumption that the residual process  $\{Z(s|s_0)\}$  is independent at all distances; and evaluate a sequence of “free” pairwise parameter estimates (Wadsworth and Tawn, 2019) for the normalising functions and those functions that describe the marginal characteristics of  $\{Z(s|s_0)\}$ . That is, we fit individual parameters, i.e.,  $\alpha_{s_i}^{(s_0)}$  etc. for  $i = 1, \dots, d$  with  $s_i \neq s_0$ , rather than a spatial function  $\alpha\{h(s, s_0)\}$ , and we do this for seven different conditioning sites  $s_0$  sampled randomly over  $\mathcal{S}$ . This approach can be used to assess the stationarity of  $\{X(s)\}$ ; if we observe clear disagreement in the parameter estimates for the different conditioning sites, then the assumption of stationarity of  $\{X(s)\}$  is unlikely to be appropriate. We find no evidence for non-stationarity in the parameter estimates presented in Figure 2; while we observe some volatility in the free parameter estimates, the general patterns appear to be the same regardless of the choice of conditioning site. We use the spatial structure in the free estimates to motivate our choice for the forms of the parameter functions detailed in Sections 2.2.1 and 2.2.2.

Using the sampling method described in Section 3.2, the full spatial fit uses  $d_s = 5000$  triples of sites with each sampling location being used as a conditioning site at least once. As the estimates of  $\alpha$  and  $\beta$  in Figure 2 decay quickly with increasing spatial distance, i.e., for any distance greater than  $25\text{km}$ ,  $\alpha \approx 0$  and  $\beta < 0.5$ ; this illustrates that the underlying process  $Y(s)$  exhibits fairly localised strong extremal dependence. This suggests that we should focus on modelling extremal dependence locally, as this will be the driving factor of the aggregate behaviour. A distance of  $25\text{km}$  in the anisotropic setting corresponds a distance of at most  $28\text{km}$  in the original setting, and so we set  $h_{max} = 28\text{km}$ . Although 5000 triples of sites represents a very small proportion of all possible triples, we observe a good model fit from Figure 2, which shows that, even at distances greater than  $28\text{km}$ , the fitted parametric functions for the dependence parameters correspond well to the sequences of free estimates. We fitted a variant of (9) for  $\delta$ , without the max operator constraint, and found the best fit was achieved with this  $\delta$  function. That gave values for  $\delta\{h(s, s_0)\} < 1$  at the two closest sites




FIG 2. Estimates of parameters that determine the extremal dependence structure plotted against inter-site distance  $h$ , which is calculated under the anisotropy transformation for the full spatial model. Estimates from the free fits described in Section 4.3 are given by the black points, parametric spatial functions are given in red (asymptotically independent model) and blue (asymptotically dependent model). Bottom right: estimates from model for  $\chi_q(s, s_O)$  in (2) with  $q = 1/(24 \times 90)$ . Distances (km) are given in the spatial anisotropy setting.

to  $s_O$ . We found that this had negligible effect on the analysis, but clearly form (9) would be more ideal.

Figure 2 can be used to make further inference about the underlying dependence structure of the precipitation process. For example, we find that  $\Delta$  in (6) can be taken to be zero without restricting the quality of the fit and similarly we can set  $\kappa_{\delta_4} = \kappa_{\beta_3} = 1$ . The estimate  $\Delta = 0$  suggests that the process is asymptotically independent at even the closest distances as asymptotic dependence requires both  $\alpha(h) = 1$  and  $\beta(h) = 0$  for all  $h$ , which the estimates in Figure 2 suggest is not the case; a fit imposing asymptotic dependence is discussed later. Furthermore, we found that incorporating spatial anisotropy into the dependence model improved the overall fit; stronger extremal dependence was found along an approximate  $-10^\circ$  bearing, reducing by at most 7% over different directions. Parameter estimates (and standard errors) are provided in Table S1 in the Supplementary Materials (Richards, Tawn and Brown, 2021). Although not illustrated in Figure 2,  $\rho$  decays quickly with distance, with  $\rho(100) \approx 0.2$ .

To further support the choice of  $h_{max}$ , we estimate  $\chi_q(s, s_O)$  in (2) for  $s_O \in \mathcal{S}$  in the centre of  $\mathcal{S}$ , taking  $q$  corresponding to a one-year return level probability, and look to see how far away  $s$  must be for  $\chi_q(s, s_O)$  to be less than  $\gamma$  for small  $\gamma > 0$  (see Section 3.4). We estimate  $\chi_q(s, s_O)$  by simulating  $5 \times 10^4$  replications, using Algorithm 1, from the fitted model and this is illustrated in Figure 2, bottom-right panel; for  $\gamma = 0.1$ , we find that a distance of  $h_{max}$  is sufficient and so we set  $\tau = h_{max}$  in Section 3.4, discussed further in Section 4.4.

Figure 3 illustrates six extreme fields: three realisations from the model defined in (15) and three observations from the data. Fields are chosen such that the site in the centre of  $\mathcal{S}$  exceeds its 99.9%-quantile but the maximum over the entire field does not exceed 30mm/hr; this is to make it easier to compare the spatial structure in the fields. Realisations from the model appear to replicate the roughness in observed events well. Furthermore, in Figure 3 we observe that replications from the model are able to exhibit some of the different physical properties of extreme precipitation. For example, the top-left panel displays a spatially

flat event whilst the other two illustrate localised extreme events; multiple sub-events in the top-middle and a single sub-event in the top-right. Further diagnostics for fitted extremal dependence models have been proposed in the literature, e.g., by [Blanchet and Davison \(2011\)](#), [Huser and Wadsworth \(2020\)](#); however, our focus lies on inference for aggregates and so we present diagnostics for these variables only.

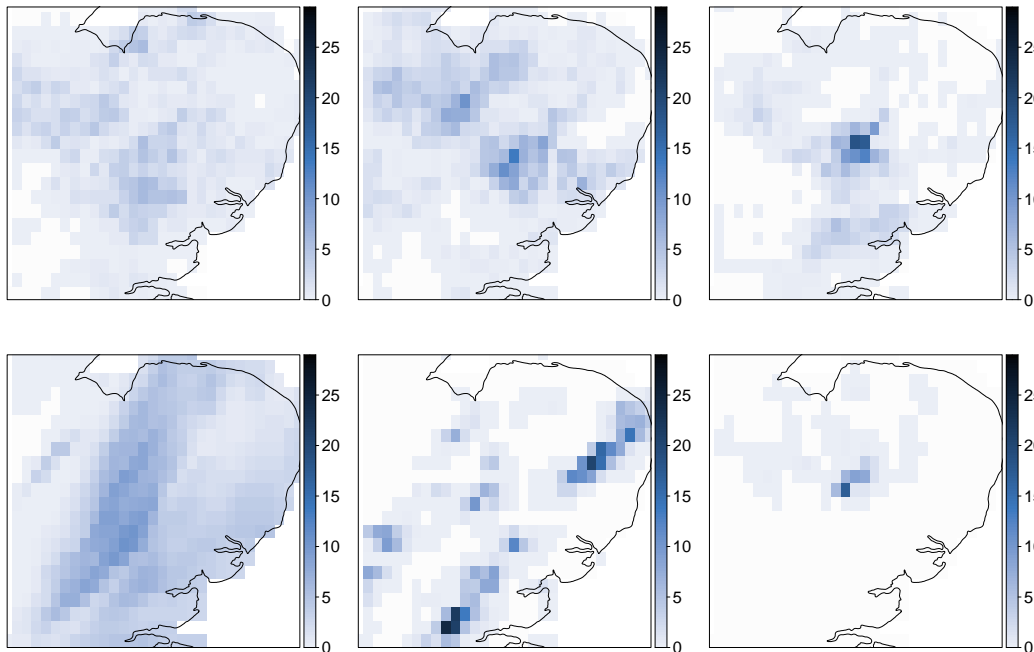


FIG 3. *Extreme precipitation fields (mm/hr). Top-row: realisations from the fitted model described in (15). Bottom-row: observed fields from the data.*

As  $\alpha$  quickly goes to zero with distance, all extremal dependence is instead exhibited through  $\beta$ . This is atypical of fits of this model for other applications ([Wadsworth and Tawn, 2019](#); [Shooter et al., 2020](#); [Simpson and Wadsworth, 2021](#)), where the  $\alpha$  function drives the extremal behaviour of their modelled processes, e.g., temperature and sea wave heights. Having  $\beta$  controlling extremal dependence would suggest that the process that generates the extreme precipitation we are modelling is somewhat rough; this concurs with the observed fields containing an extreme value shown in Figure 3, and consistent with the spatial nature of strong convective rainfall. To illustrate this, we note that, for small  $h = h(s, s_O)$ , we have  $\alpha(h) \approx 0$  and  $\beta(h) \approx 1$ , then  $\mathbb{E}[X(s)|X(s_O) = x] \approx \mu(h)x$  and  $\text{Var}(X(s)|X(s_O) = x) \approx x^2\sigma^2(h)$ , and so the largest events at  $s_O$  are the most variable. This has not been observed in other applications as the extremal dependence in these processes is typically quite smooth with  $\text{Var}(X(s)|X(s_O) = x) \approx \sigma^2(h)$  as  $\beta(h) \approx 0$  when  $h$  is small. Even at the largest  $h$  value possible in  $\mathcal{S}$ , the process  $\{X(s)\}$  does not exhibit independence; although  $\alpha$  and  $\beta$  tend to zero the residual process does not attain standard Laplace margins with  $\delta(h) = 1$ .

Existing literature for approaches that rely on modelling the underlying process to make inference the extremal behaviour of spatial aggregates of precipitation typically use models that only allow for asymptotic dependence ([Coles, 1993](#); [Coles and Tawn, 1996](#); [Buisson, de Haan and Zhou, 2008](#)). We fit such a model to illustrate that imposing asymptotic dependence may lead to poor inference for the tails of spatial aggregates. We term this the ‘‘AD model’’ and the model described above as the ‘‘AI model’’. To specify the

AD model, we fix  $\alpha(h) = 1$  and  $\beta(h) = 0$  for all  $h$  and we change  $\sigma(h)$  in (9) to  $\sigma(h) = \kappa_{\sigma_3} (1 - \exp\{-(h/\kappa_{\sigma_1})^{\kappa_{\sigma_2}}\})$  with  $\kappa_{\sigma_1}, \kappa_{\sigma_2}, \kappa_{\sigma_3} > 0$ , as we no longer require that  $\sigma(h) \rightarrow \sqrt{2}$  as  $(h) \rightarrow \infty$ . The corresponding  $\mu(\cdot)$  and  $\delta(\cdot)$  functional forms remain the same and the spatial anisotropy setting described in (12) is still used. To fully capture the behaviour of  $\mu(\cdot)$ , we found we had to take  $h_{max} = 75km$ . The estimated spatial functions for the AD model are illustrated in Figure 2. With  $\alpha$  and  $\beta$  fixed, we observe that the other parameters are forced to compensate for this misspecification. For example, we observe a strictly negative  $\mu$  function; this is to compensate for fixing the  $\alpha$  value too large for the data. Given this, we re-estimated the free parameters with  $\alpha = 1$  and  $\beta = 0$  fixed and observed good agreement between the spatial functions and these new estimates. However, this does not imply that the model as a whole fits well, this is emphasised in Section 4.4 where spatial aggregates of simulated fields  $\{Y(s)\}$  are studied.

4.4. *Diagnostics and tails of spatial aggregates.* Q-Q plots, presented in Figure 4, assess how well the tails of the simulated distributions compare against the tails of the empirical distribution of the spatial averages. **Confidence intervals given for the simulated quantiles are derived by applying to the observations the stationary bootstrap with expected block size of 48 hours.** We evaluate a sample of 50 parameter estimates by applying the procedure described in Section 3.2 to stationary bootstrap samples of the data; note that we keep  $h_{max} = 28km$ , but use different, randomly selected triples of sites for each sample. Only 50 estimates were derived due to the computational resources required for fitting. For each parameter set estimate, we draw  $5 \times 10^5$  realisations of  $\{Y(s) : s \in \mathcal{S}\}$  using the regime described in Section 3.3. Then  $\bar{R}_{\mathcal{A}}$  is calculated for each sample and for each region  $\mathcal{A}$ ; per the discussion in Section 4.3, each  $\mathcal{A}$  is at least  $\tau = 28km$  away from the boundaries of  $\mathcal{S}$ . Figure 4 illustrates generally very good fits for the tails of  $\bar{R}_{\mathcal{A}}$  with nested regions  $\mathcal{A}$ . The AI model appears to slightly underestimate the true magnitude of the largest aggregates for the largest regions; while this may suggest that the model is not capturing dependence at further distances from the conditioning site, Figure 2 suggests that the model fits well even at the furthest distances. This leads us to suspect that there is a mixture of events present in the data when we consider large spatial regions for aggregation, and that the model is not flexible enough to capture these mixtures, see Section 5. To illustrate the benefits of using our approach, Figure 4 also illustrates the same diagnostics for the AD model described in Section 4.3 for the smallest and largest regions. The AD model provides much poorer fits than the AI models, as it always overestimates the quantiles; this suggests that the AD model overestimates the dependence within the original process even for the smallest aggregation regions. A similar plot for non-overlapping regions is illustrated in Figure S6 of the Supplementary Materials (Richards, Tawn and Brown, 2021); here we observe some underestimation in the largest estimated return levels for the regions closest to the east coast, which is a possible indication of non-stationarity along this coast.

As discussed in Section 1, obtaining physically consistent return level estimates of spatial aggregates is essential. We compare two methods for achieving this: (i) performing a long-run simulation from our model, deriving empirical estimates of return-levels from these replicates; (ii) fitting a GPD to the observed aggregate tails and extrapolating to the desired return-level. Ideally, we want to use only the former approach as this mitigates the potential issues with using method (ii) discussed in Section 1; however, for computational efficiency, we perform a shorter run for method (i) with  $5 \times 10^5$  realisations and use a fitted GPD to extrapolate to the largest return-levels. Figure 5 presents estimates of return level curves for  $R_{\mathcal{A}}$  over the nested regions, illustrated in Figure 4, using methods (i) and (ii), top-left and top-right panels, respectively. For each region  $\mathcal{A}$ , a GPD is fitted to exceedances of the respective sample  $R_{\mathcal{A}}$  above the 99.9% quantile and return level curves are estimated from these fits. In



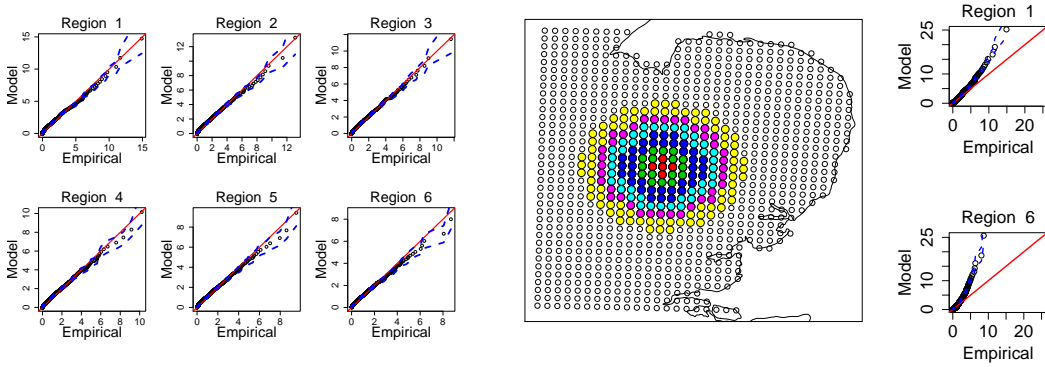


FIG 4.  $Q$ - $Q$  plots for model, and empirical,  $\bar{R}_A$  of regions of increasing size. Left: AI model, right: AD model. Probabilities range from 0.7 to a value corresponding to the 20 year return level. 95% confidence intervals are given by the blue dashed lines.  $Q$ - $Q$  plots for all six regions for the AD model are given in Figure S5 of the Supplementary Materials (Richards, Tawn and Brown, 2021). Centre: aggregate regions  $\mathcal{A}$  with corresponding areas (125, 525, 1425, 2425, 3350, 5425) –  $\text{km}^2$ . Regions 1-6 are coloured red, green, blue, cyan, purple, yellow; regions include both the coloured and interior points.

the top-right panel of Figure 5, we observe intersection in the return level curves estimated for the two smallest regions using method (ii). This problem does not arise using the computationally efficient version of approach (i), e.g., in the top-left panel of Figure 5, where we have  $5 \times 10^5 \times 20/43200 \approx 231$ , and 20, years of data for inference, respectively. Furthermore, the confidence intervals produced by method (i) are tighter, as more data are used for extrapolation; we illustrate this in the bottom-left panel of Figure 5, where we overlay return level estimates  $R_A$  using both methods, for a single region  $\mathcal{A}$ . Confidence intervals are derived for both methods by fitting a GPD to 50 bootstrap samples of  $R_A$ : in method (i), these are the samples as described at the top of this section; for (ii), we perform a simple bootstrap of the observed data, assuming temporal independence. A higher exceedance threshold for approach (i) was considered, but we found that the difference in estimates was negligible; to support the use of the 99.9% quantile, we illustrate a pooled  $Q$ - $Q$  plot in Figure 5, transforming exceedances from all 50 bootstrap samples onto standard Exponential margins using their respective GPD fits and observe an excellent overall fit.

A further point of interest for practitioners is inference on the joint behaviour of  $(\bar{R}_A, \bar{R}_B)$  for different regions  $\mathcal{A}, \mathcal{B} \in \mathcal{S}$ . We investigate this joint behaviour for the different aggregate regions given in Figure 4. Figure 6 illustrates realisations of pairwise  $(\bar{R}_A, \bar{R}_B)$  for both model and empirical estimates with nested regions  $\mathcal{A}, \mathcal{B}$ , showing that the model captures the joint distributions well; a similar figure is given for non-overlapping regions in Figure S7 of the Supplementary Materials (Richards, Tawn and Brown, 2021), in which we observe that extreme events do not typically occur together. This suggests that the extremal behaviour of the aggregates is driven by spatially-localised events. Further evidence for this can be found in Figure 6 for aggregates over nested regions; if the extremes of  $R_A$  were driven by events with a spatial profile that encompassed  $\mathcal{S}$ , then we would not observe weakening extremal dependence between aggregates over the smallest, and increasingly larger, regions. Instead, we would observe no change in the strength of extremal dependence.

**5. Discussion.** We have presented extensions of the Heffernan and Tawn (2004) and Wadsworth and Tawn (2019) models for modelling the extremal dependence for precipitation data. As illustrated in Section 2.2, this model provides flexibility over existing models for extreme precipitation as it can capture asymptotic independence. Simulating from this model

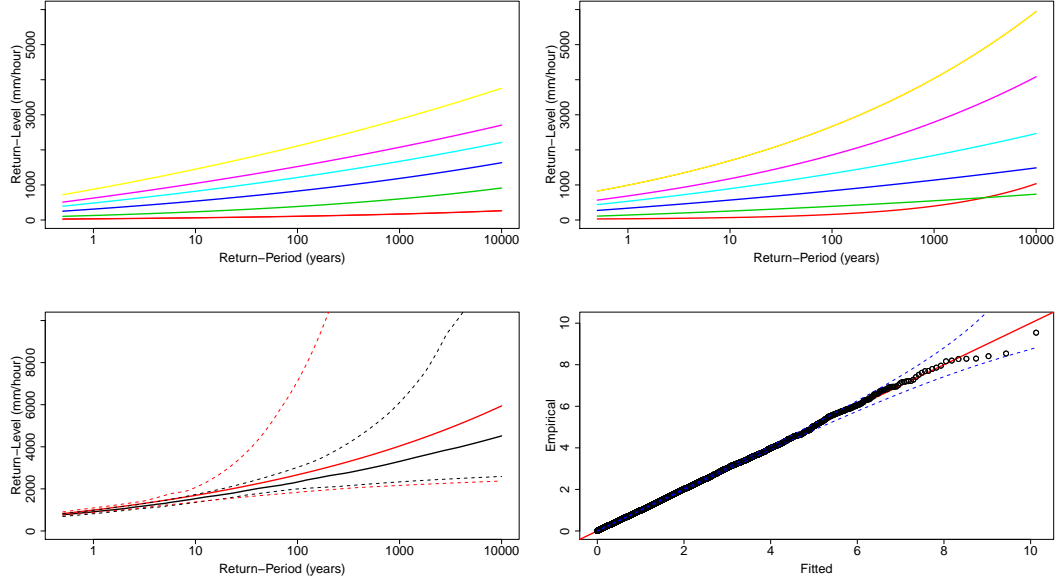


FIG 5. *Top: Estimated return level curves of  $R_A$  using the model (left) and observations (right). Colours correspond to the regions illustrated in Figure 4. Bottom-left: return level estimates for Region 6 in Figure 4 using methods (i) and (ii) in black and red, respectively. 95% confidence intervals for the methods are given by the coloured dashed lines. Bottom-right:  $Q-Q$  plot for pooled GPD fit for approach (i), over all 50 bootstrap samples, on standard Exponential margins. 95% tolerance bounds are given by the dashed lines.*




FIG 6. *Plots of  $2 \times 10^4$  realisations of pairwise  $(\bar{R}_A, \bar{R}_B)$  for nested regions  $A, B$ , illustrated in Figure 4. Black points are model estimates, red points are from the data. The regions  $A, B$  are labelled on the respective panels.*

is simple, and replications can be used to make reliable inference about the tail behaviour of spatial aggregates of the underlying process once issues linked to edge effects are addressed. This approach circumvents an issue that is common with independent inference on the tails of spatial aggregates over different regions, namely that they run the risk of making inference that is inconsistent with the physical properties of the process.

A particular drawback of our approach is that inference using the full likelihood is computationally infeasible. To overcome this issue, we proposed methods for model fitting and assessing parameter uncertainty that are based on a pseudo-likelihood approach which requires specification of a hyper-parameter  $h_{max}$ . We found that these methods worked well in our application, as we were able to choose a suitable  $h_{max}$  quite low for which the model fits well in a reasonable time-frame, see Figure 2. This is because our data exhibits fairly

localised extreme events; in applications where this is not the case, a larger  $h_{max}$  will be required which could potentially lead to more samples being required for fitting.

Data used in Section 4 are from a climate model, which means that sampling locations are comprised of non-overlapping grid-boxes, rather than point locations. In our application, we take  $R_{\mathcal{A}}$  to be the corresponding summations, rather than the integrals defined in (1); however, this is not to say that our approach cannot be used if we require inference on the tail of an integral. We have detailed a fully spatial model for both the dependence and marginal behaviour of  $\{Y(s)\}$ , and so it is possible to create a sample of  $\{Y(s) : s \in \mathcal{S}\}$  where  $\mathcal{S}$  is not necessarily the sampling locations. We can then approximate  $R_{\mathcal{A}}$  by specifying  $\mathcal{S}$  as a fine-grid and taking the sum of  $\{Y(s) : s \in \mathcal{A}\}$ .

When considering spatial aggregates over the largest regions  $\mathcal{A}$ , we find that our approach slightly underestimates the largest events, see Figure 4. Whilst this may be caused by boundary effects, it could also be caused by a complexity of the data generating process that is not captured by the model. As the size of  $\mathcal{A}$  increases, it becomes less likely that the tail behaviour of  $R_{\mathcal{A}}$  is driven by a single type of extreme event. There are two possible areas of complexity that are missed for regions that are sufficiently large: (i) multiple occurrences of localised high-intensity convective events (Schroeder, Kirchengast and O, 2018), whereas we modelled single occurrences in Section 4; (ii) events consisting of a mixture of localised high-intensity convective and widespread low-intensity non-convective, events. Our data appears to exhibit these; recall in Section 4.3 we remarked that we considered a lower threshold  $u$  in (4) for modelling, but we found that this was not feasible as the data exhibits mixtures of dependence. A higher threshold had to be specified to remove observed fields that exhibited long-range spatial dependence to improve model fitting. Further improvements can be made to inference on the tails of  $R_{\mathcal{A}}$  by modelling frontal events. To illustrate this, consider that we model  $R_{\mathcal{A}} | (\max_{s \in \mathcal{S}} X(s) > v)$ , i.e.,  $R_{\mathcal{A}}$  given an extreme event somewhere in  $\mathcal{S}$ , and undo said conditioning using the data. We do not model  $R_{\mathcal{A}}$  given that there is no extreme event anywhere in  $\mathcal{S}$ , i.e., such as caused by a frontal event. As the size of  $\mathcal{A}$  grows, we will increasingly find that these events will drive the extremal behaviour of  $R_{\mathcal{A}}$ ; this could be further explanation behind the underestimation in Figure 4, and so should be incorporated into the model. We also considered another measure of the extremal dependence in Appendix C of the Supplementary Materials (Richards, Tawn and Brown, 2021), which suggests improvement may be possible using mixture modelling. A possible approach to this problem is to incorporate covariates on precipitation field type into the model.

**Acknowledgements.** Jordan Richards and Jonathan Tawn gratefully acknowledge funding through the STOR-i Doctoral Training Centre and Engineering and Physical Sciences Research Council (grant EP/L015692/1). Simon Brown was supported by the Met Office Hadley Centre Climate Programme funded by BEIS and Defra. The authors are grateful to Robert Shooter of the Met Office Hadley Centre, UK, and to Jennifer Wadsworth and Emma Simpson of Lancaster University and University College London, respectively, for helpful discussions.

## SUPPLEMENTARY MATERIAL

### Supplement:

(doi: **COMPLETED BY THE TYPESETTER; .pdf**). Supporting material is provided for methodology proposed in the main text. Additional figures are provided as diagnostics for model fits in our application. The data can be downloaded from the CEDA data catalogue, see [Met Office Hadley Centre \(2019\)](#). Code that supports our findings can be found in the [Jbrich95/scePrecip](#) repository on GitHub.

## REFERENCES

- BENOIT, L., ALLARD, D. and MARIETHOZ, G. (2018). Stochastic rainfall modeling at sub-kilometer scale. *Water Resources Research* **54** 4108–4130.
- BLANCHET, J. and DAVISON, A. C. (2011). Spatial modeling of extreme snow depth. *The Annals of Applied Statistics* 1699–1725.
- BREIMAN, L. (1965). On some limit theorems similar to the arc-sin law. *Theory of Probability & Its Applications* **10** 323–331.
- BROWN, S. J. (2020). Future changes in heatwave severity, duration and frequency due to climate change for the most populous cities. *Weather and Climate Extremes* **30** 100278.
- BUISSHAND, T. A., DE HAAN, L. and ZHOU, C. (2008). On spatial extremes: With application to a rainfall problem. *Ann. Appl. Stat.* **2** 624–642.
- CASTRO-CAMILO, D. and HUSER, R. (2020). Local likelihood estimation of complex tail dependence structures, applied to US precipitation extremes. *Journal of the American Statistical Association* **115** 1037–1054.
- MET OFFICE HADLEY CENTRE (2019). UKCP Local Projections on a 5km grid over the UK for 1980-2080. <https://catalogue.ceda.ac.uk/uuid/e304987739e04cdc960598fa5e4439d0>. Accessed 24/01/2022.
- COLES, S. G. (1993). Regional modelling of extreme storms via max-stable processes. *Journal of the Royal Statistical Society. Series B (Methodological)* **55** 797–816.
- COLES, S. G. (2001). *An Introduction to Statistical Modeling of Extreme Values*. Springer series in Statistics. Springer-Verlag, London.
- COLES, S. G. and TAWN, J. A. (1996). Modelling extremes of the areal rainfall process. *Journal of the Royal Statistical Society. Series B (Methodological)* **58** 329–347.
- COOLEY, D. and SAIN, S. R. (2010). Spatial hierarchical modeling of precipitation extremes from a regional climate model. *Journal of agricultural, biological, and environmental statistics* **15** 381–402.
- DAVISON, A. C. and SMITH, R. L. (1990). Models for exceedances over high thresholds (with discussion). *Journal of the Royal Statistical Society: Series B (Methodological)* **52** 393–425.
- DE FONDEVILLE, R. and DAVISON, A. C. (2021). Functional peaks-over-threshold analysis. To appear in the *Journal of the Royal Statistical Society: Series B (Methodological)*.
- ENGELKE, S., DE FONDEVILLE, R. and OESTING, M. (2018). Extremal behaviour of aggregated data with an application to downscaling. *Biometrika* **106** 127–144.
- FASIOLO, M., WOOD, S. N., ZAFFRAN, M., NEDELLEC, R. and GOUDE, Y. (2021). Fast calibrated additive quantile regression. *Journal of the American Statistical Association* **116** 1402–1412.
- FERREIRA, A., DE HAAN, L. and ZHOU, C. (2012). Exceedance probability of the integral of a stochastic process. *Journal of Multivariate Analysis* **105** 241–257.
- HARRISON, D., DRISCOLL, S. and KITCHEN, M. (2000). Improving precipitation estimates from weather radar using quality control and correction techniques. *Meteorological Applications* **7** 135–144.
- HEFFERNAN, J. E. and TAWN, J. A. (2004). A conditional approach for multivariate extreme values (with discussion). *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **66** 497–546.
- HUSER, R. and DAVISON, A. C. (2013). Composite likelihood estimation for the Brown-Resnick process. *Biometrika* **100** 511–518.
- HUSER, R. and GENTON, M. G. (2016). Non-Stationary dependence structures for spatial extremes. *Journal of Agricultural, Biological, and Environmental Statistics* **21** 470–491.
- HUSER, R. and WADSWORTH, J. L. (2020). Advances in statistical modeling of spatial extremes. *Wiley Interdisciplinary Reviews: Computational Statistics* e1537.
- JOE, H. (1997). *Multivariate Models and Multivariate Dependence Concepts*. CRC Press.
- KENDON, E. J., ROBERTS, N. M., FOWLER, H. J., ROBERTS, M. J., CHAN, S. C. and SENIOR, C. A. (2014). Heavier summer downpours with climate change revealed by weather forecast resolution model. *Nature Climate Change* **4** 570–576.
- KENDON, E., FOSSER, G., MURPHY, J., CHAN, S., CLARK, R., HARRIS, G., LOCK, A., LOWE, J., MARTIN, G., PIRRET, J. et al. (2019). UKCP Convection-permitting model projections: Science report.
- LOWE, J., BERNIE, D., BETT, P., BRICHENO, L., BROWN, S., CALVERT, D., CLARK, R., KAREN, EAGLE, EDWARDS, T., FOSSER, G., MAISEY, P., MCINNES, R., MCSWEENEY, C., YAMAZAKI, K. and BELCHER, S. (2018). UKCP18 Science Overview Report. *Met Office Hadley Centre: Exeter, UK*.
- NADARAJAH, S., ANDERSON, C. W. and TAWN, J. A. (1998). Ordered multivariate extremes. *Journal of the Royal Statistical Society: Series B (Methodology)* **60** 473–496.
- NORTHROP, P. J. and JONATHAN, P. (2011). Threshold modelling of spatially dependent non-stationary extremes with application to hurricane-induced wave heights. *Environmetrics* **22** 799–809.
- PADOAN, S. A., RIBATET, M. and SISSON, S. A. (2010). Likelihood-based inference for max-stable processes. *Journal of the American Statistical Association* **105** 263–277.

- PALACIOS-RODRÍGUEZ, F., TOULEMONDE, G., CARREAU, J. and OPITZ, T. (2020). Generalized Pareto processes for simulating space-time extreme events: an application to precipitation reanalyses. *Stochastic Environmental Research and Risk Assessment* **34** 2033–2052.
- POLITIS, D. N. and ROMANO, J. P. (1994). The stationary bootstrap. *Journal of the American Statistical Association* **89** 1303–1313.
- REICH, B. J. and SHABY, B. A. (2012). A hierarchical max-stable spatial model for extreme precipitation. *The Annals of Applied Statistics* **6** 1430–1451.
- RICHARDS, J. and TAWN, J. A. (2021). On the tail behaviour of aggregated random variables. *arXiv e-prints* **arXiv:2105.11917**.
- RICHARDS, J., TAWN, J. A. and BROWN, S. (2021). Supplement to “Modelling extremes of spatial aggregates of precipitation using conditional methods”.
- RICHARDS, J. and WADSWORTH, J. L. (2021). Spatial deformation for nonstationary extremal dependence. *Environmetrics* e2671.
- RODRIGUEZ-ITURBE, I., COX, D. R. and ISHAM, V. (1987). Some models for rainfall based on stochastic point processes. *Proceedings of the Royal Society of London Series A* **410** 269–288.
- SCHROEER, K., KIRCHENGAST, G. and O, S. (2018). Strong dependence of extreme convective precipitation intensities on gauge network density. *Geophysical Research Letters* **45** 8253–8263.
- SHARKEY, P. and WINTER, H. C. (2019). A Bayesian spatial hierarchical model for extreme precipitation in Great Britain. *Environmetrics* **30** e2529.
- SHOOTER, R., TAWN, J. A., ROSS, E. and JONATHAN, P. (2020). Basin-wide spatial conditional extremes for severe ocean storms. *Extremes* 1–25.
- SIMPSON, E. S., OPITZ, T. and WADSWORTH, J. L. (2021). High-dimensional modeling of spatial and spatio-temporal conditional extremes using INLA and the SPDE approach. *arXiv e-prints* **arXiv:2011.04486**.
- SIMPSON, E. S. and WADSWORTH, J. L. (2021). Conditional modelling of spatio-temporal extremes for Red Sea surface temperatures. *Spatial Statistics* **41** 100482.
- THOMASSEN, E. D., SØRUP, H. J. D., SCHEIBEL, M., EINFALT, T. and ARNBJERG-NIELSEN, K. (2020). Data-driven distinction between convective, frontal and mixed extreme rainfall events in radar data. *Hydrology and Earth System Sciences Discussions* **2020** 1–26.
- WADSWORTH, J. L. and TAWN, J. A. (2019). Higher-dimensional spatial extremes via single-site conditioning. *arXiv e-prints* **arXiv:1912.06560**.
- WESTRA, S. and SISSON, S. A. (2011). Detection of non-stationarity in precipitation extremes using a max-stable process model. *Journal of Hydrology* **406** 119–128.
- WOOD, S. (2006). *Generalized Additive Models: An Introduction with R*. Chapman & Hall/CRC Texts in Statistical Science. Taylor & Francis.
- YOUNGMAN, B. D. (2019). Generalized additive models for exceedances of high thresholds with an application to return level estimation for U.S. wind gusts. *Journal of the American Statistical Association* **114** 1865–1879.