Journal of Applied Statistics
Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/cjas20

Estimating the probability of simultaneous rainfall extremes within a region: a spatial approach

Lee Fawcett[a] & David Walshaw[a]
[a] School of Mathematics & Statistics, Newcastle University, Newcastle NE1 7RU, UK
Published online: 14 Nov 2013.

To cite this article: Lee Fawcett & David Walshaw , Journal of Applied Statistics (2013): Estimating the probability of simultaneous rainfall extremes within a region: a spatial approach, Journal of Applied Statistics, DOI: 10.1080/02664763.2013.856872

To link to this article: http://dx.doi.org/10.1080/02664763.2013.856872

Journal of Applied Statistics, 2013
http://dx.doi.org/10.1080/02664763.2013.856872

Taylor & Francis
Taylor & Francis Group

Estimating the probability of simultaneous rainfall extremes within a region: a spatial approach

Lee Fawcett* and David Walshaw

School of Mathematics & Statistics, Newcastle University, Newcastle NE1 7RU, UK

(*Received 11 July 2013; accepted 15 October 2013*)

In this paper we investigate the impact of model mis-specification, in terms of the dependence structure in the extremes of a spatial process, on the estimation of key quantities that are of interest to hydrologists and engineers. For example, it is often the case that severe flooding occurs as a result of the observation of rainfall extremes at several locations in a region simultaneously. Thus, practitioners might be interested in estimates of the joint exceedance probability of some high levels across these locations. It is likely that there will be spatial dependence present between the extremes, and this should be properly accounted for when estimating such probabilities. We compare the use of standard models from the geostatistics literature with max-stables models from extreme value theory. We find that, in some situations, using an incorrect spatial model for our extremes results in a significant under-estimation of these probabilities which – in flood defence terms – could lead to substantial under-protection.

Keywords: extreme value theory; Gaussian processes; geostatistics; max-stable processes; rainfall extremes

1. Background and data

Over the last few decades, research in the field of extreme values has grown rapidly – both in terms of the mathematical development of the subject and its applications. Mathematical accounts and reviews can be found in, for example, [13,23,28,29]; more statistical treatments may be found in, amongst others, [2,5,16]. In recent years, the need for accuracy and precision when estimating the extremes of environmental processes has been a primary driver for the development of increasingly sophisticated applications – especially in the modelling of multivariate and spatial extremes. For example, understanding the behaviour of rainfall or sea-surge extremes is crucial in flood protection, and data on such variables are often multivariate and/or spatial in nature. Such applications often require inferences to be made well beyond the range of observed data, requiring a certain degree of faith in the applicability of the mathematics underpinning the

*Corresponding author. Email: lee.fawcett@ncl.ac.uk

extrapolation; one consequence of this lack of data is highly uncertain tail estimates, often giving extremely wide confidence intervals – a problem that a spatial/multivariate analysis can at least attempt to address through the potential to pool information across sites or regions. There are many useful references to work that has been important in the development of multivariate and spatial extremes: see, for example, [7,19,24,27] for multivariate extremes and [4,10,30,31,34] for work with a more spatial flavour.

In the last 15 years in the UK there has been a notable increase in the frequency, and severity, of extreme precipitation events leading to localised, and sometimes more widespread, flooding. For example: flash-flooding in Glasgow in 2002, resulting in the evacuation of over 200 people from their homes, landslides and the contamination of drinking water; the Boscastle floods in the southwest of England in 2004, resulting in the destruction of over 100 homes and business properties; and flooding across the entire UK in 2007, 2008 and 2012, resulting in numerous deaths and millions of pounds worth of damage to property and infrastructure. The year 2012 was the second wettest year in the UK since records began, and six years since 1998 are in the top 10 wettest years (1998, 1999, 2000, 2002, 2008 and 2012). Observational studies and climate models indicate that the occurrence and magnitude of such extreme rainfall events will continue to increase in the future.

In many of the worst cases listed above, rivers burst their banks as a result of the observation of simultaneous rainfall extremes at a number of sites across a region. Thus, of particular interest here might be the estimation of:

(i) joint exceedance probabilities at a collection of sites within a region, or perhaps
(ii) exceedance probabilities/high quantiles of the aggregate rainfall observed across *all* sites within a region.

By definition, such quantities are both multivariate and spatial in nature. Although (as mentioned above) working within a multivariate or spatial setting allows for the potential to share information between sites, reducing the uncertainty of estimates of quantities such as those given in (i) and (ii) [18], it is important that the models used characterise appropriately the observed extremes, to lend confidence when extrapolating into high-dimensional space.

To date, models for spatial extremes have been based on max-stable processes (see Section 2). A rather restrictive assumption of this class of models is that the dependence structure of the observed extremes is comparable to that for a limiting model where dependence holds for all events more extreme than those that have already occurred. In reality, dependence may be observed for data above levels of practical interest, with these data being independent in the limit. This issue of non-convergence of the dependence structure, and *asymptotic independence*, is discussed in [19,20,24]. The most well-known spatial model giving such non-negligible dependence at observable levels, whilst being independent in the limit, is probably a Gaussian process [9]. Although diagnostic tools have been developed to help determine whether a data set should be modelled using an asymptotically independent or asymptotically dependent model (see Section 4.1 of this paper), the interpretation of these diagnostics can be rather subjective and this can be a difficult task. The aim of this paper is to compare estimates of quantities like those in (i) and (ii) under model mis-specification; that is, how are our estimates affected when using a spatial model that assumes genuine extremal dependence when, in fact, the extremes of the spatial process are asymptotically independent, or vice-versa?

Figure 1 shows a map of locations in the UK for which we have daily aggregate rainfall totals (measured in mm). These sites are colour-coded according to geographically defined climate regions representing areas of different physiological character, within which the rainfall climate might be regarded as reasonably coherent. For further details regarding this partition, see [35]. For some regions, data are available for up to 40 years (1961–2000 inclusive). To focus on

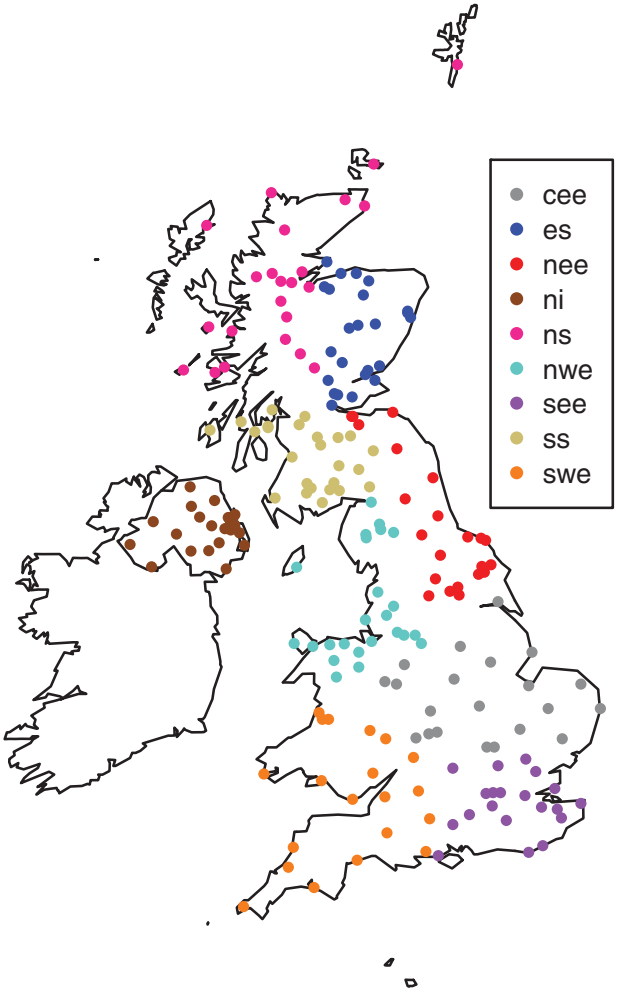

Figure 1. (Colour online) The 204 rainfall measurement locations by region. The points shown in grey represent the central and eastern England (cee) region, as studied in this paper.

extreme rainfall, we filter out a set of annual maxima for each site; that is, in our analysis we discard all but the largest daily rainfall measurement in each calendar year, reducing the number of observations per site to a maximum of 40. Within each region, and for each calendar year, we treat the set of annual maxima from the sites within that region as a single multivariate observation (in terms of the simulated data in Section 3, this is referred to as a single *replication*).

This paper is set out as follows. In Section 2, we provide an outline of some models used for analysing spatial data. We include a brief description of Gaussian processes as used in the field of geostatistics, before giving an overview of spatial models that have been developed as an extension to the standard models from extreme value theory; for demonstration purposes, we then apply these models to the rainfall extremes in central and eastern England, as shown in Figure 1. In Section 3 we demonstrate a recently published R library for the simulation of spatial extremes; we use this library to simulate extremes with known dependence structure (using spatial models giving both asymptotic dependence and asymptotic independence) and investigate the consequences of model mis-specification in terms of the estimation of quantities

such as those given in (i) and (ii) above. In Section 4, we briefly review some diagnostics for assessing the nature of dependence present in real-life spatial extremes, with a demonstration of the application of one of these techniques to rainfall extremes observed in central and eastern England. As with the simulated data in the previous section, we then investigate the effects of model mis-specification for the rainfall process over this region of the UK.

2. Models for spatial data

2.1 *Gaussian processes*

By far the most simple and well-understood approach to modelling spatial data is to assume our spatial process $\{Y(x)\}$ is Gaussian; that is, $\{Y(x)\}$ is a stochastic process for which any finite linear combination of samples has a joint Gaussian distribution. This is often written as $\{Y(x)\} \sim \mathcal{GP}(\mu, \rho)$, where μ and ρ are the mean and correlation functions, respectively.

In the geostatistics literature [9,11,33], the correlation function of the Gaussian process is usually defined by its range, nugget effect and sill/partial sill, as shown in Figure 2. For a pair of sites, the range can be thought of as the distance beyond which observations on the variable of interest are no longer spatially correlated. At an infinitesimally small separation distance, there is often a correlation 'nugget effect', attributed to measurement errors or spatial sources of variation at distances smaller than the sampling interval; the partial sill can be thought of as the spatial correlation after the nugget effect, and the sill itself is the sum of the partial sill and the nugget effect. There are a number of standard correlation functions available that will achieve the desired effect of a decay in dependence between the measured variable at a pair of sites with distance, as illustrated in Figure 2. For example, the exponential correlation function, with scale parameter $\lambda > 0$, is given by

$$\rho(h) = \exp\left\{\frac{-h}{\lambda}\right\}, \tag{1}$$

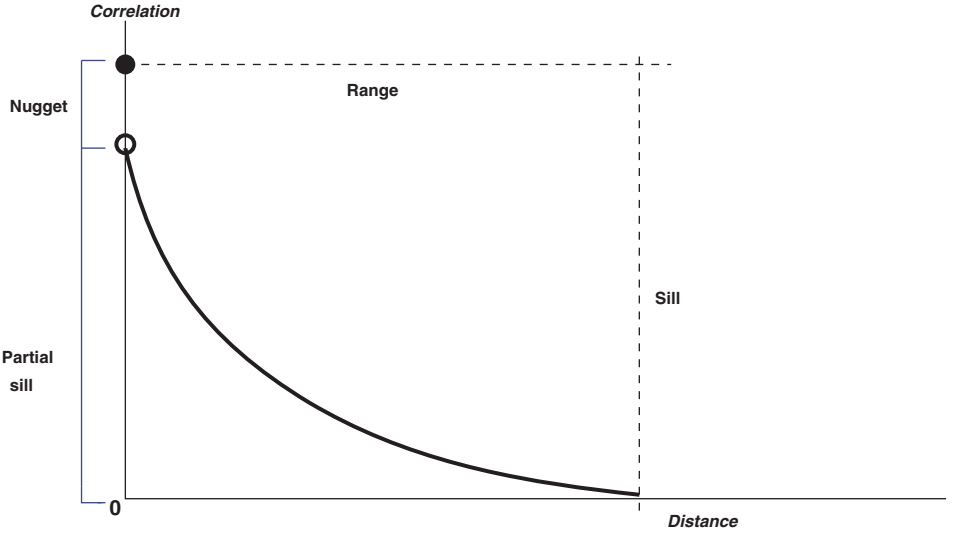

Figure 2. Anatomy of a typical correlation function.

where h is the distance between a pair of sites; a generalisation of this is the powered exponential function with additional shape parameter ν, for $0 < \nu \leq 2$:

$$\rho(h) = \exp\left\{-\left(\frac{h}{\lambda}\right)^{\nu}\right\}. \tag{2}$$

The Whittle–Matérn correlation function, with shape and scale parameters $\nu > 0$ and $\lambda > 0$, respectively, is given by

$$\rho(h) = \frac{(h/\lambda)^{\nu}}{2^{\nu-1}\Gamma(\nu)} L_{\nu}\left(\frac{h}{\lambda}\right), \tag{3}$$

where L_{ν} is the modified Bessel function of order ν and $\Gamma()$ is the gamma function. Correlation functions such as those given by Equations (1)–(3) are *isotropic*, that is, they depend only on the length of h and not its orientation; *anisotropy* can be added by replacing h with $(h^{\mathrm{T}}Ah)^{0.5}$, where A is a positive definite matrix with determinant 1; see, for example, [11]. The exponential function, as given in Equation (1), is often used for its simplicity; the powered exponential and Whittle–Matérn functions, given in Equations (2) and (3), respectively, are more flexible owing to the addition of the shape parameter ν, and thus also widely used in practice. Other commonly used functions include the Cauchy and Bessel correlation functions; see [10] for more details.

Unlike max-stable models (see below), such Gaussian-based geostatistical models are widely understood, and extremely flexible owing to the range of correlation functions that are available. Thus, in the context of the *extremes* of spatial processes, it is rather trivial to place the sets of block maxima on the Gaussian scale via the probability integral transformation, henceforth using methods available from the standard geostatistical toolbox to model the dependence structure. However, the requirement that the model for the original data should be max-stable restricts the structure of the correlation function that can be used, even after transformation to the Gaussian scale. Also, although the correlation function ρ governs the strength of dependence between observations at a pair of sites across the entire range of positive dependence, all Gaussian processes (unlike the max-stable models outlined below) bear the property of asymptotic independence; a property that may/may not be desirable for the spatial extremes of an environmental process (see the discussion in Section 1).

2.2 *Limiting models for extremes*

The aim of this section is to summarise how standard models for extremes have been extended to provide a spatial interpretation. Before detailing these so-called *max-stable processes*, we will briefly outline the standard framework for modelling univariate/multivariate extremes within the context of the traditional block/componentwise maxima setting (respectively).

2.2.1 *Univariate considerations*

Suppose $Y_1, Y_2, \ldots, Y_n$ are independent and identically distributed random variables with distribution function F. Then, if there exists sequences of constants $\{a_n\} > 0$ and $\{b_n\}$ such that

$$\Pr\{a_n^{-1}(\max[Y_1, Y_2, \ldots, Y_n] - b_n) \leq x\} \to G(x)$$

as $n \to \infty$, G must be max-stable; that is,

$$G^n(b'_n + a'_n x) = G(x)$$

must hold for sequences $\{a'_n\} > 0$ and $\{b'_n\}$, where $n \in \mathbb{N}$. It turns out that the complete family of non-degenerate limiting distributions with this property is the generalised extreme value (GEV)

distribution, with distribution function

$$\mathcal{G}(y;\mu,\sigma,\xi) = \begin{cases} \exp\left[-\left(1+\dfrac{\xi(y-\mu)}{\sigma}\right)_+^{-1/\xi}\right], & \xi \neq 0, \\ \exp\left[-\exp\left(\dfrac{-(y-\mu)}{\sigma}\right)\right], & \xi = 0, \end{cases} \tag{4}$$

where $-\infty < \mu < \infty$, $\sigma > 0$ and $-\infty < \xi < \infty$ are location, scale and shape parameters (respectively) and u_+ denotes $\max(u, 0)$. The fact that the support of the GEV depends on ξ would lead some to prefer to think of the GEV as a family of distributions indexed by ξ. However, our characterisation above as a single three-parameter distribution with a functional form for the support is standard throughout the extreme value literature. In practice, Equation (4) is often used to model a set of block maxima; typically, n is set to be the number of observations per year, giving the usual analysis of annual maxima. Estimates of an extreme quantile z_r can be obtained by equating the right-hand side of Equation (4) to $1 - r^{-1}$ and then solving for $y = z_r$, where z_r is the r-observation *return level* associated with return period r. With block length equal to one year, this is equivalent to the *r-year* return level, typically interpreted as the value that is exceeded once on average every r years. Specifically,

$$\hat{z}_r = \begin{cases} \mu + \dfrac{\sigma}{\xi}[(-\log(1-r^{-1}))^{-\xi} - 1], & \xi \neq 0, \\ \mu - \sigma \log[-\log(1-r^{-1})], & \xi = 0. \end{cases} \tag{5}$$

where μ, σ and ξ are replaced with their estimates $\hat{\mu}$, $\hat{\sigma}$ and $\hat{\xi}$, perhaps obtained by maximising the GEV log-likelihood. Indeed, in the maximum-likelihood setting, confidence intervals for $\hat{z}_r$ are usually obtained via profile log-likelihood owing to the severe asymmetry of the surface of the log-likelihood often encountered for return levels.

2.2.2 *Extension to the multivariate setting*

Now suppose we have independent sequences of D-variate random variables $(Y_{1i}, Y_{2i}, \ldots, Y_{Di})$, $i = 1, 2, \ldots, n$, with componentwise maxima

$$\max[Y_{11}, Y_{12}, \ldots, Y_{1n}], \max[Y_{21}, Y_{22}, \ldots, Y_{2n}], \ldots, \max[Y_{D1}, Y_{D2}, \ldots, Y_{Dn}].$$

If it exists (and is non-degenerate), the limiting joint distribution of these maxima, after marginal transformation to standard Fréchet with distribution function $\exp(-1/z)$, for $z > 0$, can be written as

$$\Pr\{Z_1 \leq z_1, Z_2 \leq z_2, \ldots, Z_D \leq z_D\} = \exp\{-V(z_1, z_2, \ldots, z_D)\}, \tag{6}$$

$z_1, z_2, \ldots, z_D > 0$, where the exponent measure $V(z_1, z_2, \ldots, z_D)$ can be written according to its so-called spectral representation:

$$V(z_1, z_2, \ldots, z_D) = \int_{\mathcal{S}_D} \max\left(\frac{\omega_1}{z_1}, \frac{\omega_2}{z_2}, \ldots, \frac{\omega_D}{z_D}\right) \mathrm{d}M(\omega_1, \omega_2, \ldots, \omega_D),$$

where M is a measure on the D-dimensional unit simplex $\mathcal{S}_D$ (see [27]) satisfying the constraint

$$\int_{\mathcal{S}_D} \omega_j dM(\omega_1, \omega_2, \ldots, \omega_D) = 1, \quad j = 1, \ldots, D, \tag{7}$$

in order to achieve unit Fréchet margins. Unlike the univariate setting, there is no simple parametric form for the limiting distribution; V can take any form subject to Equation (7). Commonly

used choices include the logistic, negative logistic, bilogistic and Dirichlet models – see Chapter 8 in [5] for more details.

2.2.3 *Extending further: the infinite-dimensional case*

The infinite-dimensional analogues of multivariate extreme value distributions, as provided by Equation (6), are *max-stable processes* [12]. More precisely, consider a stochastic process $\{Y(x)\}$, $x \in \mathbb{R}^D$, having continuous sample paths. Then the limiting process is given by

$$\{a_n^{-1}(x)(\max[Y_1(x), Y_2(x), \ldots, Y_n(x)] - b_n(x))\} \to \{Z(x)\},$$

as $n \to \infty$, where Y_i, $i = 1, \ldots, n$ are independent replications of Y, $a_n(x) > 0$ and $b_n(x) \in \mathbb{R}^D$ are sequences of continuous functions and the limiting process Z is assumed to be non-degenerate. Schlather [30] shows that $Z(x)$ can be written, according to its spectral characterisation, as:

$$Z(x) = \max[\zeta_1 Y_1(x), \zeta_2 Y_2(x), \ldots],$$

where $\zeta_1, \zeta_2, \ldots$ are the points of a Poisson process on $(0, \infty]$ with intensity $d\Lambda(\zeta) = \zeta^{-2} d\zeta$ and $\mathbb{E}[Y(x)] = 1$ for all $x \in \mathbb{R}^D$. Different choices for the process Y give some useful max-stable processes, such as the *Gaussian extreme value process*, sometimes referred to as the Smith model after the (famously) unpublished University of Surrey technical report by Smith [31]; the *extremal Gauss process* [30], arrived at by allowing the $\{Y(x)\}$ to be a (stationary) standard Gaussian process with some correlation function ρ (see Section 2.1); and *Brown–Resnick processes*, obtained by taking $\{Y(x)\}$ as $\exp\{U(x) - \sigma^2(x)/2\}$, where $U(x)$ is a (stationary) Gaussian process with variance function $\sigma^2(x)$ [22]. The extremal Gauss process cannot account for independence between extremes when the distance h increases indefinitely; recently, Brown–Resnick processes have been cited as an important alternative, as they possess the property of independence in the limit as $h \to \infty$. de Haan and Pereira [14] give additional examples; a comprehensive review is provided by Davison *et al*. [10].

2.3 *Application to rainfall data: central and eastern England*

As an illustration of the spatial models discussed in this section so far, we demonstrate the use of a recently published R library – namely the `CompRandFld` library [26] – to fit a Gaussian process, and two max-stable processes, to the annual maximum rainfall measurements at locations in the central and eastern England region shown in Figure 1. We have rainfall measurements for 21 sites in this region. In some years, the annual maximum rainfall observation is missing at some sites; when this occurs, we regard the multivariate observation for that year as missing entirely. For central and eastern England this leaves us with 26 years worth of complete componentwise maxima. Although wasteful of data, this approach for dealing with missingness is adequate for the purpose of the work in this section; [1] demonstrate a data augmentation procedure to generate artificial data to replace the missing components. The top row of plots in Figure 3 shows the distribution of the annual maxima for the 21 sites in this region.

To fit a Gaussian process to our spatial data, we first need to transform the margins to standard Normal. As discussed in Section 2.2.1, the GEV (4) is the limiting model for our annual maxima at each site. Denoting the marginal GEV distribution functions by $\mathcal{G}_j(y; \mu_j, \sigma_j, \xi_j)$, $j = 1, \ldots, 21$, we transform to standard Normal using

$$\Phi^{-1}\{\mathcal{G}_j(y; \hat{\mu}_j, \hat{\sigma}_j, \hat{\xi}_j)\},$$

where Φ is the distribution function of the standard Normal, and $\hat{\mu}$, $\hat{\sigma}$ and $\hat{\xi}$ are the maximum-likelihood estimates of the GEV location, scale and shape respectively. Estimates of the marginal

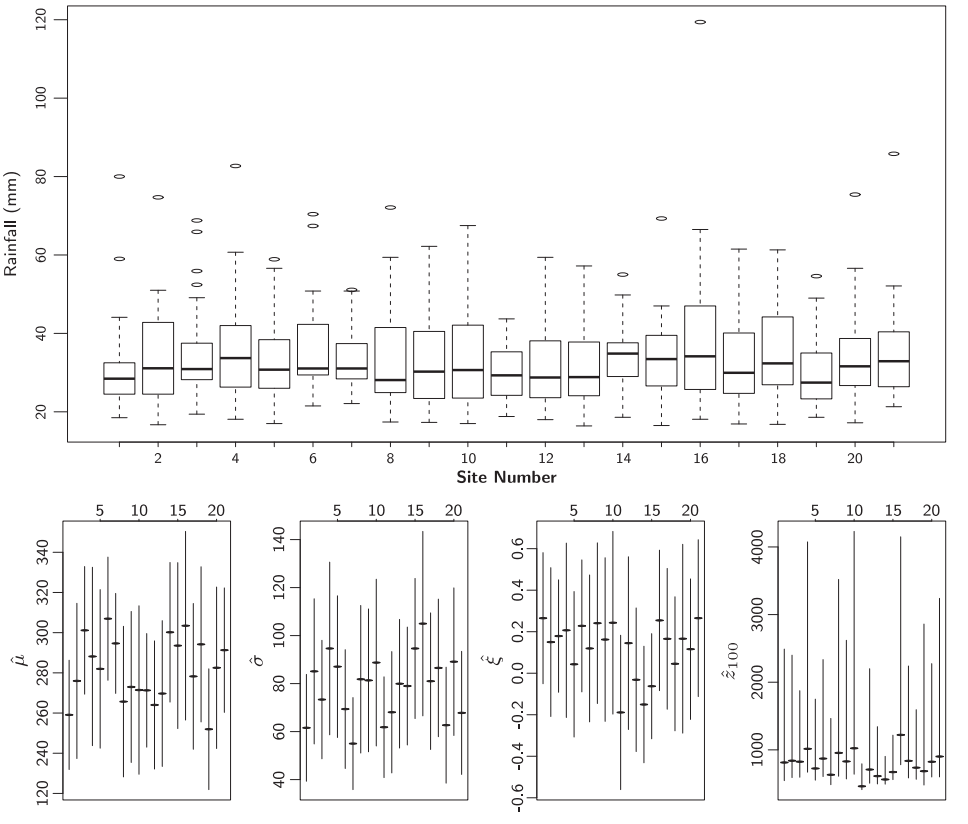

Figure 3. Top row: Boxplots showing the distribution of annual maxima for each site in central/eastern England; bottom row: maximum-likelihood estimates of marginal GEV parameters, and 100 year return levels, with 95% confidence intervals (obtained via profile likelihood for the return levels).

GEV parameters and 100-year return levels, with 95% confidence intervals, are shown in the bottom row of plots in Figure 3. These estimates have been obtained using conventional maximum-likelihood estimation. Coles and Dixon [6] suggest that the performance of this method may be poor for small samples, and they propose a penalised likelihood alternative. However when we implemented the penalised likelihood approach for our 21 sites, using the penalty function recommended by Coles and Dixon [6], we found almost no change in our parameter estimates or standard errors, and the resulting plots of parameter estimates were visibly indistinguishable from those shown in Figure 3. The `FitComposite` function in `CompRandFld` can then be used to fit a Gaussian process to the sample of transformed spatial extremes by maximising the full log-likelihood in the usual way. Using the exponential, powered exponential and Whittle–Matérn correlation functions, as given by Equations (1)–(3), respectively, gives the results shown in the top row of Table 1; parameter estimates and their standard errors are shown, along with the values of the maximised log-likelihood.

Transforming the marginal annual maxima to standard Fréchet, by applying

$$-\log\{\mathcal{G}_j(y;\hat{\mu}_j,\hat{\sigma}_j,\hat{\xi}_j)\}^{-1},$$

we can also use `FitComposite` to fit various max-stable models. However, rather than attempting to maximise the full log-likelihood, we now use composite likelihood methods [8,25,32], as it would be computationally infeasible to obtain the full joint density function from

Table 1. Estimates of the parameters in the correlation functions (with standard errors in parentheses) for each of the spatial models fitted to rainfall annual maxima for sites in central and eastern England.

Model	Correlation	Scale (λ)	Shape (ν)	$\ell_{\max}$	CLIC
Gaussian process	Exponential	56.444 (44.937)	–	−706.52	–
	Powered exponential	44.390 (141.503)	0.689 (1.607)	−705.45	–
	Whittle–Matérn	78.569 (220.167)	0.328 (1.785)	−706.18	–
Extremal Gauss process	Exponential	168.025 (64.105)	–	−23575.98	47034
	Powered exponential	168.344 (64.724)	1.075 (0.188)	−23572.12	47020
	Whittle–Matérn	168.424 (83.844)	0.158 (0.078)	−23572.48	47010
Brown–Resnick process	Exponential	22.230 (3.889)	–	−23457.71	46947
	Powered exponential	20.342 (4.003)	0.532 (1.251)	−23441.60	46701
	Whittle–Matérn	167.879 (52.316)	1.000[a]	−23423.18	46523

Notes: The maximised log-likelihood $\ell_{\max}$ is also shown (the maximised *composite* log-likelihood for max-stable models), along with the composite likelihood information criterion (where composite likelihood has been used).
[a]The parameter was held fixed.

Equation (6) unless D was small. Specifically, we assume independence between each of our multivariate observations (i.e. independence across years); *within* years, the log-likelihood is formed by summing the logged pairwise marginal densities, and a sandwich estimator is used to estimate the variance matrix of the parameter vector. For more details, see [10]. Thus, Table 1 also shows parameter estimates for an extremal Gauss process and a Brown–Resnick process. The maximised composite log-likelihood is also shown, along with the composite log-likelihood information criterion; [10] demonstrate how these figures can be used to aid model choice.

By allowing the marginal GEV parameters to depend on the latitude/longitude of the sites at which rainfall data has been collected, Davison *et al*. [10] also show how the spatial models can be interpolated to provide return level estimates at locations for which no observations have been collected; hence, they produce return level maps for rainfall over a region in Switzerland. In our work, we are interested in quantities such as (i) and (ii), as defined in Section 1, and how the estimation of such quantities can be sensitive to the assumptions regarding the spatial structure in our process. As discussed in Section 1, identifying the type of spatial dependence in our extremes (i.e. asymptotic/sub-asymptotic) can be difficult, and we will return to this issue in Section 4.1.

3. Simulation study

In this section, we use the `CompRandFld` package to fit, and simulate from, Gaussian and max-stable processes, as outlined in Sections 2.1 and 2.2, respectively. The function `RFSim` within this library simulates one or more replications of a random field – either Gaussian or max-stable – with a given correlation model and parameters therein. In the context of our rainfall extremes, we use the term replication to correspond to a single multivariate observation across several sites within a region. The primary aim of this section is to investigate the consequences of model mis-specification; in particular, the effects of assuming max-stability when in fact the true process is asymptotically independent, and vice-versa, on estimates of various quantities of practical interest.

3.1 *Study design*

In each main arm of the study, we simulate a ‘master’ random field, this field being either Gaussian (and so possessing the property of asymptotic independence), or max-stable. For our

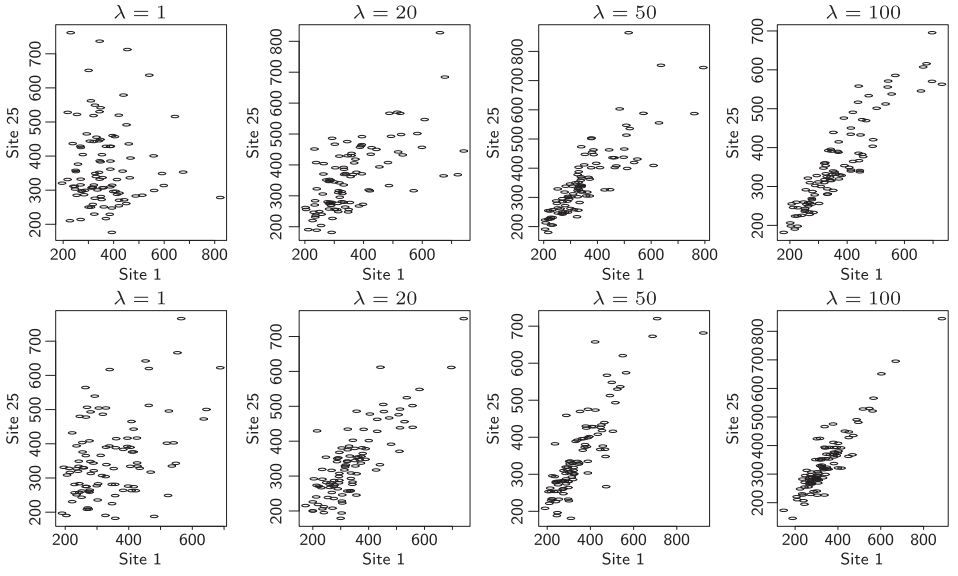

Figure 4. Plots of data for two sites simulated from (1) a Gaussian process (top row), and (2) an extremal Gauss process (bottom row), for the four values of scale parameter λ used in the correlation function.

max-stable random fields, we simulate from both the extremal Gauss process and a Brown–Resnick process. For each of our simulated master data sets, we use the exponential correlation function, as given by Equation (1), considering a range of values of the scale parameter λ to reflect the various levels of spatial dependence that might be encountered with real life environmental extremes; specifically, we use $\lambda = 1, 20, 50$ and 100, where the strength of spatial correlation increases with λ. For each random field we generate, we assume a zero nugget effect. As with the rainfall data explored in Section 2.3, our simulated spatial data will have GEV marginal distributions; we transform from standard Normal in the Gaussian process, and from standard Fréchet in the max-stable processes, to GEV with $\mu = 300$, $\sigma = 80$ and $\xi = 0.1$ – values similar to those estimated for the rainfall data in central and eastern England shown in Figure 3. Each master random field will consist of $N = 100$ replications across $n = 50$ 'sites'; for each site, Cartesian co-ordinates are generated randomly from a Uniform $\mathcal{U}(0, 10)$ distribution. Plots of the data from two randomly chosen sites, after transformation to GEV margins, are shown in Figure 4 for both the Gaussian and extremal Gauss processes. The effect of λ on the dependence is obvious, as is the effect of asymptotic dependence as provided by the extremal Gauss process.

To each master random field we then fit the correct spatial process, and an incorrect spatial process. For example, if the master data set has been simulated from a max-stable process, we will use the `FitComposite` function demonstrated in Section 2.3 to fit the correct max-stable process – but we will also fit, incorrectly, a Gaussian process. The resulting estimated parameters from both the correct and incorrect fits will then be used to simulate $K = 10^6$ replications from both models, giving a $(K \times n)$ matrix of realisations $y_{k,i}$, $k = 1, \ldots, K$ and $i = 1, \ldots, n$, for both the correct and incorrect models. As well as comparing realisations from a Gaussian process to those from the max-stable processes, we will also consider the consequences of specifying the incorrect max-stable process; that is, for example, the effects of fitting a Brown–Resnick process when in fact the master data set has been simulated from an extremal Gauss process.

3.2 *Quantities of interest*

Using each of our $(K \times n)$ matrices of simulated data, we can investigate the consequences of model mis-specification on the estimation of quantities such as those outlined in (i) and (ii) in Section 1. For example, we discussed in Section 1 that recent flooding events in the UK have often occurred as a result of the observation of simultaneous rainfall extremes across several locations within a region. Thus, from our K replications, we can obtain, empirically, estimates of the joint exceedance probabilities of the marginal r-year return levels z_r for S sites of interest. In this study, we use $r = 10$, 50, 200, 1000, 5000 and 10,000, and we obtain estimates of the joint exceedance probabilities P_r for $S = 3$ and $S = 10$ sites of interest. Of the 50 sites for which we simulate data, we select these sites of interest in two ways: (i) as the S sites which are closest together according to the minimum polygonal perimeter connecting them, giving $\{s_t^{\text{NEAR}}\}$, $t = 1, \ldots, S$; (ii) as the S sites which are furthest apart according to their polygonal perimeter, giving $\{s_t^{\text{FAR}}\}$.

Then, for each of the correct and incorrect models for spatial dependence, we estimate P_r using

$$\hat{P}_r = \sum_{k=1}^{K} \frac{I_k}{K}, \tag{8}$$

where

$$I_k = \begin{cases} 1 & \text{if } y_{k,s_t^*} > z_r, \quad t = 1, \ldots, S; \\ 0 & \text{otherwise,} \end{cases}$$

at each replication $k = 1, \ldots, K$, where * is used generically for NEAR and FAR. From our $(K \times n)$ matrix of realisations from each of the correct and incorrect models, we also obtain estimates of return levels for total maxima across all sites. For each row k, $k = 1, \ldots, K$, we find

$$T_k = \sum_{i=1}^{n} y_{k,i}, \tag{9}$$

which – in the context of the rainfall data analysed in Section 2.3 – would correspond to total maximum daily rainfall accumulations across all sites within a region, a surrogate for the total maximum daily rainfall accumulation within the region as a whole. Denoting by $T^{[k]}$ the kth order statistic of $\{T_1, T_2, \ldots, T_K\}$, we then estimate empirically the r-year return level for total maxima across all replications as

$$\hat{z}_{r(\text{TOT})} = T^{[K(1-r^{-1})]};$$

as for estimates of $\hat{P}_r$, we consider $r = 10$, 50, 200, 1000, 5000 and 10,000.

3.3 *Bootstrapping*

Our estimates $\hat{P}_r$ and $\hat{z}_{r(\text{TOT})}$, using Equations (8) and (9), respectively, are one-off estimates obtained empirically from the simulated data. To assess the variability of these estimates, we could repeat the simulation procedure multiple times, each time with $K = 10^6$. However, computationally this would be burdensome. Rather, we use bootstrap methods to estimate $\mathsf{var}[\hat{P}_r]$, $\mathsf{var}[\hat{z}_{r(\text{TOT})}]$ and confidence intervals for both. Specifically, for each bootstrap replication b, $b = 1, \ldots, B$, we randomly sample (with replacement) K rows from each $(K \times n)$ matrix of

simulated data. For each bootstrap replication of the simulated spatial process, we find $\hat{P}_r$ and $\hat{z}_{r(\mathrm{TOT})}$, yielding a collection of estimates

$$\{\hat{P}_r^{(1)}, \hat{z}_{r(\mathrm{TOT})}^{(1)}, \ldots, \hat{P}_r^{(B)}, \hat{z}_{r(\mathrm{TOT})}^{(B)}\},$$

from which we can estimate variances or construct confidence intervals. Here, we use $B = 1000$. We estimate $\mathsf{var}[\hat{P}_r]$ and $\mathsf{var}[\hat{z}_{r(\mathrm{TOT})}]$ by finding the variance of the bootstrap replicates $\hat{P}_r^{(b)}$ and $\hat{z}_{r(\mathrm{TOT})}^{(b)}$, respectively, $b = 1, \ldots, B$. For confidence intervals, rather than record the 2.5%- and 97.5%-iles from the bootstrap samples (the standard *percentile* method), we use *bias-corrected accelerated* (BC_a) intervals as proposed in [15]. This method corrects for bias owing to non-normality (particularly useful for $\hat{z}_{r(\mathrm{TOT})}$); it also accelerates convergence to a solution by correcting for the rate of change of the normalised standard errors of $\hat{P}_r$ and $\hat{z}_{t(\mathrm{TOT})}$, relative to their true values, in constructing the confidence bounds of the percentile method. For a detailed demonstration of the practical implementation of these BC_a intervals, see [17].

3.4 *Results*

Table 2 shows some results from one arm of the simulation study. Here, the master random field has been simulated from an Extremal Gauss process with an exponential correlation model, with scale parameter $\lambda = 20$. The correct form of model has then been fitted, yielding an estimated correlation parameter of $\hat{\lambda} = 21.043$; a million replications of this fitted random field have then been simulated, and the first column of Table 2 shows estimates of P_r for $S = 3$ and $S = 10$, and also $z_{r(\mathrm{TOT})}$. The next column shows estimates of these quantities of interest when an *incorrect* spatial model is fitted: specifically, we show estimates for the standard Gaussian process (after marginal transformation to standard Normal), results again being based on a million replications of this fitted process.

Figure 5 summarises some of these results: the top row of plots are densities of the bootstrap samples for $\hat{P}_{10}$, $\hat{P}_{50}$, $\hat{P}_{200}$ and $\hat{P}_{1000}$ for simulations from the (correct) extremal Gauss process (red), and the (incorrect) Gaussian process (blue), for the $S = 3$ sites closest together; the bottom row shows the same information but for the $S = 3$ sites furthest apart.

Both the table of results and the plots reveal significant differences in estimates of P_r, when the sites of interest lie far apart (FAR), for the different spatial structures assumed. For example, estimates of P_r are consistently larger when (correctly) assuming the max-stable extremal Gauss model, than when (incorrectly) assuming asymptotic independence with the standard Gaussian process, and significantly so. In practical terms this might suggest that using a model which assumes asymptotic independence when, in fact, there is asymptotic dependence, could lead to significant under-estimation of the likelihood of simultaneous extremes at multiple locations.

Similarly, estimates of P_r when the sites of interest lie within close proximity of each other (NEAR) are almost always larger when we have simulated from the (correct) extremal Gauss random field, relative to the (incorrect) standard Gaussian process; however, any differences are hardly ever significant when comparing the corresponding bootstrapped confidence intervals. This might be as we would expect: for sites very close together, dependence between extremes from a standard Gaussian random field is more likely to persist than for sites further apart, giving less distinction between estimates of P_r when simulating from the extremal Gauss random field and the standard Gaussian process.

Considering differences between estimates of $z_{r(\mathrm{TOT})}$ when assuming the correct form of spatial structure and when assuming asymptotic independence, for all but the shortest return period estimates are larger when we (correctly) use the extremal Gauss process. Again, in practical terms, this might suggest that an incorrect assumption of asymptotic independence could lead to significantly under-estimated return levels and perhaps substantial under-protection in flood

Table 2. Estimates of: (i) the joint probability, P_r, of simultaneously exceeding r-year return levels in $S = 3$ margins and $S = 10$ margins, and (ii) r-year return levels for total rainfall maxima across all margins, $z_{r(\text{TOT})}$. Bootstrapped 95% confidence intervals are shown in parentheses.

True spatial process: Extremal Gauss ($n = 50, N = 100$) *Correlation model*: Exponential, $\lambda = 20$			*Simulated process* ($n = 50$, $K = 10^6$)			
			Extremal Gauss process ($\hat{\lambda} = 21.043$)		Standard Gaussian process ($\hat{\lambda} = 117.367$)	
$\hat{P}_r \times 100 (S = 3)$	$r = 10$	NEAR	9.272	(8.690, 9.821)	8.202	(7.652, 8.740)
		FAR	4.392	(4.001, 4.780)	2.227	(1.951, 2.520)
	$r = 50$	NEAR	1.735	(1.481, 1.990)	1.683	(1.450, 1.932)
		FAR	0.798	(0.620, 0.982)	0.191	(0.110, 0.282)
	$r = 200$	NEAR	0.449	(0.330, 0.590)	0.512	(0.371, 0.670)
		FAR	0.207	(0.120, 0.303)	0.029	(0.000, 0.071)
	$r = 1000$	NEAR	0.110	(0.051, 0.181)	0.101	(0.040, 0.171)
		FAR	0.049	(0.010, 0.101)	0.000	–
	$r = 5000$	NEAR	0.021	(0.000, 0.050)	0.020	(0.000, 0.050)
		FAR	0.020	(0.000, 0.050)	0.000	–
	$r = 10,000$	NEAR	0.010	(0.000, 0.030)	0.010	(0.000, 0.031)
		FAR	0.010	(0.000, 0.030)	0.000	–
$\hat{P}_r \times 100 (S = 10)$	$r = 10$	NEAR	5.896	(5.440, 6.370)	4.208	(3.790, 4.601)
		FAR	2.627	(2.320, 2.931)	0.827	(0.650, 1.001)
	$r = 50$	NEAR	1.056	(0.861, 1.271)	0.708	(0.541, 0.880)
		FAR	0.429	(0.311, 0.550)	0.039	(0.010, 0.080)
	$r = 200$	NEAR	0.308	(0.210, 0.420)	0.179	(0.100, 0.271)
		FAR	0.098	(0.040, 0.161)	0.000	–
	$r = 1000$	NEAR	0.060	(0.020, 0.111)	0.029	(0.000, 0.071)
		FAR	0.029	(0.000, 0.070)	0.000	–
	$r = 5000$	NEAR	0.010	(0.000, 0.030)	0.000	(0.000, 0.030)
		FAR	0.006	(0.000, 0.011)	0.000	–
	$r = 10,000$	NEAR	0.000	–	0.000	–
		FAR	0.000	–	0.000	–
$\hat{z}_{r(\text{TOT})}$ (Thousand)	$r = 10$		24.182	(24.096, 24.267)	24.336	(24.254, 24.423)
	$r = 50$		32.299	(32.051, 32.520)	31.763	(31.570, 31.973)
	$r = 200$		40.308	(39.745, 40.966)	38.738	(38.331, 39.123)
	$r = 1000$		50.948	(49.826, 51.972)	47.651	(46.452, 49.302)
	$r = 5000$		64.583	(60.050, 70.194)	58.668	(54.897, 61.809)
	$r = 10,000$		71.626	(66.681, 75.694)	63.732	(59.630, 71.285)

defence terms. Differences between estimates of our quantities of interest were negligible when comparing results from the extremal Gauss process and a Brown–Resnick process, although it is noted that estimates assuming a Brown–Resnick process were consistently smaller.

Another arm of the study (for which results are not shown here) revealed that using a max-stable process when in fact our extremes are asymptotically independent (simulated from a standard Gaussian process) can result in significantly *over*-estimated P_r and $z_{r(\text{TOT})}$, especially when using the extremal Gauss process. In practical terms, such *over*-protection, perhaps in terms of flood defence, could have huge financial implications. Of course, these findings can only be extended to real data if we are able to clarify the nature of the dependence present between our observed spatial extremes. As discussed in Section 1, this can be challenging.

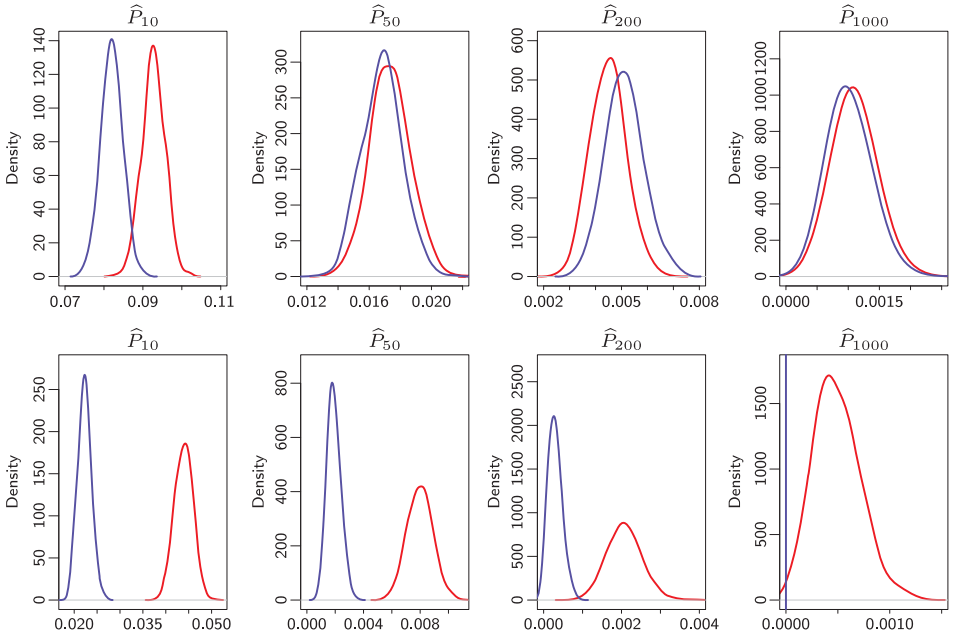

Figure 5. (Colour online) Bootstrap sampling distributions of the joint probability of exceeding the marginal 10, 50, 200 and 1000 year return levels at $S = 3$ sites of interest when fitting the correct max-stable model (extremal Gauss process, results in red) and an incorrect asymptotically independent model (standard Gaussian process, results in blue). The top row of plots show results for the $S = 3$ sites closest together (NEAR), the bottom row for the $S = 3$ sites furthest apart (FAR).

4. Rainfall extremes: central and eastern England revisited

4.1 *Nature of spatial dependence*

As in Section 2.2.3, let $Z(x)$ be a stationary max-stable process with standard Fréchet margins. Then it can be shown [3] that

$$\Pr\{Z(x) < z, Z(x+h) < z\} = \exp\left\{\frac{-\theta(h)}{z}\right\},$$

for some separation distance h. The function $\theta(h)$ is known as the *extremal coefficient function* where, generally, $1 \leq \theta(h) \leq D$, where $\theta(h) = 1$ and $\theta(h) = D$ represent full dependence, and complete independence, respectively. The analogous function in the asymptotically independent class is the *coefficient of tail dependence function* $\eta(h)$, where $\eta(h) = 1/\theta(h)$; thus, $\eta(h) = 1$ would suggest asymptotic dependence between a pair of sites separated by a distance h. A Hill estimator [21] can be used to estimate $\eta(h)$, as in [34].

Figure 6 shows Hill estimates of $\eta(h)$ for all pairs of sites in the central and Eastern England region shown in Figure 1, including (block) bootstrapped 95% confidence intervals. This plot illustrates the difficulties often encountered when attempting to assess the nature of spatial dependence across a region. For pairs of sites that are closer together, that is, for small h – perhaps $h < 100$ km – it seems plausible that $\eta(h) = 1$, suggesting genuine extremal dependence and the suitability of a max-stable model. However, for larger values of h asymptotic independence may be more plausible since here it is more likely that $\eta(h) < 1$.

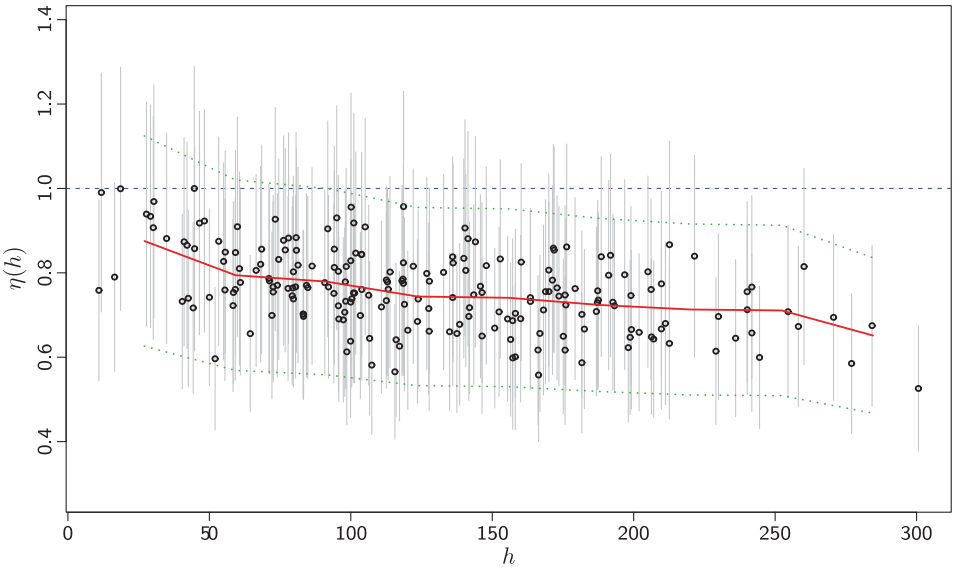

Figure 6. (Colour online) Hill estimates of the coefficient of tail dependence (black circles) for every pair of sites in the central and eastern England region, with 95% bootstrapped confidence intervals. The red and green lines show loess curves fitted through the Hill estimates, and the lower/upper 95% confidence bounds, respectively. The separation distance is in km.

4.2 *Estimation of P_r and $z_{r(\mathrm{TOT})}$*

Although the plot in Figure 6 is informative, the mix of dependence structures, depending on the separation distance h, makes it difficult for us to choose between, for example, a standard Gaussian process and an extremal Gauss process, when attempting to estimate P_r and $z_{r(\mathrm{TOT})}$. As we discovered in the simulation study, model mis-specification, in terms of the dependence structure, may not be important when estimating P_r for sites of interest that are close together (see the top row of plots in Figure 5). However, if our sites of interest are further apart, using a Gaussian process when, in fact, the true process is max-stable (for example) could lead to significant under-estimation (and possibly substantial under-protection in flood defence terms; see the bottom row of plots in Figure 5).

Table 3 shows estimates of P_r for the $S = 3$ sites closest together, and furthest apart, in the central and eastern England region; estimates of $z_{r(\mathrm{TOT})}$ are also shown. An exponential correlation function was assumed (see Equation (1)) with scale parameter $\lambda = 56.444/\lambda = 168.025$ for a standard Gaussian/extremal Gauss process (respectively), as obtained in Section 2.3 and shown in Table 1. The fitted values of the nugget, range and sill were then used alongside these values of λ (and the fitted marginal GPD parameters as shown in Figure 3) to simulate $K = 10^6$ replications from the fitted Gaussian/extremal Gauss random field, and estimates of P_r and $z_{r(\mathrm{TOT})}$ were obtained as in the simulation study in Section 3; results, with associated bootstrapped 95% confidence intervals, are shown in Table 3.

As observed in the simulation study, we see that for the $S = 3$ sites *furthest apart*, there are significant differences in estimates of P_r given by the two fitted random fields for all return periods r. This is because extremal dependence persists as h increases in the extremal Gauss random field, giving significantly higher estimates of P_r here. Interestingly, we also see significant differences in estimates of P_r for the two spatial models when looking at the three sites *closest together* for return periods $r = 10$, 50 and 200 years, with those from the max-stable model being significantly higher. The problem with our real data, of course, is that it is difficult to tell – from

Table 3. Estimates of: (i) the joint probability, P_r, of simultaneously exceeding r-year return levels in $S = 3$ margins, and (ii) r-year return levels for total rainfall maxima across all margins, $z_{r(\text{TOT})}$, for the rainfall extremes observed across the central and eastern England region. Bootstrapped 95% confidence intervals are shown in parentheses.

Rainfall data: Central/Eastern England			Extremal Gauss process ($\lambda = 168.025$)		Standard Gaussian process ($\lambda = 56.444$)	
$\hat{P}_r \times 100\%(S = 3)$	$r = 10$	NEAR	9.199	(8.610, 9.740)	5.218	(4.800, 5.660)
		FAR	8.518	(7.960, 9.070)	0.577	(0.430, 0.730)
	$r = 50$	NEAR	1.757	(1.480, 2.020)	0.771	(0.600, 0.950)
		FAR	1.568	(1.320, 1.820)	0.031	(0.000, 0.070)
	$r = 200$	NEAR	0.490	(0.360, 0.640)	0.199	(0.120, 0.300)
		FAR	0.450	(0.330, 0.580)	0.000	–
	$r = 1000$	NEAR	0.079	(0.030, 0.140)	0.041	(0.010, 0.090)
		FAR	0.069	(0.020, 0.120)	0.000	–
	$r = 5000$	NEAR	0.010	(0.000, 0.030)	0.020	(0.000, 0.050)
		FAR	0.010	(0.000, 0.030)	0.000	–
	$r = 10,000$	NEAR	0.000	–	0.000	–
		FAR	0.000	–	0.000	–
$\hat{z}_{r(\text{TOT})}$ (Thousand mm)	$r = 10$		10.378	(10.270, 10.507)	9.719	(9.642, 9.806)
	$r = 50$		14.123	(13.788, 14.380)	12.023	(11.858, 12.202)
	$r = 200$		18.023	(17.087, 19.131)	14.209	(13.666, 14.661)
	$r = 1000$		21.815	(20.651, 24.002)	16.613	(15.606, 17.490)
	$r = 5000$		25.998	(22.740, 29.153)	18.574	(17.457, 20.768)
	$r = 10,000$		27.635	(24.002, 31.265)	19.271	(17.490, 20.968)

Notes: An extremal Gauss process, and a standard Gaussian process, have been used to allow for an asymptotic dependence/independence dependence structure.

Figure 6 – which of these two models best captures the spatial dependence present in the rainfall extremes in this region.

5. Concluding remarks

In this paper, we have shown that significant estimation bias can occur for some key parameters of hydrological interest, if the nature of the dependence in our spatial extremes is not considered carefully. In particular, we have investigated the effect of model mis-specification on the estimation of joint exceedance probabilities P_r at several sites within a region, as well as high quantiles of total annual maximum rainfall $z_{r(\text{TOT})}$ across all sites. We have shown that – for a selection of sites that might be considered geographically 'far apart' – assuming asymptotic independence when in fact there exists genuine extremal dependence can result in significant under-estimation of the probability of exceeding marginally high quantiles at all sites simultaneously (see Figure 5). Conversely, substantial *over*-estimation of such quantities is observed when fitting a max-stable model to extremes which, spatially, exhibit extremal independence. In the context of flood defence, this is worrying; the former situation could lead to dangerous under-protection, whilst the latter to unjustified financial investment. However, estimates of such key quantities are much less sensitive to the choice of dependence structure for a collection of sites considered to be geographically 'close together'. For example, the spatial dependence implied by a standard Gaussian process for a set of such sites is strong enough, above levels of practical interest, to give estimates of P_r that are not significantly different to those obtained from a max-stable model.

Work has already been done on how to assess the nature of spatial dependence for such real-life extremes; see, for example [3,34]. The difficulty of such diagnostics, as shown in Figure 6 in this paper, is that it is plausible that our data possess a mix of dependence structures (both asymptotic/sub-asymptotic) depending on the separation distance between a pair of sites. Future work, then, will investigate the potential of *mixture models* to model such rainfall extremes across a region, allowing for both forms of spatial dependence simultaneously.

Acknowledgements

We would like to thank Simone Padoan for useful discussions about this work, and for demonstrating the `CompRandFld` package in R. We would also like to thank two referees for their time and suggestions.

References

[1] J. Atyeo and D. Walshaw, *A region-based hierarchical model for extreme rainfall over the UK, incorporating spatial dependence and temporal trend*, Environmetrics 6 (2012), pp. 509–521.
[2] J. Beirlant, Y. Goegebeur, J. Teugels, and J. Segers, *Statistics of Extremes: Theory and Applications*, Wiley, Chichester, 2004.
[3] L. Bel, J.N. Bacro, and Ch. Lantuéjoul, *Assessing extremal dependence of environmental spatial fields*, Environmetrics 19 (2008), pp. 163–182.
[4] T.A. Buishand, L. de Haan, and C. Zhou, *On spatial extremes: With application to a rainfall problem*, Ann. Appl. Stat. 2 (2008), pp. 624–642.
[5] S.G. Coles, *An Introduction to Statistical Modeling of Extreme Values*, Springer, London, 2001.
[6] S.G. Coles and M.J. Dixon, *Likelihood-based inference for extreme value models*, Extremes 2 (1999), pp. 5–23.
[7] S.G. Coles and J.A. Tawn, *Modelling extreme multivariate events*, J. Roy. Statist. Soc. Ser. B 53 (1991), pp. 377–392.
[8] D.R. Cox and N. Reid, *A note on pseudolikelihood constructed from marginal densities*, Biometrika 91 (2004), pp. 729–737.
[9] N.A.C. Cressie, *Statistics for Spatial Data*, Wiley, New York, 1993.
[10] A.C. Davison, S.A. Padoan, and M. Ribatet, *Statistical modeling of spatial extremes*, Stat. Sci. 27 (2012), pp. 161–186.
[11] P.J. Diggle and P.J. Ribeiro, *Model-Based Geostatistics*, Springer, New York, 2007.
[12] L. de Haan, *A spectral representation for max-stable processes*, Ann. Probab. 12 (1984), pp. 1194–1204.
[13] L. de Haan and A. Ferreira, *Extreme Value Theory: An Introduction*, Springer, New York, 2006.
[14] L. de Haan and T.T. Pereira, *Spatial extremes: Models for the stationary case*, Ann. Statist. 34 (2006), pp. 146–168.
[15] B. Efron, *Better bootstrap confidence intervals*, J. Amer. Statist. Assoc. 82 (1987), pp. 171–185.
[16] P. Embrechts, C. Klüppelberg, and T. Mikosch, *Modelling Extremal Events: For Insurance and Finance. Applications of Mathematics (New York)*, Springer, Berlin, 1997.
[17] L. Fawcett and D. Walshaw, *Estimating return levels from serially dependence extremes*, Environmetrics 23 (2012), pp. 272–283.
[18] L. Fawcett and D. Walshaw, *A hierarchical model for extreme wind speeds*, J. Roy. Statist. Soc. Ser. C 55 (2006), pp. 631–646.
[19] J.E. Heffernan and J.A. Tawn, *A conditional approach for multivariate extreme values*, J. Roy. Statist. Soc. Ser. B 66 (2004), pp. 497–546.
[20] J.E. Heffernan and S.I. Resnick, *Limit laws for random vectors with an extreme component*, Ann. Appl. Probab. 17 (2007), pp. 537–571.
[21] B. Hill, *A simple general approach to inference about the tail of a distribution*, Ann. Statist. 3 (1975), pp. 1163–1174.
[22] Z. Kabluchko, M. Schlather, and L. de Haan, *Stationary max-stable fields associated to negative definite functions*, Ann. Probab. 37 (2009), pp. 2042–2065.
[23] M.R. Leadbetter, G. Lindgren, and H. Rootzén, *Extremes and Related Properties of Random Sequences and Series*, Springer, New York, 1983.
[24] A.W. Ledford and J.A. Tawn, *Statistics for near-independence in multivariate extreme values*, Biometrika 83 (1996), pp. 169–187.
[25] B.G. Lindsay, *Composite likelihood methods*, In *Statistical Inference from Stochastic Processes (Ithaca, New York, 1987)*, N.U. Prabhu, ed., Contemp. Math. 80, 1988, pp. 221–239.
[26] S.A. Padoan and M. Bevilacqua, *CompRandFld: Composite-Likelihood based Analysis of Random Fields*, R Package Version 1.0.3, 2013. Available at http://CRAN.R-project.org/package=CompRandFld

[27] J. Pickands, *Multivariate extreme value distributions*, In *Proceedings of the 43rd session of the International Statistics Institute, Vol. 2*, Bull. Inst. Internat. Statist. 49, 1981, pp. 859–878.
[28] S.I. Resnick, *Heavy-Tail Phenomena: Probabilistic and Statistical Modeling*, Springer, New York, 2007.
[29] S.I. Resnick, *Extreme Values, Regular Variation, and Point Processes*, Springer, New York, 1987.
[30] M. Schlather, *Models for stationary max-stable random fields*, Extremes 5 (2002), pp. 33–44.
[31] R.I. Smith, *Max-stable processes and spatial extremes*, unpublished report, 1990.
[32] C. Varin, *On composite marginal likelihoods*, AStA Adv. Stat. Anal 92 (2008), pp. 1–28.
[33] H. Wackernagel, *Multivariate Geostatistics: An Introduction with Applications*, Springer, New York, 2003.
[34] J.L. Wadsworth and J.A. Tawn, *Dependence modelling for spatial extremes*, Biometrika 99 (2012), pp. 253–272.
[35] T.M.L. Wigley, J.M. Lough, and P.D. Jones, *Spatial patterns of precipitation in England and Wales and a revised, homogeneous England and Wales precipitation series*, J. Climatol. 4 (1984), pp. 1–25.