Author for correspondence:
Benjamin D. Youngman
e-mail: b.youngman@exeter.ac.uk

A geostatistical extreme value framework for fast simulation of natural hazard events

Benjamin D. Youngman and David B. Stephenson

College of Engineering, Mathematics and Computer
Science
University of Exeter
Laver Building
North Park Road
Exeter, EX4 4QE, UK

We develop a statistical framework for simulating
natural hazard events that combines extreme value
theory and geostatistics. Robust generalised additive
model forms represent generalised Pareto marginal
distribution parameters while a Student-t process
captures spatial dependence and gives a continuous-
space framework for natural hazard event simulations.
Efficiency of the simulation method allows many
years of data (typically over 10,000) to be obtained
at relatively little computational cost. This makes
the model viable for forming the hazard module of
a catastrophe model. We illustrate the framework
by simulating maximum wind gusts for European
windstorms, which are found to have realistic
marginal and spatial properties, and validate well
against wind gust measurements.

1. Introduction

Natural hazard events can have devastating and
widespread effects on society. In 2014 natural hazards
are estimated to have caused USD 106 and 29 billion
of economic and insured loss, respectively. Over the
past few decades, European windstorms, for example,
have been the second biggest cause of insured loss
from natural hazards globally, after US hurricanes.
Windstorms Christian, Xavier, Dirk and Tini, which
struck over the 2013/2014 winter, caused insured losses
totalling USD 3.3 billion. To ensure their solvency,
insurance companies must have accurate understanding
of their potential losses.

Events that cause significant loss are rare, resulting in a lack of vital data and other relevant knowledge. One method for overcoming data scarcity is to simulate events. This can help build a probabilistic view of loss, in addition to providing information on types and strengths of defences required to offer sufficient protection.

One strategy for producing probabilistic estimates of losses from natural hazard events is known as catastrophe modelling; see Grossi and Kunreuther (2006) for an overview. Such models are usually formed by linking hazard, vulnerability, damage and loss modules. When combined these characterise the extent of the hazard event, the property susceptible to damage, the damage caused to the property, given the hazard, and the subsequent loss. Often catastrophe models are used to estimate loss distributions, of which extreme quantiles are typically important. (For example, the Solvency II directive[1] is based on the 99.5% quantile of a company's annual loss distribution.) Hazard modules can be used to produce arbitrarily many synthetic events from which losses can be calculated, to improve precision of extreme quantile estimates. Various contrasting approaches to the hazard module exist: translations, distortions or parsimonious parameterisations of historical events, which can fail to capture the full population variation in events; physically-based simulation models similar to climate models, which can under-represent processes, such as wind gust speeds (Haas et al., 2014; Roberts et al., 2014), and in turn underestimate losses; or multivariate statistical models that incorporate extreme value statistics, which have extended from bivariate dependencies (Bonazzi et al., 2012; Brodin and Rootzén, 2009) to max-stable processes (Davison and Gholamrezaee, 2012).

The area affected by a natural hazard event can vary considerably in size: heatwaves might affect entire continents, whereas flooding events may only span a few metres. To capture events entirely, simulation domains may be very different in size. Local variation of natural hazards can also vary considerably: relative variations in temperature over a small domain are typically much less than those of rainfall. European windstorms, however, can affect large areas and have high local variability; adequate simulations of these must represent a large domain at high resolution. A robust simulation model for natural hazard events must allow various different combinations of domain size, resolution and variability amounts on different spatial scales.

To meet these simulation criteria, we propose a framework that couples extreme value and geostatistical methods. Works by Casson and Coles (1999), Cooley et al. (2007) and Sang and Gelfand (2009) consider a similar coupling, although here we focus on a geostatistical model for residuals—an approach sometimes referred to as *anamorphosis*. Our approach allows excesses corresponding to exceedances of a high threshold to be simulated, and provides an alternative benchmark for catastrophe models for the following reasons: geostatistical models can give highly efficient simulations of high-resolution random fields, compared to fully multivariate models; various forms for dependence over space and time exist for geostatistical models, which have well represented various types of environmental phenomena (see, for example, Diggle and Ribeiro (2007)); and statistically sound marginal estimates for extremes are used. Furthermore, we can quickly implement the method on a personal computer, thus requiring less computational time and resource than other types of hazard module, such as those built similarly to climate models.

The following section describes a spatial, extreme value framework for simulating natural hazard events. Section 3 presents a model for simulating extreme windstorm events for a large part of Europe together with some simulations from the model. Section 4 provides a summary of the framework presented and of its performance for simulating windstorm events.

2. Method

This section gives details of the framework proposed to give realistic simulations of natural hazard events. Throughout let $Y(s,t)$ represent values of a natural hazard process at location $s \in S$ and time $t = 1, \ldots, T$. The simulation model generates excesses above a high threshold,

[1]EU directive 2009/138/EC, available from http://ec.europa.eu/finance/insurance/solvency/solvency2/index_en.htm

which are assumed to follow a generalised Pareto distribution (GPD). The threshold can either be chosen or estimated, but must be sufficiently high that the GPD assumption is valid. For loss estimation it is useful if the threshold is below any value of the natural hazard process above which damage can occur; for example, around 25ms^{-1} for wind gusts. The simulation model must also represent spatiotemporal dependence realistically and include an estimate of the threshold exceedance rate to allow simulations to represent specific time periods.

(a) Marginal threshold excess model

We assume that excesses of some spatially-varying threshold $u(s)$, for location s, follow the generalised Pareto distribution (GPD), so that

$$Y(s,t) - u(s) \,|\, Y(s,t) > u(s) \sim GPD(\psi_u(s), \xi(s)), \qquad (2.1)$$

where $\psi_u(s) > 0$ and $\xi(s)$ are scale and shape parameters, respectively, for $\{y : 1 + \xi(s)y/\psi_u(s) > 0\}$; the scale parameter may vary with $u(s)$. We consider generalised additive model (GAM, Chavez-Demoulin and Davison (2005); Hastie and Tibshirani (1990)) forms for GPD parameters, summarised as

$$\log \psi_u(s) = \psi_u + f_0(lon(s), lat(s)) + \sum_{p=1,\ldots,q} f_p(z_p(s)) + \sum_{p<p'} f_{pp'}(z_p(s), z_{p'}(s)) \qquad (2.2)$$

where $lon(s)$ and $lat(s)$ represent longitude and latitude, and $z_p(s)$, $p = 1, \ldots, q$, represent covariates, at location s. The log link in equation (2.2) ensures $\psi_u(s) > 0$ for all $s \in S$; a transformation for $\xi(s)$ may not be necessary. Parametric or nonparametric forms can be considered for $f_0()$, $f_p()$ and $f_{pp'}(,)$. We focus on GAM forms and propose specific types that are suitable for simulating many different hazard types. These are found to be more flexible than parametric forms, such as those described in Coles (2001, Ch. 6), and simpler to fit than non-parametric forms, such as Casson and Coles (1999) and Cooley et al. (2007). We represent $f_0(,)$ as a thin plate regression spline (Wood, 2003), as longitude and latitude have the same units; for this reason we distinguish longitude and latitude from other covariates. For other covariates, where interactions may occur and scales and units may differ, we propose additive and tensor-product forms (De Boor, 1978).

GPD models with parameters with spline forms can be fitted in various ways. Chavez-Demoulin and Davison (2005) propose an approach based on maximum likelihood that combines Newton-Raphson and generalised ridge regression steps. This can be used to simultaneously estimate both regression and smoothing parameters. Their orthogonal parameterisation guarantees convergence. For interpretability, we use spline forms for $\log \psi_u(s)$ and $\xi(s)$, which are not orthogonal. We estimate smoothing parameters using generalised (approximate) cross-validation (Wahba, 1990), where at each iteration we find GPD parameters by numerically maximising the penalised likelihood (Green, 1987; Pauli and Coles, 2001).

(b) Threshold estimation

Here we consider models for the spatially-varying threshold, $u(s)$. A constant threshold may not suit large domains, especially if marginal distributions exhibit large variation. We use quantile regression to estimate $u(s)$, which is assumed to have GAM form. Put briefly, $u(s)$ is chosen to minimise the so-called tilted loss function, subject to some additional roughness penalty terms, which depend on the spline choices; see Koenker (2005, Ch. 7) for fuller details. This approach to modelling and estimating $u(s)$, coupled with a GPD model for excesses of $u(s)$, was proposed by Northrop and Jonathan (2011). We extend the additive spatial GAM forms for longitude and latitude used in Northrop and Jonathan (2011) to thin plate regression splines forms. Link functions may also be considered to relate $u(s)$ to GAM terms.

For simulations to represent given time periods, the rate of upcrossing of the threshold must be taken into account, which we denote $\zeta(s) = \Pr(Y(s,t) > u(s))$. If the quantile used in quantile

regression is fixed at ζ, then $\zeta(s) = \zeta$ for all $s \in S$, $t = 1, \ldots, t$. Alternatively the threshold can be fixed and logistic regression used to estimate $\zeta(s)$, using a similar additive form to equation (2.2); see Wood (2006) for inference details in this case.

(c) Residual dependence model

By considering marginal and dependence models separately, the model can be related to copula-based approaches. Disadvantages to this separation, identified in Mikosch (2006), are potential bias in stochastic dependence estimates, a lack of supporting statistical theory and potentially poor representation of multivariate extremes or, by virtue of being static, temporal dependence. Here the advantages, given in §1, of fast, high-resolution simulations, robust dependence forms, and extreme value margins are seen to outweigh these disadvantages. To efficiently achieve high-resolution simulations, we restrict attention to residual models based on spatially-continuous stochastic processes, which are widely used in geostatistics.

Gaussian processes are commonly used in geostatistics, and may be used for anamorphosis. They are robust and can give efficient simulations on high-resolution grids using circulant embedding methods (Wood and Chan, 1994), or for irregularly-spread locations. The Gaussian process model imposes asymptotic independence between different locations s and s' at time t: if y^+ is the upper endpoint of the distribution of $Y(s,t)$ and $\chi = \lim_{y \to y_+} \Pr\{Y(s,t) > y \,|\, Y(s',t) > y\}$, then $\chi = 0$ (Sibuya, 1959); when $\chi > 0$ asymptotic dependence occurs. This assumption may be inappropriate for some natural hazards; see Coles et al. (1999) for environmental examples. The Student-t process can be seen as a generalisation. It imposes asymptotic dependence when its degrees of freedom, ν, are finite, and asymptotic independence when infinite, as the latter can be parameterised as a Gaussian process. We treat ν as an unknown parameter.

We consider a Student-t process tail model: a Student-t and continuous space extension of the multivariate Gaussian tail model of Bortot et al. (2000). The marginal model of §2(a) can be used to give uniformly distributed residuals on $[1 - \zeta(s), 1]$,

$$e^*(s,t) = 1 - \zeta(s) \left[1 + \xi(s) \frac{Y(s,t) - u(s)}{\psi_u(s)} \right]^{-1/\xi(s)} \qquad \text{if } Y(s,t) > u(s), \qquad (2.3)$$

which may be transformed to the final residuals $e(s,t) = G_\nu^{-1}(e^*(s,t))$, where G_ν^{-1} is the inverse Student-t distribution function with ν degrees of freedom. Residuals are modelled with a tail Student-t process, which is written

$$e(s,t) \sim \textit{tail-tProcess}_\nu \Big(0, \, c\big((s,t), \big) \Big), \qquad (2.4)$$

where $c((s,t),)$ is a spatiotemporal correlation function and ν is the degrees of freedom of the process. A suitable form for the correlation function will depend on the application, although inclusion of nugget term is likely to be necessary to allow local measurement error. A form for $c((s,t),)$ suitable for extreme windstorm events is presented in §3. Parameters of the tail Student-t process can be estimated by maximising the censored log-likelihood of the corresponding finite-dimensional multivariate tail Student-t distribution.

It is natural to consider max-stable processes (Schlather, 2002; Smith, 1990) as alternative residual models. These fit readily into the proposed framework by converting the uniform residuals of equation (2.3) to have unit Frechét distribution, ie. $e(s,t) = -1/\log e^*(s,t)$. Options for max-stable models include Schlather, extremal t and Brown-Resnick processes (Davis and Resnick, 1984; Opitz, 2013; Schlather, 2002). We focus on anamorphosis-based models for their interpretability, relative to widely-used geostatistical methods, and simulation efficiency. Related findings are presented in §3.(b).iii and §3.(d).iii.

(d) Event simulation

Values of a natural hazard process are simulated for a fixed number of locations $\mathcal{S} = \{s_1, \ldots, s_{S^*}\}$, which may be regularly or irregularly spread. Let $\mathcal{T}$ be the time domain of the event. For example, if daily values are simulated for a 3-day event centred at time t^*, $\mathcal{T} = \{t^* - 1, t^*, t^* + 1\}$. The following algorithm details simulation of a single event, and assumes that the threshold, marginal excess and residual dependence models have been fitted.

   (i) Simulate $e(s, t)$ from $tail\text{-}tProcess_\nu\big(0,\, c\big((s, t),\, \big)\big)$ for locations $s \in \mathcal{S}$, times $t \in \mathcal{T}$.
   (ii) Obtain uniform residuals $e^*(s, t) = G_\nu(e(s, t))$.
   (iii) Set

$$y(s, t) = \left\{ \begin{array}{ll} u(s) + \frac{\psi_u(s)}{\xi(s)}[(1 - e^*(s, t))^{-\xi(s)} - 1] & \text{if } e^*(s, t) > 1 - \zeta(s), \\ \text{`censored at } u(s)\text{`} & \text{otherwise}, \end{array} \right.$$

   where $\zeta(s) = \Pr(Y(s, t) > u(s))$; recall §2(b).

Repeat steps 1–3 to give the required number of simulated events. Extensions for the case where simulated values of non-exceedances of the threshold $u(s)$ are also required are proposed in §4.

3. Simulation of extreme windstorm events

This section illustrates the framework by simulating windstorm events for a large part of Europe.

(a) Data

The windstorm data are maximum daily surface wind gust speeds (in ms⁻¹), measured by anemometer, and extracted from the National Climatic Data Center global summary of the day database[2]. These are supplemented with MIDAS land surface station observations (Met Office, 2012), to improve coverage for Norway and the UK, and ECA&D data (Klein Tank et al., 2002) to improve coverage for Spain. The period 1st January 1994 to 31st December 2014 is studied. Only winter (defined as December, January, February) storms are modelled. Locations of data stations used are shown on Figure 1 together with a plot representing their elevations.

   Figure 2 shows plots of wind gust speeds for seven stations. These are chosen as they are the nearest two to London and Paris, and the nearest to Madrid, the Ruhr and Milan (which are the five most populated metropolitan areas within the domain studied, in descending order[3]). Whether the gusts speeds are asymptotically dependent or independent is assessed. These probabilities are estimated empirically and uncertainty quantified by bootstrap re-sampling the data. Two stations are chosen near to London and Paris to help assess asymptotic dependence for these proximate and highly populated station pairs; see also §3.(b).iii. Between the London and Paris stations, dependence tends to be higher across the range of gust speeds, compared with between stations from different areas. However, the rarity of concomitant exceedances for the highest gust speeds, such as above 25ms⁻¹, makes it inconclusive from Figure 2 whether the gust speeds are asymptotically dependent or independent. Further analysis is given in §3.(b).iii by considering extremal index estimates.

(b) Model specification and estimates

This section gives details of the model specifications used to estimate and then simulate extreme wind gust speeds.

---

[2] https://data.noaa.gov/dataset/global-surface-summary-of-the-day-gsod
[3] Source https://en.wikipedia.org/wiki/List_of_metropolitan_areas_in_Europe; accessed March 28, 2016

Station locations

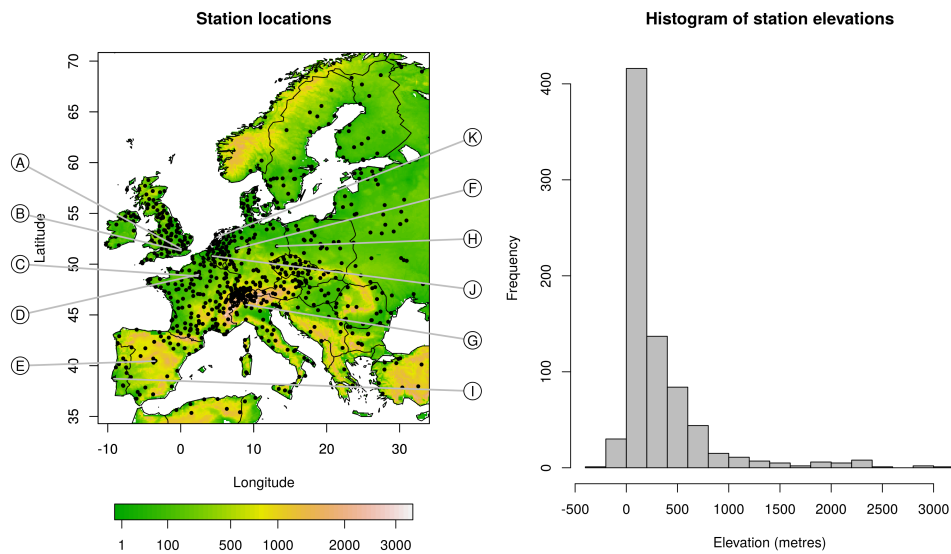**Histogram of station elevations**

Figure 1. Left: station locations overlaid on a map of elevation. Seven sites used for initial data analysis (Figure 2) are identified as A–G. Nine sites used for model checking (Figure 6) are identified as A, C, E–K. Right: histogram of station elevations.

(i) Threshold estimation

We use a spatially-varying threshold. A constant threshold did not suit the large domain, primarily because it was not possible to find a single threshold for which the GPD assumption was valid for all stations.

The high proportion of missing data, which is typical of European wind gust measurements, presents a significant obstacle to estimating the threshold. Estimating the threshold at a pre-specified quantile using only available measurements leads to bias. This is because of a tendency towards the start of the study period for stations to only record high wind gust measurements; ignoring this, and assuming that measurements are missing at random, leads to severe overestimation of high quantiles. We overcome this with an infilling procedure based on a regression model: its response is Box-Cox transformed wind gust measurements using $\lambda = 0.2$ (Box and Cox, 1964); its covariates are taken from temporally complete ERA-Interim reanalysis data (Dee et al., 2011), such as gust and wind speeds and mean sea-level pressure. A multivariate Gaussian model is then fitted to the model's residuals. The infilled data are obtained by simulating residuals for the missing data, conditional on residuals derived from available data. Residuals and predictions from the ERA-Interim regression model are combined, and the infilled gust speeds obtained by inverting the Box-Cox transformation. These infilled values are only used to estimate the threshold, and not the excess model.

As described in §2(b), and proposed by Northrop and Jonathan (2011), the threshold is estimated by nonparametric quantile regression. The 98th percentile is estimated for both practical and theoretical reasons. Klawa and Ulbrich (2003) found this to perform well for estimating financial loss from European windstorms over Germany, while it is sufficiently high that tests show the GPD assumption for threshold excesses to be valid, as supported by figures 6 and 7. We find estimation of both the threshold and excesses models unreliable at higher thresholds, due to too few data, and that GPD estimates incur systematic spatial bias for thresholds below the 95th percentile. Empirical estimates of the 98th percentile for each
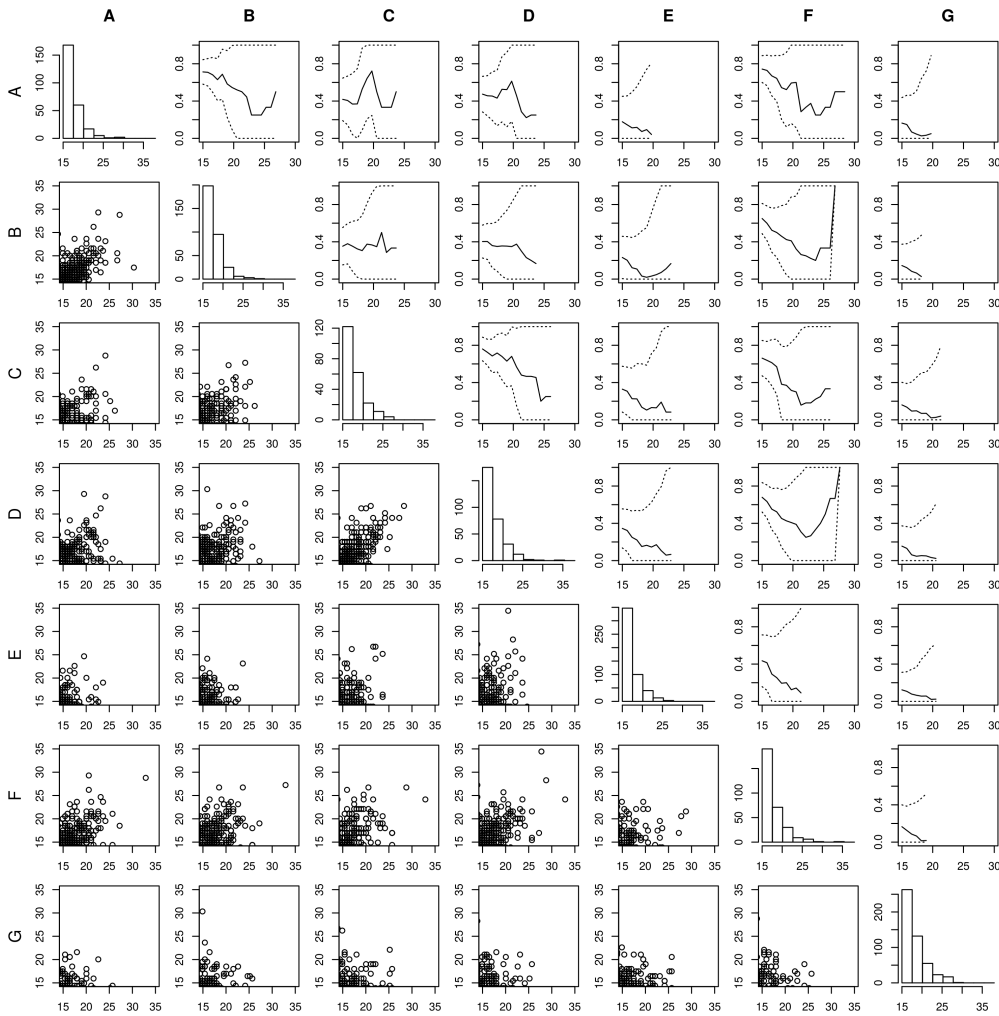
Figure 2. Plots of wind gust speeds exceeding 15ms^{-1} for stations A–G. Diagonal plots show histograms of wind gust speeds. Lower off-diagonal plots show daily wind gust speeds at two stations. Upper off-diagonal plots show empirical estimates of $\Pr\{Y(s,t) > y \mid Y(s',t) > y\}$ for station pairs; see §2(c).

station show variation with elevation and mean winter wind speed[4]. Consequently the quantile regression model for $u(s)$ takes the form

$$u(s) = f_0(lon(s), lat(s)) + f_{12}(elev(s), mws(s)), \qquad (3.1)$$

where $f_0(\,,\,)$ is a thin plate regression spline, $elev(s)$ and $mws(s)$ are elevation and mean winter wind speed at location s, respectively, and $f_{12}(\,,\,)$ is formed by a tensor-product of P-splines (which negates the need for $f_p(\,)$ terms).

The estimate of the 98th percentile, shown in Figure 3, ranges from approximately 20ms^{-1}, which includes densely populated cities, such as Paris and London, to over 35ms^{-1} for the Alps and Pyrenees (although there are very few stations at such altitudes). Figure 3 shows that elevation influences the threshold estimate most; mean winter wind speed has a small, but much lesser, effect. The thin plate spline captures that windstorms tend to track north of the UK and

---

[4]Mean wind speed data obtained from https://crudata.uea.ac.uk/cru/data/hrg/; see New et al. (2002).
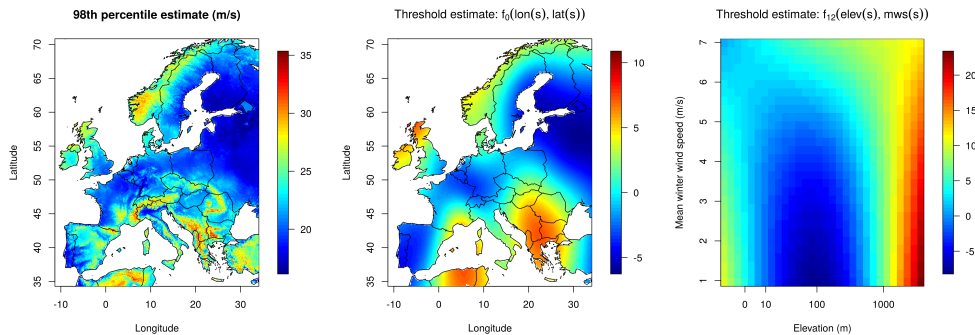
Figure 3. Estimate of 98th percentile of the distribution of wind gust speeds (ms^{-1}) from nonparametric quantile regression (left) and contributions to estimate of thin-plate regression spline (centre, $f_0(lon(s), lat(s))$) and tensor product spline of elevation and mean winter wind speed (right, $f_{12}(elev(s), mws(s))$ are shown.

Norway. Relative to the rest of the region studied, wind gust data for the Balkans are scarce; their increased estimate is accompanied by large uncertainty.

(ii) Excess model

We then model excesses of the estimated 98th percentile as GPD. The following GAM forms are used:

$$\log \psi_u(s) = f_0(lon(s), lat(s)) + f_1(elev(s)) + f_2(mws(s)),$$

$$\xi(s) = f_0(lon(s), lat(s)) + f_1(elev(s)),$$

where $f_1(\,)$ and $f_2(\,)$ are P-splines. Estimates of $\log \psi_u(s)$ and $\xi(s)$ are shown in Figure 4, with contributions from each of the GAM components shown in Figure 5. Elevation appears to have greatest effect. In general, increased elevation coincides with higher values of $\log \psi_u(s)$ and lower values of $\xi(s)$, although for $\log \psi_u(s)$ this relationship is not entirely monotonic increasing. Increased mean winter wind speed also coincides with higher $\log \psi_u(s)$ values, but to a much lesser extent than elevation. The thin plate spline contributions show that $\log \psi_u(s)$ decreases smoothly from the north-west to the south-east of the region studied; a similar—but lesser—effect can be seen for $\xi(s)$, although its increased values around the Balkan states are most apparent. Uncertainty in the $\xi(s)$ estimate around the Balkan states is, however, again large, due to few data.

Quantile plots for assessing the GPD fits are shown in Figure 6. These are achieved by omitting the data for the nine validation sites (A, C, E–K) to give estimates of a reduced model, which is used to predict sites' GPD parameters. Data for these sites are included when estimating the final model. For the nine sites considered, which coincide with the nine most populated metropolitan areas, and are highlighted in Figure 1, adequate agreement between model-based and empirical estimates of the distribution of excesses of $u(s)$ can be seen. Slight signs of higher model-based estimates than observed gust speeds are apparent. We suspect that careful quality control on measurements could offer slight improvement, as, due to a scarcity of data, we have taken a conservative approach to quality control. We have also only considered fairly simple covariate choices. Other similarly simple covariates, such as mean winter temperature, were also considered. Of those considered, we have chosen the best fitting model. However, allowing less simple covariates may improve the model's fit, in particular covariates derived from reanalysis data.

We also collectively assess the fits for the stations using the Kolmogorov-Smirnov (K-S) test. Ties in the data make the K-S test in its standard form unreliable. Instead we use a Monte
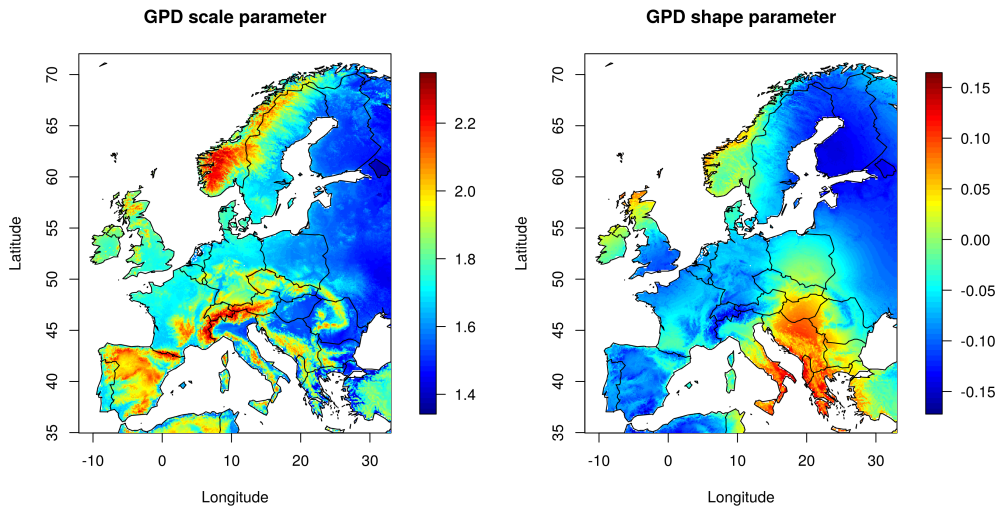
Figure 4. GPD model parameter estimates. Left: GPD scale parameter estimate (log scale, $\log(\mathrm{ms}^{-1})$). Right: GPD shape parameter.

Carlo technique to obtain p-values, which involves simulating sets of excesses from estimated GPDs, rounding the simulated excesses, and then computing the K-S test statistic for each set of simulated data. The p-values, shown for each station in Figure 7, result from comparing the K-S test statistic for the gust speed data with the test statistics from the simulated data. For 40 out of 789 stations ($\simeq 5.1\%$), the p-value is below 5%, which indicates adequate fit of the GPD across all stations. Importantly, there are no signs of systematic spatial deviations, such as large parts of the study region where the GPD fit is inadequate.

(iii) Residual model

We build a residual model based on an anisotropic correlation function. This is formed by considering a transformed space with coordinates $\tilde{s} = ARs$ where

$$A = \begin{pmatrix} 1/\phi_1 & 0 \\ 0 & 1/\phi_2 \end{pmatrix}, \qquad R = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix}$$

and $s = (s_1, s_2)^T$. This corresponds to scaling by $1/\phi_1$ and $1/\phi_2$ in longitude and latitude directions, and a counter-clockwise rotation through angle θ. We allow anisotropy as windstorms follow tracks, which typically causes dependence between gust speeds to be greatest along the track. Angle θ allows the prevailing direction of the tracks to deviate from perfectly following the longitude or latitude axes. The correlation function includes a nugget, $1 - \tau$, to allow measurement error, and is written

$$c((s,t),(s',t')) = c(d) = (1-\tau)\delta(d) + \tau\rho(d), \tag{3.2}$$

where $\delta(\ )$ is the Kronecker delta function and $d = ||\tilde{s} - \tilde{s}'||$ for locations $\tilde{s}$ and $\tilde{s}'$ in transformed space. We choose the Whittle form for $\rho(\ )$, so that

$$\rho(d) = 2^{1-\kappa}\{\Gamma(\nu)\}^{-1} d^\kappa K_\kappa(d), \tag{3.3}$$

for $\kappa > 0$ and K_κ the modified Bessel function of the second kind. The Whittle form is chosen for compatibility with the storm model of Cox and Isham (1988), which follows in §3.(c).i; otherwise Matérn or powered exponential forms might be considered. We are required to estimate τ, ϕ_1, ϕ_2, θ and κ.
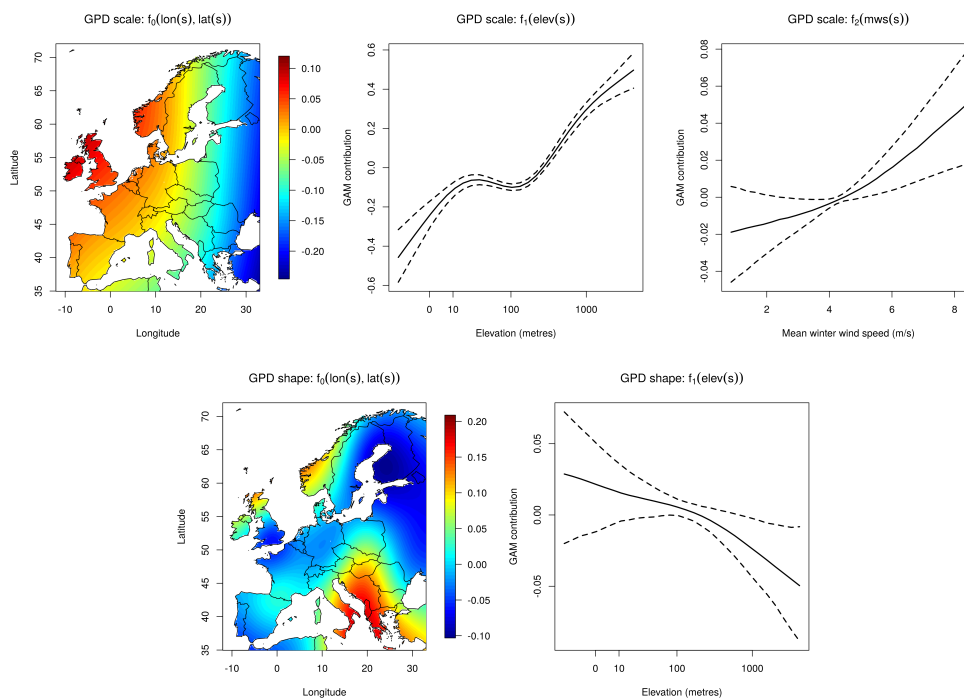
Figure 5. Row 1: additive contributions of $f_0(lon(s), lat(s))$ (left), $f_1(elev(s))$ (centre) and $f_2(mws(s))$ (right) to (log) GPD scale parameter estimate. Row 2: contributions of $f_0(lon(s), lat(s))$ (left) and $f_1(elev(s))$ (right) to GPD shape parameter estimate.

Figure 7 shows pairwise estimates of the extremal index for the nine validation sites of Figure 6 based on the censored method of Schlather and Tawn (2003), with 95% uncertainty bounds estimated by profiling the likelihood. The maximum likelihood estimates generally equal 2 (and uncertainty bounds all include 2), which indicates asymptotic independence. This is confirmed when we fit the tail Student-t process, under the assumption of no temporal dependence, and find that its likelihood is maximised for large ν (ie. $\nu > 10,000$); at this value the difference between Student-t and Gaussian processes is negligible. Consequently we focus on Gaussian processes. The estimated correlation structure of the Gaussian process is represented in Figure 8 by a semi-variogram, which compares empirical with model-based estimates of the dependence structure against distance defined on transformed space, $d_s = ||\tilde{s} - \tilde{s}'||$; the structure itself is also shown with distance defined on the original scale. Figure 8 shows a decay in dependence. The practical range, where correlation falls below 0.05, is approximately 13 degrees. The mean estimates approximately show that, once rotated counter-clockwise by 7 radians, dependence extends 2.8 times further along the longitude than the latitude axis. Figure 7 shows extremal index estimates for the nine validation sites based on simulations from the Gaussian process model. These mimic the availability of the wind gust data so that missing values occur simultaneously. Once uncertainty is taken into account, we see that simulations from the model are consistent with the extremal index estimates from the wind gust data, which supports the adequacy of the estimated dependence structure.
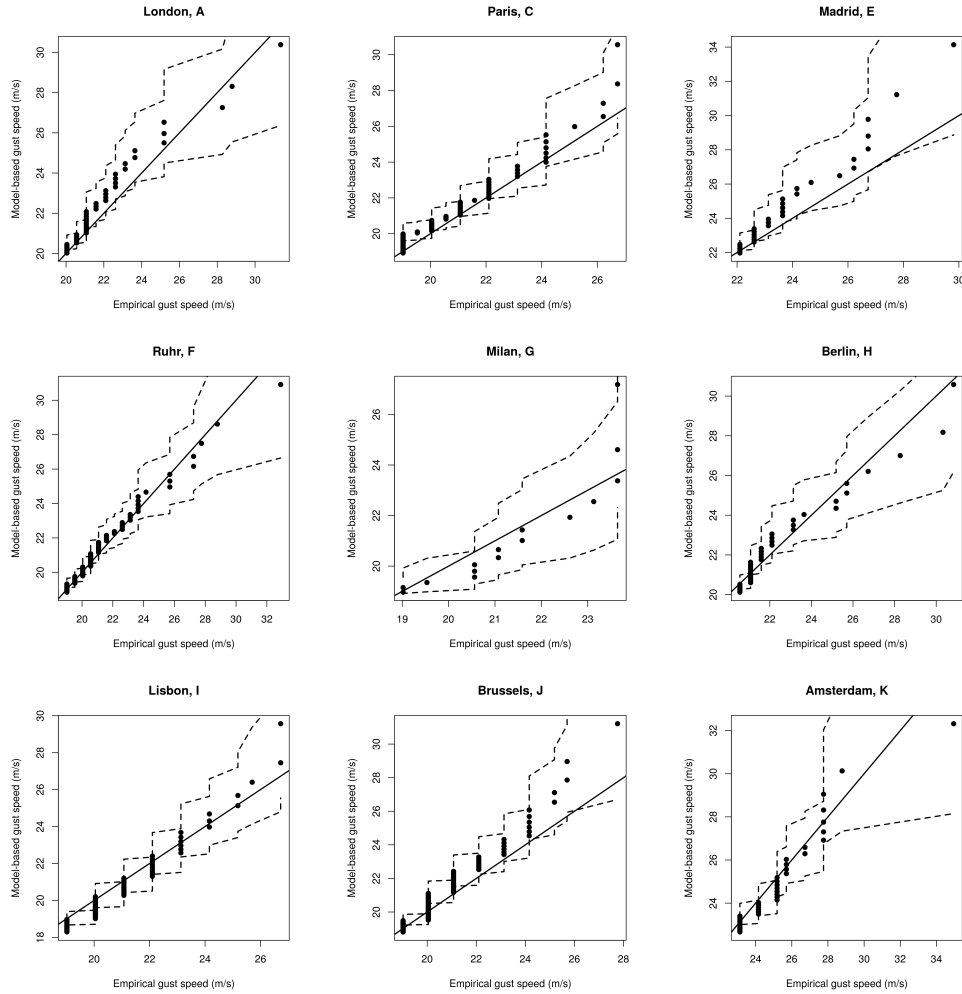
(c) Simulated windstorm events

Figure 6. GPD quantile-quantile plots for nine stations identified as A, C, E–K in Figure 1.

(i) Simulation specification

We base windstorm simulations on a spatiotemporal extension of the Whittle correlation function, developed as a storm model in Cox and Isham (1988) and later presented for Gaussian processes in Gneiting et al. (2006) and Schlather (2010). Each event, $j = 1, \ldots, J$, is assumed to move at a random bivariate normally distributed velocity $V_j \sim BVN(\mu_j, M/2)$ and has corresponding correlation function

$$c_j((s,t),(s',t')) = \frac{1}{|I_2 + d_t^2\, \widetilde{M}\,|^{1/2}} \rho\big(||(d_s - d_t\tilde{\mu}_j)^T (I_2 + d_t^2\, \widetilde{M})^{-1}(d_s - d_t\tilde{\mu}_j)||\big), \qquad (3.4)$$

for $s \neq s'$, $t \neq t'$, where $c_j((s,t),(s',t')) = 1$ otherwise, $d_t = t - t'$, $\tilde{\mu}_j = AR\mu_j$ and $\widetilde{M} = ARMR^T A^T$. The large number of windstorms, J, results in too many μ_js to reliably obtain maximum likelihood estimates of M and the μ_js alongside τ, ϕ_1, ϕ_2, θ and κ. A simpler model with $\mu_j = \mu$ for all j fails to capture the variation that is empirically evident. Instead the μ_js and M are empirically estimated from storm tracks extracted from the ERA-Interim reanalysis (Dee et al., 2011) using the tracking algorithm of Hodges (1995, 1999). Comparison against the windstorm track observations of the IBTrACS database (Knapp et al. (2010); not shown)
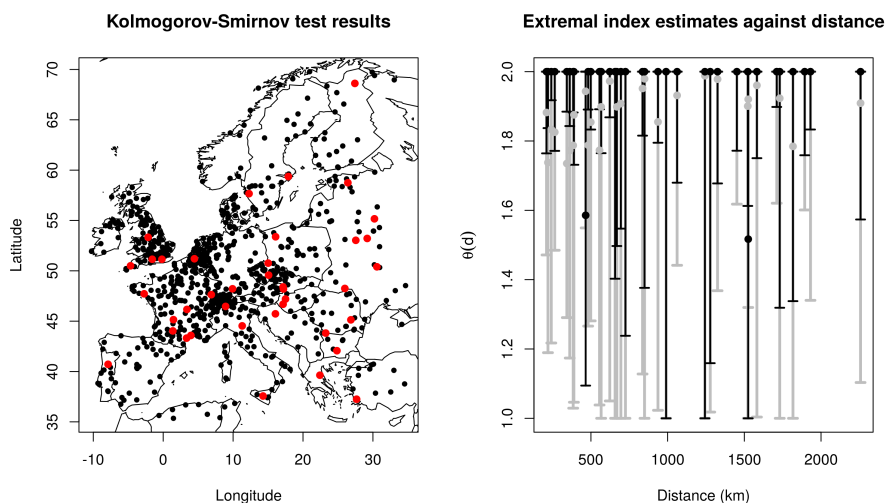
Kolmogorov-Smirnov test results

Extremal index estimates against distance

Figure 7. Left: stations with Kolmogorov-Smirnov test p-value greater than 5% (black) and less than 5% (red). Right: pairwise, censored extremal index estimates for nine stations of Figure 6 based on wind gust speed data (black) and simulations from Gaussian process model (grey). Estimates ($\bullet$, $\circ$) are shown alongside 95% uncertainty estimates based on profile likelihood (——) and variability from multiple simulations (——).

Semi-variogram for transformed distance

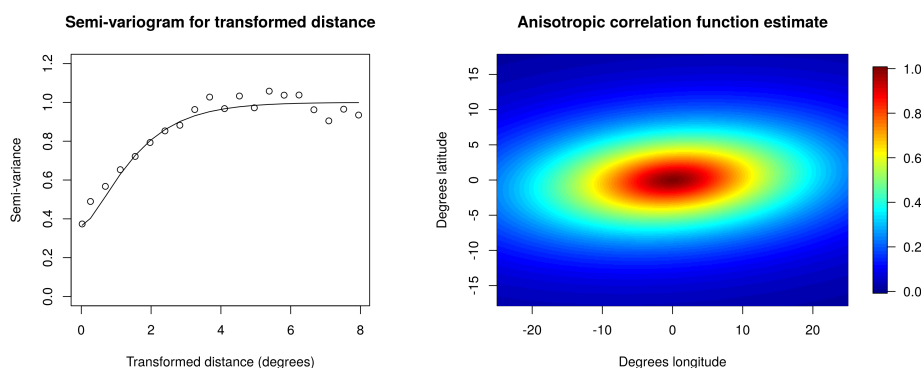**Anisotropic correlation function estimate**

Figure 8. Left: semi-variogram against distance, defined in degrees for transformed coordinates, with Whittle model superimposed. Right: estimated anisotropic correlation function, ie. $\rho(\ )$ from equation (3.3).

shows that, although the intensities of windstorms may tend to be underestimated, velocities are represented sufficiently well for the purposes of this simulation.

(ii) Results on simulated events

Events are simulated for locations s_k on a 300×300 grid $S = \cup_{k=1}^{n_S} s_k$, where $n_S = 90,000$. We simulate events equivalent to 10,000 Dec–Feb winters and analyse 3-day events because a 72-hour period is typically used in the insurance industry to define individual events. Only excesses corresponding to exceedances of $u(s)$ are simulated. Under this specification, each simulated winter takes approximately 20 seconds, but simulations are entirely paralellisable. Figure 9 shows two randomly-chosen 3-day event simulations. For both events, exceedances of the estimated 98th
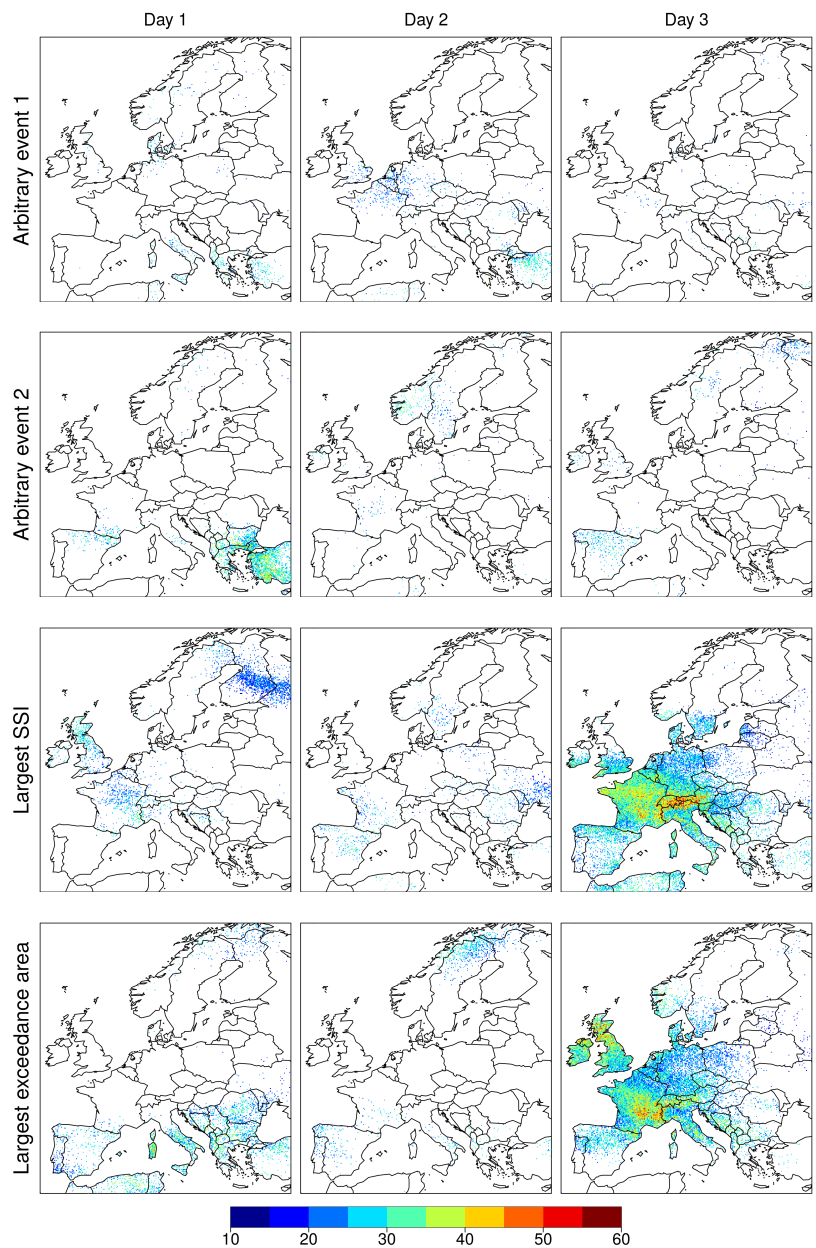
Figure 9. Plots of two randomly-chosen 3-day windstorm event simulations (rows 1 and 2) and extreme events, based on largest SSI (row 3) and exceedance area (row 4) over three days.

percentile can be seen on each of the three simulated days. Very few damage-causing gust speeds exceeding 25ms^{-1} occur. Therefore neither event is likely to be classed as extreme in terms of the financial loss that it might have caused.

To highlight potentially catastrophic events, two measures are used to quantify whether a windstorm event is extreme (Roberts et al., 2014). Each event, $j = 1, \ldots, J$, is defined for locations $s \in S$ and a time domain $\mathcal{T}_j = \{t_j - 1, t_j, t_j + 1\}$, where t_j gives the peak time of event j. Plots of the events with the highest values of each of these measures are shown in Figure 9.

Klawa and Ulbrich (2003) propose a *storm severity index* (SSI) given by

$$SSI(\mathcal{T}_j) := \sum_{s_k \in S} \sum_{t \in \mathcal{T}_j} I[s_k \in S_{SSI}, Y(s_k, t) > u(s_k)] \left(\frac{Y(s_k, t)}{u(s_k)} - 1 \right)^3,$$

which is designed to quantify an event's kinetic energy, which relates closely to financial loss, where $S_{SSI} \subset S$ is a region for which the SSI will be calculated. Compared to S, S_{SSI} excludes the Balkan states, where shape parameter estimates are relatively imprecise; corresponding simulated gusts would otherwise dominate SSI values. Windstorms with large financial losses seldom occur in the excluded regions. Figure 9 shows the 3-day event with highest SSI. We see that many threshold exceedances occur on day three in various countries, in particular France, Switzerland, Germany, Italy, Belgium, and Austria, but also, to a lesser extent, Netherlands, UK, Denmark, and Slovenia. The renowned windstorms Martin, Lothar and Kyrill affected a similar set of countries.

We define the *exceedance area* of a storm as

$$A(\mathcal{T}_j) := \sum_{s_k \in S} \sum_{t \in \mathcal{T}_j} I[Y(s_k, t) > u(s_k)].$$

This measure captures an event's spatial extent over three days; that is, the proportion of Europe affected by high wind gust speeds, relative to the local 98th percentile. Figure 9 also shows the 3-day event with the largest exceedance area. The event with largest exceedance area is very similar to that with largest SSI, as it has greatest effect on day three and affects a similar set of countries.

(d) Alternative residual models

One of the benefits of the separation of the simulation model into marginal and residual models is the ability to combine alternative models for each.

(i) The effect of measurement error

The simulations in Figure 9 could prove awkward for loss estimation if the signal-to-noise ratio (the ratio of τ to $1 - \tau$ in equation (3.2)) is seen to be small, ie. if the gust speeds have relatively large measurement error. Figure 10 shows single-day windstorm events simulated with different signal-to-noise ratios, based on the GPD estimates described in §3.(b).ii. Part of the motivation for this analysis is that increasing the signal-to-noise ratio might lead to simulated events representative of unobserved actual windstorm events, as opposed to representing measurements. The effect of increasing the signal-to-noise ratio is clear, with the spatial structure of events becoming smoother, eventually forming larger, more distinct patches affecting fewer parts of Europe.

(ii) A reanalysis-based residual model

As a further sensitivity analysis, we swap the measurement-based residual model for one based on ERA-Interim reanalysis wind gust output, which are exempt from measurement error. We estimate the ERA-Interim 3-day dependence structure using an empirical probability transformation to convert margins to Gaussian, interpolate the resulting data onto the simulation grid S and then calculate the empirical correlation matrix. We then use this to simulate Gaussian residuals, which are then converted to the original wind gust scale using the GPD estimates of §3.(b).ii. Figure 11 shows simulated events chosen according to the same criteria as Figure 9. Increased smoothness of the events is clear to see. We also note that the events with largest SSI and exceedance area coincide. As the ERA-Interim data are aggregated onto a 0.75° grid, their residual structure will inevitably be smoother than in reality, which will lead to larger-than-realistic loss estimates; residuals based on the measurements will suffer the opposite. These alternative residual models could therefore be used to place upper and lower bounds on loss estimates. Based on the preceding two sensitivity analyses for the residual model, we should
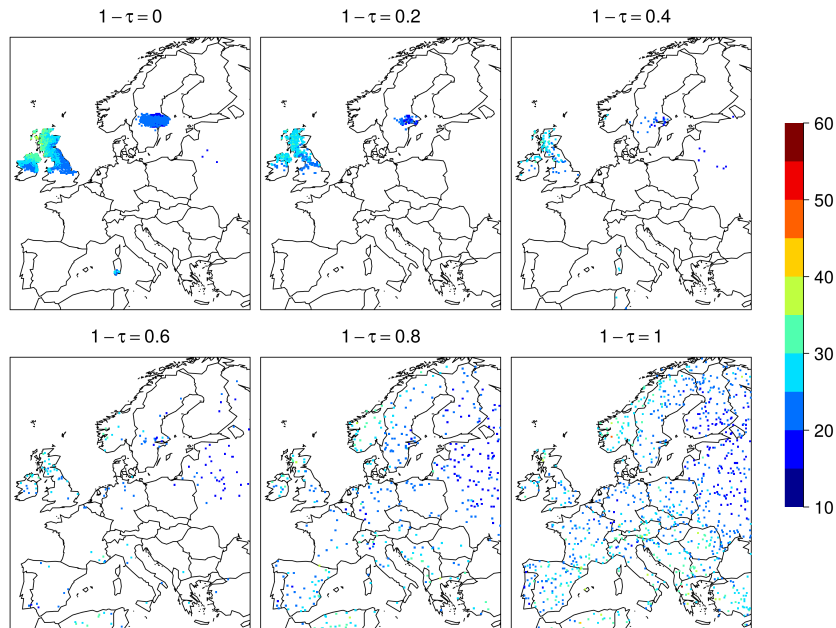
Figure 10. Simulated wind gusts for single days based on differing measurement error amounts, ie. $1 - \tau$ in equation (3.2). Simulations use the same seed to aid comparison between different signal-to-noise ratios.

hope to reduce the uncertainty in loss estimates with higher-resolution wind gust output instead of ERA-Interim and/or more precise measurements.

(iii) A max-stable process residual model

To fit a max-stable process residual model we obtain unit Frechét residuals as described in §2(c), ie. $e(s,t) = -1/\log e^*(s,t)$ with $e^*(s,t)$ defined in equation (2.3). We fit the max-stable models by maximising pairwise censored likelihoods (Huser and Davison, 2014; Ledford and Tawn, 1996; Lindsay, 1988; Padoan et al., 2010) and hence present the two-dimensional case. For two locations $\tilde{s}$ and $\tilde{s}'$, defined on the transformed space of §3.(b).iii, we assume that

$$\Pr(e(\tilde{s},t) < e_1, e(\tilde{s}',t') \leq e_2) = \exp\{-V(e_1, e_2)\},$$

where $V(,)$ is the exponent measure function. As asymptotic independence is indicated by Figure 7, we consider the extension of the Schlather model (Schlather, 2002) proposed in Davison and Gholamrezaee (2012) where

$$V(e_1, e_2) = \left(\frac{1}{e_1} + \frac{1}{e_2}\right)\left[1 - \frac{\alpha(d)}{2}\left(1 - \sqrt{1 - 2\frac{\{c(d)+1\}e_1 e_2}{(e_1 + e_2)^2}}\right)\right],$$

with $c( )$ as in equation (3.2), where $d = ||\tilde{s} - \tilde{s}'||$. We take $\alpha(d) = \{1 - |d|/(2r)\}_+$, which gives asymptotic independence when $|d| > 2r$, and find that $r \to \infty$. The resulting estimate $c( )$ has a relatively large nugget estimate ($\tau \simeq 0.63$); its $\rho( )$ component is shown on the original coordinate scale in Figure 12, which shows a very short range of dependence. Extremal index estimates for the max-stable model with distance defined on the transformed space, which are shown in Figure 12, highlight a lack of agreement between the max-stable model and the wind gust data. Simulations from the model (not shown), which take on average 90 times longer than those for Gaussian process model, have qualitatively similar characteristics to the $\tau = 1$ case presented in Figure 10. We have, however, only considered a fairly simple form for $\alpha( )$. More flexible forms
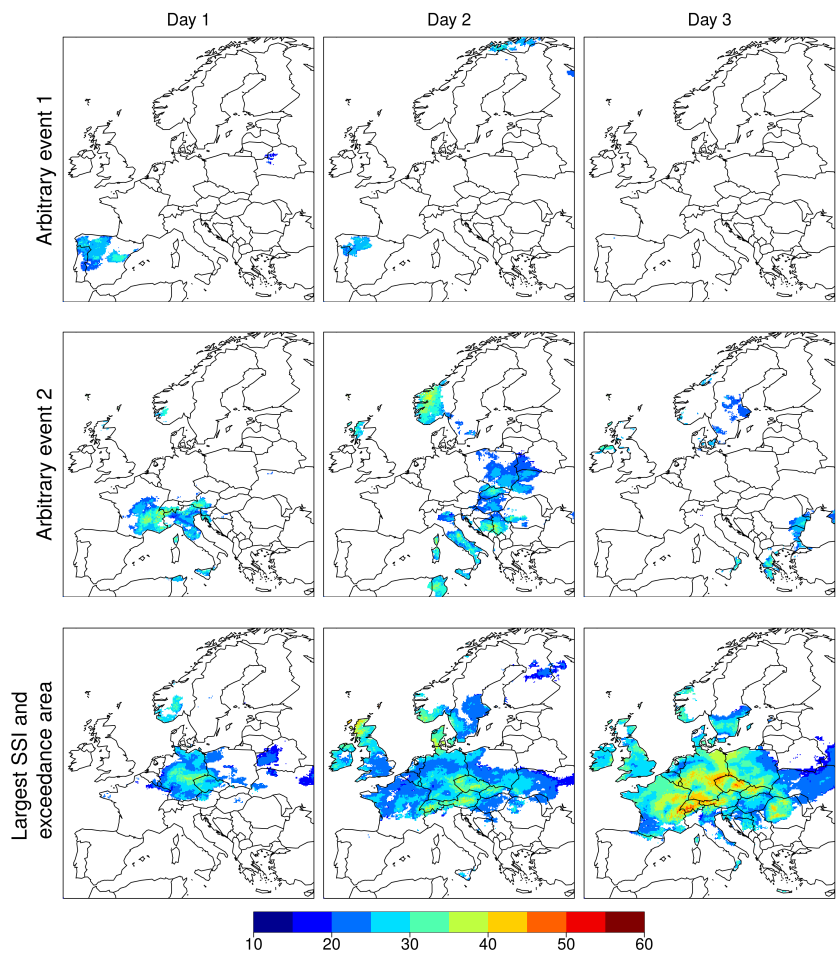
Figure 11. Events simulated using ERA-Interim residual model. Two randomly-chosen 3-day windstorm event simulations (rows 1 and 2) and an extreme event (row 3) based on largest SSI or exceedance area (both happen to coincide) over three days.

are presented in Davison and Gholamrezaee (2012, §2(e)) and Huser and Davison (2014) use a further extension when modelling extreme hourly rainfall.

4. Summary

We have presented a statistical framework for simulating natural hazard events that combines extreme value theory, to accurately capture event magnitude, and geostatistics, to robustly incorporate spatial dependence. The framework can be used to quickly give high-resolution simulations that can be formally validated. The framework has been used to produce realistic simulations of European windstorm events. By virtue of its generality, speed and that it can be implemented using a standard computer, the framework could be readily used as a hazard module in a catastrophe model for other types of hazard, or as a tool for assessing other modules. Various aspects of the proposed framework that could be changed or extended to bring potential improvements are discussed in the remainder of this section.

The Student-t process residual model includes the Gaussian process as a special case. It can be fitted quickly using maximum likelihood, based on finite-dimensional counterparts of
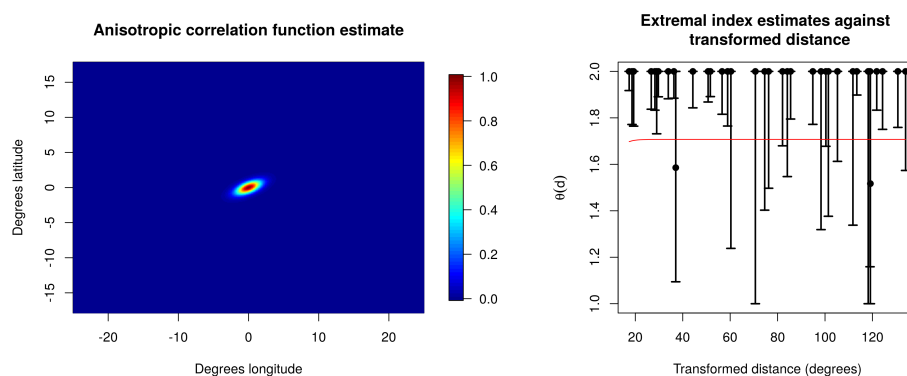
Figure 12. Left: plot of correlation function estimate for asymptotically independent parameterisation of Schlather model. Right: extremal index estimates from model (red) against empirical estimates (black).

the processes, and subsequently gives fast simulations. The finite-dimensional representation—as opposed to continuous-space—could better suit certain types of hazard phenomena. One example, as in studied Keef et al. (2009), is river flows, where complex dependencies between flow gauges may benefit from a finite specification. Its finite-dimensional representation loses certain computational benefits, such as those gained by circulant embedding, but remains relatively efficient. Student-t and Gaussian processes lack max-stability. Although in §3.(d).iii we have looked into using max-stable processes to simulate extreme windstorm events, we note that they may better represent other types of natural hazard, or other temporal resolutions. For example, max-stable processes have been used in Huser and Davison (2014) to model hourly rainfall, and for modelling annual maximum snow depths (Blanchet and Davison, 2011) and temperatures (Davison and Gholamrezaee, 2012). Estimates of the extremal index, as in Figure 7, may help reveal which, if any, residual model is most suitable. If the model is used as a hazard module in a catastrophe model, so that loss estimates are produced, it is important that, in addition to formal statistical validation of the residual model, validation against an appropriate loss function for the application is also performed. The anamorphosis and max-stable models can also be extended to capture dependencies between multiple processes; although a trade-off between resolution and increased dimensionality may be needed to maintain computational feasibility.

We have used thin plate spline forms to capture spatial smoothness in marginal distribution parameters, which appear to capture the expected parameter variation reliably, upon appropriate choice of basis dimension and cross-validation of smoothing parameters. We see spline forms as a compromise between parametric forms, which can be inflexible, and stochastic forms, such as Gaussian processes, for which inference might require computer-intensive Markov chain Monte Carlo (MCMC) techniques. We have seen little to justify this extra computational demand and also think it could hinder use of the model by end-users. The GAM forms can be extended to incorporate temporal variability, such as a model for the entire year by modelling each season separately or by allowing temporal variation in parameter estimates. For extreme gust speeds it is not valid to assume a constant distribution across the year, which is likely to be true for most hazard phenomena. Unless a fairly homogeneous period can be identified when events tend to occur, such as heatwaves in summer, a temporal model should be used. Similarly, temporal variability may benefit corresponding residual models.

For some hazard types, such as modelling flooding due to extreme rainfall, simulated values of non-exceedances of the threshold $u(s)$ may be required. To model marginal distributions in this case, a GPD tail model can be coupled with a model for non-exceedances, which can be empirical (Keef et al., 2009), nonparametric (Carreau et al., 2009), or parametric (Frigessi et al., 2002). To give spatial simulations, empirical estimates of spatial dependence can be used together with a

suitable interpolation scheme. A geostatistical model can be used; although a different spatial dependence structure between exceedances and non-exceedances of $u(s)$ should be considered for sufficient flexibility. Alternatively, flooding can be studied by modelling extreme river flows (Keef et al., 2009), for which the proposed framework is better suited.

We have allowed for measurement error by including a nugget term in correlation functions. This is computationally convenient, but may lack interpretability when compared to the model $Y(s,t) \sim N\big(W(s,t), \omega^2\big)$, where $W(s,t)$ is an actual—but unobservable—value of the hazard process at location s and time t, as it may be possible to specify ω^2 based on knowledge of the measurement process. The $W(s,t)$s can then be modelled as the $Y(s,t)$s have been previously. As it is not possible to integrate out the $W(s,t)$s analytically within the present framework, numerical procedures, such as MCMC, will be required for inference. Integration would also need to carry over the threshold excess and residual models, negating much of the simplicity offered by GAM forms; this might make stochastic processes preferable for capturing spatial variation in parameters. This model can also be formulated as a hierarchical model and, if informative prior knowledge is available, implemented from a Bayesian perspective.

One aspect of the model where scope exists to add generality is the correlation function, which we have assumed is the same across the study region. This assumption can be relaxed; for example, by allowing parameters of the correlation function to vary in space. We considered replacing ϕ_1 with $\phi_1(s)$ and using simple forms such as $\phi_1(s) = \exp\{\phi_{1,0} + \phi_{1,1} lon(s) + \phi_{1,2} lat(s)\}$. Although such model extensions could not be reliably estimated here, they may benefit other applications. Correlation can also be defined on an alternative (potentially higher-dimensional) space, such as climate space used in Cooley et al. (2007). When modelling gust speeds over larger domains, it should be noted that storms have highs and lows, which suggests an oscillating correlation function might be more suitable, as illustrated in Lindgren et al. (2011). Windstorms are also known to behave differently over land from sea; therefore partitioning distance into over-land and over-sea components might better represent dependencies. For hazard types for which events tend to affect smaller areas, the proposed model may be built based on a smaller region, which might improve the validity of using spatially-constant correlation functions.

References

Blanchet, J. and A. C. Davison (2011). Spatial modeling of extreme snow depth. *The Annals of Applied Statistics*, 1699–1725.

Bonazzi, A., S. Cusack, C. Mitas, and S. Jewson (2012). The spatial structure of European wind storms as characterized by bivariate extreme-value copulas. *Natural Hazards and Earth System Science 12*(5), 1769–1782.

Bortot, P., S. Coles, and J. Tawn (2000). The multivariate Gaussian tail model: an application to oceanographic data. *Journal of the Royal Statistical Society: Series C (Applied Statistics) 49*(1), 31–049.

Box, G. E. P. and D. R. Cox (1964). An analysis of transformations. *Journal of the Royal Statistical Society. Series B (Methodological) 26*(2), 211–252.

Brodin, E. and H. Rootzén (2009). Univariate and bivariate GPD methods for predicting extreme wind storm losses. *Insurance: Mathematics and Economics 44*(3), 345–356.

Carreau, J., P. Naveau, and E. Sauquet (2009). A statistical rainfall-runoff mixture model with heavy-tailed components. *Water Resources Research 45*(10).

Casson, E. and S. Coles (1999). Spatial regression models for extremes. *Extremes 1*(4), 449–468.

Chavez-Demoulin, V. and A. C. Davison (2005). Generalized additive modelling of sample extremes. *Journal of the Royal Statistical Society: Series C (Applied Statistics) 54*(1), 207–222.

Coles, S. (2001). *An Introduction to Statistical Modeling of Extreme Values*. Springer Series in Statistics. Springer London.

Coles, S., J. Heffernan, and J. Tawn (1999). Dependence measures for extreme value analyses. *Extremes 2*(4), 339–365.

Cooley, D., D. Nychka, and P. Naveau (2007). Bayesian spatial modeling of extreme precipitation return levels. *Journal of the American Statistical Association 102*(479), 824–840.

Cox, D. R. and V. Isham (1988). A simple spatial-temporal model of rainfall. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences 415*(1849), 317–328.

Davis, R. and S. Resnick (1984). Tail estimates motivated by extreme value theory. *Ann. Statist. 12*(4), 1467–1487.

Davison, A. C. and M. M. Gholamrezaee (2012). Geostatistics of extremes. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, Volume 468, pp. 581–608. The Royal Society.

De Boor, C. (1978). *A Practical Guide to Splines*. Number v. 27 in Applied Mathematical Sciences. Springer-Verlag.

Dee, D. P., S. M. Uppala, A. J. Simmons, P. Berrisford, P. Poli, S. Kobayashi, U. Andrae, M. A. Balmaseda, G. Balsamo, P. Bauer, P. Bechtold, A. C. M. Beljaars, L. van de Berg, J. Bidlot, N. Bormann, C. Delsol, R. Dragani, M. Fuentes, A. J. Geer, L. Haimberger, S. B. Healy, H. Hersbach, E. V. Hólm, L. Isaksen, P. Kållberg, M. Köhler, M. Matricardi, A. P. McNally, B. M. Monge-Sanz, J.-J. Morcrette, B.-K. Park, C. Peubey, P. de Rosnay, C. Tavolato, J.-N. Thépaut, and F. Vitart (2011). The ERA-Interim reanalysis: configuration and performance of the data assimilation system. *Quarterly Journal of the Royal Meteorological Society 137*(656), 553–597.

Diggle, P. and P. Ribeiro (2007). *Model-based Geostatistics*. Springer series in statistics. Springer.

Frigessi, A., O. Haug, and H. Rue (2002). A dynamic mixture model for unsupervised tail estimation without threshold selection. *Extremes 5*(3), 219–235.

Gneiting, T., M. Genton, and P. Guttorp (2006). *Statistical Methods for Spatio-Temporal Systems*, Chapter Geostatistical space-time models, stationarity, separability and full symmetry, pp. 151–175. CRC Press.

Green, P. J. (1987). Penalized likelihood for general semi-parametric regression models. *International Statistical Review/Revue Internationale de Statistique 55*, 245–259.

Grossi, P. and H. Kunreuther (2006). *Catastrophe Modeling: A New Approach to Managing Risk*. Huebner International Series on Risk, Insurance and Economic Security. Springer Science+Business Media.

Haas, R., J. G. Pinto, and K. Born (2014). Can dynamically downscaled windstorm footprints be improved by observations through a probabilistic approach? *Journal of Geophysical Research: Atmospheres 119*(2), 713–725.

Hastie, T. J. and R. J. Tibshirani (1990). *Generalized additive models*, Volume 43. CRC Press.

Hodges, K. I. (1995). Feature tracking on the unit-sphere. *Monthly Weather Review 123*(12), 3458–3465.

Hodges, K. I. (1999). Adaptive constraints for feature tracking. *Monthly Weather Review 127*(6), 1362–1373.

Huser, R. and A. C. Davison (2014). Space-time modelling of extreme events. *Journal of the Royal Statistical Society: Series B (Statistical Methodology) 76*(2), 439–461.

Keef, C., J. Tawn, and C. Svensson (2009). Spatial risk assessment for extreme river flows. *Journal of the Royal Statistical Society: Series C (Applied Statistics) 58*(5), 601–618.

Klawa, M. and U. Ulbrich (2003). A model for the estimation of storm losses and the identification of severe winter storms in Germany. *Natural Hazards and Earth System Science 3*(6), 725–732.

Klein Tank, A., J. Wijngaard, G. Können, R. Böhm, G. Demarée, A. Gocheva, M. Mileta, S. Pashiardis, L. Hejkrlik, C. Kern-Hansen, et al. (2002). Daily dataset of 20th-century surface air temperature and precipitation series for the European climate assessment. *International Journal of Climatology 22*(12), 1441–1453. Data from `http://www.ecad.eu`.

Knapp, K. R., M. C. Kruk, D. H. Levinson, H. J. Diamond, and C. J. Neumann (2010). The international best track archive for climate stewardship (IBTrACS) unifying tropical cyclone data. *Bulletin of the American Meteorological Society 91*(3), 363–376.

Koenker, R. (2005). *Quantile Regression*. Econometric Society Monographs. Cambridge University Press.

Ledford, A. W. and J. A. Tawn (1996). Statistics for near independence in multivariate extreme values. *Biometrika 83*(1), 169–187.

Lindgren, F., H. Rue, and J. Lindström (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology) 73*(4), 423–498.

Lindsay, B. (1988). *Statistical Inference from Stochastic Processes (Ithaca, NY, 1987). Contemp. Math.*, Chapter Composite likelihood methods, pp. 221–239. Amer. Math. Soc.

Met Office (2012). Met Office integrated data archive system (MIDAS) land and marine surface stations data (1853-current). NCAS British Atmospheric Data Centre, 4 February 2015., `http://catalogue.ceda.ac.uk/uuid/220a65615218d5c9cc9e4785a3234bd0`.

Mikosch, T. (2006). Copulas: Tales and facts. *Extremes 9*(1), 3–20.

New, M., D. Lister, M. Hulme, and I. Makin (2002). A high-resolution data set of surface climate over global land areas. *Climate research 21*(1), 1–25.

Northrop, P. J. and P. Jonathan (2011). Threshold modelling of spatially dependent non-stationary extremes with application to hurricane-induced wave heights. *Environmetrics 22*(7), 799–809.

Opitz, T. (2013). Extremal t processes: Elliptical domain of attraction and a spectral representation. *Journal of Multivariate Analysis 122*, 409–413.

Padoan, S. A., M. Ribatet, and S. A. Sisson (2010). Likelihood-based inference for max-stable processes. *Journal of the American Statistical Association 105*(489), 263–277.

Pauli, F. and S. Coles (2001). Penalized likelihood inference in extreme value analyses. *Journal of Applied Statistics 28*(5), 547–560.

Roberts, J. F., A. J. Champion, L. C. Dawkins, K. I. Hodges, L. C. Shaffrey, D. B. Stephenson, M. A. Stringer, H. E. Thornton, and B. D. Youngman (2014). The XWS open access catalogue of extreme European windstorms from 1979 to 2012. *Natural Hazards and Earth System Science 14*(9), 2487–2501.

Sang, H. and A. Gelfand (2009). Hierarchical modeling for extreme values observed over space and time. *Environmental and Ecological Statistics 16*(3), 407–426.

Schlather, M. (2002). Models for stationary max-stable random fields. *Extremes 5*(1), 33–44.

Schlather, M. (2010). Some covariance models based on normal scale mixtures. *Bernoulli 16*(3), 780–797.

Schlather, M. and J. A. Tawn (2003). A dependence measure for multivariate and spatial extreme values: Properties and inference. *Biometrika 90*(1), 139–156.

Sibuya, M. (1959). Bivariate extreme statistics, I. *Annals of the Institute of Statistical Mathematics 11*(2), 195–210.

Smith, R. L. (1990). Max-stable processes and spatial extremes. *Unpublished manuscript*.

Wahba, G. (1990). *Spline Models for Observational Data*. CBMS-NSF Regional Conference Series in Applied Mathematics. Society for Industrial and Applied Mathematics (SIAM, 3600 Market Street, Floor 6, Philadelphia, PA 19104).

Wood, A. T. and G. Chan (1994). Simulation of stationary Gaussian processes in $[0, 1]^d$. *Journal of computational and graphical statistics 3*(4), 409–432.

Wood, S. N. (2003). Thin-plate regression splines. *Journal of the Royal Statistical Society (B) 65*(1), 95–114.

Wood, S. N. (2006). *Generalized Additive Models: An Introduction with R*. Chapman & Hall/CRC Texts in Statistical Science. Taylor & Francis.