Nat. Hazards Earth Syst. Sci. Discuss.,
https://doi.org/10.5194/nhess-2018-102-AC1, 2018

Natural Hazards and Earth System Sciences
Discussions
Open Access

Interactive comment on "On the relevance of extremal dependence for spatial statistical modelling of natural hazards" *by* Laura C. Dawkins and David B. Stephenson

Laura C. Dawkins and David B. Stephenson

l.c.dawkins@exeter.ac.uk

Received and published: 10 August 2018

Key:

- **Reviewer's comment**
- *Our response*
- Additional/edited text in the manuscript

1 Reviewer 1

1.1 Specific comments

- **Comparison of Gumbel, Gaussian and power law copulas to the empirical estimate is shown only for a very few pairs of sites. At the end of Section 3.3 it is suggested that the results found at these sites are representative of results for other pairs of locations. How many other pairs of locations were tested? How confident are you that asymptotic independence is the dominant dependence structure across all pairs of sites? Have you considered ways in which you could formally test this over all pairs of sites?**

  This is a very good point, and important to demonstrate. We have addressed this within the paper at the end of Section 3.3 by adding a few paragraphs and some additional analysis:

  When aiming to develop a statistical model for high dimensional spatial data over a large geographical domain, it is essential to systematically explore the dominant extremal dependence class across all locations. Here, we present an approach for doing so, which uses this quick-to-calculate coefficient of tail dependence diagnostic, demonstrated by application to our windstorm footprint data set. We first take a stratified (based on the distribution of locations over longitude and latitude) sample of 100 locations within the European domain. One such sample is shown in [Figure 1 in attachments](a). Since the extremal dependence is likely to decrease with increasing separation distance (Wadsworth and Tawn (2012)) and we hope to understand if asymptotic independence is dominant and hence present at small separation distances, for each of these 100 locations, we estimate the coefficient of tail dependence, η (and the associated

Printer-friendly version

Discussion paper

95% profile likelihood confidence interval) when paired with the 100 nearest locations within the full domain. [Figure 1 in attachments](b) demonstrates how the 100 nearest locations are geographically distributed for one such sampled location in our windstorm footprint dataset. For each pairing, the coefficient of tail dependence is calculated using w as the 0.9 quantile threshold of the structure variable, found to ensure stable estimates of η using diagnostic plots as in [Figure 1 in attachments] (c). The estimated η parameters and confidence intervals for these 100$\times$100 pairs of locations are plotted against separation distance to explore how, throughout the domain, η varies at small separation distances and changes with increasing separation distance, shown in [Figure 1 in attachments] (d). The parameter estimate related to the pair of locations in pink and blue in [Figure 1 in attachments] (b) is shown in pink. This method is repeated many times with 10 such repetitions shown in [Figure 1 in attachments] of the Supplementary Material at the end of the paper, showing very similar results.

[Figure 1 in attachments here] - (a) A stratified (based on the distribution of locations over longitude and latitude) sample of locations within the European domain, with stratified grid shown in grey; (b) a demonstration of the 100 nearest locations [turquoise] to one of these sampled locations [blue], with one such point selected at random [pink]; (c) the coefficient of tail dependence diagnostic plot (as in Fig. 4) for wind gusts at the blue location paired with the pink location; (d) the coefficient of tail dependence (estimated using w as the 0.9 quantile threshold of the structure variable) and 95% profile likelihood confidence intervals, for each of the 100 sampled locations paired with their 100 nearest locations in the full domain, plotted against separation distance in kilometres, with the estimate based on the pair of locations in (b) and (c) added in pink.

[Figure 1 in attachments] (d) shows that for small separation distances (<180 km) a proportion of pairs of locations have coefficients of tail dependence parameter, η, estimates close to 1, with $\eta = 1$ within the confidence interval, indicating asymptotic dependence. Within the range (0-50 km) 69$\%$ of pairs of locations exhibit this behaviour, however this proportion reduces rapidly as separation distance increases, to 30$\%$ for locations separated by (50-100 km), 13$\%$ for locations separated by (100-150 km) and 3$\%$ for locations separated by (150-200 km). Hence, while there is evidence of asymptotic dependence for some locations in close proximity, even at very small separation distances (50 km) a larger proportion of locations exhibit asymptotic independence. Indeed, here and in [Figure 2 in attachments] of the Supplementary Material, beyond a separation distance of approximately 200km the coefficients of tail dependence parameter estimates drop well below 1, indicating asymptotic independence. Therefore, since separation distances within the domain extend to up to 3500km, we conclude that asymptotic independence is the dominant extremal dependence structure across the spatial domain.

It is important to consider the validity of representing even this small proportion of asymptotically dependent pairs of locations incorrectly as asymptotically independent. To explore this, Bortot et al. (2000) carried out a simulation study in which they fit the Gaussian, Ledford and Tawn (1996) and Gumbel models to bivariate data simulated from three parent populations with different classes of extremal dependence. They conclude that, for asymptotically independent parent populations the Gaussian copula is able to provide accurate inferences for tail probability estimates, out performing the Gumbel copula model, and even for asymptotically dependent parent populations, the estimation error of the Gaussian copula model was deemed to be acceptably small. This suggests that, when data dimensionality prohibits the use of flexible extremal dependence models, such as Huser and Wadsworth (2018), and asymptotic independence is found to be the dominant extremal dependence structure across the spatial domain, using an asymptotically independent model, such as the Gaussian tail model, is preferable over using a model for asymptotic dependence throughout the domain. In Section 4 we present a further, natural hazards relevant, diagnostic approach for further validating this, based on an estimates of the aggregate natural hazard losses.

[Figure 2 in attachments here] - For 10 stratified samples of 100 locations within the European domain: the coefficient of tail dependence (estimated using w as the 0.9 quantile threshold of the structure variable) and 95$\%$ profile likelihood confidence intervals, for each of the 100 sampled locations paired with their 100 nearest locations in the full domain, plotted against separation distance in kilometres.

- **A thought on the justification for asymptotic independence based on the model for wind gust speeds. It is assumed in the physical model that each of the Cartesian components of wind gusts speeds $U(s)$ and $V(s)$ follows an independent Gaussian spatial process. By properties of the multivariate normal distribution, each the vector of each component $U = (U(s1), U(S2))$ and $V = (V(s1), V(S2))$ at any two sites $s1$ and $s2$ follows a bivariate normal distribution and consequently the components of each of U and V are asymptotically independent. However the Gaussian process assumption is just another modelling assumption, so the question is how accurate is it, ie. how strong is the evidence in favour of the Gaussian assumption? If actually the speed components followed some other process that had**

asymptotically dependence bivariate margins then the conclusion would be very different.

The reviewer raises a good point. This brief justification has been added to the article in the third paragraph of Section 3.4:

The distribution of each velocity component is expected by the Central Limit Theorem to be close to normally distributed since the net displacement of a particle in turbulence is the result of many irregular smaller displacements. The distribution of each component has zero skewness due to the symmetry of the fluid equations (negative deviations are as likely as positive ones) but can have slightly more kurtosis (i.e. fatter tails) than the normal distribution due to intermittency in the flow. Measurements of velocity components in the atmospheric surface layer reveal that the distributions are near to Gaussian (e.g. Chu et al. (1996)).

- **On a related point, it is not entirely obvious where the equation for χ_{max} and the expression for Pr$(u_1^2 > t, u_2^2 > t)$ comes from. Although these expressions are correct, would it be possible to put a derivation in the appendix?**

  It's reassuring that the expressions are correct and the sensible suggestion of adding a short appendix has been adopted:

  Assume the velocity components (u_1, v_1) and (u_2, v_2) at two separate locations in an isotropic turbulent flow can be represented as bivariate normally distributed vectors (u_1, u_2) and (v_1, v_2) that are independent and identically distributed with zero expectations.

The individual velocity components, (u_1, u_2) and (v_1, v_2), are both asymptotically independent because of each being bivariate normally distributed.

The squares of the individual velocity components, e.g. (u_1^2, u_2^2), are also asymptotically independent. This is proven by rewriting the joint probability of exceedance:

$$\begin{aligned} &\Pr(u_1^2 > t^2, u_2^2 > t^2) \\ &= \Pr(u_1 > t, u_2 > t) + \Pr(u_1 > t, u_2 \le t) + \Pr(u_1 \le t, u_2 > t) + \Pr(u_1 \le t, u_2 \le t) \\ &= \chi_{++}\Pr(u_1 > t) + \chi_{-+}\Pr(u_1 > t) + \chi_{-+}\Pr(u_1 \le t) + \chi_{--}\Pr(u_1 \le t) \\ &= \chi_{++}\Pr(u_1^2 > t^2) + \chi_{+-}\Pr(u_1^2 > t^2), \end{aligned}$$

which is obtained by noting that $\Pr(u_1^2 > t^2) = \Pr(u_1 > t) + \Pr(u_1 \le t)$, and conditional probabilities $\chi_{++} = \chi_{--}$ and $\chi_{+-} = \chi_{-+}$ by symmetry of the bivariate normal distribution about $(0, 0)$. Since the components are bivariate normal, χ_{++} and $\chi_{+-} \to 0$ as $t \to \infty$, and so $\Pr(u_1^2 > t^2, u_2^2 > t^2)/\Pr(u_1^2 > t^2) \to 0$. Hence, the square of the velocity component is also asymptotically independent.

Perhaps rather counter-intuitively, the sum of two independent identically distributed asymptotically independent variables is not necessarily asymptotically independent. It, therefore, remains to be proven whether or not $(u_1^2 + v_1^2, u_2^2 + v_2^2)$ is asymptotically independent.

- **The second half of Section 4 (p21 onwards) needs some re-working to clarify the points that are being made. For example, it is not clear why we might be interested in the first and second moments derived after line 326; also these are not all moments but expectation (first moment), conditional expectation (conditional first moment) and variance (function of second and first moment).**

This part of Section 4 has now been rewritten more clearly present our motivation and correctly refer to these quantities. In combination with your comment below - ***does your model enable you to look at the sizes of losses at the two sites rather than just the probability of a loss jointly occurring at each site?****, we have included a more detailed review of alternative windstorm loss functions, in which the size of the loss is represented as a function of the wind (rather than being equal to 1 as it is in our loss function). We have changed the definition of the conceptual loss function to a more generic form which quantifies the size of the losses as well as the exceedance of the loss threshold.*

Within this section we initially present the probability mass function of this generic bivariate conceptual loss function, demonstrating how the success of a given model in representing the bivariate conceptual loss for the pair (X, Y) *closely relates the its characterisation of* $\chi(p)$*, and hence present a comparison of empirical and modelled* $\chi(0.01)$ *and* $\bar{\chi}(0.01)$ *for the 2 opposing dependence models (Gaussian/Gumbel).*

We then derive the conditional expected loss for the generic loss function which is a function of $\chi(p)$ *as well as the the size of the loss at each location, therefore motivating the use of this conditional first moment for comparing how well the Gaussian and Gumbel models represent the size of the joint losses, rather than just their conditional probability of occurrence.*

After paragraph 1 in Section 4 we have edited:

Similar to other natural hazard loss models, in the absence of confidential insurance industry exposure and vulnerability information, it has become common in the literature to define conceptual windstorm loss as a function of the footprint wind gust speeds (see Dawkins et al. (2016) for a review). While these conceptual windstorm loss functions vary in the detail of their composition, it is possible to express most in a general form, for the pair (X, Y), as:

L(X,Y) = g[V(X)e(X)H{X-U(X)} + V(Y)e(Y)H{Y-U(Y)}] where V is a function the wind gust speeds characterising the magnitude of the hazard, e represents exposure (e.g. population density), U quantifies a high threshold of the wind gust speed above which losses occur, H is a Heaviside function such that $H\{m\} = 1$ if $m > 0$ and $H\{m\} = 0$ otherwise, and g is an additional function applied in some cases to reduce skewness. For example, in the widely used and rigorously validated conceptual loss function of Klawa and Ulbrich (2003), $V(X) = (X - x_{0.98})^3$, $U(X) = x_{0.98}$ (where $x_{0.98}$ is the 98% quantile of X) and $e(X)$ is represented by the population density at the location (with equivalent expressions for Y), while Cusack (2013) used a loss function in which $V(X) = (X - x_{0.99})^3$, $U(X) = x_{0.99}$, the 99% quantile of X, and $g[\cdot] = \sqrt[3]{\cdot}$. See Table 2.1 in Dawkins (2016) for a summary of previously published conceptual loss functions in terms of the components of Eqn. (1).

Then, after explaining our choice of loss function and loss threshold, we introduce the generic loss function.

Since, for a given storm event, $V(X)e(X)$ and $V(Y)e(Y)$ in Eqn. (1) are constants, this equation can be simplified to:

L(X,Y) $\propto C_X H\{X - U(X)\} + C_Y H\{Y - U(Y)\}$ where $C_X = V(X)e(X)$ and $C_Y = V(Y)e(Y)$. In our case $C_X = C_Y = 1$, and $U(X) = x_{0.99}$, $U(Y) = y_{0.99}$, the 99% quantiles of X and Y respectively. Therefore, while in this study we use just one conceptual loss function in which the magnitude of the loss is always equal to 1, due to its identified suitability for representing loss in this data set (Roberts et al. (2014)), it is simple to adapt the following analysis to accommodate alternative loss functions in which the size of the loss is included as a function of the excesses of the natural hazard, by incorporating a model for the marginal distribution of hazard at each location. This would be an interesting area of future exploration within this windstorm footprint application, beyond the scope of this analysis.

The probability mass function of the bivariate conceptual loss function can easily be obtained in terms of the Extremal Dependence Coefficient, $\chi(p)$, by considering the joint probability of (X, Y) in each of the quadrants shown in Fig. 2:

$\Pr$(L(X,Y) = C$_X + C_Y) = \chi(p)p$,
$\Pr(L(X,Y) = C_X) = \Pr(L(X,Y) = C_Y) = 2(1-\chi(p))p$,
$\Pr(L(X,Y) = 0) = 1 + p(\chi(p) - 2)$, highlighting how the success of a given model in representing the bivariate conceptual loss for the pair (X, Y) closely relates the how well it characterises $\chi(p)$, and hence the extremal dependence between X and Y.

We then present the results in Figure 7, and finally introduce the conditional expected loss:

As well as being relevant for representing the probability mass function of the bivariate conceptual loss function, $\chi(p)$ can also be shown to characterise the conditional expectation of joint loss:

E(L(X,Y)) = (C$_X + C_Y)\chi(p)p + C_X(1-\chi(p))p + C_Y(1-\chi(p))p = (C_X + C_Y)p,$ and,
$\mathbb{E}(L(X,Y)|L(X) = C_X) = (C_X + C_Y)\chi(p)p + C_X(1-\chi(p))p = p(C_Y\chi(p) + C_X).$
This conditional first moment of the loss distribution can therefore be used to compare how well the opposing dependence models represent the size of the joint losses, rather than just their conditional probability of occurrence, since the expression includes C_X and C_Y. Here, $C_X = C_Y = 1$, hence the conditional expectation of joint loss is equivalent to the conditional expectation of loss jointly occurring at both locations given a loss has occurred at one location.

Followed finally by Fig. 8

- **Why does the expected loss not depend on the extremal dependence between the two sites (line 327)?**

  We agree this need clarification. We have added this explanation to the end of the final paragraph in the previous response:

  It should be noted that the (non-conditional) expected loss, $\mathbb{E}(L(X,Y))$, does not depend on $\chi(p)$. This is because the expectation of a sum is the sum of the expectations, hence expected total loss over two or more locations is simply the sum of the expected losses at each location, and so is unaffected by the amount of dependency between sites.

- **Could you clarify what is being illustrated (line 331) in this final part of the section?**

  This sentence has been removed in the rewriting of Section 4.

- **Although the Gumbel/Gaussian discrepancy is clear in Figure 7, it might be informative to also look at spatial plots of the differences between the empirical and model-based estimates of each of chi(u) and chibar(u) to see if there is any spatial clustering in these differences, ie. do the models repre-**

sent the empirical behaviour better for some regions/distances/directions than others?

We have now change Figure 7 to show the difference between the empirical and modelled estimates of χ and $\bar{\chi}$ and changed the interpretation of the plot accordingly.

[Figure 3 in attachments here] - The difference between empirical and modelled $\chi(0.01)$ for (a) the Gaussian model and (b) the Gumbel model, and the difference between empirical and modelled $\bar{\chi}(0.01))$ for (c) the Gaussian model and (d) the Gumbel model, for London paired with all other locations over land.

[Figure 3 in attachments] demonstrates how the Gaussian model is able to represent empirical $\chi(0.01)$ well throughout the domain. Conversely the Gumbel model greatly over estimates $\chi(0.01)$ for all pairs of locations with non-zero empirical $\chi(0.01)$, bar the neighbouring grid cell. However, this neighbouring location is also well represented by the Gaussian model. The Gaussian model reproduces $\bar{\chi}(0.01)$ well for locations within a small to medium separation distance from London, with this distance being greater in the West-East direction, reflecting the common path of storms over Europe (Hoskins and Hodges (2002)). The Gaussian model over and under estimates $\bar{\chi}(0.01)$ for far away locations, with underestimation particularly in furthest away locations. This discrepancy is most likely due to the very small sample of joint extremes at these pairs of locations making estimates of $\bar{\chi}(0.01)$ highly uncertain. The Gumbel model greatly overestimates $\bar{\chi}(0.01)$, for all location, except again for those locations in very close proximity to London. This discrepancy in the Gumbel model is likely due to a misspecification of asymptotic dependence between most locations, resulting in an overestimation of the conditional dependencies in the extremes.

- **Finally, how does Figure 8 change with the choice of p?**

While this would be interesting to explore we feel that this is beyond the scope of the study. We have chosen to present the results for one conceptual loss function with the selection of this function and the the value of p justified in the text. I (Dawkins) have, however, addressed this in part in previous work published in my PhD thesis (Dawkins et al. (2016)), identifying no change in the overall results when p *is varied. We have now added a few sentences to the end of Section 4 to acknowledge this point:*

As previously mentioned, alternative windstorm loss thresholds have been implemented in other studies, for example the 98% quantile in Klawa and Ulbrich (2003), and the fixed thresholds of 20ms^{-1} in Bonazzi et al. (2012) and Dawkins et al. (2016) and 25ms^{-1} in Lamb and Frydendahl (1991) and Roberts et al. (2014). An exploration of the effect of the choice of loss threshold and, indeed loss function, on how the opposing dependence models represent joint losses would be an extremely interesting area of further investigation, however beyond the scope of this study. Dawkins (2016) goes some way in addressing this by presenting a comparison for the 98% quantile and 25ms^{-1} fixed loss thresholds in the same form as [Figure 3 in attachments]. Dawkins (2016) found that the overall suitability of the opposing models remained the same for both threshold, although the discrepancy of the Gumbel model was slightly smaller for the lower, 98% quantile, loss threshold. This was thought to be because modelled exceedances further from the upper limit of the joint distribution were less sensitive to a misspecification of the extremal dependence characteristic in the Gumbel model.

- **And does your model enable you to look at the sizes of losses at the two sites rather than just the probability of a loss jointly occurring at each site?**

  This point has been addressed in the response above in which we restructure Section 4 and explain how the analysis could be extended to explore the size of the losses.

1.2 Technical corrections

- **Throughout apposite/apposing should be opposite/opposing.**

  Thank you we have changed these.

- **Lines 71-76: these sentences are not entirely clear. On line 73 please clarify what the 'parametric representations' are representations of. May also be clearer to split the sentence on lines 73–76 into two sentences, the first to discuss what will be done (i.e. two copula dependence models fitted) and the second to explain why there two copulas were chosen.**

  Thank you we have changed this to:

  For a given pair of locations within a windstorm hazard field, we fit Gumbel and Gaussian bivariate copula dependence models, and explore how well these models represent the empirical estimates of $\chi(p)$ and $\bar{\chi}(p)$. These two copula models characterise opposing extremal dependence class, the Gaussian copula characterising asymptotic independence and the Gumbel asymptotic dependence, hence this comparison gives an indication of the form of extremal dependence within the data.

- **line 94: please could you clarify exactly which 'statistical property' is meant here. Something like 'the extremal dependence class estimated from the data'.**

  This has been removed in the restructuring of Section 4

- **line 135: based on**

  Thank you we have changed this.

- **line 137/138: doesn't quite get across the message that sites separated by different distances/directions will have different levels of dependence. Also, why is this likely to be the case? Have you looked at dependence as a function of distance/direction?**

  We have previously looked at the correlation in wind gust speeds as a function of distance and have therefore added this to the supplementary material and refer to it at the end of the first paragraph in Section 3:

  These three locations are shown in Fig. 1, and these two pairings are chosen because of their contrasting separation distances, and hence degrees of dependence (as shown in [Figure 4 in attachments] in the Supplementary Material).

Printer-friendly version

Discussion paper

[Figure 4 in attachments here] - Empirical correlation between (a) London, (b) Amsterdam and (c) Berlin and all other land locations over land, plotted against distance in (d), (e) and (f) respectively and for distance binned average correlation in (g), (h), (i) respectively.

- **RHS sign in the inequality should be reversed.**

  Thank you we have now changed this.

- **line 167: could append this paragraph to the previous one.**

  Thank you we have now changed this.

- **line 171: no need to state 'empirical exploration' as it is made clear later in the sentence that the estimate is an empirical one.**

  Thank you we have now changed this.

- **line 177-180: think this sentence can be removed as it doesn't quite fit here and is better covered in Section 3.3**

We agree and have now removed these lines and refer to Fig 3 in Section 3.3, which we now reference in the caption for Figure 3.

- **line 182: clarify that rarity of extreme events in historical data is not specific to this particular data set.**

  Yes, this is true. We have now edited the first paragraph of Section 3.3 to:

  Here, as in all datasets of environmental phenomena, the rarity of very extreme events makes it impossible to empirically quantify the asymptotic limits $\chi(0)$ and $\bar{\chi}(0)$, necessary for extremal dependence class identification.

- **line 186: not sure that 'model the asymptotic limit' is quite right, maybe 'predict' instead of 'model'. The model can only reflect the (sub-asymptotic) evidence in the data and the only extra information used in obtaining an estimate of the asymptotic limit is the assumption that the model fitted to sub-asymptotic data can be extrapolated to make predictions on higher (asymptotic) levels.**

  We have now altered the first paragraph of Section 3.3 to take this comment into account:

  To overcome this, Ledford and Tawn (1996) developed a bivariate tail model, able to characterise both classes of extremal dependence, which when fit to a bivariate random variable can be used to predict the asymptotic limit of the conditional probability measures and hence give an estimate of the class of extremal dependence, based on the sub-asymptotic evidence in the data and the assumption that the model can be extrapolated to asymptotic levels.

- **Notation: switching from X and Y to Z1 and Z2 mixes two different ways of distinguishing sites (different letters v. subscripts). Could change (X; Y) to (X1;X2) or (Y1; Y2), or change (Z1;Z2) to (ZX;ZY).**

  Thank you we have now changed this.

- **line 196: 'asymptotic' - 'extremal', as asymptotic dependence is a particular class of extremal dependence.**

  Thank you we have now changed this.

- **line 198: expression for (0) has an excess bracket.**

  Thank you we have now changed this.

- **line 201: 'model' missing after Ledford and Tawn (1996)**

  Thank you we have now changed this.

- **lines 218-220: this sentence would be clearer split into two. First describing the models for (X; Y) and then describing how model-based predictions of chi(p) and chi bar (p) are obtained from the models and are compared to the empirical estimate in Figure 3.**

  We agree and have now made this edit:

  In addition, as a comparison, alternative parametric bivariate dependence models, known as the Gaussian and Gumbel copulas, are used to model the pair (X, Y), since each copula characterises an opposing extremal dependence class. Model based predictions of $\chi(p)$ and $\bar{\chi}(p)$ for each copula are included in Fig. 3. The representation of $\chi(p)$ and $\bar{\chi}(p)$ in the limit $p \to 0$ for these opposing models then gives further indication of the extremal dependence class present.

- **Figure 6 caption ms-1.**

  Thank you we have now changed this.

References

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

Bortot, P., Coles, S., and Tawn, J. (2000). The multivariate gaussian tail model: an application to oceanographic data. *Applied Statistics*, 49:31–49.

Chu, C. R., Parlange, M. B., Katul, G. G., and Albertson, J. D. (1996). Probability density functions of turbulent velocity and temperature in the atmospheric surface layer. *Water Resources Research*, 32(6):1681–1688.

Cusack, S. (2013). A 101 year record of windstorms in the Netherlands. *Climate Change*, 116:693–704.

Dawkins, L. C. (2016). *Statistical modelling of European windstorm footprints to explore hazard characteristics and insured loss*. PhD thesis, University of Exeter, College of Engineering Mathematics and Physical Sciences.

Dawkins, L. C., Stephenson, D. B., Lockwood, J. F., and Maisey, P. E. (2016). The 21st century decline in damaging european windstorms. *Natural Hazards and Earth System Science*, 16:1999–2007.

Hoskins, B. J. and Hodges, K. I. (2002). New perspectives on the Northern Hemisphere winter storm tracks. *Journal of Atmospheric Science*, 59:1041–1061.

Huser, R. G. and Wadsworth, J. L. (2018). Modeling spatial processes with unknown extremal dependence class. *Journal of the American Statistical Association*, 0:1–11.

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

Lamb, H. and Frydendahl, J. (1991). *Historic Storms of the North Sea, British Isles and Northwest Europe*. Cambridge University Press.

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

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

Wadsworth, J. L. and Tawn, J. A. (2012). Dependence modelling for spatial extremes. *Biometrika*, 99(2):253–272.