# Modelling the spatial extent and severity of extreme European windstorms 

Paul Sharkey ${ }^{1}$, Jonathan A. Tawn ${ }^{2}$ and Simon J. Brown ${ }^{3}$<br>Email: pgshky@gmail.com, j.tawn@lancs.ac.uk, simon.brown@metoffice.gov.uk<br>${ }^{1}$ British Broadcasting Corporation, Bridge House, Salford Quays, Salford, M50 2LH, U.K.<br>${ }^{2}$ Department of Mathematics and Statistics, Lancaster University, Lancaster, LA1 4 YF, U.K.<br>${ }^{3}$ Met Office Hadley Centre, Exeter, EX1 3PB, U.K.


#### Abstract

Windstorms are a primary natural hazard affecting Europe that are commonly linked to substantial property and infrastructural damage and are responsible for the largest spatially aggregated financial losses. Such extreme winds are typically generated by extratropical cyclone systems originating in the North Atlantic and passing over Europe. Previous statistical studies tend to model extreme winds at a given set of sites, corresponding to inference in a Eulerian framework. Such inference cannot incorporate knowledge of the life cycle and progression of extratropical cyclones across the region and is forced to make restrictive assumptions about the extremal dependence structure. We take an entirely different approach which overcomes these limitations by working in a Lagrangian framework. Specifically, we model the development of windstorms over time, preserving the physical characteristics linking the windstorm and the cyclone track, the path of local vorticity maxima, and make a key finding that the spatial extent of extratropical windstorms becomes more localised as its magnitude increases irrespective of the location of the storm track. Our model allows simulation of synthetic windstorm events to derive the joint distributional features over any set of sites giving physically consistent extrapolations to rarer events. From such simulations improved estimates of this hazard can be achieved both in terms of intensity and area affected.


Keywords: Climate extremes, extratropical cyclones, extreme value analysis, Lagrangian model, spatial dependence.

## 1 Introduction

While the winter climate of the United Kingdom and northern Europe is typically associated with mild, wet weather that poses little infrastructual or societal risk, there has been an increased focus
in recent years on the impact of windstorms in this part of the world. These events are often the consequence of extratropical cyclones, and are directly linked to the occurrence of flooding, transport chaos and considerable damage to infrastructure. Roberts et al. (2014) describe a comprehensive catalogue of European windstorms in the period 1979-2012 that contains extensive information related to the meteorology and monetary impact of each storm. Storm Daria, which occurred in January 1990, is believed to be the most destructive windstorm in this period, with an estimated insured loss of $\$ 8.2 \mathrm{bn}$.

Windstorms are often a consequence of the passage of extratropical cyclones. Extratropical cyclones are synoptic-scale weather systems associated with low pressure that generally originate in the North Atlantic and progress northeasterly towards Europe. These systems can be characterised by paths of local vorticity maxima, which we refer to as tracks. Cyclones are typically formed as a result of horizontal temperature gradients and evolve according to a particular lifecycle with associated frontal systems (Shapiro and Keyser, 1990) where cold and warm air masses converge. High winds tend to occur along these boundaries (Hewson and Neu, 2015). A large body of research exists on cyclone identification, storm tracking and feature extraction in reanalysis datasets (Murray and Simmonds, 1991; Hodges, 1995), which produce good approximations of how a track develops in space and time. The study of cyclones with respect to the cyclone centre, a Lagrangian frame of reference, has provided useful insight into the evolution of cyclones and the physical processes that drive evolution (Catto et al., 2010; Rudeva and Gulev, 2011; Dacre et al., 2012). However, there has been little work on characterising the ensuing extreme winds in a robust way within such a framework.

We would like to assess the joint risk of multiple locations, in a set of fixed sites, experiencing the same windstorm event. This assessment is particularly difficult when using only the data from these sites as it is not possible to exploit knowledge that these extreme events are arising from the passage of extratropical cyclones over time and space. Consequently, inference cannot account for extreme events that, just by chance, did not pass over the observed sites and it does not incorporate the knowledge of the properties of extratropical cyclones into the estimation of probabilities of rarer events than those observed. Our paper describes an approach which incorporates this physical knowledge through a statistical model of windstorms, which can be used in practice to assess the marginal and joint risk of these weather systems over Europe while accounting for the varying probability of storm tracks over the region.

A common approach to statistical modelling of extreme wind speeds is to use techniques from extreme value analysis, which use models built on asymptotic arguments to estimate probabilities of events beyond the range of the data. In meteorological applications, such probabilities are commonly used by practitioners to design infrastructure appropriately to defend against the natural hazard being studied. The most widely-used approach in extreme value analysis is to consider excesses above a suitably high threshold. Consider a sequence of independent and identically distributed (i.i.d.) random variables $X_{1}, \ldots, X_{n}$. Under weak conditions on the $X_{i}$, the unique, non-degenerate distribution that the scaled excesses of a threshold by $X_{i}$ converges to, as the threshold tends to the upper limit $x_{F}$ of $X_{i}$, is the generalised Pareto distribution (GPD) (Pickands, 1975; Davison and Smith, 1990). We make the assumption that this limiting result holds for a large enough threshold $u$. The GPD takes the form

$$
\begin{equation*}
\operatorname{Pr}\left(X_{i}-u>x \mid X_{i}>u\right)=\left(1+\frac{\xi x}{\sigma_{u}}\right)_{+}^{-1 / \xi}, x>0 \tag{1}
\end{equation*}
$$

where $c_{+}=\max (c, 0)$ and where $\sigma_{u}>0$ and $\xi \in \mathbb{R}$ denote the scale and shape parameters respectively. The shape parameter is invariant to the choice of threshold but the scale parameter is threshold-dependent. The threshold $u$ is typically determined using some standard selection diagnostics (Coles, 2001) such as ensuring that the parameters are stable with respect to the threshold for all threshold choices larger than $u$.

There have been numerous studies using extreme value models to estimate extreme wind speeds (Coles and Walshaw, 1994; Fawcett and Walshaw, 2006; Ribatet, 2013). However, these models have no consideration of the physical processes generating the extremes. Some recent studies have, however, modelled extreme winds in the context of an extratropical cyclone. Della-Marta and Pinto (2009) use a GPD model to assess changes in extreme wind intensity under climate change scenarios, which led to results showing that the frequency of intense wind events in Europe is predicted to increase. Sienz et al. (2010) extended this approach to model the effect of the North Atlantic Oscillation (NAO) index. Bonazzi et al. (2012) modelled the tail dependence of wind speeds between locations over Europe using a bivariate extreme value copula and found dependence to be greater in the west-east direction, which is consistent with the passage of extratropical cyclone tracks over Europe. More recently, Youngman and Stephenson (2016) use extreme value analysis coupled with a geostatistical model to capture the spatio-temporal development of windstorms over Europe, but again the direct influence of the storm track is not accounted for.

The existing approaches share a common philosophy in that windstorms are modelled in an Eule-
rian frame of reference. In fluid mechanics, this approach refers to the scenario whereby an observer measures observations of a process at a fixed location while the process, e.g., a windstorm, passes over. So, it is a Eulerian framework that is being adopted when analysing data collected at a fixed set of sites over time, such as when modelling weather station data. It is the standard approach for statistical modelling of spatial data (Cressie, 1993), with the advantage that one can build large data sets, with time series at each location being observed. However, if one's concern is focused more on modelling the evolution and influence of the process itself, a Lagrangian frame of reference is required. This approach to modelling requires following the process and collecting observations as it moves through space and time. We take this approach as it is a natural framework on which to build a model for windstorms, since it allows us to model the behaviour of extreme winds relative to the storm centre.

Recent advances in climate modelling have resulted in the increased availability of high-resolution datasets that are spatially and temporally complete, allowing large-scale processes to be modelled in a Lagrangian framework (Catto et al., 2010). As the observed data record is relatively short with regard to storm tracks, and even more so with regard to windstorms, there is a need for a statistical model to provide extra information about the possible extreme, long-term characteristics of windstorms that are generated by the extratropical cyclone. In particular, we need to be able to assess the likelihood of observing more severe storms than those observed, where these might occur, and how large the spatial extent of the event might be.

Sharkey (2018) and Sharkey et al. (2019) model synthetic storm tracks in a way that replicates the climatology of extratropical cyclones in the North Atlantic. This model allows extrapolation to events that have larger vorticity than previously observed. We extend this work to model synthetic windstorm events relative to their storm tracks in a Lagrangian framework. We first describe a model for the instantaneous area, which we refer to as a footprint, affected by strong winds in the vicinity of the storm centre. We represent the footprint as an ellipse because it provides a parsimonious envelope for typical event shapes. We then present an algorithm which determines the location and shape of the footprint at any point in time. We model the evolution over time of both the characteristics of the ellipse and the magnitude and spatial distribution of the extreme winds within the footprint. These models allow us to generate a series of footprints for multiple windstorms associated with the synthetic storm tracks of Sharkey et al. (2019), providing a method for estimating the entire area affected by extreme winds from the passage of a single storm and the aggregated risk arising from extreme windstorms over the North Atlantic and Europe.

The paper is structured as follows. In Section 2, we introduce the data and our methods for extracting the features of the windstorm from the data and an exploratory analysis of these features. We discuss our modelling strategy in Section 3, introducing our approaches to modelling the evolution of the windstorm footprints and the extreme winds within the footprints. We derive a range of estimated spatial properties of windstorms over Europe in Section 4, before concluding in Section 5 with some discussion.

## 2 Windstorm definition and exploratory analysis

### 2.1 Data description

As in Sharkey et al. (2019), our work uses storm track data covering the North Atlantic and European domain. Our dataset consists of storm track locations at 3-hourly intervals with an associated vorticity measure representing the strength of the storm. Storms are identified and tracked over the period 1979-2014 from the ERA-Interim reanalysis dataset (Dee et al., 2011) using a feature extraction approach outlined in Hoskins and Hodges (2002) based on the tracking algorithm introduced in Hodges (1995). We restrict our attention to the set of storm tracks produced during an extended winter period (October-March), when storms are widely regarded to be most intense. We exclude Mediterranean extratropical cyclones as these do not produce large financial losses. We also exclude "medicanes" (Akhtar et al., 2014) as these arise from a different physical process and are not captured well by reanalysis data. We denote the longitude and latitude coordinates of the storm track at time $t$ by $\operatorname{Lon}_{t}$ and $\operatorname{Lat}_{t}$ respectively. The vorticity associated with the track at $\left(\operatorname{Lon}_{t}\right.$, Lat $\left._{t}\right)$ is denoted by $\Omega_{t}$.

Our model is based on wind speed data from the EURO4 numerical weather prediction model (Standen et al., 2017), which is a downscaled version of the ERA-Interim reanalysis dataset. Data are available on a 4 km spatial resolution over Europe and part of the North Atlantic, amounting to 1,100, 000 cells (see Figure 1). Values are obtained at hourly intervals over the period 1979-2014. We linearly interpolate the storm track locations and vorticity within each 3-hourly interval to match the hourly temporal resolution of the wind speed data. We select only the wind speed fields at times corresponding to the set of storm tracks. In particular, as we are looking to model the effect of the storm track on the spatio-temporal evolution of wind speeds in the vicinity of the track, we isolate the field of interest as a square-shaped region centred at the storm centre with sides of approximately $1,600 \mathrm{~km}$ in length (see Figure 2, which in the left panel shows such a region at a
time step when storm Daria was located over the UK). We believe that this field is large enough so that the extreme winds generated by a windstorm are sufficiently captured.

### 2.2 Marginal model

Initial investigation of the data confirms, as expected, that winds over the sea are markedly stronger than those over land (see Figure 2, left panel). This property is largely due to open water exerting significantly less drag on the atmosphere in contrast with the land surface, orography and manmade structures that impede strong winds. The contrast in scale over land and sea, and to a lesser extent, over low-lying and high-lying land, motivates a standardisation of wind speeds in each cell to have a common marginal distribution. In this sense, our approach is based on a copula modelling strategy over space (Joe, 1997).

Let $X(s, t)$ be a random variable denoting the wind speed in cell $s$ at time $t$, for $s=1, \ldots, 1,100,000$. We assume that $X(s, t)$ is identically distributed over $t$ for each cell $s$ but do not make any specific modelling assumption about its spatial or temporal dependence structure. We propose a marginal model for the distribution of $X(s, t)$, for each cell $s$, of the form

$$
F_{s}(x)= \begin{cases}\hat{F}_{s}(x) & x \leq u_{s}  \tag{2}\\ 1-\lambda_{u_{s}}\left(1+\xi_{s} \frac{x-u_{s}}{\sigma_{u_{s}}}\right)^{-1 / \xi_{s}} & x>u_{s}\end{cases}
$$

where $\hat{F}_{s}$ denotes the empirical distribution function for realisations of $X(s, t)$ over time. For realisations above $u_{s}$, the GPD in (1) is used as a conditional model for excesses above $u_{s}$, with cell-specific parameters $\left(\sigma_{u_{s}}, \xi_{s}\right)$. To undo this conditioning, a third parameter $\lambda_{u_{s}}$, denoting the probability of an exceedance of $u_{s}$, must be specified. Parameter stability plots were checked at a number of cells, which indicated that a threshold corresponding to the $98 \%$ marginal quantile in each cell would be a good choice for all cells. We therefore choose this quantile to be the cell-specific threshold in each cell. Parameter estimates are obtained using maximum likelihood techniques. We note that we do not attempt to impose spatial smoothness on the form of $F_{s}$.

Figure 1 shows the parameter estimates of the GPD corresponding to each cell in the region over Europe along with the threshold, which corresponds to the $98 \%$ quantile in each cell. This figure shows explicitly the contrast in wind speed magnitudes between locations on land and sea. This contrast is also reflected in the estimation of the scale parameter, but the shape parameter exhibits no such contrast between land and sea, with most estimates occurring in the region $(-0.2,0)$, indicating that the distribution of wind speeds has a finite endpoint in general. The numerical


Figure 1: The extreme value threshold (left) used in the specification of the GPD model, along with parameter estimates of the scale (centre) and shape parameters (right) for each cell.


Figure 2: Wind speeds, in m/s, at 3pm on January 25th, 1990 in the vicinity of Storm Daria (left) and standardised onto $\operatorname{Exp}(1)$ margins (right). The storm centre is represented by the cross. The white box contains an example of a localised convective event. Land/sea borders are not explicitly shown in the left panel, but can be seen due to the contrast in magnitude between winds over land and sea.
maximisation algorithm used to obtain the parameter estimates mostly converges, however there are certain regions that exhibit unusual behaviour. For example, the $98 \%$ quantile threshold is
a lot higher in areas over Iceland than other land locations, while the Italian Alps see unusually high estimates of the shape parameter. The challenges of representing wind over high orography in numerical weather models is well documented, as evidenced by these two particular locations, and recalibration methods have been proposed (Howard and Clark, 2007). For this analysis we exclude regions with orography above 500 m from our analysis.

We propose a marginal standardisation to unit exponential $\operatorname{Exp}(1)$ margins over space and time using a probability integral transform under the marginal model (2). The exponential scale has been found to be most ideal for studying extremal dependence (Heffernan and Tawn, 2004; Wadsworth et al., 2013). We define $X^{E}(s, t)$ to be the relative wind speed on $\operatorname{Exp}(1)$ margins in cell $s$ and time $t$, such that for all cells and all times we have that

$$
X^{E}(s, t)=-\log \left\{1-F_{s}(X(s, t))\right\}
$$

where $F_{s}$ is defined as in (2). We use the term relative wind speeds to describe $X^{E}(s, t)$ as this quantity defines wind speeds relative to the marginal characteristics of cell $s$. Figure 2 shows the effect of the standardisation on one time step of Daria. In particular, we see spatially correlated values of high relative wind speeds over both land and sea as a result of the transform and the land/sea contrast is almost entirely removed. Thus, this transformation recovers the underlying process of the windstorm. This approach reveals the shape of the windstorm event without the influence of the marginal characteristics at each location, which should allow for a simpler approach to modelling the spatial extent of the event.

### 2.3 Feature extraction

Figure 2 shows that, as well as the large band of strong relative winds clustered near the storm centre, small localised fragments of high relative wind speeds are visible on the western edge of the footprint. Such localised events are due to convection and not due to the larger scale dynamics of the storm. Since we believe them not to be directly linked with the extratropical storm system, we are not concerned with modelling these features. Our work is focused on modelling the features of an extratropical cyclone that occur on larger spatial scales and have the potential to produce much larger impacts, such as large insurance losses (Roberts et al., 2014), than these localised convective events.

We therefore require some methodology to extract the main features of interest, such as the cluster of high relative wind speeds, from the standardised fields and to see how these track through time.

Our strategy is to identify all large-scale high relative wind speeds associated with the extratropical cyclone at each time step, to then bound these by an elliptical region in space, and to find all such regions associated with each cyclone track. We call each of these regions a footprint and we refer to a windstorm as being the collection of footprints over a cyclone lifetime. If a windstorm has gaps without footprints, either before, between or after periods with footprints, we refer to these as inactive and active phases of the windstorm respectively. For a storm centre ( $\mathrm{Lon}_{t}, \mathrm{Lat}_{t}$ ) with associated vorticity $\Omega_{t}$ at the $t$-th time step of a cyclone, we identify the footprint from the associated wind field by a multi-stage process below. To aid understanding of this process Figure 3 (left and centre panels) illustrates Steps 2-4 for storm Daria and Figure 4 presents the features we extract.

## Step 1: Removal of small-scale features

We apply a spatio-temporal Gaussian filter (Nixon and Aguado, 2012) to each field, which removes the effect of the small-scale convective events. We define a threshold $v$, indicating a large relative wind speed for the filtered data, with relative wind speeds below this level temporarily masked and those that exceed it retained. If the entire filtered wind field is masked, then it is interpreted as there being no significant windstorm activity, no footprint is formed and the windstorm is deemed to be in an inactive phase.

## Step 2: Clustering of large-scale high relative wind speeds

When there is at least one exceedance of $v$ we identify clusters of exceedances using the clustering algorithm DBSCAN (Density-Based Spatial Clustering of Applications with Noise) (Ester et al., 1996), which recursively groups cells into distinct clusters based on adjacency to neighbouring cells, without the need to specify the number of clusters in advance. We then extract the largest of the identified clusters and define it by $S_{t}$ containing $m_{t}$ locations in 2-dimensional space such that $S_{t}=\left\{s_{t, 1}, \ldots, s_{t, m_{t}}\right\} \in \mathbb{R}^{2}$.

## Step 3: Identifying the bounding ellipse

We want to model the behaviour of the wind field in the region that bounds the cluster identified in Step 2. We specify all footprints to be ellipses as we found clusters tended to have this shape and that these shapes can be described parsimoniously (by their centre, semi-major and semi-minor axes, and associated angles, see Figure 4). To select the ellipse to contain all elements of $S_{t}$ we want it to be the minimum-area ellipse containing the cluster, which is achieved by using Khachiyan's algorithm (Moshtagh, 2005; Todd and Yıldırım, 2007). A set $\mathcal{E}_{t}$ contained by an ellipse can be
written as

$$
\mathcal{E}_{t}=\left\{s \in \mathbb{R}^{2}:\left(s-c_{t}\right)^{T} E_{t}\left(s-c_{t}\right) \leq 1\right\},
$$

where $c_{t}$ is the centre of the ellipse, $E_{t}$ is a positive-definite matrix, and the area of $\mathcal{E}_{t}$ is $\left\{\operatorname{det} E_{t}^{-1}\right\}^{1 / 2}$. Khachiyan's algorithm finds $c_{t}$ and $E_{t}$ to minimise det $E_{t}^{-1}$ subject to $\left(s_{t, i}-c_{t}\right)^{T} E_{t}\left(s_{t, i}-c_{t}\right) \leq 1$ for $i=1, \ldots, m_{t}$ using conditional gradient ascent methods. Once the footprint ellipse has been identified we take all the original relative wind speeds from inside the ellipse to give the footprint. At time $t$, we denote the semi-major and semi-minor axes of the ellipse by $A_{t}$ and $B_{t}$ respectively, the grid cell distance between the storm centre and the centre of the ellipse by $R_{t}^{E}$, the angle between due south and the vector between the storm centre and the centre of the ellipse by $\Theta_{t}^{E} \in[-\pi, \pi]$, and the orientation of the ellipse relative to due north by $\Gamma_{t} \in[-\pi / 2, \pi / 2]$.

## Step 4: Elimination of spurious footprints

Occasionally we found footprints not generated by the extratropical cyclone spuriously identified by Steps $1-3$ when $R_{t}^{E}$ is large or $\left\{\operatorname{det} E_{t}^{-1}\right\}^{1 / 2}$ is sufficiently small. We exclude these footprints and treat the process as in an inactive phase.

## Step 5: Selecting features from the footprint

We are mostly interested in the extreme relative wind speed and its location within footprint $\mathcal{E}_{t}$ with the maximum value being denoted by $W_{t}$, the grid cell distance between the centre of the ellipse and the location of the maximum being $R_{t}^{W}$, and the bearing between the two cells relative to due south being $\Theta_{t}^{W} \in[-\pi, \pi]$. In Section 3.4 we model the entire spatial process of the relative wind speeds in footprint $\mathcal{E}_{t}$ conditional on the value of $W_{t}$ and its location.

### 2.4 Exploratory analysis of footprint features

We undertake an extensive exploratory analysis on a number of aspects driving and influencing the evolution of windstorm activity, of which some findings are reported here. First, we investigate the dependence structure of characteristics of the ellipse representing the windstorm footprint, as shown in Figure 4. We assess the factors influencing the activation and termination of windstorm events. We also look at some quantities representing the spatial distribution of wind speeds relative to the storm centre, and how these compare with previous studies. Finally, we explore how the distribution of relative wind speeds within the footprint varies with respect to characteristics of the storm track and footprint.


Figure 3: The largest cluster of relative winds (left) after applying the spatio-temporal filter and DBSCAN clustering to the field in Figure 2, the ellipse-shaped footprint determined by Khachiyan's algorithm (centre) and an example of simulating wind speeds from our model within this ellipse. (right).

Figure 5 shows boxplots illustrating some key dependencies between variables of a footprint. The area of the ellipse, which is proportional to $\Delta_{t}=A_{t} \times B_{t}$, tends to increase as vorticity and maximum relative wind speed increases, indicating that the strongest events tend to occur on a larger spatial scale. The radius $R_{t}^{E}$ tends to decrease as $\Omega_{t}$ increases, although the effect is small, suggesting that footprints tend to occur closer to the storm centre when a large vorticity is observed. Maximum relative wind speed $W_{t}$ tends to increase as $\Omega_{t}$ increases, though this dependence is weak. We also examine partial autocorrelation plots (not shown) to determine how individual components of the ellipse depend on their lags. This analysis shows evidence of a second-order temporal dependence structure in most components reflecting the smooth evolution of footprints through the windstorm. This analysis suggests there is considerable uncertainty in this dataset.

Next consider what factors influence whether a windstorm is active or not. With regards to windstorm activation, we investigate the components of the track that may trigger an event. We find that the probability of windstorm activations tends to increase as vorticity increases, signalling a direct link between the intensity of the storm track and the occurrence of windstorm events. With regard to an inactive phase caused by termination of an active phase, we need to account for the variables of the windstorm footprint in addition to information from the track. We find that a storm is more likely to terminate if it is associated with small values of relative maximum wind


Figure 4: Graphical representation of the model variables derived from the ellipse shape used to summarise the windstorm footprint.


Figure 5: Boxplots showing the dependence between $\Omega_{t}$ and $R_{t}^{E}, \Omega_{t}$ and $\Delta_{t}^{1 / 2}, \Delta_{t}^{1 / 2}$ and $W_{t}$, and $\Omega_{t}$ and $W_{t}$.


Figure 6: The spatial density of track locations that are associated with observed (left) and simulated (right) windstorms.
speed, vorticity and area. These findings indicate a termination is more likely if the windstorm weakens both in terms of its magnitude and its spatial scale.

This exploratory analysis also shows some spatial variation in the occurrence of footprints. Figure 6 (left panel) shows the density of storm track locations when windstorms are in an active phase. This figure shows that windstorms tend to occur over the North Atlantic and western Europe, with the density decaying as one moves to the edges of the domain. The reductions in density on the western boundary are the result of the filtering of footprints and the boundary effect on the EURO4 simulation of weather. Most extratropical cyclones enter the EURO4 domain through this western boundary but their intensity is diminished due to the coarser resolution of the driving GCM. Their subsequent intensification takes a number of time steps which results in smaller footprints and lower winds towards this boundary.

For each windstorm in an active phase, i.e., with an identified footprint, we collect all fields centred on the storm's maximum vorticity, like in Figure 3, and assess the spatial distribution of these fields relative to the storm centre. Figure 7 shows the mean, $95 \%$ and $99 \%$ quantiles of wind speeds. This figure shows that relative wind speeds tend to be larger in regions southwest of the storm centre, which arise with the passage of cold fronts. This figure also shows the density of events over all footprints, illustrating that windstorms are most likely to occur southwest of the storm centre, with very few events occurring in the northern half of the field in comparison. These are consistent with the observations of Catto et al. (2010) and Rudeva and Gulev (2011).

Unusual events with large magnitudes are detected on the northwest and southeast edges of the


Figure 7: The mean, $95 \%$ quantile, $99 \%$ quantile and density of wind speeds relative to the storm centre (represented here as a cross) over all observed footprints.
domain. Despite our filtering, these events may not be generated by the extratropical cyclone; however, they are very rare and their impact minimal. The mean behaviour shows a local minimum occurring close to the storm centre, which likely arises as a result of low wind speeds occurring at pressure minima (Catto et al., 2010), which is known to be spatially adjacent to where the maximum vorticity occurs (Hoskins and Hodges, 2002).

We investigate the factors influencing the distribution of relative wind speeds within each footprint through analysis of their mean and standard deviation. This reveals, intuitively, that the mean and standard deviation of winds tend to increase as the maximum relative wind increases. In contrast, an increase in vorticity tends to slightly reduce the mean relative wind, while increasing the standard deviation. Additionally, the standard deviation of wind increases as the area of the ellipse increases, which one would expect given the increased chances of observing smaller and larger values of wind speed. By construction the strongest relative winds tend to occur near the location of the maximum, determined by $R_{t}^{W}$ and $\Omega_{t}^{W}$, while weaker relative winds are more likely closer to the perimeter of the ellipse. Additionally, the relative winds tend to exhibit anisotropic properties, which tends to manifest in a 'stretching' of the band of strongest winds oriented in the direction perpendicular to $\Theta_{t}$. As the windstorm evolves, this gives the effect of the winds 'bending' around the storm centre.

## 3 Windstorm modelling

### 3.1 Introduction

We propose an approach for modelling and simulating windstorms relative to an extratropical cyclone track that is motivated by our findings in Section 2.4. In particular, along a cyclone track we model whether the windstorm is active or not at each time step, and when it is active we model the evolution of the footprint's characteristics along with the spatial distribution of wind speeds within a footprint. Specifically if $\left[t_{S}, t_{T}\right]$ is a contiguous active interval, with inactive periods immediately before and afterwards, we call $t_{S}$ and $t_{T}$ the activation and termination times respectively. A windstorm can have repeated phases of activity and inactivity along a cyclone track. For each active phase the features of the process identified in Figure 4, denoted $\boldsymbol{Z}_{t}$ at time $t$, are jointly modelled over time for all $t_{S} \leq t \leq t_{T}$ in Section 3.2; whilst we present models for the occurrence $t_{S}$ and $t_{T}$ in Section 3.3, and in Section 3.4 we model the wind speed fields within the footprint.

### 3.2 Footprint modelling

During active phases, we model the joint temporal evolution of the footprints variables $\boldsymbol{Z}_{t}=$ $\left\{A_{t}, B_{t}, W_{t}, R_{t}^{E}, \Theta_{t}^{E}, R_{t}^{W}, \Theta_{t}^{W}, \Gamma_{t}\right\}$ of Figure 4 conditional on the storm track variables (Lon Lat $\left._{t}, \Omega_{t}\right)$. For simplification purposes for our inference, in practice we condition on where ( $\operatorname{Lon}_{t}$, Lat $_{t}$ ) in $\nabla_{t}$, where $\nabla_{t}$ is a region of size $20^{\circ} \times 14^{\circ}$ centred at $\left(\operatorname{Lon}_{t}\right.$, Lat $\left._{t}\right)$. The exploratory analysis of Section 2.4 motivates modelling $\boldsymbol{Z}_{t}$ as a $k$-th order Markov process during active phases along the cyclone track and for $\boldsymbol{Z}_{t}$ and $\boldsymbol{Z}_{t^{\prime}}$ to be conditionally independent of each other given the track information if $t$ and $t^{\prime}$ correspond to different cyclone tracks or different active phases along a single track.

By the Markov property, the distribution of the current value of a process is affected only by the previous $k$ time steps of the process. Let $\boldsymbol{Z}_{i: j}=\left\{\boldsymbol{Z}_{t}: t=i, \ldots, j\right\}$, where $i \leq j$, then it is only necessary to model the joint distribution of $\boldsymbol{Z}_{t: t+k}$, from which the conditional density function of $\boldsymbol{Z}_{t+k} \mid \boldsymbol{Z}_{t: t+k-1}$ can be derived. Empirical evidence, such as Figure 5, helps to identify which components of $\boldsymbol{Z}_{t: t+k}$ are independent or conditionally independent, which helps to identify $k$. For example, we found $k=2$ and a weakly dependent relationship between $\Omega_{t}$ and $W_{t}$ with $\Omega_{t}$ and $W_{t}$ conditionally independent given $A_{t}$ and $B_{t}$. This result allows us to simplify our model for $\operatorname{Pr}\left(\boldsymbol{Z}_{t+k} \leq z \mid \boldsymbol{Z}_{t: t+k-1}=\boldsymbol{z}_{t: t+k-1}\right)$. First, consider an initialisation time $t_{A}$, with $t_{S} \leq t_{A} \leq t_{T}$, which is determined by an algorithm in Section 3.3. We simulate the initialisation value $\boldsymbol{Z}_{t_{A}}$ from the distribution of $\boldsymbol{Z}_{t_{A}} \mid\left(\operatorname{Lon}_{t_{A}}\right.$, $\left.\operatorname{Lat}_{t_{A}}, \Omega_{t_{A}}\right)$, which we estimate using a conditional kernel density (see the online appendix). Then consider forward propagation of $\boldsymbol{Z}_{j}$, given $\boldsymbol{Z}_{t_{A}}$, for $t_{A}<j \leq t_{T}$, using the following conditional distributions:

$$
\begin{aligned}
& R_{j}^{E} \quad \mid \quad R_{(j-k))_{t_{A}}: j-1}^{E}=r_{(j-k)_{\geq t_{A}}: j-1}^{E}, \Omega_{j}=\omega_{j}, \\
& \Theta_{j}^{E} \mid \quad \Theta_{(j-k)_{\geq t_{A}}: j-1}^{E}=\theta_{(j-k)_{\geq t_{A}}: j-1}, R_{j}^{E}=r_{j}^{E} ; \\
& A_{j} \mid \quad A_{(j-k)_{\geq t_{A}}: j-1}=a_{(j-k)_{\geq t_{A}}: j-1}, \Omega_{j}=\omega_{j}, R_{j}^{E}=r_{j}^{E}, \Theta_{j}^{E}=\theta_{j}^{E} ; \\
& B_{j} \mid \quad B_{(j-k)_{\geq t_{A}}: j-1}=b_{(j-k)_{\geq t_{A}}: j-1}, \Omega_{j}=\omega_{j}, R_{j}^{E}=r_{j}^{E}, \Theta_{j}^{E}=\theta_{j}^{E} ; \\
& W_{j} \mid \quad W_{(j-k)_{\geq t_{A}}: j-1}=w_{(j-k)_{\geq t_{A}}: j-1}, R_{j}^{E}=r_{j}^{E}, \Theta_{j}^{E}=\theta_{j}^{E}, A_{j}=a_{j}, B_{j}=b_{j} ; \\
& \Gamma_{j} \mid \quad \Gamma_{(j-k)_{\geq t_{A}}: j-1}=\gamma_{(j-k)_{\geq t_{A}}: j-1}, \Theta_{j}^{E}=\theta_{j}^{E} ; \\
& R_{j}^{W} \mid \quad R_{(j-k)_{\geq t_{A}}: j-1}^{W}=r_{(j-k)_{\geq t_{A}}: j-1}^{W}, A_{j}=a_{j} ; \\
& \Theta_{j}^{W} \mid \quad \Theta_{(j-k)_{t_{A}}: j-1}^{W}=\theta_{(j-k)_{\geq t_{A}}: j-1}^{W},
\end{aligned}
$$

where we use the notation $(j-k)_{\geq t_{A}}=\max \left\{j-k, t_{A}\right\}$ and in each case the distribution is conditional on $\left(\operatorname{Lon}_{j}, \operatorname{Lat}_{j}\right)$ being in $\nabla_{j}$. Realisations for $R_{j}^{W}$ and $\Theta_{j}^{W}$ are rejected if the simulated position of the maximum occurs outside of the footprint at time $j$. We model each of these con-
ditional distributions using kernel density estimates, the formulation of which can be found in the online appendix. Backwards simulation for $t_{T} \leq j<t_{A}$ is implemented similarly, but with substituting $\boldsymbol{Z}_{j}$ for $\boldsymbol{Z}_{t_{A}-j}$.

### 3.3 Initialisation, activation and termination

We have to model the probability distributions for initialisation time $t_{A}$, activation time $t_{S}$ and termination $t_{T}$, for each active period given other covariates of the footprint and storm track process. The model formulation is identical for each of these, with $t_{A}$ being generated first and $t_{S}$ and $t_{T}$ conditional on $t_{A}$. Let $T_{t}$ be a Bernoulli random variable such that: $T_{t}=1$ if the storm is active at time $t$ and $T_{t}=0$ otherwise. We model $T_{t}$ by a Bernoulli logistic generalised additive model (Wood, 2006). So $T_{t} \sim \operatorname{Bernoulli}\left(p_{t}\right)$, where

$$
p_{t}=\frac{\exp \left\{\sum_{i=1}^{q} \beta_{i}\left(\nu_{i, t}\right)\right\}}{1+\exp \left\{\sum_{i=1}^{q} \beta_{i}\left(\nu_{i, t}\right)\right\}},
$$

where $\beta_{i}$ is a smooth non-linear function of covariate $\nu_{i}$ with $i \in(1, \ldots, q)$, where $q$ is the number of covariates and $\nu_{i, t}$ denotes the realisation of the $i$ th covariate at time $t$. The smooth functions are represented by penalised regression splines, where the smoothing parameter is determined using generalised cross validation (GCV) and the model is fitted using penalised maximum likelihood, with the choice of covariates we used discussed in Section 4.1.

Now consider how to determine the initialisation time $t_{A}$ out of the possible values of $1 \leq t \leq \ell$, where $\ell$ is the duration of the storm track. Section 2.4 demonstrated that the windstorm is typically in an active phase at the time when the maximum vorticity is observed, as this time is associated with the strongest winds over the cyclone track and hence has the highest probability $p_{t}$ of being active. We denote this time by $t_{\Omega}$, with $1 \leq t_{\Omega} \leq \ell$, If the windstorm is determined to be active at $t_{\Omega}$, i.e., $T_{t_{\Omega}}=1$, then we set $t_{A}=t_{\Omega}$. If a windstorm is inactive at $t_{\Omega}$, i.e., $T_{t_{\Omega}}=0$, we attempt to initialise the active phase successively forwards, and then backwards, in time until an active phase is first identified, i.e., the first occasion when we get a Bernoulli realisation of 1 . If the initialisation is found during the forwards procession then $t_{A}>t_{\Omega}$ and $t_{S}=t_{A}$ and we only propagate the Markov chain forwards. Similarly, initialisation on the backwards procession gives $t_{A}<t_{\Omega}$ we set $t_{T}=t_{A}$. We allow for the possibility that multiple phases of consecutive footprints can occur on the same track. If $t_{S}>1$, we proceed backwards along the track to check if the windstorm reactivates, in which the same procedure applies until $t=1$. Similarly, if $t_{T}<\ell$, we proceed forwards along the track to check if the windstorm reactivates, in which the same procedure applies until $t=\ell$.

### 3.4 Modelling wind speeds within a footprint

We also require an approach for modelling the spatial distribution of relative wind speeds within a footprint at each time step of a windstorm. In theory many different spatial processes could be used, but a natural class of models to consider is Gaussian processes, which are widely used in geostatistics (Cressie, 1993; Stein, 1999; Diggle and Ribeiro, 2007) to model spatial data as they are simple to use and parsimonious. A Gaussian process describes the joint distribution of random variables over a continuous domain such as space inside a footprint, while for any finite collection of locations in the space the variables follow a multivariate Gaussian distribution.

Let $\left\{X^{E}(s, t): s \in \mathcal{E}_{t}\right\}$ denote the field of relative wind speeds in the ellipse $\mathcal{E}_{t}$ at time $t$, where $X^{E}(s, t)$ is marginally $\operatorname{Exp}(1)$ distributed over $t$ for each $s$. To model using Gaussian processes we first need to transform the variables marginally to be Gaussian. Let $D_{t}$ be the distribution function of $X^{E}(s, t)$ for all $s \in \mathcal{E}_{t}$ at time $t$. As the values in the footprint are not typical of the wind field, $D_{t}$ is not the distribution function of an $\operatorname{Exp}(1)$ variable. Based on findings from Section 2.4, we use a weighted nonparametric estimate $\tilde{D}_{t}$ of the distribution function $D_{t}$. Using a kernel, we weight observed relative wind speeds in $\mathcal{E}_{t}$ conditional on the corresponding values of $W_{t}, \Delta_{t}$ and $\Omega_{t}$. We use a probability integral transform to convert to a Gaussian field, which we denote by $X^{G}(s, t)$, with

$$
\begin{equation*}
X^{G}(s, t)=\Phi^{-1}\left[D_{t}\left(X^{E}(s, t)\right)\right], \tag{3}
\end{equation*}
$$

for all $s \in \mathcal{E}_{t}$ and all $t$, where $\Phi$ denotes the standard Gaussian distribution function.

We make the assumption that for each $t,\left\{X^{G}(s, t): s \in \mathcal{E}_{t}\right\}$ follows a Gaussian process with zero mean and unit variance. The correlation between sites $\left(s_{i}, s_{j}\right) \in \mathcal{E}_{t}$ is modelled by $\rho\left(\left(s_{i}, s_{j}\right) J_{t} \Psi_{t}\right)$ where $\rho(\cdot)$ denotes an isotropic correlation function and the time varying anisotropic effects are described by the matrices

$$
\Psi_{t}=\left[\begin{array}{cc}
\cos \psi_{t} & -\sin \psi_{t} \\
\sin \psi_{t} & \cos \psi_{t}
\end{array}\right] \quad J_{t}=\left[\begin{array}{cc}
1 & 0 \\
0 & \zeta_{t}
\end{array}\right],
$$

where $\psi_{t}$ is the time-varying anisotropy angle representing the counter-clockwise rotation of the space and $\zeta_{t}>1$ is the time-varying anisotropy ratio, which controls the degree of stretching along the angle where correlation decays most slowly with increasing distance. Supported by the analysis of within-footprint winds and their anisotropic properties, we fix $\psi_{t}$ to be the angle perpendicular to $\Theta_{t}$ and set $\zeta_{t}=A_{t} / B_{t}$.

The correlation function is typically chosen so that the correlation between $X^{G}\left(s_{1}, t\right)$ and $X^{G}\left(s_{2}, t\right)$ decreases as the distance $\left|s_{2}-s_{1}\right|$ between the sites increases. A common choice of correlation function is the Matérn family, which has the form

$$
\begin{equation*}
\rho(u)=\left\{2^{\kappa-1} \Gamma(\kappa)\right\}^{-1}\left(u / \alpha_{t}\right)^{\kappa} K_{\kappa}\left(u / \alpha_{t}\right), \tag{4}
\end{equation*}
$$

where $\kappa>0$ is a shape parameter that determines the smoothness of the underlying process, $K_{\kappa}$ denotes a modified Bessel function of the second kind of order $\kappa$, and $\alpha_{t}>0$ is a time-varying scale parameter, with dimensions of distance, with increasing $\alpha_{t}$ corresponding to stronger correlations. The Matérn family is a generalisation of other common choices of correlation functions including the exponential $(\kappa=0.5)$ and Gaussian $(\kappa \rightarrow \infty)$, with larger $\kappa$ giving smoother fields. We fix $\kappa=0.6$, which gives fields of similar smoothness to the observed fields. We estimate the parameter $\alpha_{t}$ for each footprint using variogram methods. We avoid the use of likelihood methods due to the computational difficulties that arise with large spatial datasets like ours. Our investigations (not shown here) suggest that $\alpha_{t}$ is constrained by the value of $\Delta_{t}$, which is proportional to the area of the ellipse. We therefore model $\alpha_{t} \mid \Delta_{t}$ using a conditional kernel density, the formulation of which is outlined in the online appendix. At time $t$, we generate a realisation of $\alpha_{t}$ conditional on the simulated realisation of $\Delta_{t}$ determined by the model in Section 3.2.

We incorporate information about the physical structure of the footprint in determining the structure of the Gaussian field $\left\{\tilde{X}^{G}(s, t): s \in \mathcal{E}_{t}\right\}$ by using conditional simulation (Diggle and Ribeiro, 2007). We impose three conditions on the simulated fields: that the maximum relative wind speed is simulated at the position determined by $R_{t}^{W}$ and $\Theta_{t}^{W}$; that the lower limit of $\tilde{D}_{t}$ occurs on the outer perimeter of the ellipse; and that the lower limit of $\tilde{D}_{t}$ occurs everywhere in the region corresponding to the local minimum of wind speeds (see Section 2.4) near the storm centre when a footprint is simulated in the vicinity of this region. The first and second conditions create a two-dimensional pseudo-Brownian bridge between the position of maximum and the perimeter of the footprint. To impose the third condition, we specify a second ellipse centred at the local minimum with random dimensions for size; specifically we fix the maximum length of the semimajor and semi-minor axes of this ellipse to be 40 and 35 units of grid-cell distance respectively, with random perturbations modelled as $\operatorname{Exp}(0.05)$ random variables. These values were found to replicate well the average behaviour of wind speeds in the region at which the local minimum occurs.

After simulation of $\left\{\tilde{X}^{G}(s, t): s \in \mathcal{E}_{t}\right\}$, we transform this field to obtain a field of relative wind
speeds conditional on the characteristics of the footprint such that

$$
\tilde{X}^{E}(s, t)=\tilde{D}_{t}^{-1}\left(\Phi\left[\tilde{X}^{G}(s, t)\right]\right) .
$$

An example of this is shown in Figure 3, in which we have simulated a field of relative wind speeds conditional on the footprint of storm Daria at this particular time step. Our model captures the correlation structure of the field quite well, with the weakest winds occurring on the outer perimeter of the ellipse and the large winds occurring in similar locations to the observed footprint. For this simulated field, the decay from the maximum relative wind speed appears more isotropic than the observed field. When viewed in an Eulerian framework, this level of spatial difference is not too important as these footprints move over space with the cyclone track, so blur out this distinction. Having obtained the simulated relative wind fields, we can then transform these onto the observed margins, such that for each $s \in \mathcal{E}_{t}$

$$
\tilde{X}(s, t)=F_{s}^{-1}\left(\tilde{X}^{E}(s, t)\right),
$$

where $F_{s}$ denotes the marginal model for cell $s$, as defined in (2). With this formulation, we are assuming the relative wind fields at consecutive time steps are conditionally independent given temporally correlated realisations of $W_{t}$ and $\alpha_{t}$. While this assumption gives good results in practice, further investigation may be necessary to assess whether performance could be improved by specifying a spatio-temporal structure in the correlation function $\rho(\cdot)$.

## 4 Results

We examine the performance of the windstorm model in terms of simulated footprint characteristics and then the wind speeds within the footprint. The joint risk from extreme windstorms at locations in northern England and eastern Germany is then explored by combining the windstorm model presented here with the track model of Sharkey et al. (2019) to produce estimates of joint event probabilities through simulation.

### 4.1 Validation of footprint model

We explore first whether the characteristics of windstorm footprints are being captured through an assessment of the marginal distributions of the individual components and their dependence structure. QQ plots based on the simulation of footprints using the model described in Section 3.2, applied to 2,944 synthetic storm tracks, the same number as in the observed record, are assessed both for the observed tracks and tracks generated by the model of Sharkey et al. (2019). Both


Figure 8: QQ plots, with $95 \%$ tolerance intervals, comparing the observed and simulated marginal distributions of $R_{t}^{E}, \Theta_{t}, \Delta_{t}^{1 / 2}$ and $W_{t}$. Simulated values are based on footprints relative to 2,944 synthetic tracks from the model of Sharkey et al. (2019).
showed similar positive results, so we illustrate only the latter (see Figure 8). They show that the marginal distributions of radius, bearing, proportional area and relative maximum wind speed are being captured well by the model. Figure 9 shows that, when compared with Figure 5, the dependence structure of these components is also consistent with the observations. We can thus conclude that the physical structure of the observed windstorm footprints is replicated sufficiently by the model.


Figure 9: Boxplots showing dependence structure of some aspects of the simulated footprints. Simulated values are based on footprints relative to 2,944 synthetic tracks from the model of Sharkey et al. (2019).

We examine the components of a windstorm influencing the activation and termination models that were outlined in Section 3.3. In both cases, we investigate multiple combinations of covariates and compare model fit using AIC. The best fitting activation model includes functions of vorticity, longitude and latitude. Figure 10 (left panel) shows the estimated smooth function $\beta_{i}$ associated with vorticity, which shows that $\beta_{i}$ tends to increase approximately linearly as vorticity increases.


Figure 10: The smooth functions $\beta_{i}$ (with $95 \%$ confidence intervals) showing the effect of $\Omega_{t}$ on the probability of windstorm activation (left) and of $\Delta_{t}^{1 / 2}$ (centre) and $W_{t}$ (right) on the probability of termination.

This relationship has the effect that the probability of activation tends to increase as vorticity increases, which reflects our findings from Section 2.4. The best fitting termination model includes functions of $\Delta_{t}^{1 / 2}$ and $W_{t}$. Figure 10 (centre and right panels) shows that $\beta_{i}$ tends to decrease nonlinearly as both $\Delta_{t}^{1 / 2}$ and $W_{t}$ increase, which means that the probability of termination also tends to decrease. The effect of $\Delta_{t}^{1 / 2}$ tends to level off at high values, though wide confidence intervals suggest that this effect is highly uncertain. This analysis confirms our belief that weakening events in terms of spatial extent and maximum relative wind speed are consistent with the termination of an active phase of a windstorm. Figure 6 (bottom panel) shows the spatial density of windstorm occurrence in the simulations. We see that the large-scale spatial variation of windstorms from the model reflects that of the observations in Figure 6 (top panel).

### 4.2 Validation of model for wind speeds within a footprint

We check that the model replicates well the physical structure and location of winds relative to the storm centre. For this task we simulate the wind speeds within the footprints generated in Section 4.1 using the wind model of Section 3.4. Figure 11 shows the mean, $95 \%$ quantile, $99 \%$ quantile and spatial density of simulated winds from all simulated footprints relative to the centre of the storm, quantities that were previously studied for the observed data in Figure 7. The comparison of observed and simulated winds is very good; in particular it replicates the feature
that the largest magnitudes tend to be found in the region southwest of the storm centre. The local minimum that occurs near the storm centre as a result of small pressure gradients is also captured. The upper quantiles of the spatial distribution are slightly more dispersed compared to the observed characteristics; however, we are satisfied that the large-scale features have been captured by the model. Figure 11 also shows that high magnitude events can be generated by the model on the outer edges of the domain. These are rare occurrences, but we believe these to be attributed to the detection of spurious events in the feature extraction algorithm. Section 5 discusses possible improvements to the algorithm so that the detection of spurious events is minimised.


Figure 11: The mean, $95 \%$ quantile, $99 \%$ quantile wind speed and density of within-footprint winds relative to the storm centre (represented by the cross) over a set of simulated windstorms.

By simulating windstorms relative to synthetic tracks, our model also allows us to perform an


Figure 12: QQ plots, with $95 \%$ tolerance intervals, at six locations comparing the distribution of wind speeds (in $\mathrm{m} / \mathrm{s}$ ) from the observed and simulated windstorms.

Eulerian analysis at different locations over the North Atlantic and Europe. We assess how our model captures the distribution of extreme winds at a number of locations to examine whether our approach succeeds in generating physically realistic synthetic values at these different sites. We compare our simulated values with winds that are contained within the observed footprints at each site. QQ plots for six locations, three on land and three at sea, are shown in Figure 12. The 100-year return level (not shown), estimated from our model for the spatial regions discussed in Section 4.3, compares favourably with estimates from the marginal model in (2), with an average percentage error of $2.4 \%$ over all cells. This result demonstrates that our model can be used to assess marginal risk at different locations for events beyond the range of the observational record.

### 4.3 Joint risk from windstorms

We use our approach to estimate quantities related to joint risk, that is, the probability that multiple locations are affected by extreme wind speeds simultaneously. As our model captures the spatial extent of these meteorological events, we should therefore capture the risk of multiple locations experiencing extreme wind speeds from the same storm. The results are based on 50,000 extratropical cylcone tracks generated using the track model of Sharkey et al. (2019). For each
track, a windstorm is simulated using the windstorm model of Section 3. This dataset represents approximately 600 years of data, under the assumption that there are on average 811 windstorms per year as found in the observed record.

Figure 13 (left panel) compares the joint behaviour on common $\operatorname{Exp}(1)$ margins between wind speeds in Lancaster and Manchester, two cities in northern England 73km apart. The joint correlation structure is largely captured by the model, which has the added benefit of being able to simulate joint events of magnitudes beyond the range of the observation record. One way of summarising the joint extremal behaviour of a process at arbitrary locations $s_{1}$ and $s_{2}$ is to estimate the quantity $\chi\left(q ; s_{1}, s_{2}\right)$ (Coles et al., 1999), defined as

$$
\chi\left(q ; s_{1}, s_{2}\right)=\operatorname{Pr}\left(X^{E}\left(s_{2}, t\right)>x^{q} \mid X^{E}\left(s_{1}, t\right)>x^{q}\right),
$$

where $x^{q}$ is the $100 q \%$ quantile of the common $\operatorname{Exp}(1)$ distribution. Estimates of the quantity $\chi\left(q ; s_{1}, s_{2}\right)$ obtained from the observed and simulated data are shown for a range of $x^{q}$ in Figure 13 (right panel), where $s_{1}$ and $s_{2}$ are chosen to be sites in Lancaster and Manchester respectively. Estimates from the data and the large simulated sample from the model are obtained as conditional proportions. Figure 13 shows that extremal dependence tends to decrease as the magnitude of the event increases. Here, $95 \%$ binomial confidence intervals are used to assess the uncertainty for the observed and the Monte Carlo uncertainty for the simulated data. To obtain these intervals, we use an effective sample size $n \theta\left(x^{q}\right)$ defined in terms of the sample size $n$ and a threshold-based extremal index $\theta\left(x^{q}\right)$ (Ferro and Segers, 2003; Eastoe and Tawn, 2012) to account for temporal dependence. For the model-based estimates, the confidence intervals do not represent the uncertainty due to the model parameter estimation. A fuller assessment of model uncertainty can be obtained using a parametric bootstrap, which would have the effect of widening the model-based confidence intervals. Despite not representing the full uncertainty in the model-based estimates, it is clear from the overlapping of the confidence intervals that there is little difference between data and model-based estimates of $\chi\left(q ; s_{1}, s_{2}\right)$ here over $q$ within the range of the observed data. Critically though, this figure also illustrates how our model allows estimation of $\chi\left(q ; s_{1}, s_{2}\right)$ beyond the range of the observational record, indicating that $\chi\left(q ; s_{1}, s_{2}\right)$ continues to decay to 0 beyond the observed data.

Estimating $\chi\left(q ; s_{1}, s_{2}\right)$ at a fixed critical level $x^{q}$ at a set of sites $s_{2}$ can allow us to explore the spatial extent of extreme events. Figure 14 (top panels) shows $\chi\left(q ; s_{1}, s_{2}\right)$ calculated across a number of locations in northern England, with $s_{1}$ being Lancaster in this instance. We explore two cases where $x^{q}$ is chosen to be the $90 \%$ quantile and the 10 -year return level at each site. In particular, we see that the probability surface decays more steeply as $\left|s_{2}-s_{1}\right|$ increases for the more extreme events.


Figure 13: Scatter plot (left) showing observed (red) and simulated (black) wind speeds on an $\operatorname{Exp}(1)$ scale in Manchester and Lancaster. Estimates of $\chi\left(q ; s_{1}, s_{2}\right)$ measuring extremal dependence between these locations as a function of $x^{q}$ (right) using the observed (red) and simulated (black) data, with $95 \%$ binomial confidence intervals using an effective sample size. The vertical line denotes the maximum observed wind speed on the $\operatorname{Exp}(1)$ scale.

We also see that Liverpool, Manchester and Leeds are more likely than not to experience an event on the $90 \%$ quantile simultaneously to Lancaster; however, this scenario is less likely for an event corresponding to the 10-year return level. Similarly, in Germany, as shown in the bottom panels in Figure 14, the probability of experiencing an extreme windstorm event simultaneously with Berlin decreases as the event becomes more extreme. The spatial extremal dependence is slightly stronger for Berlin than Lancaster, as might be expected given Berlin is more inland on a large land mass. In both regions, there is little evidence of anisotropy in the extremal dependence structure, with perhaps some indication of stronger dependence in the northwest-southwest direction centred at Berlin. The results in Figures 13 and 14 illustrate how the spatial extent of an extreme windstorm event becomes more localised as the magnitude increases. This result implies that in the limit, extreme values at each location tend not to occur simultaneously, which corresponds to the property of asymptotic independence. Models for asymptotic dependence, that is, when $\chi\left(q ; s_{1}, s_{2}\right) \rightarrow c>0$ as $q \rightarrow 1$ for $s_{1} \neq s_{2}$, lead to extreme values tending to occur simultaneously, are well-established but tend to over-estimate the probability of extreme joint events given that the underlying process is asymptotically independent, i.e., when $c=0$. Models that capture asymptotic independence are less well-established; see Ledford and Tawn (1996), Heffernan and Tawn (2004), Wadsworth and Tawn (2012) and Winter et al. (2016) for some examples. We have shown that our model captures the property of asymptotic independence over space, while accounting for the complex non-stationarity of the extratropical cyclone system.


Figure 14: Estimates of $\chi\left(q ; s_{1}, s_{2}\right)$ for northern England (top) and eastern Germany (bottom) conditioning on a critical value $x^{q}$, where $s_{1}$ is the cell where Lancaster (top) and Berlin (bottom) are located. In the left panels, $x^{q}$ is taken to be the $90 \%$ quantile, while $x^{q}$ is taken to be the 10 -year return level in the right panels. Both regions are of equal size.

## 5 Discussion

We have presented a novel approach to modelling windstorms in a Lagrangian framework. We described two models; first, for the evolution and development of the footprint relative to the storm track, and second, for the spatial distribution of extreme winds within the footprint. The Lagrangian framework allows us to pool information regarding events over the spatial domain being studied, which allows extrapolation of the characteristics of windstorms over space. The model provides a mechanism for generating synthetic windstorm events, the analysis of which allows improved estimation of joint risk associated with extreme windstorms over Europe.

There are, however, opportunities to improve the performance of the model. While the feature
extraction approach introduced in Section 2.3 appears to extract the main features of a windstorm, steps could be taken to improve the robustness of this procedure. Firstly, one could conduct a sensitivity study on the choice of threshold $v$, which controls the level under which the wind speed fields are masked. Ideally, one should choose a high enough value of $v$ such that convective events and non-extreme winds are masked, but small enough so that the localised features of the windstorm are not excluded. "Sting jets", meteorological phenomena associated with rapidly developing storms, can produce damaging winds on very small spatial scales (Baker, 2009; Hewson and Neu, 2015) and therefore it is important that the extraction algorithm should not exclude such features.

The feature extraction algorithm could also be improved to minimise the detection of spurious footprints that may not be generated by the extratropical cyclone. We see examples of this through high magnitude events occurring on the outer edges of the domains in Figure 7. After the spatiotemporal filtering step, our algorithm selects the largest cluster in size to define the footprint. One could alternatively define a score function that incorporates multiple criteria in a detection strategy, which could be motivated by physical intuition regarding the structure of a windstorm. For example, the score function could give a higher weight to clusters with a bearing relative to the storm centre closest to southwest, where most footprints seem to occur. To induce some smoothness between consecutive time steps, one could assign a higher weight to a cluster such that the Euclidean distance between the position of this cluster and the selected cluster at the previous time step is minimised. Exploration of different score functions could add valuable improvements to our feature extraction algorithm, and ultimately, the model.

Our conditional kernel strategy for modelling $\left\{\boldsymbol{Z}_{t}\right\}$ appears to perform well. However, we could alternatively model the extremal behaviour of $W_{t}$ using the approaches described in Winter and Tawn (2017) and Sharkey et al. (2019), whereby a GPD model is defined above a high threshold and the extremal temporal dependence structure is modelled using a $k$ th order extremal Markov process. This approach stems from the conditional multivariate extreme value methodology of Heffernan and Tawn (2004). We did not implement this approach as part of this study due to the additional complexity involved and that extrapolation occurred naturally through the features of our model. In addition, the observed values of $W_{t}$ correspond to upper tail values of $X_{i}$, the distribution of wind speeds in cell $i$. Because we have a large dataset of observations in the upper tail, we felt that standard statistical modelling approaches were sufficient in this case. However, investigation of the benefits of imposing an extremal temporal dependence structure on the upper tail of $W_{t}$ represents an interesting avenue of future research.

The potential future risk associated with extreme windstorms is of pressing concern in addition to how their characteristics might be affected by a changing climate. For the North Atlantic, previous studies have indicated the winter storm track will potentially increase over the UK and northern Europe but eliciting this signal is difficult (Zappa et al., 2013). This uncertainty, in combination with the low probability of a cyclone producing extreme winds makes estimating future windstorm risk very challenging. The windstorm model presented here and the track model of Sharkey et al. (2019) together provide a new tool that can be applied to future climate simulations to potentially provide improved estimates of such risk.

## Acknowledgements

The authors are thankful for the contributions of two referees that vastly improved the quality of the paper. The authors gratefully acknowledge the support of the EPSRC funded EP/H023151/1 STOR-i Centre for Doctoral Training, the Met Office and EDF Energy. We extend our thanks to Hugo Winter for helpful discussions and support. We thank Kevin Hodges for the storm track data. Simon Brown was supported by and Paul Sharkey partially supported by the Joint UK BEIS/Defra Met Office Hadley Centre Climate Programme (GA01101).

## References

Akhtar, N., Brauch, J., Dobler, A., Béranger, K., and Ahrens, B. (2014). Medicanes in an ocean-atmosphere coupled regional climate model. Natural Hazards and Earth System Sciences, 14:2189-2201.

Baker, L. (2009). Sting jets in severe northern European wind storms. Weather, 64(6):143-148.
Bonazzi, A., Cusack, S., Mitas, C., and Jewson, S. (2012). The spatial structure of European wind storms as characterized by bivariate extreme-value copulas. Natural Hazards and Earth System Sciences, 12(5):1769-1782.

Catto, J. L., Shaffrey, L. C., and Hodges, K. I. (2010). Can climate models capture the structure of extratropical cyclones? Journal of Climate, 23(7):1621-1635.

Coles, S. G. (2001). An Introduction to Statistical Modeling of Extreme Values. Springer.
Coles, S. G., Heffernan, J. E., and Tawn, J. A. (1999). Dependence measures for extreme value analyses. Extremes, 2(4):339-365.

Coles, S. G. and Walshaw, D. (1994). Directional modelling of extreme wind speeds. Journal of the Royal Statistical Society: Series C (Applied Statistics), 43:139-157.

Cressie, N. (1993). Statistics for Spatial Data. John Wiley \& Sons.
Dacre, H. F., Hawcroft, M. K., Stringer, M. A., and Hodges, K. I. (2012). An Extratropical Cyclone Atlas: A tool for illustrating cyclone structure and evolution characteristics. Bulletin of the American Meteorological Society, 93(10):1497-1502.

Davison, A. C. and Smith, R. L. (1990). Models for exceedances over high thresholds (with discussion). Journal of the Royal Statistical Society. Series B (Methodological), 52(3):393-442.

Dee, D. P., Uppala, S. M., Simmons, A. J., Berrisford, P., Poli, P., Kobayashi, S., Andrae, U., Balmaseda, M. A., Balsamo, G., and Bauer, P. (2011). The ERA-Interim reanalysis: Configuration and performance of the data assimilation system. Quarterly Journal of the Royal Meteorological Society, 137(656):553-597.

Della-Marta, P. M. and Pinto, J. G. (2009). Statistical uncertainty of changes in winter storms over the North Atlantic and Europe in an ensemble of transient climate simulations. Geophysical Research Letters, 36(14).

Diggle, P. J. and Ribeiro, P. J. (2007). Model-based Geostatistics. Springer.
Eastoe, E. F. and Tawn, J. A. (2012). Modelling the distribution of the cluster maxima of exceedances of subasymptotic thresholds. Biometrika, 99(1):43-55.

Ester, M., Kriegel, H.-P., Sander, J., and Xu, X. (1996). A density-based algorithm for discovering clusters in large spatial databases with noise. In Proceedings of the Second International Conference on Knowledge Discovery and Data Mining (KDD-96), pages 226-231. AAAI Press.

Fawcett, L. and Walshaw, D. (2006). A hierarchical model for extreme wind speeds. Journal of the Royal Statistical Society: Series C (Applied Statistics), 55(5):631-646.

Ferro, C. A. T. and Segers, J. (2003). Inference for clusters of extreme values. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 65(2):545-556.

Heffernan, J. E. and Tawn, J. A. (2004). A conditional approach for multivariate extreme values (with discussion). Journal of the Royal Statistical Society: Series B (Statistical Methodology), 66(3):497-546.

Hewson, T. D. and Neu, U. (2015). Cyclones, windstorms and the IMILAST project. Tellus A, $67(1): 27128$.

Hodges, K. I. (1995). Feature tracking on the unit sphere. Monthly Weather Review, 123(12):34583465.

Hoskins, B. J. and Hodges, K. I. (2002). New perspectives on the Northern Hemisphere winter storm tracks. Journal of the Atmospheric Sciences, 59(6):1041-1061.

Howard, T. and Clark, P. (2007). Correction and downscaling of NWP wind speed forecasts. Meteorological Applications, 14(2):105-116.

Joe, H. (1997). Multivariate Models and Multivariate Dependence Concepts. CRC Press.

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

Moshtagh, N. (2005). Minimum volume enclosing ellipsoid. Technical Report.

Murray, R. J. and Simmonds, I. (1991). A numerical scheme for tracking cyclone centres from digital data. Australian Meteorological Magazine, 39(3):155-166.

Nixon, M. S. and Aguado, A. S. (2012). Feature Extraction \& Image Processing for Computer Vision. Academic Press.

Pickands, J. (1975). Statistical inference using extreme order statistics. The Annals of Statistics, $3(1): 119-131$.

Ribatet, M. (2013). Spatial extremes: Max-stable processes at work. Journal de la Société Française de Statistique, 154(2):156-177.

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

Rudeva, I. and Gulev, S. K. (2011). Composite analysis of North Atlantic extratropical cyclones in NCEP-NCAR reanalysis data. Monthly Weather Review, 139(5):1419-1446.

Shapiro, M. A. and Keyser, D. A. (1990). Fronts, jet streams, and the tropopause. US Department of Commerce, National Oceanic and Atmospheric Administration, Environmental Research Laboratories, Wave Propagation Laboratory.

Sharkey, P. (2018). Statistical models for extreme weather events. PhD thesis, Lancaster University.
Sharkey, P., Tawn, J. A., and Brown, S. J. (2019). A stochastic model for the lifecycle and track of extreme extratropical cyclones. arXiv:1905.08840 [stat.AP].

Sienz, F., Schneidereit, A., Blender, R., Fraedrich, K., and Lunkeit, F. (2010). Extreme value statistics for North Atlantic cyclones. Tellus A, 62(4):347-360.

Standen, J., Wilson, C., Vosper, S., and Clark, P. (2017). Prediction of local wind climatology from Met Office models: Virtual Met Mast techniques. Wind Energy, 20(3):411-430.

Stein, M. L. (1999). Interpolation of Spatial Data: Some Theory for Kriging. Springer Science \& Business Media.

Todd, M. J. and Yıldırım, E. A. (2007). On Khachiyan's algorithm for the computation of minimumvolume enclosing ellipsoids. Discrete Applied Mathematics, 155(13):1731-1744.

Wadsworth, J., Tawn, J., et al. (2013). A new representation for multivariate tail probabilities. Bernoulli, 19(5B):2689-2714.

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

Winter, H. C. and Tawn, J. A. (2017). kth-order Markov extremal models for assessing heatwave risks. Extremes, 20(2):393-415.

Winter, H. C., Tawn, J. A., and Brown, S. J. (2016). Modelling the effect of the El Nino-Southern Oscillation on extreme spatial temperature events over Australia. The Annals of Applied Statistics, 10(4):2075-2101.

Wood, S. (2006). Generalized Additive Models: an Introduction with R. CRC press.
Youngman, B. D. and Stephenson, D. B. (2016). A geostatistical extreme-value framework for fast simulation of natural hazard events. Proceedings of the Royal Society of London Series A, 472(2189).

Zappa, G., Shaffrey, L. C., Hodges, K. I., Sansom, P. G., and Stephenson, D. B. (2013). A multimodel assessment of future projections of North Atlantic and European extratropical cyclones in the CMIP5 climate models. J. Climate, 26(16):5846-5862.

