University of Exeter
Department of Mathematics

Statistical modelling of European windstorm footprints to explore hazard characteristics and insured loss

Laura Claire Dawkins

February 2016

Supervised by
David B. Stephenson
Christopher A. T. Ferro
Ken Mylne

Submitted by Laura Claire Dawkins, to the University of Exeter as a thesis for the degree of Doctor of Philosophy in Mathematics , February 2016.

I certify that all material in this thesis which is not my own work has been identified and that no material has previously been submitted and approved for the award of a degree by this or any other University.

(signature) .....

Abstract

This thesis uses statistical modelling to better understand the relationship between insured losses and hazard footprint characteristics for European windstorms (extra-tropical cyclones). The footprint of a windstorm is defined as the maximum wind gust speed to occur at a set of spatial locations over the duration of the storm. A better understanding of this relationship is required because the most damaging historical windstorms have had footprints with differing characteristics. Some have a large area of relatively low wind gust speeds, while others have a smaller area of higher wind gust speeds. In addition, this insight will help to explain the surprising, sharp decline in European wind related losses in the mid 1990's. This novel exploration is based on 5730 high resolution model generated historical footprints (1979-2012) representing the whole European domain.

Functions of extreme footprint wind gust speeds, known as storm severity measures, are developed to represent footprint characteristics. Exploratory data analysis is used to compare which storm severity measures are most successful at classifying 23 extreme windstorms, known to have caused large insured losses. Summarising the footprint using these scalar severity measures, however, fails to capture different combinations of spatial scale and local intensity characteristics. To overcome this, a novel statistical model for windstorm footprints is developed, initially for pairs of locations using a bivariate Gaussian copula model; subsequently extended to represent the whole European domain using a geostatistical spatial model. Throughout, the distribution of wind gust speeds at each location is modelled using a left-truncated Generalised Extreme Value (GEV) distribution. Synthetic footprints, simulated from the geostatistical model, are then used in a sensitivity study to explore whether the local intensity or spatial dependence structure of a footprint has the most influence on insured loss. This contributes a novel example of sensitivity analysis applied to a stochastic natural hazards model.

The area of the footprint exceeding 25ms^{-1} over land is the most successful storm severity measure at classifying extreme loss windstorms, ranking all 23 within the top 18% of events. Marginally transformed wind gust speeds are identified as being asymptotically independent and second-order stationary, allowing for the spatial dependence to be represented by a geostatistical covariance function. The geostatis-

tical windstorm footprint model is able to quickly (∼3 seconds) simulate synthetic footprints which realistically represent joint losses throughout Europe. The sensitivity study identifies that the left-truncated GEV parameters have a greater influence on insured loss than the geostatistical spatial dependence parameters. The observed decline in wind related losses in the 1990's can therefore be attributed to a change in the local intensity rather than the spatial structure of footprint wind gust speeds.

Thank you to my supervisors, David Stephenson, Ken Mylne and Chris Ferro, for their guidance and support. Thank you for keeping me on track while allowing me to make this project my own.

Thanks also to Ben Youngman and Theo Economou for their help and support throughout. Their enthusiasm and encouragement has been invaluable, next round is on me!

Thank you to all of my collaborators on the NERC CREDIBLE project, particularly Francesca Pianosi and Thorsten Wagener for their guidance in the final chapter of this thesis, and Jonty Rougier and Jon Tawn for the extremely useful statistical discussions at group meetings. Thanks also to my collaborators on the XWS catalogue project, especially Julia Roberts for her help and guidance when working with the Met Office data.

A huge thank you to my family; Mum, Dad, Emma and Megan, for being so supportive through the ups and downs of the last 3 years. Their never-ending encouragement and praise has got me through! Thank you Mum, for the many discussions we had about my work, both over the phone and on the sofa with a cup of tea, they always helped me to organise my thoughts.

Thank you to the "Genei Table", for exuding such infinite power and wisdom, allowing me to concentrate and resist the

temptation of procrastination for at least twice as long as normal while writing up. Many thanks to my constant "Genei Table" companion, Tiggie the cat. I still find cat hairs in my laptop keyboard from time to time!

Thank you to all of my amazing friends who have kept me sane, mostly by distracting me with coffee breaks, phone calls from Philadelphia, Cheesy Tuesdays, and an endless stream of funny YouTube videos and gifs. Thank you to my housemates Caroline, Jenny and Anna and everyone I have shared an office with, particularly Nathan, James and Rachel, who have always been so supportive and ready to listen. Thanks also to Nina for continually being there when I needed to vent (and just every afternoon) with a pack of biscuits and words of encouragement. We both got our code working in the end!

Finally I'd like to thank the great brown liquid wonderment that is coffee! None of this would have been possible without you.

Contents

List of Tables

List of Figures

Publications

Material from Chapter 3 has been published in

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

1. Introduction

Weather related natural disasters cause a great deal of damage and economic loss globally, currently estimated to cost $221 billion annually (Mildenhall et al., 2014). Extratropical cyclones, also known as windstorms, are the 2nd largest cause of global insured loss (Martínez-Alvarado et al., 2014), and are a major contributor to losses in Europe. For example windstorm Kyrill (17th – 19th January 2007) and the Great Storm of October '87 (15th – 17th October 1987) caused insured losses of $6.7 billion and $6.3 billion respectively (indexed to 2012). In addition, windstorms often arrive in quick succession, increasing the risk of large aggregate losses. For example the cluster of European windstorms in December 1990, Anatol (2nd – 4th), Lothar (25th – 27th) and Martin (26th – 28th) caused an aggregate insured loss of $13.9 billion (see Table 3.1).

The severity of a windstorm hazard event is often summarised in terms of its footprint, defined as the maximum wind gust speed to occur at a set of spatial locations over the duration of the storm. Figure 1.1 shows the footprints for windstorm Kyrill and the Great Storm of October '87, over the European domain. These two windstorms incurred a similar insured loss, however their associated footprints have very different characteristics. Windstorm Kyrill has a large area of high wind gust speeds and a low peak intensity while the Great Storm of October '87 has a smaller area of high wind gust speeds but a higher peak intensity. So what is the relative importance of these characteristics in determining the insured loss that will be incurred?

One method for relating windstorm footprints to insured loss is to use a storm severity measure, a function of extreme footprint wind gust speeds. Storm severity measures represent multiple characteristics of the footprint, and can therefore be used to explore which characteristics are most damaging, e.g. the area or peak intensity. Klawa and Ulbrich (2003) developed a storm severity measure based on the assumption that losses occur locally for wind gust speeds exceeding the local 98th percentile of daily maximum wind gust speeds and grow according to a cubic function between wind and loss. Klawa and Ulbrich (2003) showed that this storm severity measure, multiplied by population density at each location, is a reasonable indicator of insured loss in Germany and has subsequently been applied in numerous studies which explore windstorm losses in Germany (Leckebusch et al., 2007; Pinto

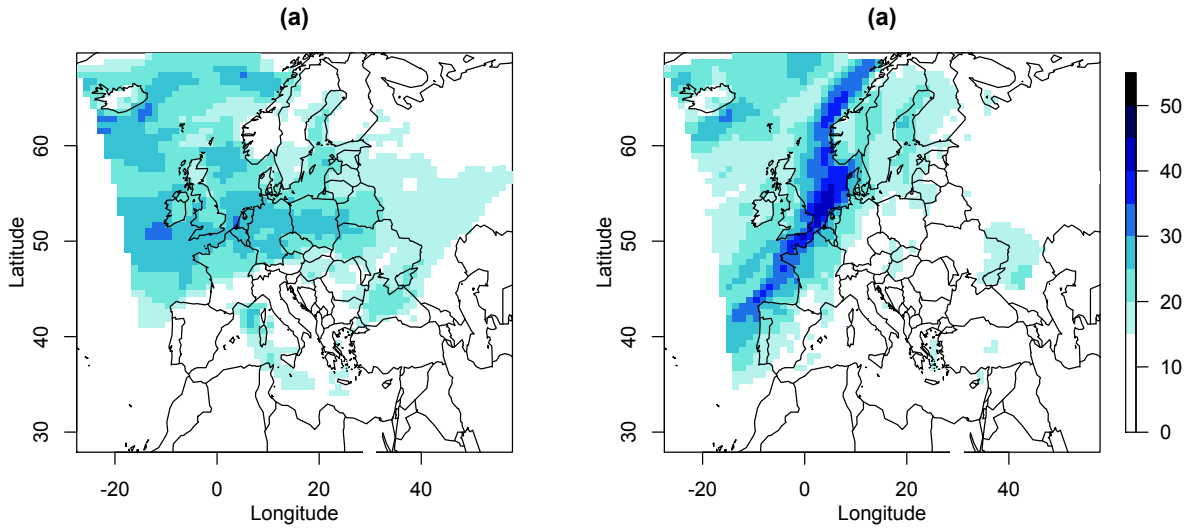

Figure 1.1. Footprints for (a) windstorm Kyrill and (b) the Great Storm of October '87. Here defined as the maximum wind gust speed (in ms^{-1}) over a 72 hour period centred on the time at which the maximum 925hPa wind speed occurs over land, estimated by downscaling ERA-Interim reanalysis using the Met Office unified model (see Section 2.3.1 for more information about the data).

et al., 2007, 2010; Donat et al., 2010, 2011a,b). A number of other storm severity measures have also been developed, such as Della-Marta et al. (2009), Haylock (2011) and Lamb and Frydendahl (1991). In all cases, the relationship between the storm severity measures and insured loss is only explored based on either small regions of Europe, annual aggregate loss data or using low spatial resolution wind gust speed data. No storm severity measure, and hence combination of footprint characteristics, has therefore been identified as being the best representation of insured loss for the whole of Europe based on high resolution wind gust data.

During an extreme windstorm event insured losses occur at multiple locations, therefore exploring the spatial structure of individual storms is also relevant to understanding how windstorm footprints relate to total insured loss. Condensing information into a scalar storm severity measure, which incorporates multiple features of the footprint into one value, inevitably fails to capture the full spatial details of storm intensity and cannot represent different combinations of spatial scale and local intensity. This motivated Bonazzi et al. (2012) to propose a model for the bivariate distribution of footprint wind gust speeds at pairs of locations, in terms of the local intensity and spatial dependence, to explore which of these characteristics has the most influence on insured loss, approximated by a storm severity measure.

There are a number of limitations to the investigation carried out by Bonazzi et al. (2012). The model is only able to represent two locations at once and is developed

based on a small set of 135 extreme historical windstorm footprints, using wind gust speeds separated by up to 100km, not covering the whole European domain. In developing this model Bonazzi et al. (2012) assumed that wind gust speeds at pairs of locations are extremally dependent, but did not test this assumption. In addition, the sensitivity of the storm severity measure to changes in local intensity and spatial dependence was only investigated using combinations of the model fitted to windstorm footprints in positive and negative North Atlantic Oscillation phases, therefore not encompassing a wide range of variation in the model parameters. A better understanding of the relationship between the footprint characteristics, local intensity and spatial dependence, and insured loss is therefore required, using better model development based on a more complete data set and a more complete sensitivity study.

This improved understanding will also help to explain the recent surprising decline in windstorm related insured losses. As explained by catastrophe modelling company RMS, windstorms Christian and Xaver in the winter of 2013-2014 are two outliers during a general lull of European windstorm activity that has lasted about 20 years (Mark, 2013). This lead RMS to release an updated European windstorm model that offered adjustments for the recent decline in windstorm severity. These adjustments ensured recent and historic wind history were represented by including the most up-to-date long-term historical wind record, going back 50 years. This

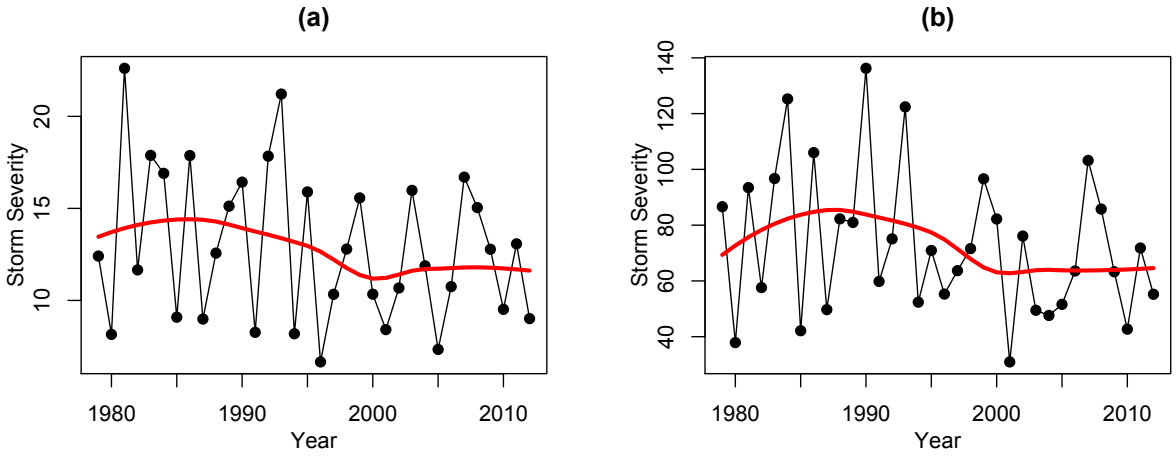

Figure 1.2. The annual (January-December) mean of (a) the storm severity measure introduced by Klawa and Ulbrich (2003), and (b) the area of the footprint exceeding 25ms^{-1}, a component of the storm severity measure developed by Lamb and Frydendahl (1991). The red curve is a locally weighted scatterplot smoothing curve, fitted to the annual means.

decrease in insured losses is also reflected in the annual averages of a range of storm severity measures. Two such storm severity measures are presented in Figure 1.2, calculated from 5730 windstorm footprints over extended winters (October-March) 1979-2012 (see Section 2.3.1). Figure 1.2 shows how both the storm severity mea-

sure developed by Klawa and Ulbrich (2003) and the area of the footprint exceeding 25ms^{-1} declined in the mid 1990's. This decrease in windstorm severity in the past two decades was also identified by Cusack (2013) when applying the Klawa and Ulbrich (2003) storm severity measure to wind gust speeds observed at stations in the Netherlands (1910-2011).

1.1. Aims and Questions

The main objective of this thesis is to develop and critically apply a statistical methodology to explore the relationship between windstorm footprint characteristics and insured loss. The modelling takes advantage of a recently developed, very large data set of 5730 reanalysis model generated historical windstorm footprints at a downscaled grid cell horizontal resolution of 25km (kindly modelled by J. Standen and provided by J. Roberts of the Met Office; see Section 2.3.1). Donat et al. (2010) identify a significant advantage in using higher resolution wind gust speed data for the predictability of insured loss.

This investigation is carried out in two stages. Firstly, by developing a number of storm severity measures, based on existing measures in the literature, and exploring which is best at representing European wide insured loss. Secondly, by developing a spatial geostatistical model for windstorm footprints over the whole European domain, which will be used in a sensitivity study to explore whether the footprint local intensity or spatial dependence has most influence on insured loss, approximated by the most successful storm severity measure.

The following questions are addressed:

- How can windstorm footprints best be summarised using storm severity measures?
- How well do storm severity measures represent insured loss?
- How should the local intensity of damaging wind gust speeds be modelled?
- How should the spatial dependence in damaging wind gust speeds be modelled?
- Can a windstorm footprint model realistically represent historical storm severity across the European domain?
- Which windstorm footprint characteristic, spatial dependence or local intensity, is most influential in causing the recent decline in storm severity?

1.2. Structure of this thesis

Chapter 2 presents a review of previously published research on windstorm footprints. Firstly the definition for the footprint is explored. Existing methods for relating the windstorm footprint to insured loss will then be reviewed, including catastrophe modelling and existing storm severity measure based explorations. Windstorm footprints can also be related to insured loss using a statistical model, hence, previously published statistical models of extreme wind speeds will also be considered. Lastly, the large data set of high resolution historical windstorm footprints used throughout the thesis will be introduced and evaluated.

In Chapter 3, the relationship between windstorm footprints and insured loss is explored using storm severity measures. The way in which the windstorm footprint can be summarised by different storm severity measures is first explored. The storm severity measures are then applied to historical windstorm footprint data to establish which are most successful at representing insured loss by exploring how well they classify notable windstorms which caused large insured losses. Some of this work has already been published in Roberts et al. (2014).

Chapter 4 presents the development and validation of a bivariate model for footprint wind gust speeds at pairs of locations. The emphasis of this chapter is the investigation of extremal dependence between locations throughout the European domain to ensure the correct class of dependence model is used. The exploration of this property is novel, therefore improving upon our current understanding of the statistical characteristics of windstorm footprints. Extreme value theory is used to model the distribution of wind gust speeds at each location.

In Chapter 5, the bivariate model is extended to a geostatistical model to represent the whole European domain. The spatial dependence is modelled using an anisotropic covariance function of separation distance. This involves resolving challenges that arise in model fitting to the large footprint data set. The resulting statistical windstorm footprint model will allow for the quick and computationally inexpensive simulation of realistic synthetic footprints.

In Chapter 6, the relationship between windstorm footprints and insured loss is explored in a sensitivity study, identifying which footprint statistical characteristic, local intensity or spatial dependence, has greatest influence on the storm severity measure that best represents insured loss. This is done by stochastically simulating synthetic footprints with variations on the model parameters, which specify the footprint statistical characteristics, and calculating the resulting storm severity. This

is the first known application of of formal sensitivity analysis methodology to a stochastic windstorm model.

2. Background

This chapter reviews previous research relevant to the methods used within this thesis. Firstly, the definition of a windstorm footprint will be discussed. Existing methods for relating the windstorm footprint to insured loss will then be reviewed, including catastrophe modelling and storm severity measures. Windstorm footprints can also be related to insured loss using a statistical model, as in Bonazzi et al. (2012). Therefore, previously published statistical models of extreme wind speeds will also be discussed. Finally, the windstorm footprint data used throughout the rest of this thesis will be described.

2.1. What is a Windstorm Footprint?

European windstorm footprints have commonly been used to represent the extremity of a windstorm event, either in the current climate, (e.g., Klawa and Ulbrich 2003, Heneka et al. 2006, Donat et al. 2010, Della-Marta et al. 2010, Donat et al. 2011b, Haylock 2011, Bonazzi et al. 2012, Cusack 2013, Roberts et al. 2014), or under future climate change conditions, (e.g., Dorland et al. 1999, Leckebusch et al. 2006, Leckebusch et al. 2007, Pinto et al. 2007, Leckebusch et al. 2008a, Leckebusch et al. 2008b, Schwierz et al. 2010, Donat et al. 2011a, Pinto et al. 2012).

There is no unique definition of a windstorm footprint. Unlike tropical cyclones, windstorms do not have a simple spatial-temporal structure; they can vary hugely in size, strength and duration. Small intense storms can be very damaging, as can large less intense storms. So how should a footprint be defined in order to appropriately summarise the intensity of a windstorm?

In the introduction, the windstorm footprint is defined as the maximum wind gust speed to occur at a set of spatial locations over the duration of the storm. The exact definition of the windstorm footprint in terms of the meteorological variable used to represent the maximum wind intensity, the event duration and the spatial domain size and resolution varies a great deal in previous studies, depending on the aim of the study and the available data. These variations will now be explored.

2.1.1. Meteorological variable

The maximum wind gust speed has been shown to have a robust relationship with storm damage (Klawa and Ulbrich, 2003) and is therefore most commonly used as the footprint meteorological variable. Seregina et al. (2014) defined wind gust speeds as the short-term exceedances of mean wind speeds for a certain time range. Wind gust speeds are therefore determined for a given averaging period. This period length is an important consideration and is often related to the intended application. Some applications require information about extreme values of the shortest wind gusts and therefore use a one second averaging period, while in other cases, the damaging wind gusts are those that affect an entire structure and a longer five-ten second averaging period is used. For windstorm footprints, Klawa and Ulbrich (2003) used a two second wind gust averaging period, while other studies such as Bonazzi et al. (2012) and Roberts et al. (2014) used a three second averaging period.

The maximum wind gust speed can be provided by observational weather stations or estimated from numerical atmospheric models. Observational station data is generally used when the region of interest is relatively small. For example Dorland et al. (1999) investigated storm damage under climate change in only one country, the Netherlands, and used station data from the Royal Dutch Meteorological Office (KNMI) for 30 locations. Similarly Bonazzi et al. (2012) focused on a small region of North West Europe. They used a large number of weather station data sources to obtain a fairly homogenous spread across Europe with typical distance between stations of around 50 to 100 km in the domain, with higher station densities in more urbanised areas such as in Germany, the UK, and France. Della-Marta et al. (2009), however, argued that wind gust speed observations generally have limited spatial resolution and/or spatial representativeness and contain temporal inhomogeneities. For example Klawa and Ulbrich (2003) developed a windstorm loss model for Germany and were only able to use 24 of the 90 German stations due to inhomogeneities found in the data at the other 66 stations. The resulting locations of the station data therefore under represented the North East of Germany.

To overcome potential inhomogeneities in station data, a number of studies have used simulated output from numerical atmospheric models, e.g. Regional Climate Models (RCMs) or Global Climate Models (GCMs). Using atmospheric model data also allows for projection under future climate scenarios. Della-Marta et al. (2009) discussed how atmospheric models are required for the temporal and spatial homogeneity needed for a continental-scale overview of the European storm climate. A drawback of using these atmospheric models is that the physical processes related to wind gust speeds are generally not explicitly resolved and must therefore be calculated as a model diagnostic using gust parameterisation schemes (Donat et al.,

2011a). There are several physical mechanisms that may lead to strong near-surface wind gusts, reflected in different parameterisations of numerical models. Most approaches have in common that the gust parameterisation is influenced by the large-scale wind, the surface roughness length, and the atmospheric stability (Seregina et al., 2014).

Gust parameterisation can lead to uncertainties in modelled wind gust speeds. As explained by Haas et al. (2014), model simulated wind gust speeds usually differ from observations because regional climate models have biases and cannot capture all local effects. Leckebusch et al. (2006) used an ensemble of four GCMs to explore the frequency and intensity of European windstorms under climate change conditions. They found that the model that contained a gust parameterisation tended to over estimate the number of extreme events in the historical period compared to the other three models which only provided maximum wind speeds. Della-Marta et al. (2009) used the ERA-40 reanalysis dataset to explore return periods of windstorms over Europe. They illustrated that even in reanalysis data, which is thought to provide physical consistency with in situ wind speed observations through data assimilation, wind gust speeds, calculated using parameterisation schemes, still exhibited unrealistic behaviour near coasts and steep topography.

Born et al. (2012) showed that different parameterisation schemes can sometimes lead to differences of up to 10-20 ms^{-1} in the estimated wind gust speeds at particular sites. Roberts et al. (2014) explained that a commonly used alternative method for predicting wind gust speeds is to use the maximum wind speed at the vertical levels from which momentum may be transported to the surface (e.g. Brasseur 2001). Roberts et al. (2014) argued that this method is more physically based than using a gust paramerisation, although it is unclear whether this method adds a significant improvement to the wind gust speed estimates (Sheridan, 2011).

Haas et al. (2014) developed a dynamical and statistical downscaling method to adjust RCM simulations of wind gust speeds toward observations to help correct for the bias in the wind gust parameterisation. Statistical downscaling is a method which relates large-scale variables to local variables via statistical transfer functions, while dynamical downscaling combines large-scale reanalyses or GCM data with RCMs resulting in high-resolution simulations over a region of interest (Haas et al., 2014). The statistical-dynamical downscaling method used by Haas et al. (2014) was a combination of these two techniques. They showed an improvement for 88 of 100 events considered, at about 64% of test sites when the downscaling method was applied, indicating some potential in this method for improving the wind gust speed fields used to create windstorm footprints.

Alternatively, another wind-related variable can be used in place of wind gust speed. Della-Marta et al. (2010) used geostrophic wind speed at the 850hPa pressure level. They, however, acknowledged that 850hPa geostrophic wind speed is not fully representative of the wind gust speeds and hence wind-related loss, and in using 850hPa geostrophic wind speed they lost information about real surface winds affected by, for example, surface roughness. However their justification for this choice of meteorological variable was that they gained a wind climatology that was independent of gust parameterisations, which could be used in model inter-comparison.

Some GCMs do not include a gust parameterisation. As a result Pinto et al. (2007), Donat et al. (2010), Donat et al. (2011a) and Donat et al. (2011b) used the maximum instantaneous 10 metre wind speed as an alternative footprint meteorological variable. Similarly, due to the 100 year long historical record required, Cusack (2013) analysed daily maximum wind speeds, rather than wind gust speeds since wind gust speeds have only recently been recorded. Della-Marta et al. (2009), however, argued that instantaneous wind fields do not necessarily sample maximum storm intensity. To overcome this problem, Seregina et al. (2014) presented a statistical methodology for predicting wind gust speeds from wind speed data using a transfer function between the distribution parameters of wind and wind gust velocities. They found that, despite uncertainties caused by the short length of the observational records, the method lead to consistent results, enabling a wide range of possible applications.

2.1.2. Spatial domain and resolution

The choice of spatial domain and resolution largely depends on the application and aims of the given windstorm footprint investigation. Applications that require the use of GCMs for climate change studies are restricted by their relatively low resolution. For example Leckebusch et al. (2008b) used a GCM with a low horizontal grid point spacing of 2.5° and Leckebusch et al. (2008a) used a 1.87° horizontal resolution model. Similarly the four GCMs used by Leckebusch et al. (2006) had resolutions ranging between 2° to 3.75°.

Donat et al. (2010) explored the benefits and limitations of dynamical downscaling for windstorm loss calculations and identified a significant advantage in using higher resolution wind gust speed data for the predictability of insured loss. In addition, Della-Marta et al. (2010) commented on how estimates of European windstorm climate and their associated losses are often compromised by either relatively short records, coarse resolution or inhomogeneous datasets. They used a dynamically downscaled GCM with 0.5° horizontal resolution to improve estimates of European

windstorm climate and the risk of insured loss. Similarly, Schwierz et al. (2010) dynamically downscaled GCM wind gust speed data to a 30km horizontal resolution and Haylock (2011) to a 0.22° resolution.

The spatial domain of the footprint, again varies depending on the application. A number of studies are based solely on Germany, for example Klawa and Ulbrich (2003), Leckebusch et al. (2007) and Donat et al. (2010), while some extend to west and central Europe, for example Pinto et al. (2007), Leckebusch et al. (2008a), Leckebusch et al. (2008b), Haylock (2011) and Bonazzi et al. (2012), and others to the whole of Europe, for example Schwierz et al. (2010) and Della-Marta et al. (2010).

2.1.3. Event duration

Studies use different definitions of the windstorm event duration and methods for identifying the maximum wind intensity at each location over the duration period. Leckebusch et al. (2008b) noted that the typical timescale of a synoptic situation is from 2.5 up to 6 days and therefore concluded that 3 days is a reasonable time duration for footprints. The 3 day window is the most commonly used in previous studies and is consistent with the reinsurance industry definition of a windstorm event (Haylock, 2011). Haylock (2011) argued that 72 hours is enough time for a storm to completely pass over western Europe, however also acknowledged that there is a possibility that a single 72 hour footprint may result from more than one windstorm crossing the domain. They explained how choosing a shorter period, however, increased the risk of one windstorm generating more than one footprint, violating the independence between events required to carry out statistical analysis. Since Haylock (2011) used daily maximum wind gust speed data, the maximum of these daily maxima over 4 calendar days was used to ensure that any 72 hour period centred around a particular day was captured. Della-Marta et al. (2009) also used a 72 hour windstorm event duration, and centred the 3 day period on the documented date of the event in the windstorm catalogue developed by Lamb and Frydendahl (1991). Lamb and Frydendahl (1991) did not explain how the date of each event was chosen, however Della-Marta et al. (2009) suggested they are based on the time during which the greatest impacts were experienced. Haas and Pinto (2012) followed the same approach, defining the 72 hour period as one day before and after the event date, where the event dates were calculated as the top 100 days in a storm severity measure in the time period of interest. They noted that this approach resulted in some storms being counted twice, for example two of the top 100 days was associated with windstorm Kyrill.

Bonazzi et al. (2012) also identified event days by selecting the top N days in a storm severity measure. However, they chose the top N days for each country separately, e.g. the top 25 for major countries such as Germany and the top 10 for smaller countries such as Slovakia, to ensure that their sample of footprints represented storms that affected all countries in their European domain. Rather than taking a day either side of the event day, they used a range of sources to identify the peak wind gust speed for each event: (i) daily peak wind gust speed time-series at stations, (ii) Reanalysis data of surface pressure, (iii) German Weather Service (DWD) operational analyses of surface pressure and (iv) reports of damage in the public domain such as media and the internet. This meant that the event duration was not fixed and varied between windstorm events. Cusack (2013) also used a flexible event duration, taking all consecutive days with a non-zero value of a storm severity measure to be one event. They did this to reduce the possibility of double counting events, but based on the above discussion, this method seems likely to result in consecutive or overlapping events being represented in one footprint.

Donat et al. (2011b) used a longer, 5 day event duration. They argued that this is a more reasonable time interval to allow for typical travelling velocities and for sequences of storms that hit Germany within a few days. They however, noted that the most significant losses related to a windstorm usually occur within 1 or 2 days and suggested that there is greater consistency between windstorm footprints if a 3 day event duration is used rather than 5 days. They found that a 5 day event duration resulted in some consecutive storms being considered as one event, but that a 3 day event duration resulted in some windstorms being excluded because they defined individual storm events as being separated by at least one day. Pinto et al. (2012) acknowledged that, while it may take 3 days for a windstorm to cross the whole of Europe, a 2 day event duration was appropriate for their study which was based on a smaller region in North-West Europe which would therefore take less time for the storm to pass through.

In contrast, Leckebusch et al. (2007), Pinto et al. (2007) and Donat et al. (2010) studied the effect of climate change on annual aggregate wind related losses in Germany. They therefore considered maximum wind gust speed footprints for each single day, and summed the resulting storm severity over the year.

The way in which windstorm footprints are defined and constructed within this thesis, in terms of the meteorological variable, region and event duration used, is described in Section 2.3.1.

2.2. Relating windstorm footprints to insured loss

This thesis aims to explore how windstorm footprints relate to insured loss. Previous approaches for exploring this relationship will be reviewed here. Firstly in the context of catastrophe modelling, which is carried out by insurers to estimate windstorm losses and price insurance by combining windstorm footprints with insured exposure and vulnerability data. Storm severity measures will then be discussed, these are commonly used to explore historical and future trends in windstorm losses when exposure is unavailable and loss is estimated as a function of extreme footprint wind gust speeds. Finally, previously published univariate and multivariate statistical models of extreme wind speeds will be reviewed.

2.2.1. Catastrophe modelling

Catastrophe models combine mathematical representations of the natural occurrence patterns and characteristics of natural hazards, such as windstorms, with information on property values, construction types and occupancy classes to provide information about the potential for large losses before they occur. The purpose of catastrophe modelling is therefore to estimate the chance of potential future catastrophic events so companies can prepare for the financial impact. Catastrophe modelling is used in the insurance and reinsurance industries for insurance pricing, risk selection, loss mitigation, reinsurance decision making and overall risk portfolio management for natural hazards (Clark, 2002).

The hazard module of the catastrophe model should represent where the natural hazard events are likely to occur, how severe they are likely to be and their relative frequency. In the case of European windstorms, the hazard module is represented by the windstorm footprint (Schwierz et al., 2010). Using windstorm footprints from historical events alone does not give a representative or comprehensive selection of all possible events that could occur. Therefore, catastrophe modellers generate synthetic events to represent all types of possible, yet realistic, scenarios. Large catalogues of synthetic footprints are used to determine the frequency, magnitude and any other characteristics of potential catastrophe events needed within the hazard module (Clark, 2002).

These synthetic event sets are traditionally created based on low-order parametric stochastic models (Beven et al., 2016). More recently, however, numerical weather and climate models are used to simulate large sets of artificial hazard events. The biases in the model output are adjusted using statistical models, and stochastic models are used for simulating losses from the artificial windstorm events, e.g. compound-

Poisson event-loss table models (Beven et al., 2016).

A number of assumptions are made when developing these models which leads to a great deal of uncertainty in the resulting estimated potential losses. Stochastic hazard and loss models often use highly idealised non-physical descriptions of complex storm processes (Beven et al., 2016). There is also the risk that the models are over-fit to data sets that cover a relatively short historical period. For example, the RMS European windstorm catastrophe model overestimated recent losses due to the data period used to develop the model (Mark, 2013). In addition, numerical weather and climate models show biases in storm properties, such as underestimation of extreme winds (Roberts et al., 2014), incorrect locations of storm tracks over Western Europe (Zappa et al., 2013) and missing processes such as sting jets (Catto et al., 2010). When relating windstorm footprints to insured loss, re/insurers, however, benefit from the availability of insured exposure and vulnerability data at a high spatial resolution. Contrary to this, most commonly scientists in academia only have access to very limited records of annual aggregate or event based insured losses. In these studies, insured loss must be approximated using a storm severity measure, or loss function (Klawa and Ulbrich, 2003), calculated from the footprint wind gust speeds, and in some cases population density, presenting a further source of uncertainty when relating footprints to insured loss.

2.2.2. Storm severity measures

Storm severity measures are descriptive statistical summaries of the severity of a windstorm event, making idealised assumptions about the hazard, vulnerability and exposure in order to estimate insured loss. It is therefore important to validate the skill of a storm severity measure before using it as an approximation for loss. This skill depends on many factors such as the footprint data used, the availability of insured loss data for validation and the time period and region studied. As a consequence previously published storm severity measures vary in composition, for example in terms of the characteristics used to summarise the footprint, e.g. damage area and duration. Exploring how well these different measures represent insured loss therefore gives insight into which characteristics of the footprint are most damaging and how the windstorm footprint is best related to insured loss. The composition of previously published storm severity measures will be reviewed, followed by their skill in representing insured loss. This review will inform the development of storm severity measures in Chapter 3 of this thesis.

Storm severity measure composition

Although storm severity measures vary in the detail of their composition, it is possible to express most in a general form. For a given windstorm event, a severity measure m can be expressed as:

$$m = g\left(\sum_{j=1}^{J} V(x(s_j), t)e(s_j, t)\right) \tag{2.1}$$

where $x(s_j)$ is the footprint wind gust speed at location $s_j, j = 1, ..., J$. The functions V and e represent the hazard and exposure respectively, which both depend on the wind gust exceeding a damage threshold t. These two functions together, summed over all locations, represent the severity of a given windstorm event. In some cases an additional function g is applied to the severity measure to reduce skewness. All previously published storm severity measures are special cases of this general form. These severity measures will first be introduced and discussed, and finally summarised based on this general form in Table 2.1.

The first storm severity measure was derived by Lamb and Frydendahl (1991) while compiling a catalogue of extreme storms over the North Sea, British Isles and north-west Europe. Lamb and Frydendahl (1991) included storms which had "acquired historical note" and used the severity measure to rank these storms within the catalogue. For a given storm event, their severity measure took the form

$$\text{loss} \approx V_{max}^3 AD \tag{2.2}$$

where V_{max}^3 is the cubed maximum surface wind speed throughout the domain of interest, A is the greatest area affected by damaging winds and D is the overall duration of the damaging winds.

From a theoretical point of view, the cube of the wind speed is proportional to the damaging power (Klawa and Ulbrich, 2003), i.e. the rate of kinetic energy delivered by advection. Lamb and Frydendahl (1991) explained how the dynamic pressure of the wind is proportional to the square of the wind speed, while the wind power is a matter of work done and therefore involves the dynamic pressure and the run of the wind, hence the cube of the wind speed. Empirically this relationship was supported by MunichRe (1993), who found that the loss extent of windstorms increased with almost the cube of the maximum wind gust speeds ($x^{2.7} \sim$ loss). Later explorations carried out by Munich Re, following the windstorm series of December 1999, generally confirmed this relationship, however it was suggested that higher exponents such as x^4 and x^5 were needed to represent the losses in areas of very extreme wind

gust speeds (MunichRe, 2002).

The area and duration parameters, A and D, used within this severity measure require the definition of a damage threshold. A threshold of 50 knots (25.7ms^{-1}) was chosen, since, as explained by Lamb and Frydendahl (1991), wind speeds of 38-44 knots (19.5-22.6ms^{-1}) were known, in the UK, to damage chimney pots and branches of trees and wind speeds of 45-52 knots (23.1- 26.8ms^{-1}) to uproot trees and cause severe damage to buildings.

Following this, Klawa and Ulbrich (2003) developed a storm severity measure, or loss function, for the estimation of windstorm losses and the identification of severe winter windstorms in Germany:

$$\text{loss} \approx \sum_{j=1}^{J} p(s_j) \left(\frac{x(s_j)}{c(s_j)} - 1 \right)^3 \quad \text{for } x(s_j) > c(s_j) \tag{2.3}$$

where $x(s_j)$ is the footprint wind gust speed at location s_j, $c(s_j)$ is the damage threshold for location s_j, the 98$^{\text{th}}$ percentile of the wind climatology (daily maximum wind gust speed during the analysis period), and $p(s_j)$ is the population density at location s_j.

Similarly to Lamb and Frydendahl (1991), Klawa and Ulbrich (2003) used the cubic relationship between wind gust speed and loss. However, rather than the cube of the maximum wind intensity over the whole domain, Klawa and Ulbrich (2003) used the cube of the excess wind gust speed, above a damage threshold, summed over all locations. They noted that the use of the cube of the excess implied that a small increase in wind speed above the relative threshold would have a large effect on the severity measure, making their approach different from Lamb and Frydendahl (1991). The excess wind has most commonly been used in subsequent studies due to the success of the Klawa and Ulbrich (2003) measure at representing insured loss values in Germany (discussed further in Section 2.2.2).

Klawa and Ulbrich (2003) used a relative rather than absolute damage threshold, different for each location. They argued that both buildings, due to enforced building regulations, and nature (e.g. forests) adapt to the local wind conditions and, as a result, the damage threshold should vary throughout the domain. Klawa and Ulbrich (2003) noted that insurance companies in Germany pay for storms when maximum gusts are above 20ms^{-1} and consequently found that this value coincided with the 98th percentile of daily maximum gust wind speed in German flat land stations. They therefore used this relative 98$^{\text{th}}$ percentile of the climatology in their storm severity measure, meaning that damages occurred in each location in the top

2% of all days. They explained how this is in line with the argument used by Palutikof and Skellern (1991) for their definition of storm vulnerability areas, based on return periods of wind extremes at British weather stations, that at any location (irrespective of the environmental characteristics) storm damages are assumed to occur at 2% of days (Klawa and Ulbrich, 2003). In addition, Klawa and Ulbrich (2003) noted that the same excess wind gust speed will lead to less/more damage when the damage threshold value is high/low. They therefore used winds scaled with the relative damage threshold value to ensure all locations had the same importance within the severity measure.

In addition, Klawa and Ulbrich (2003) included a parameter for population density as a proxy for exposure in areas affected by damaging winds. In a later study Donat et al. (2011a) explained how population density is required as a proxy for insured values because information about the spatial distribution of total insured values is usually seen as confidential by insurance companies. Donat et al. (2011a) argued that insured values are proportional to population density for developed countries in Central and Western Europe because most people are insured in these regions.

The area of damaging winds is indirectly represented in the Klawa and Ulbrich (2003) measure since the severity is a sum over all damaging locations. Klawa and Ulbrich (2003) discussed the inclusion of a duration parameter, as used by Lamb and Frydendahl (1991), which has been shown to have a positive correlation with loss by Swiss Re (Klawa and Ulbrich, 2003). Klawa and Ulbrich (2003), however, concluded that, while it is possible to define storm duration it is not easy to estimate the effect of the duration on the loss sum. In most studies the footprint wind gust speed used within the storm severity measure is defined as the maximum wind intensity over a fixed duration (e.g. 72 hours). Therefore, much like area, the duration is indirectly included in this storm severity measure.

Klawa and Ulbrich (2003) also discussed the inclusion of precipitation as a parameter, noting that this variable might increase losses in single cases. They concluded, however, much like storm duration, it is hard to define the effect of precipitation on loss in an objective way. Klawa and Ulbrich (2003) noted that the number of parameters used within a storm severity measure should be minimised since the use of too many parameters may lead to overfitting due to the limited datasets used, resulting in a meaningless severity measure.

In a majority of subsequent studies, storm severity measures follow the same composition as Klawa and Ulbrich (2003), namely Leckebusch et al. (2007), Pinto et al.

(2007), Leckebusch et al. (2008b) and Pinto et al. (2012), who used storm severity measures to explore how windstorm related losses will change under climate change conditions for Germany and regions of Europe. In addition, Donat et al. (2010) used the same storm severity measure to explore the benefits and limitations of regional multi-model ensembles for storm loss estimations, Donat et al. (2011a) to investigate future changes in windstorm losses and extreme wind speeds inferred from GCM and RCM multi-model simulations and Donat et al. (2011b) to explore the high-resolution refinement of a storm loss model and estimation of return periods of loss-intensive storms over Germany.

Cusack (2013) developed a 101 year historical record of windstorms for the Netherlands using a variation on the Klawa and Ulbrich (2003) storm severity measure, which took the cube root of the original measure, divided by the number of locations used

$$loss \approx \sqrt[3]{\frac{\sum_{j=1}^{J}\left(\frac{x(s_j)}{c(s_j)}-1\right)^3}{J}} \quad \text{for } x(s_j) > c(s_j)$$

where $c(s_j)$ is the 99^{th} percentile of the climatology wind gust speed at location s_j.

Cusack (2013) did not include the population density as a measure of exposure since assuming uniform exposure allowed for better identification of the variability in storminess due to meteorological factors. Cusack (2013) used the relative 99^{th} percentile of the climatology wind gust speed, arguing that the 98^{th} percentile threshold resulted in too many weak storms with uncertain damage characteristics being included in their analysis. The 99^{th} percentile was therefore used to increase confidence in the identification of damaging storms. Cusack (2013) found, however, that the conclusions of their study were insensitive to this change in threshold.

Contrary to Klawa and Ulbrich (2003), Haylock (2011) argued that there was little evidence of variation in country wide wind climatology and, since their storm severity measure was used to find the largest storms, decided that this variation was not critical within a storm severity measure. Consequently Haylock (2011), and similarly Bonazzi et al. (2012), used absolute damage thresholds, constant throughout the domain. Haylock (2011) used a threshold of 11ms^{-1}, equivalent to the mean 90^{th} percentile wind speed for all windstorm footprints in their data set. Bonazzi et al. (2012) used a threshold of 20ms^{-1}, previously identified by Klawa and Ulbrich (2003) as a suitable threshold for German insurance property losses and used in the measure that constituted the basis of the first European windstorm parametric index Catastrophe Bond undertaken by RMS in 2000 (Bonazzi et al., 2012).

The storm severity measures developed by Haylock (2011) and Bonazzi et al. (2012) respectively are defined as

$$loss \approx \sum_{j=1}^{J} \cos(\theta(s_j))(x(s_j) - 11)^3 \quad \text{for } x(s_j) > 11$$

and

$$loss \approx \sqrt[3]{\frac{\sum_{j=1}^{J}(x(s_j) - 20)^3 \cdot A(s_j)}{\sum_{j=1}^{J} A(s_j)}} \quad \text{for } x(s_j) > 20 \tag{2.4}$$

Rather than applying the storm severity measure to spatially uniform data, Bonazzi et al. (2012) used regions of different sizes and hence multiplied the severity within each region by the area of that region, $A(s_j)$. Haylock (2011) applied their severity measure to gridded atmospheric model data. Therefore, in a similar way, they multiplied the severity at each location by its latitude, $\theta(s_j)$, since the grid cell size will depend on its proximity to the poles.

Della-Marta et al. (2009) presented a number of footprint storm severity measures to explore the return periods of extreme windstorms. These measures ranged in complexity, with the most simple being the mean and 95% quantile footprint wind gust speed over all locations for a given event. The simple severity measures provide an interesting comparison with the more complex measures defined above.

These previously published storm severity measures can be summarised based on the general form introduced in Eqn. (2.1), shown in Table 2.1. The Della-Marta et al. (2009) 95% quantile measure is not included because it is not easily expressed in this general form.

Validation of storm severity measures

The storm severity measures in previously published studies characterise different properties of the footprint. Therefore exploring how well these measures represent insured loss identifies which characteristics are most damaging and therefore the best way to summarise the footprint to approximate insured loss.

Some studies such as Haylock (2011) and Lamb and Frydendahl (1991), did not validate how well their severity measure represents insured loss because the measure was used to rank extreme historical storms and not to represent loss. Bonazzi et al. (2012) showed no validation results, only stating that their storm severity measure, applied to station data covering North-West Europe, correlated closely

Table 2.1.. Definitions of previously published storm severity measures, summarised based on the general form introduced in Eqn. (2.1), where $H(n)$ is a Heaviside/indicator function such that $H(n) = 1$ if $n > 0$ and $H(n) = 0$ otherwise.

Study	g	V	e	t	other parameters
Lamb and Frydendahl (1991)	-	$H(x(s_j) - t)$	-	25.7ms^{-1}	$\max_{j=1,\ldots,J}(x_j)^3$, D
Klawa and Ulbrich (2003)	-	$H(x(s_j) - t)\left(\frac{x_i(s_j)}{t} - 1\right)^3$	$p(s_j)$	$c(s_j)$	-
Cusack (2013)	$\sqrt[3]{V/J}$	$H(x(s_j) - t)\left(\frac{x_i(s_j)}{t} - 1\right)^3$	-	$c(s_j)$	-
Haylock (2011)	-	$H(x(s_j) - t)\{\cos(\theta(s_j))(x(s_j) - t)^3\}$	-	11ms^{-1}	-
Bonazzi et al. (2012)	$\sqrt[3]{V/\sum_{j=1}^{J} A_j}$	$H(x(s_j) - t)(x(s_j) - t)^3 \cdot A_j$	-	20ms^{-1}	-
Della-Marta et al. (2009)	-	$x(s_j)$	J^{-1}	-	-

with aggregated storm damages. Similarly, Della-Marta et al. (2009) commented on how their spatial 95% quantile wind gust speed severity measure, applied to ERA-40 reanalysis data, was more representative of loss figures published in MunichRe (2001) compared to the other measures they introduce, without explicitly showing the results. Other studies, such as Cusack (2013), used their storm severity measure without validation due to the lack of available loss data, justifying the use of the measure by its similarity to other successful measures previously published.

The most rigorous validation exists for the storm severity measure proposed by Klawa and Ulbrich (2003) (Equation 2.3) since it is used in a number of other studies and has therefore been applied and validated multiple times.

The initial validation was carried out by Klawa and Ulbrich (2003) by both comparing annual aggregate insured loss values with the storm severity measure summed over each corresponding year, and insured loss values for a small selection of 11 extreme windstorm events with the storm severity measure for the corresponding event days. This validation was based solely on Germany due to the lack of available insured loss data for other countries. The storm severity measures were calculated using wind gust speed data from 24 stations. The event based insured losses were taken from MunichRe (1993) and MunichRe (1999) and the annual aggregate insured loss values were provided by the German Insurance Association (GDV) for the years 1980-1997. Klawa and Ulbrich (2003) commented that the event losses are based on the market insight of Munich Re and therefore may not have been representative of the whole insurance market. They also noted that the GDV annual aggregate loss data may have included indirect storm losses associated with hail and rain.

Klawa and Ulbrich (2003) found that the time series of modelled and true annual aggregate insured losses had a correlation of 0.96, demonstrating the general suitability of the model for annual aggregate losses in Germany in the period considered. They identified a small long term trend in the modelled annual insured losses based on the storm severity measure, which implied that the measure underestimated insured loss at the beginning of the period but overestimated it at the end. Klawa and Ulbrich (2003) suggested that this could have been due to the uncertainty in de-trending the insurance loss data, done to remove the effect of inflation. They found, however, that removing this bias in the trend resulted in an underestimation of some well-documented extreme windstorm events.

Klawa and Ulbrich (2003) also found a general agreement between individual event losses and those estimated using the storm severity measure for Germany. They con-

cluded that their storm severity measure is generally capable of giving a reasonable estimation of losses produced by individual storms, in spite of the simplifying assumptions made. Again, however, this analysis was solely based on Germany. The storm severity measure also greatly overestimated the loss associated with windstorm Verena (13th January 1993). Klawa and Ulbrich (2003) speculated that, since the severity measure was able to estimate the 1993 annual aggregate loss reasonably well, this was because the losses published by Munich Re were too small. A comparison of the true and modelled annual aggregate losses, taken from Klawa and Ulbrich (2003), is presented in Figure 2.1, showing how, in general, the severity measure estimates the annual aggregate losses well.

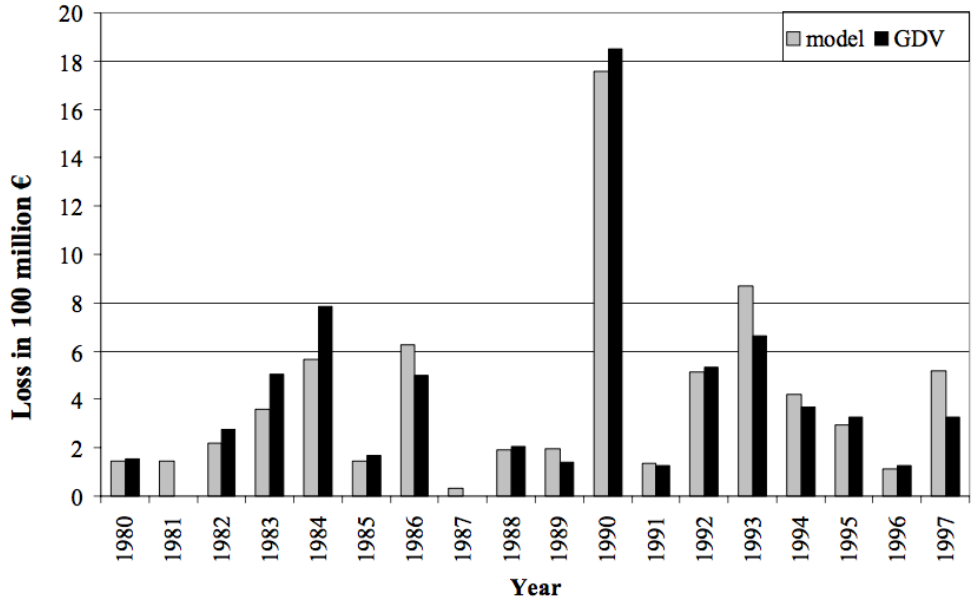

Figure 2.1. Insured annual accumulated losses in Germany. Comparison between detrended loss data (basic year = 1990) reported by the German Insurance Association (GDV), and loss estimations by the storm severity measure. Source: Klawa and Ulbrich (2003) Figure 4.

Insured losses associated with extreme storms such as Lothar (26th Decemeber 1999) and Anatol (3rd Decemeber 1999) were found to be reasonably well approximated by the storm severity measure with Lothar being overestimated by 15% and Anatol underestimated by 30%. Klawa and Ulbrich (2003) explained how this level of agreement is based on the cubic relationship between wind and loss and may have been improved further if a higher exponent was used, as was suggested by MunichRe (2002). They also identified that the loss estimation for Lothar could have been improved by including additional wind gust speed station measurements to better represent it's relatively narrow footprint.

Pinto et al. (2007) applied this storm severity measure to reanalysis and GCM data

with horizontal resolutions of 1.875° (∼ 200 km). Similarly to Klawa and Ulbrich (2003), Pinto et al. (2007) validated the estimated losses from the storm severity measure by comparing annual aggregate losses from GDV for Germany. However, rather than using the raw losses, annual aggregate insured loss values were given as the ratio between insured claims and total insured values (unit: Euros per 1000 Euros, i.e., in ‰). An advantage of this measure is that inflation, a factor that Klawa and Ulbrich (2003) found to be a potential issue, can be neglected as it is included both in insured values and in the loss (Donat et al., 2011a). Before comparing the estimated annual losses with the GDV loss ratios, Pinto et al. (2007) applied a linear regression calibration of the storm severity measure to adjust the estimates toward the German loss values. Pinto et al. (2007) identified a strong correlation between the aggregate annual loss time series of 0.87, suggesting that, after calibration, this model for insured loss is relatively successful even for low resolution GCM wind gust speed data. Donat et al. (2010) applied the Klawa and Ulbrich (2003) storm severity measure to RCM data of different horizontal resolutions ranging from 120km down to 25km to explore the benefits of regional multi-model ensembles for storm loss estimations. In the same way as Pinto et al. (2007), Donat et al. (2010) calibrated the estimated annual aggregate losses calculated by the storm severity measure using a linear regression calibration method based on German GDV annual loss ratios. They found large deviations in modelled losses, in comparison to observations, for the years 1972 and 1993, due to over or under estimations of specific storms such as Verena. Donat et al. (2010) found that for all RCM resolutions the mean annual loss ratio for Germany is around 0.15‰, as in the GDV data. However, as explained by Donat et al. (2010), this agreement was to be expected since raw severity measures were calibrated towards observed loss ratios for each individual RCM. Donat et al. (2010) identified that the inter-annual variability in terms of the standard deviation of annual loss ratios was lower compared to observed loss variability in almost all models. They found that the agreement between observed and estimated loss ratios was slightly higher for the 25km resolution RCM compared to the coarser resolution models, again based on Germany only.

Donat et al. (2011b) explored how well the Klawa and Ulbrich (2003) measure represented 34 individual windstorm events in Germany. The loss ratios of the 34 events were provided by German Verbundene Wohngebäude Versicherung (VGV). Donat et al. (2011b) applied the severity measure to two reanalysis data sets with horizontal resolutions of 79km × 79km and 180km × 275km over Germany. They applied a linear regression calibration to estimate the loss of an event based on the other 33 events in the data set. They found that, in general, the losses calculated from the different reanalysis models reproduced the insurance loss data well. However, for windstorm Kyrill, the losses were underestimated by both wind datasets by

about 15 to 20% in comparison to the observed loss records. Donat et al. (2011b) explained how this underestimation may have been due to the regulation practises of insurance companies for which loss adjustments may be applied. Donat et al. (2011b) acknowledged how these company-specific and socio-economic effects were not accounted for by the storm severity measure approach. Donat et al. (2011b), however, found that the spatial distribution of losses in the different administrative districts were reproduced well by the storm severity measure. They concluded that, in general, the results confirmed the ability of the measure for realistically calculating both country-wide accumulated losses and the spatial distribution of losses for severe storm events in Germany, particularly when higher resolution reanalysis data was used.

The ability for the presented storm severity measures to represent insured loss has therefore either been validated in a non-transparent way, validated against a very small number of windstorm events or, in the case of the Klawa and Ulbrich (2003) measure, validated rigorously, but based on only German insurance loss data. No storm severity measure has therefore been identified as being the best representation of insured loss for the whole of Europe based on high resolution wind gust data. In the same way, the best footprint characteristics for representing European wide insured loss are not known, although maximum wind gust speed cubed, population density, damage area and duration seem to be the most popular parameters for representing insured loss in practise. This insight will be used to develop storm severity measures to explore the relationship between windstorm footprints and insured loss in Chapter 3 of this thesis.

2.2.3. Statistical models of extreme wind speeds

Within this thesis a novel spatial geostatistical model for windstorm footprints will be developed. This model will allow for the simulation of synthetic footprints over the whole European domain. The parameters of the statistical model will represent the footprint characteristics of interest and will be varied within a rigorous sensitivity study to explore their relationship with insured loss. Existing approaches for modelling extreme wind speeds will now be explored to inform the model development within this thesis.

Univariate modelling of extreme wind speeds

Justus et al. (1977) explored methods for estimating wind speed distributions and proposed the use of the Weibull distribution. Following this, several other studies suggested adopting the Weibull distribution for this purpose (Haas et al., 2014).

For example, Haas et al. (2014) used the Weibull distribution to model both daily maximum wind speed and wind gust speeds associated with selected 3-day windstorm events, when developing a probabilistic approach for downscaling windstorm footprints. Similarly Seregina et al. (2014) used the Weibull distribution to model hourly mean and maximum wind speeds at 123 measurement stations.

Friederichs et al. (2009), however, pointed out that, by classical Extreme Value Theory (EVT) (Coles, 2001), the distribution of maximum wind speeds should follow the Generalised Extreme Value (GEV) distribution and exceedances of a large threshold should follow a Generalised Pareto Distribution (GPD). Friederichs et al. (2009) compared several probabilistic approaches for modelling wind gust speeds in Germany. They modelled wind gust speed data at 139 stations and concluded that the GEV is the most appropriate and theoretically consistent statistical model. Supporting this, Perrin et al. (2006) found that using the Weibull distribution leads to incorrect estimates of the tails of the distributions of extreme wind speeds and that using the GEV distribution overcomes this issue. The GEV model was also used by Payer and Küchenhoff (2004) to successfully model annual maximum wind speeds at stations in Germany, and Bonazzi et al. (2012) used the GPD to model the marginal distribution of extreme excesses in wind gust speeds at pairs of locations within their bivariate windstorm footprint model.

In addition, EVT has also been applied to model the extremes of other windstorm related variables. For example Lionello et al. (2008) fitted a GEV distribution to monthly pressure minima derived from three different climate models over the North Atlantic domain; Della-Marta and Pinto (2009) fitted a GPD to extreme central pressure and maximum vorticity, in three non-overlapping regions of Europe; and Sienz et al. (2010) used a GPD to fit a tail model to extreme values of geopotential height, mean horizontal gradient, cyclone depth and relative vorticity, for storms within a large Atlantic region.

Multivariate modelling of extreme wind speeds

Literature on modelling wind speeds beyond the univariate setting is extremely limited, with most applications of multivariate extreme value modelling based in other areas of environmental research, for example hydrology (Renard and Lang, 2007; Coles and Tawn, 1996), oceanography (Bortot et al., 2000) and surface-level ozone (Eastoe, 2009; Heffernan and Tawn, 2004).

Extremal dependence is a specific mathematical feature of all multivariate probability distributions. For pairs of random variables, it describes the joint behaviour

of two variables as they become large (Coles et al., 1999). When modelling multivariate extremes, correctly modelling the extremal dependence is very important to ensure the relationship between variables in the upper tail of the joint distribution is represented correctly. Extremal dependence can take one of two forms, asymptotic dependence, where large values of the variables tend to occur simultaneously, or asymptotic independence, where the largest values rarely occur together (Coles et al., 1999). Eastoe et al. (2013) explained how making an incorrect assumption on the form of the asymptotic dependence can lead to invalid inferences. Wadsworth et al. (2015) went on to explain how most available statistical models for multivariate extremes are suitable for either one case or the other, but not both. This highlights the importance of correctly identifying this property.

Bonazzi et al. (2012) proposed a modelling approach that assumed asymptotic dependence between windstorm footprint wind gust speeds at pairs of locations, without showing any validation of this modelling choice. An important part of developing the windstorm footprint model within this thesis will therefore be to explore the extremal dependence between locations.

Davison et al. (2012) explained how a variety of statistical tools have been used to model multivariate and spatial extremes, including Bayesian Hierarchical Models (BHMs), also known as latent variable models, max-stable random fields, copula models and geostatistics.

Within BHMs, or latent variable models, dependence is introduced by integration over latent variables or processes (Davison et al., 2012). Most commonly this is done by introducing spatial variation in the parameters of the model. Due to the complexity of the integrations involved in this approach, it is most naturally performed in a Bayesian setting (Davison et al., 2012). Economou et al. (2014) used the BHM approach to spatially and temporally model extreme mean sea level pressure associated with windstorms. Economou et al. (2014) used this approach because it allows for spatial random effects and time-dependent covariates to be included in the model parameters, allowing for a great deal of flexibility and the natural inclusion of physical mechanisms. Economou et al. (2014) however noted that a major issue with this method, also pointed out by Davison et al. (2012), is the assumption of independence of the extremes, conditional on the latent process, which may yield unrealistic estimates from the model because the spatial clustering of rare events is not properly accounted for.

A common approach for modelling spatial extremes is to use max-stable processes. For example, Coles and Walshaw (1994) used a max-stable model for the dependence

in maximum wind speeds in different directions. Max-stable processes (de Hann, 1984) are an infinite dimensional generalisation of multivariate extreme value distributions. A stochastic process $\{Z(x)\}_{x\in\mathcal{X}}$ observed within the spatial or temporal domain $\mathcal{X}$, is max-stable if there exists continuous functions $a_n(x) > 0$ and $b_n(x) \in \mathbb{R}$ such that

$$\max_{i=1,\ldots,n} \frac{Z_i(\cdot) - b_n(\cdot)}{a_n(\cdot)} \stackrel{D}{=} Z(\cdot) \quad \text{for } n \geq 1 \tag{2.5}$$

where $Z_1, ..., Z_n$ are independent replications of Z, which is assumed to be non degenerate, and "$\stackrel{D}{=}$" represents being equal in distribution. Blanchet and Davison (2011) used the max-stable approach to spatially model extreme snow depth and argued that this method should be used because it accounts for the spatial dependence of extremes in a way that is consistent with classical extreme value theory. However, Economou et al. (2014) preferred not to use the max-stable models because they had issues relating to complexity of implementation and lack of flexibility. In addition, Segers (2012) highlighted how max-stable models are too coarse to accurately describe tails of multivariate distributions with asymptotic independence and can therefore only be used to model asymptotic dependence or complete independence.

The refinement of max-stable models, to overcome this issue, is a topic of ongoing research, originating from Ledford and Tawn (1996). Ledford and Tawn (1996, 1997) presented a refinement of the max-stable approach, in the form of a bivariate tail model which also includes the case of asymptotic independence. Within this model the joint survivor function of an arbitrary random pair (X, Y) with unit Fréchet marginal distributions satisfies:

$$\Pr(X > t, Y > t) \sim \mathcal{L}(t)p^{1/\eta} \quad \text{for large} t.$$

where $\mathcal{L}(t)$ is slowly varying as $t \to \infty$, $p = \Pr(X > t) = \Pr(Y > t)$ and η is the coefficient of tail dependence, in the range (0,1]. The parameter η describes the type of asymptotic dependence between X and Y (dependence or independence) and $\mathcal{L}(t)$ the relative strength given a particular value of η (Ledford and Tawn, 1996). This model is discussed further in Section 4.4.1.

This model has been explored and developed in a number of other studies (Bortot and Tawn 1998, Coles et al. 1999, Bortot et al. (2000), Heffernan 2000, Ferro 2007). Bortot et al. (2000) explained how the difficulty in working with the Ledford and Tawn (1996, 1997) model lies in its generality since the class of $\mathcal{L}(t)$ is infinite dimensional and restrictions must therefore be made by representing it in a parametric form. In particular, Bortot et al. (2000) showed how, in using a common formulation for $\mathcal{L}(t)$, the bivariate Gaussian distribution is excluded from the range

of distributions this model can represent. The Ledford and Tawn (1996, 1997) model is also restricted to the bivariate setting and is only applicable when both variables are large.

To overcome this, Heffernan and Tawn (2004) suggested a conditional representation, in which, given that one variable is large, the distribution of the remaining variables is modelled. This approach extends easily to higher dimensions, and incorporates both classes of extremal dependence. Keef et al. (2013) argued that the Heffernan and Tawn (2004) model provides the most flexible current approach for modelling this conditional distribution. They explained how the Heffernan and Tawn (2004) model encapsulates a broad range of asymptotic independence and asymptotic dependence forms of extremal dependence structure allowing for different types of dependence between different pairs of variables. Keef et al. (2013), however, identified a number of problems with this approach including complications that arise with modelling variables when some components are positively associated and others negatively associated.

Wadsworth and Tawn (2012), however, explained how, although these new models improve the flexibility of modelling extremal dependence classes, they are all based on modelling a finite dimensional variable and current theory and methodology for infinite-dimensional spatial extremes are still restricted to the max-stable class.

Wadsworth et al. (2015) presented an alternative limit representation for bivariate extremes, to that of Ledford and Tawn (1996, 1997), that encompasses a wide variety of dependence scenarios, and, like Heffernan and Tawn (2004), is applicable when at least one variable is large. They explained how their model permits inference across both extremal dependence classes, with a smooth transition between them. Wadsworth et al. (2015) also commented on how this new bivariate model could be extended to a multivariate setting.

An alternative, again when modelling in finitely many dimensions, is the use of extreme value copulas. Bonazzi et al. (2012) used an extreme value copula approach in their bivariate model for windstorm footprint wind gust speeds at pairs of locations. This copula approach represents the joint distribution of wind gust speeds in terms of the marginal distribution of wind gust speeds at each location and their mutual dependence, represented by the copula function. This approach is easily extendable to higher dimensions. Extreme value copulas satisfy the max-stability property, but, like max-stable processes, only allow for asymptotic dependence or complete independence between variables.

To overcome this, asymptotically independent multivariate spatial processes can be modelled using a Gaussian dependence structure, or Gaussian copula, which assumes independence in the extremes. The spatial process can be marginally transformed to Gaussian at each location, a method known as Gaussian anamorphosis (Davison et al., 2012), and the dependence between locations is then determined by the correlation matrix of the associated multivariate normal distribution. The Gaussian copula is not an extreme value copula and does not satisfy max-stability, however examples exist of its application to model the dependence between asymptotically independent extremes. For example Renard and Lang (2007) used the Gaussian copula for multivariate extreme value analysis in a number of hydrology case studies, demonstrating the usefulness and the relative simplicity of this approach. They emphasised the need to check the adequacy of the asymptotic independence assumption before using the model for computation beyond the observed range of the data. Similarly, Bortot et al. (2000) developed a multivariate Gaussian tail model for the dependence between oceanographic surge height, wave period and wave height. They showed that, when modelling data that are asymptotically independent, the multivariate Gaussian model is robust, has simple diagnostics, easily interpretable parameters and extends straightforwardly to higher dimensions.

In addition, if asymptotic independence is identified, the Gaussian dependence structure can be simplified by using a geostatistical approach, in which the dependence in the spatial field is modelled as a Gaussian process with a specified covariance function which represents the dependence between locations as a function of separation distance and direction. Luo et al. (2008) use a geostatistical approach to estimate wind speed surfaces using irregularly distributed data from England and Wales and Shoji (2006) carry out a geostatistical analysis of wind velocities in Japan. No examples of the application of geostatistics for modelling extreme wind gust speeds were found.

2.3. Data used in this thesis

2.3.1. Windstorm footprint data

The windstorm footprint data used in this thesis will now be introduced, specifying the chosen definition of the footprint in terms of the meteorological variable, region and event duration used.

Following Donat et al. (2011b), Haas and Pinto (2012) and Haas et al. (2014),

windstorm footprints are created based on a reanalysis dataset known as ERA Interim (Dee et al., 2011). Reanalysis is a scientific method for developing a consistent estimate of past weather by combining observations with a forecast model output using a data assimilation scheme. The ERA Interim reanalysis is produced with a sequential data assimilation scheme, advancing forward in time using 12-hourly analysis cycles. In each cycle, available observations are combined with prior information from a forecast model to estimate the evolving state of the global atmosphere (Dee et al., 2011).

ERA Interim reanalysis has a relatively low horizontal resolution of 0.75° (∼ 80km). A much higher resolution dataset was created by scientists at the Met Office by dynamically downscaling the ERA Interim reanalysis data to a horizontal resolution of 0.22 °, equivalent to ∼ 25km at the model's equator. The high resolution dataset was created in 2012 and therefore covers the ERA Interim period, up to the date at which it was created (1979-2012).

Downscaling was performed using the Met Office Unified Model (MetUM) version 7.4 (Davies et al., 2005). The MetUM has 70 irregularly spaced vertical levels (top at 80km) and uses a rotated pole with a longitude of 177.5 ° and latitude 37.5 °, and in the rotated coordinate frame it extends from -9.36 ° to 29.58 ° in longitude, and -17.65 ° to 16.89 ° in latitude, so the grid spacing does not vary substantially over the domain (∼ 0.22 ° or ∼ 25km). The non-hydrostatic dynamical equations are solved using semi-Langrangian advection and semi-implicit time stepping (Davies et al., 2005). The downscaled region covers most of Europe and the eastern North Atlantic as shown in Figure 2.2, roughly representing the region 15 °W to 25 °E in longitude and 35 °N to 70 °N in latitude. This domain is therefore appropriate for a European wide investigation, as is proposed for this thesis.

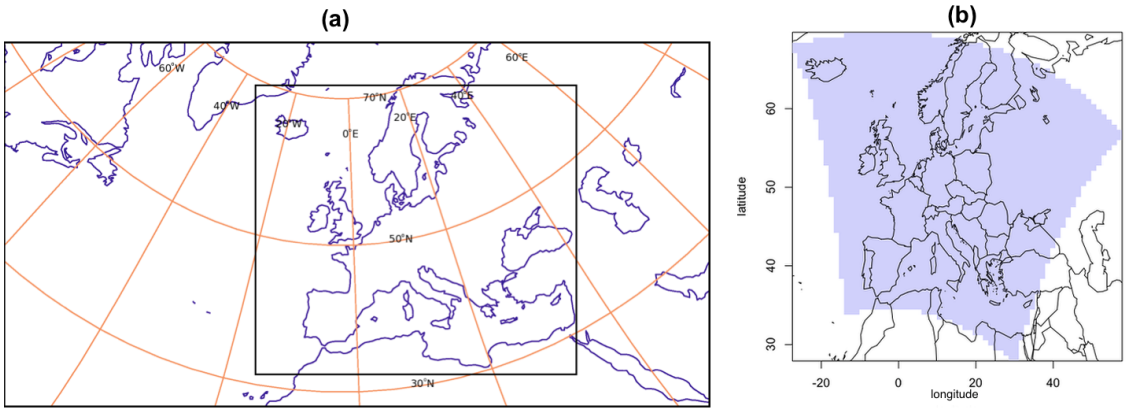

Figure 2.2. Modelling domain used to generate the footprints (a) on the rotated pole (inner rectangle) provided by J. Roberts, (b) on the non-rotated pole (shaded region)

The MetUM was initialised every day at 18Z (i.e. 6pm) using initial conditions from reconfigured ERA Interim analysis at that time. The 0.22 ° MetUM runs for 30 model hours, using lateral boundary conditions and initial conditions for the domain generated from the ERA Interim 6 hourly analyses. The first six hours of model output are disregarded due to spin up, allowing the model to adjust from the ECMWF Integrated Forecast System (ECMWF, 2006) initial conditions, leaving data for 00Z (midnight) to 00Z on each day. This results in daily 24 hour forecasts which are combined to create a new, higher resolution dataset for the ERA Interim period. Reinitialising the 0.22 ° MetUM runs every 24 hours ensures deviations from ERA Interim in the centre of the model domain are minimised.

The footprints represent the maximum 3 second wind gust speed at a height of 10 metres produced by the 0.22 ° MetUM. The wind gust speeds are derived from wind speeds, based on a gust parameterisation. The MetUM estimates gusts using the relationship $U_{gust} = U_{10m} + C\sigma$, where U_{10m} is the wind speed at 10 metre altitude and σ is the standard deviation of the horizontal wind, estimated from the friction velocity using the similarity relation of Panofsky et al. (1977). The roughness constant C is determined from universal turbulence spectra and is larger over rough terrain.

Within the 34 year period covered by the data set (1979-2012), 5730 windstorm events are identified as occurring within the specified domain during 33 extended winters (October - March 1979/80 - 2011/12). This also corresponds to the period of the year in which most windstorms occur over central Europe (Klawa and Ulbrich, 2003). Windstorm events are identified using the approach in Hoskins and Hodges (2002), based on the Hodges (1995) storm tracking algorithm. This storm tracking algorithm is applied to the ERA Interim reanalysis dataset and is an objective approach based on identifying maxima in 850 hPa relative vorticity, a measure of the local average circulation of a fluid (Hoskins and Hodges, 2002). Identification entails determining the vorticity maxima by steepest ascent maximisation (Hodges, 1995). These maxima are linked together, initially using a nearest neighbour search, and then refined by minimising a cost function for track smoothness, subject to adaptive constraints on the displacement distance and track smoothness (Hodges, 1995). Storm tracks that pass through the specified domain and last longer than 2 days are retained. Historically 6 hourly reanalysis data are used in storm tracking. Here, a new data set is used, created for the development of the eXtreme Wind Storms (XWS) catalogue presented in (Roberts et al., 2014). In this new dataset, 3 hourly data is created by combining the 03Z and 09Z outputs from the 00Z forecast and the 15Z and 21Z outputs from the 12Z forecast with the 6 hourly analyses. Roberts et al. (2014) explain how this 3 hourly data produces more reliable storm tracks

since some extreme European windstorms have very fast propagation speeds.

A windstorm footprint is created for each of the 5730 events that are identified by this method. The event duration is specified as being a 72 hour period, centred on the time at which the maximum 925hPa wind speed over land occurs within a 3° radius of the storm track, which is output every 3 hours as part of the tracking dataset introduced above. This definition gives good flexibility in the event start and end time because the wind speed data is available every 3 hours. This flexibility is much greater than some other studies such as Della-Marta et al. (2009) and Haas and Pinto (2012) who use daily data and therefore must use a 72 hour period beginning and ending at 00Z.

Evaluation of windstorm footprint data

Most previous studies, for example Schwierz et al. (2010), Haylock (2011) and Bonazzi et al. (2012) are based on a relatively small subset of very extreme events which may not be representative of the general characteristics of windstorm footprints, while here a large set of 5730 events will be used. The footprints used here are also at a high horizontal resolution (25km) compared to the data used in other studies, with only Haylock (2011) using an equivalent resolution. The higher the resolution of the model used to produce the windstorm footprints, the more sub grid scale processes can be resolved resulting in more realistic wind gust speeds. The footprints also span the whole of Europe for a reasonably long period of time (1979-2012) meaning that conclusions made here about the relationship between windstorm footprint characteristics and insured loss are representative of the whole continent and not just specific countries as is the case for studies such as Klawa and Ulbrich (2003) who focus on Germany.

The windstorm footprint data used in this thesis is evaluated in Roberts et al. (2014), in the development of the XWS catalogue. They compared the windstorm footprint wind gust speeds with observational wind gust data extracted for the MIDAS database (Roberts et al., 2014), a mixture of 1-, 3- and 6-hourly maximum wind gust speeds at sites covering most of North-West and central Europe. Roberts et al. (2014) found that modelled and observed wind gust speeds were generally in close agreement. They noted how this result was particularly impressive because the modelled wind gust speeds were simply interpolated from a 25km grid cell to a specific location without applying any corrections. However, they found that there is less agreement when considering stations with altitude greater than $\sim$ 500m. This

is a common issue in atmospheric models (Donat et al., 2010), caused by the use of an effective roughness parameterisation, needed to estimate the effect of sub-grid scale orography on the synoptic scale flow, causing unrealistically slow wind speeds at 10m (Roberts et al., 2014). Roberts et al. (2014) applied a correction proposed by Howard and Clark (2007), to windstorm Kyrill, which involved using wind speed at a height which is unaffected by the surface and assuming a log-profile to interpolate back down to 10m. They found that this correction gave a clear improvement at high-altitudes, however, they concluded that since this correction required the extraction of a great deal of archived data on all model levels, it was too time consuming to apply in practice.

Roberts et al. (2014) also found that the model slightly underestimates extreme wind gust speeds greater than $\sim 25\text{ms}^{-1}$. Roberts et al. (2014) explored this issue in great detail and concluded that this problem is due to several mechanisms, such as the underestimation of convective effects and strong pressure gradients. In particular they suggested that this happens when a storm develops outside of the MetUM domain and is therefore only being driven at the boundaries, meaning that the windstorm does not develop to become as extreme as in the observations. When the MetUM is reinitialised, which happens every 24 hours, the storm is corrected and the extremity can be reached, therefore this is seen to be more of an issue in windstorms that move through the domain quickly and therefore travel further before being reinitialised (Roberts et al., 2014). Roberts et al. (2014) concluded that no universal correction can be applied to all storms. This is reflected in the statistical recalibration method, which was adapted for each individual windstorm footprint (Roberts et al., 2014). The proposed recalibration method was based on polynomial regression between transformed wind gust speeds (Roberts et al., 2014), and its application was shown to successfully increase the under predicted wind gusts $> 25\text{ms}^{-1}$. Only 50 windstorms are selected for the XWS catalogue, making this individual recalibration a feasible proposition for Roberts et al. (2014), however recalibrating all 5730 windstorm footprints would be very time consuming and will therefore not be attempted.

Roberts et al. (2014) explained how over the 33 extended winters in the data set, 2.5 windstorms pass through the domain on average in any given 72 hour period. This means that windstorm events that pass through the domain within the same 72 hour period will be represented in both of the associated footprints, based on the definition of the footprint used in this thesis. This results in some extreme storms being represented twice because they overlap with weaker events. This could create problems when trying to attribute damage to a particular event when exploring the relationship between storm severity measures and insured loss.

To overcome this, Roberts et al. (2014) developed a footprint decontamination method which can be quickly and easily applied with the aim of removing all wind gust speeds associated with any other overlapping windstorm. They assume that all gusts within a 1000km radius of the track position at each 6 hour time step are associated with that particular windstorm and all other data in the domain is rejected. The "decontaminated" footprint is then derived by taking the maximum wind gust speed in all grid cells where there are data remaining, within the 72 hour period of the storm (Roberts et al., 2014). Various other radius sizes were explored and Roberts et al. (2014) concluded that 1000km is generally large enough to capture the high wind gust speeds associated with each windstorm, but small enough to avoid contamination. Figures 2 and 3 in Roberts et al. (2014) show how this decontamination method is able to isolate 4 individual windstorm events that occur in the same 72 hour period as windstorm Kyrill. This decontaminated dataset will be used in the final part of Chapter 3 which presents the methodology developed for selecting extreme storms for the XWS catalogue.

2.3.2. Population density data

Following Klawa and Ulbrich (2003), population density data will be used to approximate the exposure in some storm severity measures in this thesis. The population density data is based on the derived portfolio in Haylock (2011). This portfolio was created based on the 2005 Gridded Population of the World, Version 3 (GPWv3) estimates of human population by 2.5 arc-minute (2.5/60 of a degree) grid cells. A proportional allocation gridding algorithm, utilising more than 300,000 national and sub-national administrative units, is used to assign population values to grid cells. The population density grids are derived by dividing the population count grids by the land area grid, so representing persons per square kilometre. GPWv3 is produced by the Columbia University Center for International Earth Science Information Network (CIESIN) in collaboration with Centro Internacional de Agricultura Tropical (CIAT). This gridded data set was mapped onto the locations supported by the PartnerRe European windstorm model to create the geographic distribution of values presented in Haylock (2011). For use within this thesis, population density is required to be on the same grid as the footprint data and is therefore regridded once more. The latitude-longitude coordinates for the centres of the grid cells for the footprint are calculated. The population density data is then bi-linearly interpolated to these latitude-longitude coordinates. The interpolated values are then used as the population density values for the footprint grid cells (see Figure 2.3).

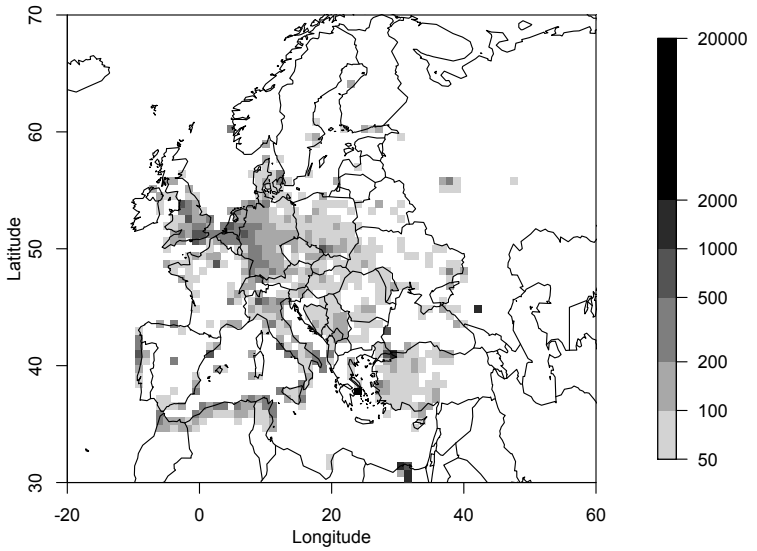

Figure 2.3. Population density (people km^{-2}) within the European domain, regridded to the same resolution as the windstorm footprints, based on the 2005 Gridded Population of the World, Version 3 data set. Locations with less than 50 people km^{-2} are shown in white.

2.3.3. Notation

The notation used to represent the windstorm footprint wind gust speeds throughout this thesis will now be defined.

Let $\{X(s) : s \in S\}$ be a continuous spatial random process that represents the maximum 3 second wind gust speed in a 72 hour period at locations s in the region of interest S. A windstorm footprint is considered to be a discrete realisation of this random process, i.e. footprints $i = 1, 2, ..., I$ are independent realisations $x_i(s)$ of the $X(s)$ process, observed at spatial locations $\{s_j; j = 1, ..., J\}$ within the region, at grid point locations. The high resolution Met Office data contains $I = 5730$ footprints over the period 1979-2012, recorded at $J = 14872$ land grid point locations over Europe (see Figure 2.2).

3. Exploring the relationship between storm severity measures and insured loss

3.1. Introduction

The footprint of a windstorm is a representation of a complex multidimensional spatial and temporal process and therefore varies from storm to storm depending on many factors. Such factors may include the state of large scale modes of variability (e.g. the North Atlantic Oscillation), cyclone location and the cyclone deepening rate (Gulev et al., 2001). These varying factors will have an impact on the resulting windstorm footprint and, in turn, on the insured loss incurred. Figure 1.1 showed how the Great storm of October '87 and windstorm Kyrill, which incurred a similar insured loss of $6.7 billion and $6.3 billion respectively, have visually very different footprints. So how does this variation in windstorm footprints relate to insured loss, and do these two storms in fact have a common damaging characteristic?

Within this chapter windstorm footprints will be related to insured loss using storm severity measures, functions that incorporate multiple parameters relating to footprint characteristics such as area and peak intensity, to scalar values of footprint extremity. A review of existing storm severity measures is presented in Section 2.2.2. The storm severity measures developed within this chapter will be based on those in these previous studies, ranging in complexity and which footprint characteristics are represented. Exploratory data analysis is used to explore the relationship between the storm severity measures and insured loss, based on the historical data set of windstorm footprints introduced in Section 2.3.1.

Due to the unavailable and therefore incomplete record of European wide insured losses associated with windstorm events, the relationship between storm severity measures and insured loss will be quantified in terms of their ability to classify a set of 23 extreme insurance loss windstorms, identified through expert elicitation with reinsurers. These 23 windstorms will first be presented for comparison to explore

any common characteristics that may be visually present before applying the storm severity measures.

This investigation will be applied to a large data set of high resolution historical windstorm footprints over the European domain, therefore providing insight into which footprint characteristics are most damaging for the whole of Europe, identified as a gap in existing knowledge in Section 2.2.2.

The final part of this chapter will present an application of this investigation; a methodology that has been developed for selecting extreme windstorms for the eXtreme Wind Storms (XWS) catalogue, also published in Roberts et al. (2014). The XWS catalogue is a freely available database of storm tracks and footprints for 50 of the most extreme windstorms to occur in Europe in recent years (1979-2012), available at www.EuropeanWindstorms.org.

This chapter will address four questions:

- How can the extremity of windstorm footprints be summarised into storm severity measures?
- How do storm severity measures relate to one another?
- How do storm severity measures relate to insured loss?
- How can storm severity measures be used to select extreme windstorms for the XWS catalogue?

3.2. Extreme insurance loss windstorms

Of the 5730 historical windstorms (introduced in Section 2.3.1), 23 have been identified as causing large insured losses by colleagues at Willis Research Network. These particular 23 windstorms were selected based on their notable socio-economic impacts and presence in literature and documentation associated with windstorm damage, for example the IMILAST project (Neu et al., 2013) and the PERILS event loss database (perils.org).

These 23 extreme insurance loss windstorm events are first compared to investigate whether they share any common characteristics. Table 3.1 contains information about each of these 23 windstorms including the insured loss incurred (where available) from Sigma (2004, 2006, 2007, 2009, 2011, 2012, 2013). The footprints for these 23 windstorm and histograms of the associated wind gust speeds are presented

Table 3.1.. Name, date of occurrence, maximum 925hPa windspeed over land (taken from the storm track), footprint area exceeding 25ms^{-1} over land and insured loss incurred (where available) for the 23 extreme insurance loss windstorms. Losses have been indexed to 2012. Missing values of insured loss represent where this information could not be attained.

Name	Date of maximum 925hPa windspeed over land	Maximum 925hPa windspeed over land (ms^{-1})	Footprint area over land (25km grid boxes)	Insured loss (Billion USD)
Daria	25 Jan 1990	39.53	882	8.2
Lothar	26 Dec 1999	36.72	822	8.0
Kyrill	18 Jan 2007	36.38	1257	6.7
Great Storm of '87	16 Oct 1987	39.53	681	6.3
Vivian	26 Feb 1990	35.16	1059	5.6
Klaus	24 Jan 2009	37.23	565	3.5
Martin	27 Dec 1999	37.18	689	3.3
Xynthia	27 Feb 2010	32.62	669	2.9
Anatol	3 Dec 1999	39.86	742	2.6
Erwin	8 Jan 2005	39.22	643	2.2
Herta	3 Feb 1990	33.16	726	1.5
Emma	29 Feb 2008	25.12	779	1.4
Wiebke	28 Feb 1990	32.24	1294	1.4
Gero	11 Jan 2005	39.13	309	0.6
Ulli	3 Jan 2012	36.32	400	0.2
Dagmar-Patrick	26 Dec 2011	30.08	116	0.04
Fanny	4 Jan 1998	34.60	366	-
Jeanette	27 Oct 2002	36.92	1499	-
Lore	28 Jan 1994	31.60	629	-
Oratio	30 Oct 2000	38.45	713	-
Stephen	26 Dec 1998	39.53	345	-
Xylia	28 Oct 1998	26.72	408	-
Yuma	24 Dec 1997	39.92	205	-

in Figures 3.1 and 3.2 respectively.

Table 3.1 indicates that both the maximum 925hPa windspeed over land and footprint area exceeding 25ms^{-1} over land vary for different storms, with neither measure decreasing with insured loss, which itself varies greatly ($0.04 billion - $8.2 billion). In addition, Figure 3.1 shows a great deal of visual variation within these 23 footprints. The position of the extreme wind gusts vary, with some located more northerly, e.g. Emma, and some more southerly, e.g. Klaus. In addition, the area over which wind gusts exceed 15ms^{-1} varies between storms, with some having a relatively small exceedance area, e.g. Xynthia, and some much larger, e.g. Daria. The range of gust intensities also varies, with windstorm Wiebke reaching 31.5ms^{-1} while windstorm Stephen's peak gust is 49.6ms^{-1}. The histograms in Figure 3.2 show that the distribution of footprint wind gust speeds vary with some peaking in the interval 5-10ms^{-1}, e.g. Yuma, some in the interval 10-15ms^{-1}, e.g. Lothar and some the interval 15-20ms^{-1}, e.g. Xylia. Variation also exists in the upper tail of

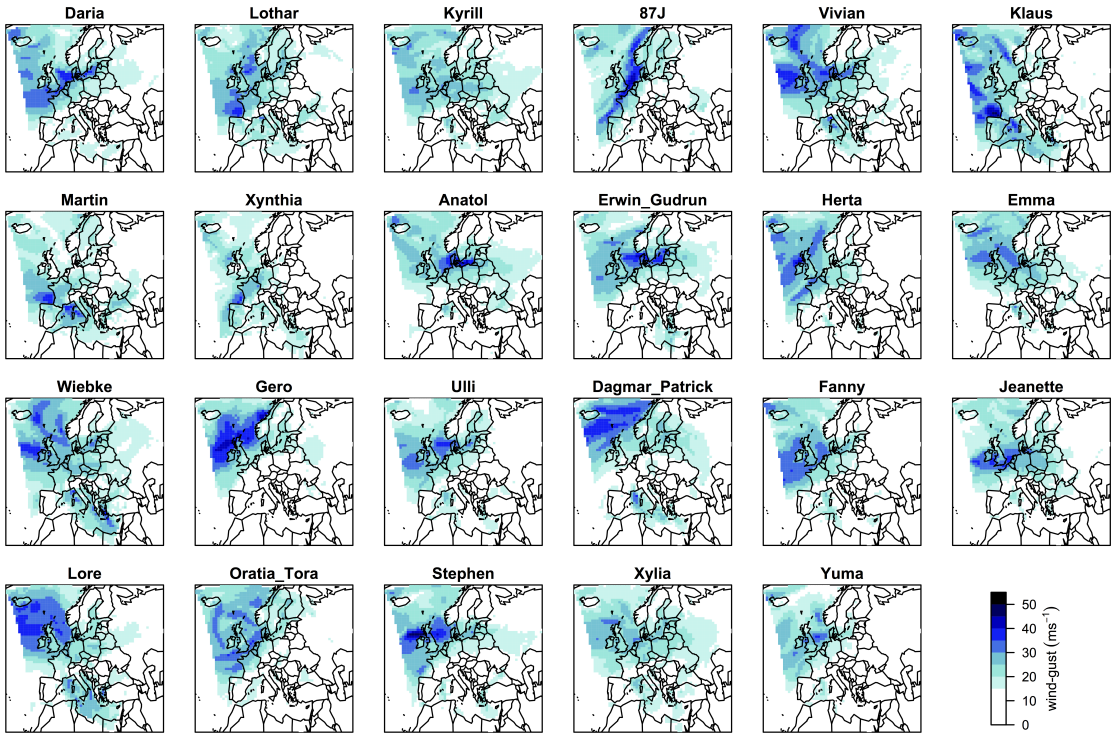

Figure 3.1. Windstorm Footprints showing the maximum 3 second wind gust speed (ms^{-1}) over the 72 hour life time of the storm, for the 23 extreme insurance loss windstorms, in the same order as in Table 3.1.

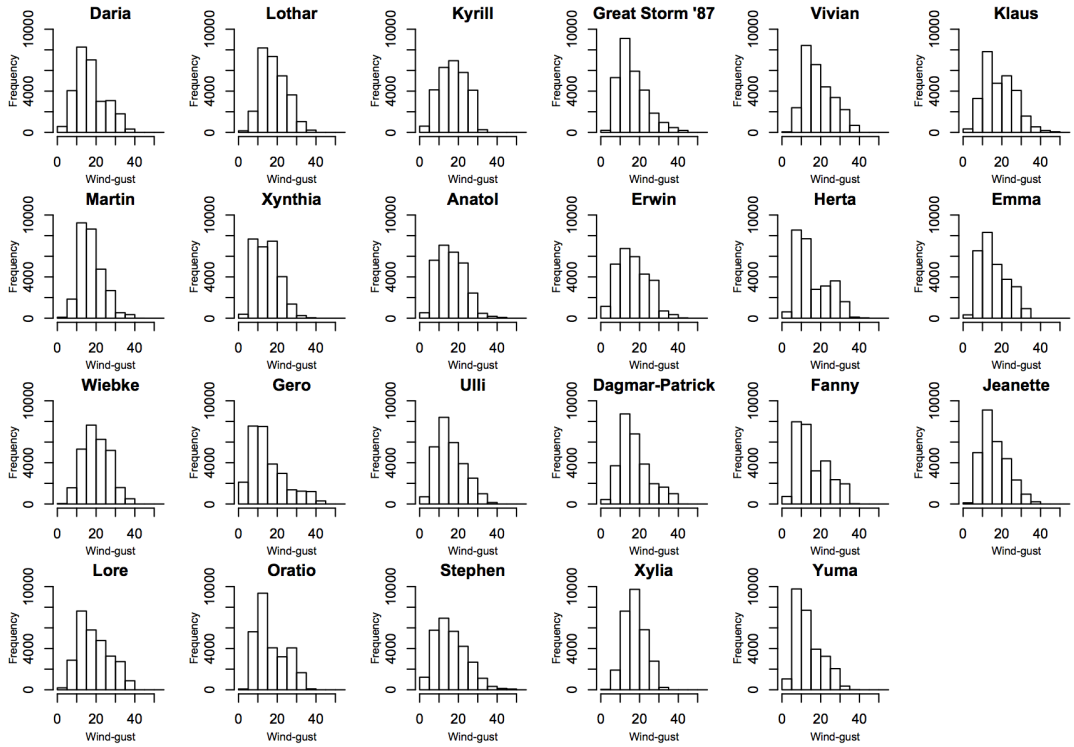

Figure 3.2. Histograms of the footprint wind gust speeds (ms^{-1}) for the 23 extreme insurance loss windstorms.

the distribution, with some storms having a heavy upper tail, e.g. the Great Storm '87 and Gero, while others have a short upper tail, e.g. Oratio and Kyrill.

No common characteristic is identified from exploring these extreme insurance loss windstorm footprints in isolation. They should therefore be explored within the context of all historical windstorms in order to identify common characteristics.

3.3. Defining storm severity measures

There are many ways in which a windstorm footprint can be summarised using a storm severity measure. Several such measures have been developed in previous studies (reviewed in Section 2.2.2) with their uses ranging from the estimation of the return period of windstorms over Europe (e.g. Della-Marta et al. 2009), to understanding how windstorms will change under anthropogenic climate change (e.g. Leckebusch et al. 2008a), to the ranking of storms according to severity in catalogues of extreme historical storms (e.g. Lamb and Frydendahl 1991). These indices also range in complexity from simple summary statistics such as the event mean and 95^{th} percentile wind gust speed (e.g. Della-Marta et al. 2009), to loss functions which quantify severity in terms of the cube of damaging wind gust speeds and population density (e.g. Klawa and Ulbrich 2003). Here, nine such severity measures will be defined. These measures will range in complexity, allowing for the comparison of simple and more complex measures; measures using wind gust speeds in exceedance of relative or absolute thresholds; and measures that do and do not include population density as an approximation of exposure.

As discussed in Section 2.2.2, most such severity measures can be represented in a general form. For windstorm footprint i a severity measure m_i can be expressed as:

$$m_i = g\left(\sum_{j=1}^{J} V(x_i(s_j), t)e(s_j, t)\right) \tag{3.1}$$

where $x_i(s_j)$ is the wind gust for footprint i at location $s_j, j = 1, ..., J$. As previously explained, the functions V and e represent the hazard and exposure respectively, which both depend on the wind gust exceeding a damage threshold t. These two functions together, summed over all locations, represent the severity of a given windstorm event. In some previously published measures an additional function g is applied to the severity measure, e.g. the cube root (Cusack, 2013). Three types of storm severity measures will be used in this study: simple footprint summary statistics, loss functions and composite measures which are a combination of other measures. These measures are intended to represent the severity of a windstorm defined by its socio-economic impact and are therefore only applied to wind gust speeds over land. The land grid-cells over Iceland are excluded since the relatively

high wind gust speeds in this region give too much weight to storms that pass over Iceland but do not have a large effect on the rest of Europe.

This form of storm severity measure is common in previous studies as explained in Section 2.2.2 and is shown to summarise the footprint characteristics that are important for representing insured loss, such as the area of damaging winds, the cubed excess wind gust speed and exposure. In contrast, some of the simple and composite measures do not take this common form. Comparing the relationship between these types of severity measures will therefore identify whether this general form in Eqn. (3.1) is indeed the most appropriate for representing insured loss.

Table 3.2 defines the nine storm severity measures used in this investigation, where measures (1 - 3) are the simple footprint summary statistics, measures (4 - 7) are the loss functions, measure (8) is an intensity measure taken from the windstorm track rather than the footprint, used in composite measure (9). In Table 3.2, define $c(s_j)$ as the 98$^{\text{th}}$ percentile of climatology wind gust speed at location s_j, i.e. the 98$^{\text{th}}$ percentile of daily maximum wind gust speed, calculated for the full reference period (October 1979 to March 2012), shown in Figure 3.3, and $p(s_j)$ as the 2005 population density (people km^{-2}) at location s_j (introduced in Section 2.3.2). Define $H(n)$ as a Heaviside or indicator function $H(n) = 1$ if $n > 0$ and $H(n) = 0$ otherwise.

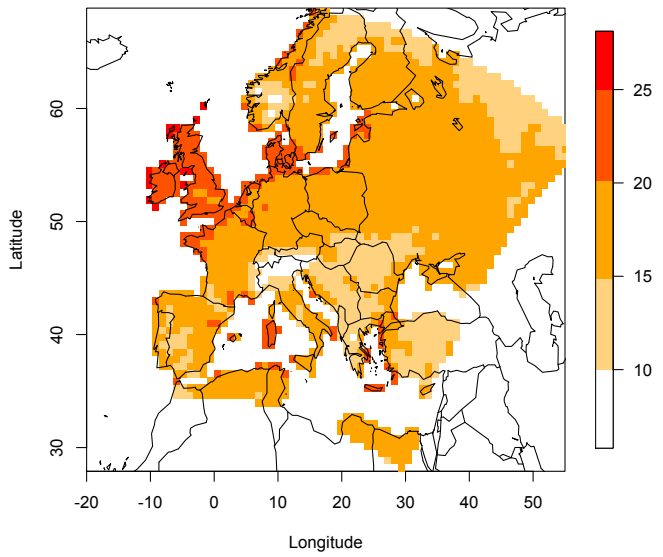

Figure 3.3. The 98$^{\text{th}}$ percentile of climatology wind gust speed (ms^{-1}) October 1979 to March 2012, for all land locations in the domain, $c(s_j)$ for $j = 1, ..., J$.

This particular range of indices is used to give a diverse representation of the characteristics of windstorm footprints. Simple footprint summary statistics $\bar{x}$ and x_{90},

Table 3.2.. Definition of the nine storm severity measures used in this investigation in the form of Equation 3.1, for windstorm footprint i.

Measure	Label	$V(x_i(s_j),t)$	$e(s_j,t)$	t
1. Area (number of grid cells) exceeding 25ms^{-1}	A_i	$H(x_i(s_j)-t)$	1	25
2. Mean wind gust speed	$\bar{x}_i$	$x_i(s_j)$	J^{-1}	-
3. Ninetieth percentile wind gust speed	x_{90_i}	-	-	
4. Absolute threshold loss function	L_{a_i}	$H(x_i(s_j)-t)(x_i(s_j)-t)^3$	1	25
5. Relative threshold loss function	L_{r_i}	$H(x_i(s_j)-t)\left(\frac{x_i(s_j)}{t}-1\right)^3$	1	$c(s_j)$
6. Population weighted, absolute threshold loss function	L'_{a_i}	$H(x_i(s_j)-t)(x_i(s_j)-t)^3$	$p(s_j)$	25
7. Population weighted, relative threshold loss function	L'_{r_i}	$H(x_i(s_j)-t)\left(\frac{x_i(s_j)}{t}-1\right)^3$	$p(s_j)$	$c(s_j)$
8. Maximum 925hPa wind speed over land within a 3° radius of the storm track	U_{max_i}	-	-	
9. Composite severity measure	LF_i	-	-	

similar to those measures introduce by Della-Marta et al. (2009), are included for comparison with the more complex indices. This gives insight into whether complex measures, most common in the literature, are indeed better at representing insured loss. Severity measure A characterises the area, or number of 25km grid cells, in which damaging wind gust speeds (>25ms^{-1}) occur. This measure is included to identify whether the area of damaging winds, used in the severity measure developed by Lamb and Frydendahl (1991), is an important characteristic for representing insured loss. The damage threshold of 25ms^{-1} has been recognised in previous studies as being the wind speed at which damage starts to occur, for example Lamb and Frydendahl (1991) used a loss threshold of 50 knots (25.7ms^{-1}) and Bonazzi et al. (2012) and Klawa and Ulbrich (2003) noted that the German reinsurance industry use a damage threshold of 20ms^{-1}.

The more complex, loss function severity measures (4 - 6), L_r, L_a and L'_a, are variations of the loss function introduced by Klawa and Ulbrich (2003) which is equivalent to measure (7), L'_r. This loss function has been used in a number of other studies to approximate windstorm loss (e.g Leckebusch et al. 2007, Leckebusch et al. 2008a, Pinto et al. 2007) and is explained in detail in Section 2.2.2. As discussed in Section 2.2.2, the Klawa and Ulbrich (2003) loss function is found to well represent annual aggregate and event based losses in Germany. This investigation will explore how successful this loss function is at representing insured loss for the whole of Europe. For comparison severity measure L_r, is included to investigate the importance of including population density in the exposure component of the measure when representing insured loss.

The loss function developed by Klawa and Ulbrich (2003) assumed that storm dam-

ages occur when a certain, locally adapted, relative threshold is exceeded. Other studies, such as Lamb and Frydendahl (1991), Haylock (2011) and Della-Marta et al. (2009), use absolute exceedance threshold, constant throughout the domain. Severity measures L'_a and L_a, which use an absolute exceedance threshold of 25ms^{-1}, are therefore included as a comparison, to explore which type of exceedance threshold is most appropriate for representing insured loss.

Lastly, a composite severity measure is included, defined as $LF = AU^3_{max}$, where A is the area of damaging wind gust speeds (severity measure 1 in Table 3.2), and U_{max} is the maximum 925hPa windspeed over land within a 3° radius of the storm track. The measure U_{max} is taken from the windstorm track rather than the footprint but is included in order to create LF, which is equivalent to the storm severity measure used by Lamb and Frydendahl (1991), defined in Section 2.2.2, Eqn 2.2. Here, however, the duration is assumed to be 72 hours for all storms, rather than being included as an additional parameter. The measure U_{max} is also included in the investigation as a comparison with x_{90}, the severity measure related to the peak intensity from the footprint.

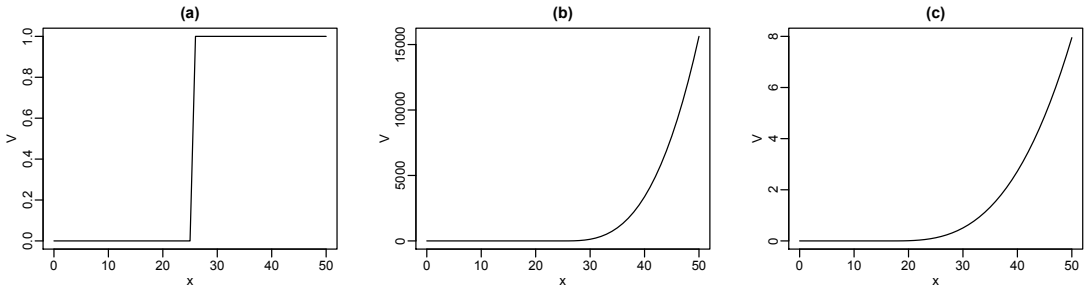

Figure 3.4. Diagrams of V for x in the range 0-50ms^{-1} for severity measures (a) A, (b) L_a and (c) L_r with absolute threshold 25ms^{-1} and relative threshold $c = 16$ms^{-1} (selected arbitrarily).

Figure 3.4 demonstrates the forms that the hazard function V takes in the different severity measure types. For severity measure A the hazard function is a simple step function. This form of hazard function is simplistic and may not be physically realistic since no damage occurs below the damage threshold of 25ms^{-1} and maximum damage occurs above this threshold. The hazard functions for measures L_a and L_r increase as a function of wind gust, x. This is more physically plausible, however there is no upper limit for the hazard function and therefore neither for the loss, which is not physically realistic. The hazard function for L_r is less steep because the excess wind gust speed is normalised.

3.4. Relating storm severity measures to one another

These storm severity measures are calculated for each of the 5730 historical windstorm footprints (see Section 2.3.1). The relationship between the storm severity measures will be quantified using two approaches. Firstly, paired scatter plots of the logarithm of each measure will be examined and the association between measures will be quantified in terms of the Spearman rank correlation coefficient. This measure of correlation is preferred because it represents monotonic rather than just linear association. Secondly, paired scatter plots of the rank of windstorms in each severity measure will be presented focusing on the most extreme events.

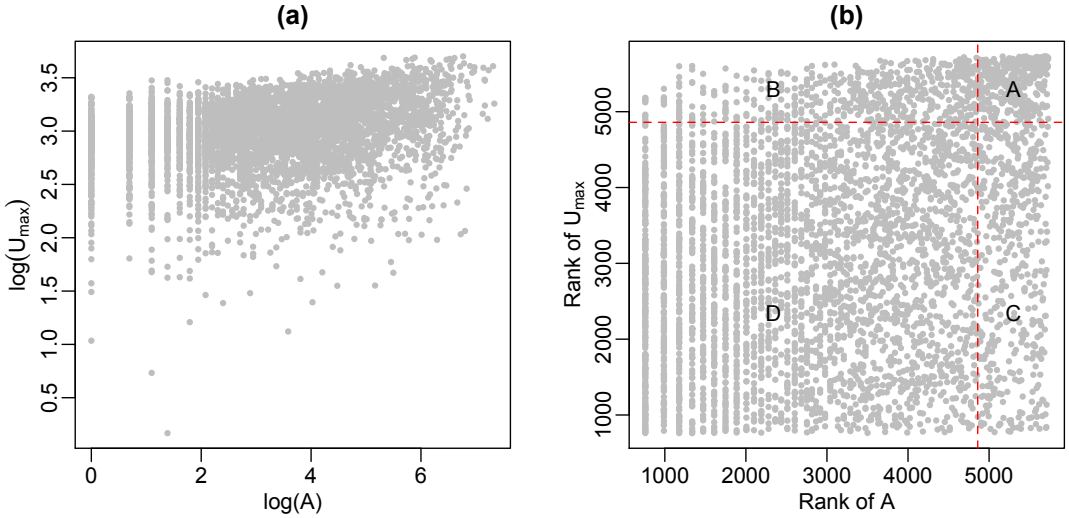

Figure 3.5. Scatter plots of (a) the logarithm of A and U_{max} (b) the rank of A and U_{max} with dashed red lines showing how the plot can be divided into four categories to calculate the extremal dependence measure.

Relationships between extreme values of the severity measures are of particular interest because of the relevance to damage and insured loss. These extremal relationships will be investigated using the extremal dependence measure χ_t, based on the Extremal Dependence Coefficient introduced by Coles et al. (1999). This measure quantifies the conditional probability of the rank of a windstorm being extreme, or above some threshold t, in one severity measure ($m_{(2)}$), given it is extreme in another severity measure ($m_{(1)}$). The extremal dependence measure is defined as:

$$\begin{aligned} \chi_t &= Pr(\text{rank}(m_{(2)}) > t | \text{rank}(m_{(1)}) > t) \\ &= \frac{Pr(\text{rank}(m_{(1)}) > t, \text{rank}(m_{(2)}) > t)}{Pr(\text{rank}(m_{(1)}) > t)} \\ &= \frac{a}{a+c} \end{aligned} \tag{3.2}$$

where n is the total number of windstorms, a is the number of windstorms that are extreme in both measures, $m_{(1)}$ and $m_{(2)}$, (section **A** in Figure 3.5 (b)) and c is the number of windstorms that are only extreme in $m_{(1)}$ (section **C** in Figure 3.5 (b)), therefore $t = n - (a + c)$. The extremal dependence measure χ_t is a conditional probability and is therefore in the range [0,1]. The larger the value of χ_t, the greater the extremal dependence between the severity measures.

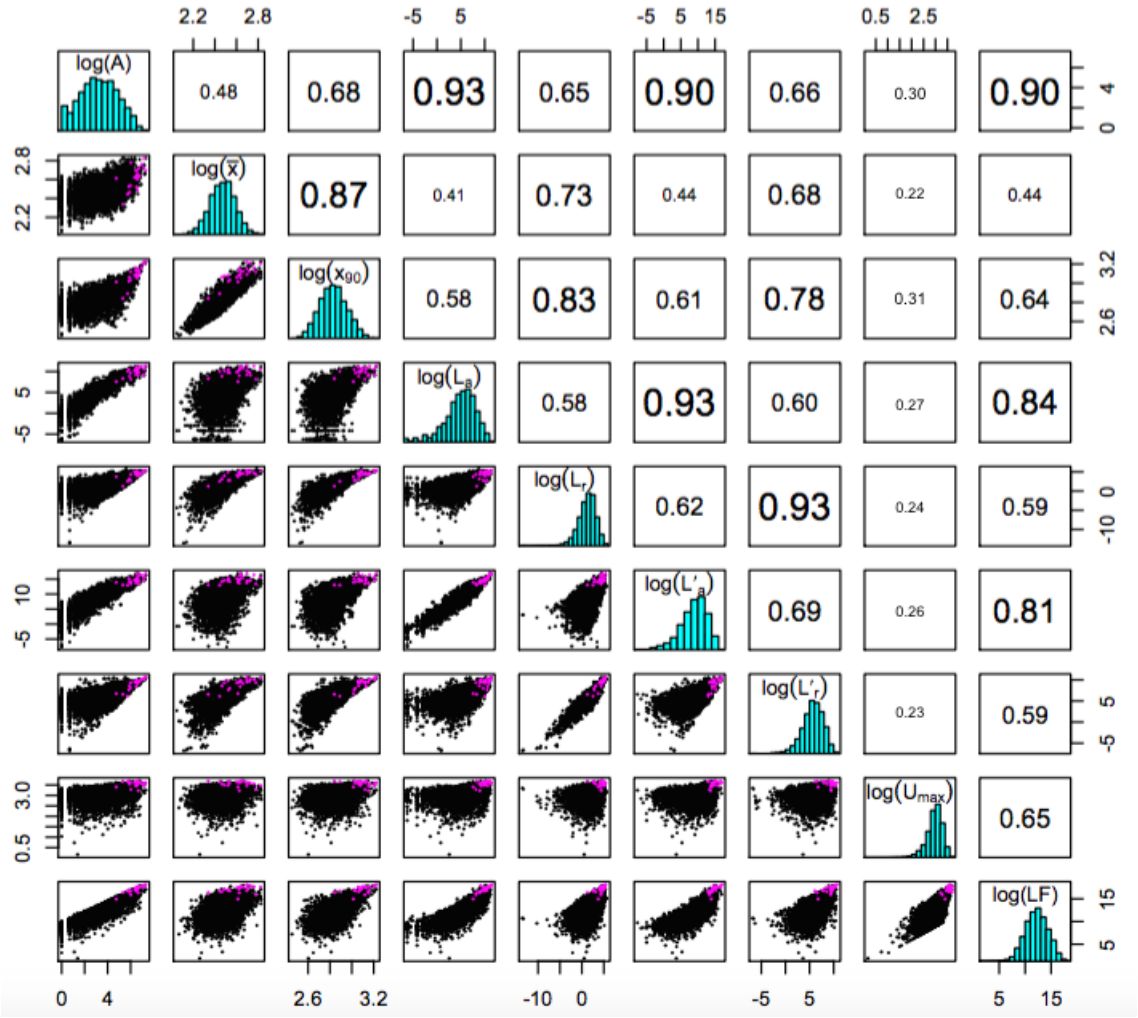

Figure 3.6. Paired scatter plots (lower panel) and the Spearman rank correlation coefficient (upper panel) for all pairs of severity measures (on natural logarithm scale). The panels on the diagonal show the histogram of the natural logarithm of each measure. The 23 extreme insurance loss storms (Table 3.1) are shown in purple.

Figure 3.6 shows the overall relationship between the storm severity measures defined in section 3.3. The logarithm of the indices is used to reveal relationships more clearly and the Spearman rank correlation coefficient is invariant under monotonic transformations such as the logarithm so it will take the same value for measures on the logarithmic scale as on the original scale.

The strongest relationships ($\rho = 0.93$) exist between A and L_a, L_a and L'_a, and

L_r and L'_r. A relationship between A and L_a might be expected since L_a is a measure of the intensity of the footprint wind gust speeds within the area A. However, the very strong relationship between these two measures suggests that including magnitude of the exceedance of the damage threshold within the severity measure gives little extra information about the severity of a windstorm. In addition, the strong positive relationships between L_a and L'_a and L_r and L'_r imply that little information is gained from including the population density as a measure of exposure within the loss functions. This is a very interesting result since population density has commonly been used as an approximation to exposure in loss functions in previous studies, e.g. Klawa and Ulbrich (2003), without any quantification of how informative it is in representing loss.

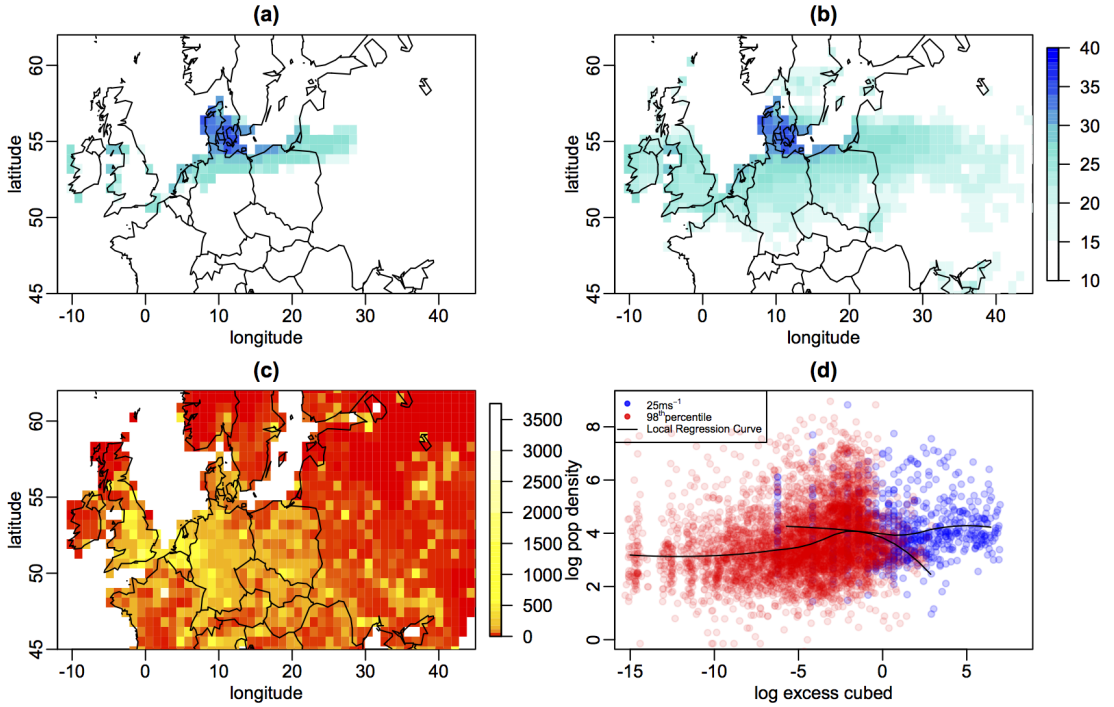

Figure 3.7. Anatol example: Exceedance footprints for windstorm Anatol (2nd-4th December 1999) using a threshold of (a) 25ms^{-1} and (b) $c(s_j)$ (c) Population density over the exceedance region (people km^{-2}) and (d) the relationship between the excess cubed wind gust speed and the population density in grid cells that exceed 25ms^{-1} (blue) and $c(s_j)$ (red).

This strong positive relationship between loss functions with and without population density included is found to be because of the difference in spatial scale of the wind gust speed excesses and the population density. To illustrate this, Figure 3.7 shows the exceedance footprints for windstorm Anatol (2nd-4th December 1999) for exceedance thresholds of 25ms^{-1} and $c(s_j)$. The spatial distribution of population density is more varied than the smoother exceedance footprint, meaning that for any magnitude of cubed excess wind gust speed a variety of population densities,

encompassing the whole range of populations, is sampled. This can be seen in Figure 3.7 (d), where for each magnitude of cubed excess wind gust speed, using the absolute threshold (blue) and the relative threshold (red), a similar range of population densities is sampled, resulting in an approximately proportional relationship between L_a and L'_a, and L_r and L'_r, as seen in Figure 3.6. A local regression curve is added to demonstrate the lack of relationship between population density and cubed excess wind gust speed. This can be further demonstrated by supposing that the population density at location s_j, denoted $p(s_j)$, can be thought of in terms of a mean population density plus some variation, $p(s_j) \approx \bar{d} + \epsilon(s_j)d$, then,

$$\begin{aligned}\sum_j V(s_j)p(s_j) &= \sum_j V(s_j)(\bar{d} + \epsilon(s_j)d) \\ &\approx \sum_j V(s_j)\bar{d} \\ &\text{since } \sum_j V(s_j)\epsilon(s_j)d \approx 0 \quad \text{if there is no correlation between } V(s_j) \text{ and } \epsilon(s_j) \\ &\propto \sum_j V(s_j)\end{aligned}$$

where $V(s_j)$ is the hazard at location s_j (see equation 3.1).

Other relationships of note include the strong positive correlation between x_{90} and the relative threshold loss function measures L_r and L'_r, indicating that this simple measure could be used to represent the more complex loss function measures.

The composite severity measure LF, defined as AU^3_{max}, has a stronger dependence on A than U_{max} and U_{max} has a weak dependency with all other severity measures. In addition, the relatively weak dependency between absolute and relative threshold loss functions (e.g. L_a and L_r) shows that these measures characterise different features of the footprints and both types may therefore be useful in determining the severity of a windstorm.

Figure 3.8 shows the relationship between the extremes of storm severity measures where the extremal dependence measure χ_t is calculated for $t = 5630$, meaning a footprint is classed as being extreme if it is ranked within the top 100 for a given severity measure.

The same strong dependencies exist between the loss functions with and without population density, meaning that the difference in spatial scale of the wind gust speed excesses and the population density holds even for the most extreme footprints. The strong dependency between x_{90} and relative threshold loss functions, L_r and L'_r, is also still present in the extremes of the measures. The strong dependence

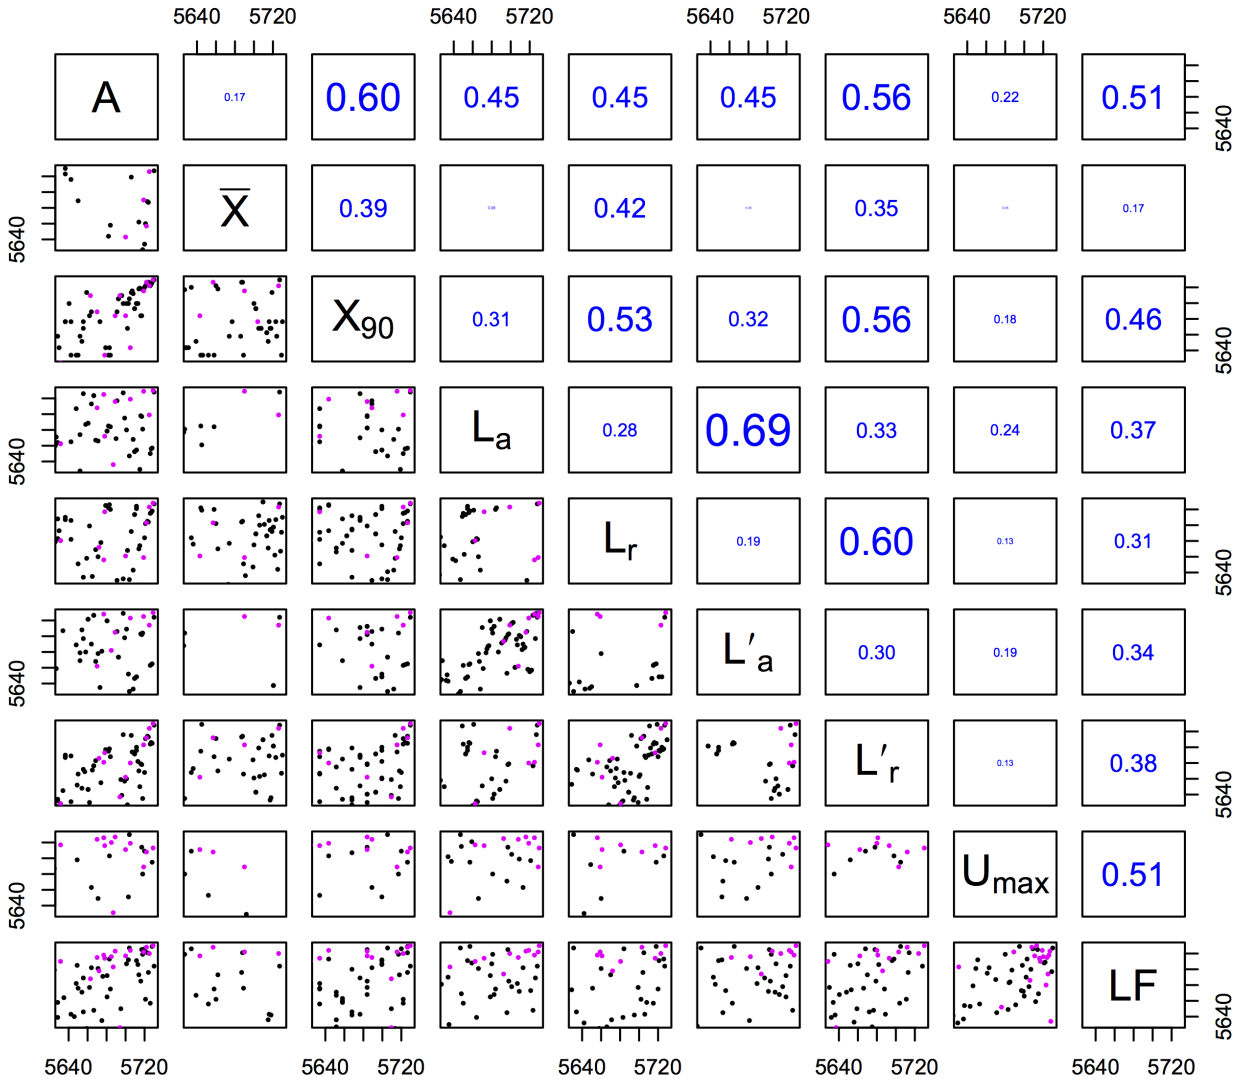

Figure 3.8. Paired scatter plots of the rank of severity measures for rank∈(5630,5730), with the 23 extreme insurance loss storms (Table 3.1) shown in purple, equivalent to section **A** of Figure 3.5 (lower panel) and extremal dependence measure χ_t for extremal threshold t=5630 (upper panel).

between composite measure LF and A holds for extreme windstorms. The same strong dependence now also exists between LF and U_{max} for extreme windstorms meaning that both the peak wind intensity and damage area are equally important within the composite measure for extreme windstorms.

3.5. Relating storm severity measures to insured loss

The relationship between the severity measures and insured loss will be quantified in terms of their ability to classify the set of 23 extreme insurance loss windstorms, introduced in Table 3.1. In other words, storm severity measures in which the 23 extreme insurance loss windstorms rank highly, are identified as being able to represent insured loss well. This is due to the unavailability of windstorm loss data, meaning that all observed storms cannot be ranked according to loss amount. The 23 extreme insurance loss storms, shown in purple in Figures 3.6 and 3.8, have generally high values in most of the 9 severity measures. They are in the top 18%, 87%, 47%, 18%, 53%, 11%, 45%, 18% and 8% of windstorms for measures A, $\bar{x}$, x_{90}, L_a, L_r, L'_a, L'_r, U_{max} and LF, respectively. This indicates that measures using the 25ms^{-1} absolute excedance threshold: A, L_a, L'_a, and also U_{max} are relatively good at representing insured loss, with LF being the best. The other measures, including the loss functions using the relative exceedance threshold are less successful at representing insured loss.

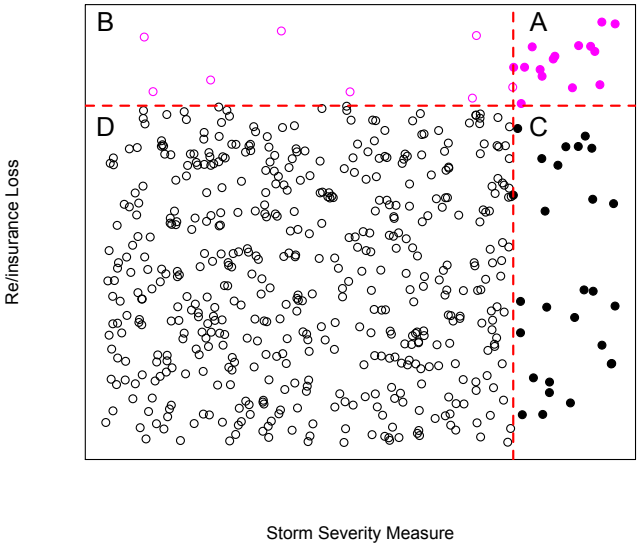

Figure 3.9. Conceptual diagram of storms ranked according to a severity measure and according to insured loss. The red dashed lines represent the division of the plot into the four categories A, B, C and D where the counts of storms in each category are denoted a, b, c and d respectively.

Alternatively, this can be thought of in terms of the extremal dependence between storm severity measures and insured loss. A conceptual scatter plot of the rank of each storm according to a storm severity measure and insured loss is shown in

Figure 3.9. The storms cannot be ranked according to the insured loss they incurred because these values are not available. Rather, the 23 extreme insurance loss storms are said to be in categories A and B and the remaining 5707 in categories C and D. The storm severity measure that has the strongest relationship with insured loss will maximise the extremal dependence measure χ_t (Eqn 3.2) for a given threshold t, equivalent to maximising $\frac{a}{a+c}$ in Figure 3.9. This will be explored by comparing the number of storms that are extreme in both insured loss and the severity measure, a, for the $(a + c)$ most extreme events of the severity measure, over a range of values of $a + c$. The greater the value of a for a given $a + c$, the stronger the extremal dependence between the severity measure and insured loss.

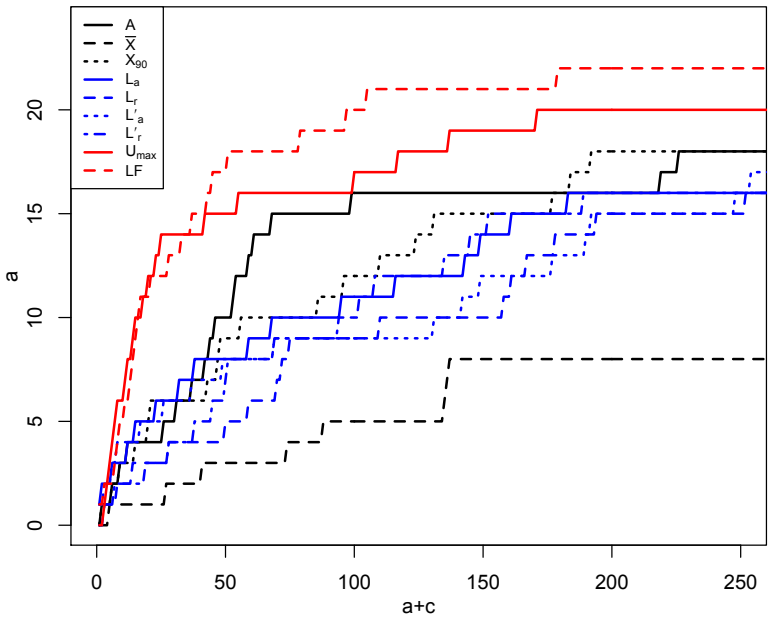

Figure 3.10. The number of storms that are extreme in both insured loss and the severity measure, a, for the $(a + c)$ most extreme values of the severity measure.

Figure 3.10 shows the relationship between a and $(a + c)$ for the 9 storm severity measures. It can be seen that measures U_{max} and particularly LF have the strongest relationships with insured loss. The measure $\bar{x}$ has the weakest relationship and the other five measures have similar relationships, with A being better related to insured loss for $(a + c) \in (50, 150)$.

This shows that the loss function measures (shown in blue), which have been used in previous studies and were found to be a good approximation of losses in Germany, are not the most successful at representing insured loss here. This approximately equal performance of the loss functions also shows that including population density

within the measure does not improve the relationship with insured loss. In addition, using a relative or absolute threshold makes very little difference. Interestingly, measures of peak intensity such as x_{90} and particularly U_{max} are more related to insured loss than these loss functions which are measures of excess wind gust intensity summed over a damage region, and may therefore be expected to be a better representation of overall loss. The measure A is the best footprint based classifier of insured loss, out performing the loss functions and x_{90} and, when combined with U_{max} to form the composite measure LF, gives the best measure for representing insured loss. This means that the historical windstorms that caused large insured losses have either a large footprint damage area, A, or a high maximum 925hPa wind intensity, U_{max}, or both.

Referring back to Table 3.1, now comparing the extreme insurance loss storms in the context of all other storms, it is seen that the two presented measures are extreme and a large area of the footprint exceeding 25ms^{-1} is a common characteristic in the footprints in Figure 3.1.

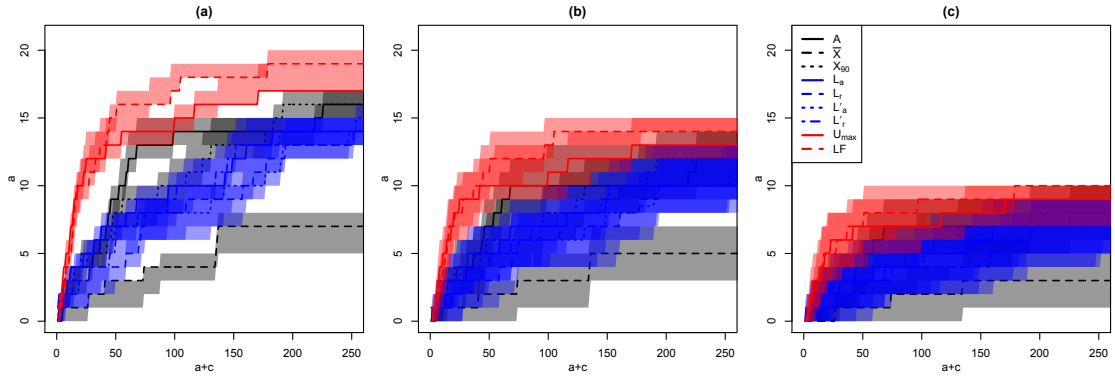

Figure 3.11. The number of storms that are extreme in both insured loss and the severity measure, a, for the $(a+c)$ most extreme values of the severity measure when using (a) 20, (b) 15 and (c) 10 extreme insurance loss storms, selected at random from the list of 23. This is repeated 1000 times for each measure. The median of these 1000 samples is plotted with shaded regions showing the 2.5% and 97.5% quantiles.

This classification method will depend on the list of 23 extreme insurance loss windstorms used. The sensitivity of the results to this list is tested by selecting 20, 15 and 10 of the 23 extreme insurance loss storms at random, 1000 times, and repeating the analysis. This bootstrap method is used to ensure the results are consistent irrespective of which particular storms are selected from the original 23. The resulting plots can be seen in Figure 3.11, showing that the relationships between the severity measures and insured loss are insensitive to the list of extreme insurance loss storms used, and LF is consistently outperforming the other measures.

These conclusions could be very different if quantitative insured loss values were used rather than this classification method. Insured loss data is, however, hard to acquire and a fully rigorous validation would necessitate the availability of European wide insured loss values for each of the 5730 windstorm events in the data set.

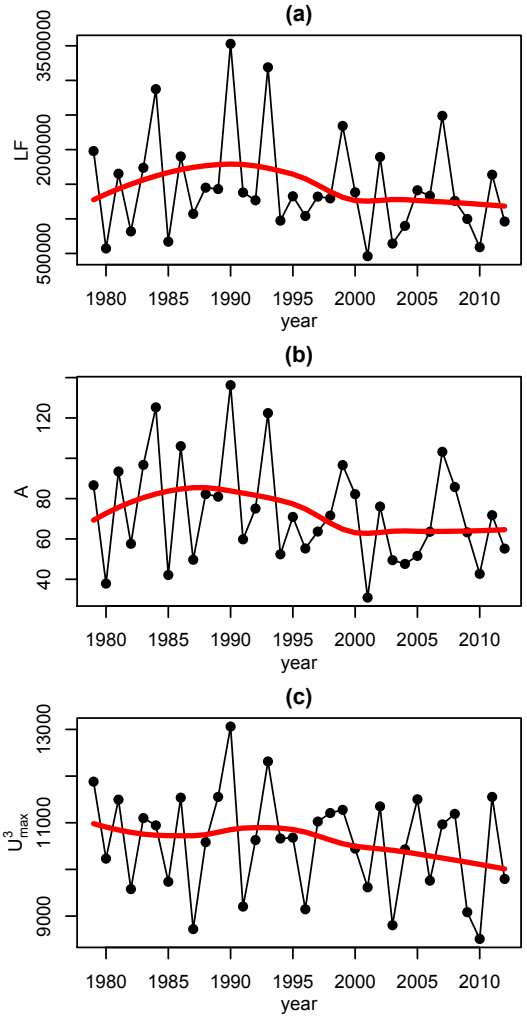

Figure 3.12. Annual means of (a) LF, (b) A and (c) U^3_{max}, based on all 5730 windstorm footprints in the data set. The red curve is the locally weighted scatterplot smoothing curve, fitted to the annual means.

As explained in the Introduction of this thesis, the re/insurance industry observed a decline in European windstorm related losses in the mid 1990's (Mark, 2013). This 20 year lull is reflected in the annual means of LF in the period 1979-2012, shown in Figure 3.12 (a), which drop at around the year 1996. This further suggests that LF is, indeed, a good representation of insured loss. In addition, Figure 3.12 shows how the sudden drop in $LF(= AU^3_{max})$ is explained by a similar drop in A, which is not reflected in the annual means of U^3_{max}, further suggesting that A is also a good representation of insured loss.

3.6. Application: selecting extreme windstorms for the XWS catalogue

In 2012, a group of scientists at the Met Office and the universities of Exeter and Reading decided to create a freely available, digitised catalogue of European windstorm tracks and footprints for 50 of the most extreme windstorms to hit Europe in recent years (1979-2012), named the eXtreme Wind Storms (XWS) catalogue. The catalogue is available at www.EuropeanWindstorms.org and is intended to be a resource for both academia and the re/insurance industry for understanding the variation and factors that influenced past windstorm events. This understanding is thought to be important when evaluating and improving weather, climate and catastrophe models (Roberts et al., 2014). Information about the 50 European windstorms selected and how the database was created are available at this web address, and in more detail in "The XWS open access catalogue of extreme European windstorms from 1979 to 2012" by Roberts et al. (2014).

The 23 extreme insurance loss storms were of particular interest to the insurance community and were therefore essential for inclusion within the catalogue. The storm severity measures were used to select the remaining 27 extreme events for the catalogue. For these additional 27 events to be of interest for the catalogue, from a damage point of view, it was concluded that they should be extreme in the severity measure that best represented insured loss, such that altogether 50 storms were classed as extreme, i.e. in categories A, B and C of Figure 3.9. The additional 27 extreme storms, selected for the catalogue, would therefore have similar characteristics to the 23 extreme insurance loss storms. Referring to Figure 3.9, this was quantified by observing which severity measure maximises $\frac{a}{a+c}$, such that $a + b + c = 50$. Graphically, this was equivalent to maximising a when the line $a = (a + c) - 27$ intercepts the plot of a against $(a + c)$.

The XWS storm selection was done using a different definition of the windstorm footprint where all wind gust speeds outside of a 1000 km radius of the track position were disregarded, reducing the overlap of different windstorm events within one footprint. This was done to ensure that no extreme event was selected twice (Roberts et al., 2014). These footprints are termed 'disaggregated' and are explained in more detail in Section 2.3.1.

Figure 3.13 shows that the severity measure LF maximises the number of extreme insurance loss storms in the top $a + c$ events such that $a + b + c = 50$. Severity measure LF was therefore used to select the 27 other extreme windstorms. The XWS catalogue is therefore made up of the 23 extreme insurance loss storms and 27

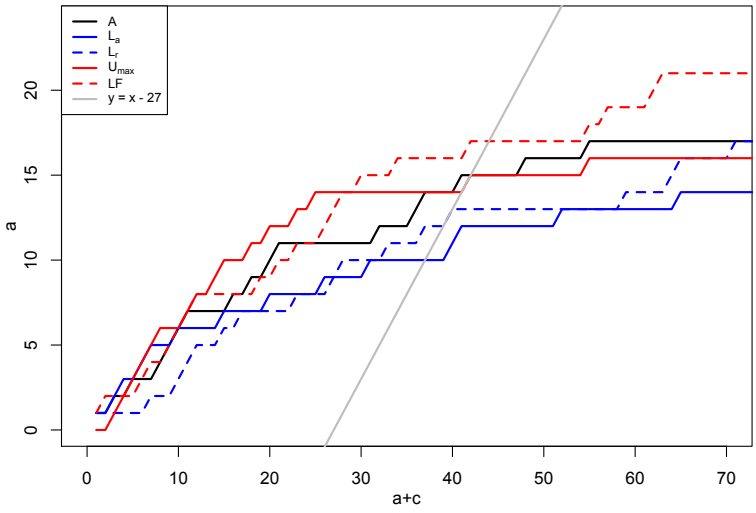

Figure 3.13. The number of storms that are extreme in both insured loss and the severity measure, a, for the $(a + c)$ most extreme values of the severity measure, for severity measures A, L_a, L_r, U_{max} and LF, using disaggregated footprints. The line $a = (a + c) - 27$ identifies which severity measure maximises $\frac{a}{a+c}$ such that 50 storms are selected for the catalogue.

others that are extreme in LF. It is interesting to note that the severity measures perform similarly when using the disaggregated footprints (Figure 3.13), compared to using the raw footprints (Figure 3.10). However measures A, L_a and L_r perform slightly better when the disaggregated footprints are used, most likely because large events overlapping smaller events are removed from the footprints, so less large events are counted twice meaning the 23 extreme insurance loss storms rank higher.

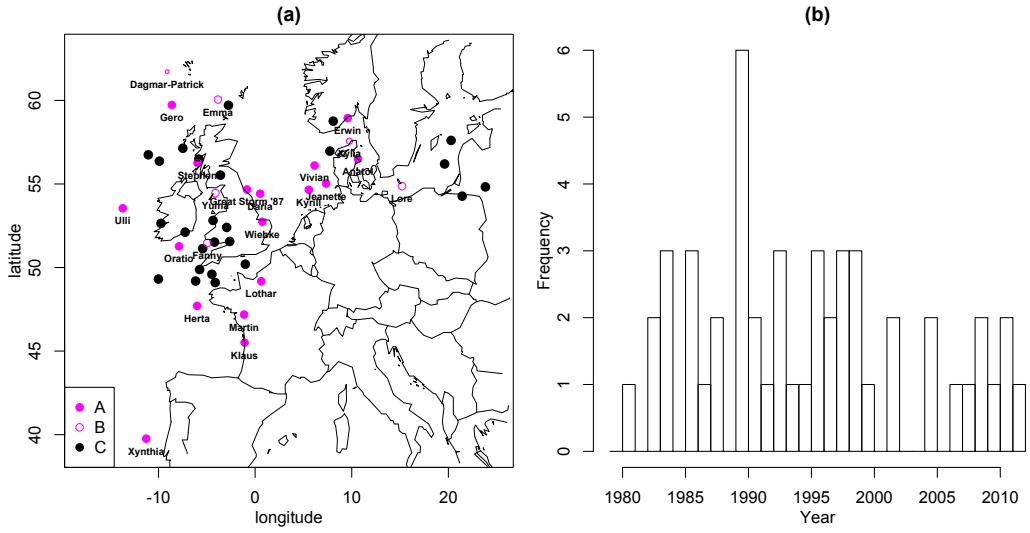

Figure 3.14. (a) The location of the centre of the 850 hPa relative vorticity when the maximum wind speed over land (U_{max}) occurs for the 50 windstorms selected for the XWS catalogue using LF. The size of the points is proportional to the magnitude of LF, and (b) the distribution of years in which the 50 windstorms selected for the XWS catalogue occurred.

The spatial and temporal distribution of the 50 windstorms selected for the XWS catalogue using severity measure LF is shown in Figure 3.14. The peak intensities are generally located around the UK and Northern Europe. These are strong Atlantic storms that are well represented in the reanalysis data used to create the footprints (Roberts et al., 2014). The 50 windstorms selected for the XWS catalogue occur in 26 of the 33 year period, well representing the period spanned by the XWS catalogue. However, more windstorms are selected in the period 1985-1995 and fewer in the period 2000-2010.

3.6.1. Why are category C storms not considered to have caused extreme insured loss?

Severity measure LF is the most successful at representing insured loss, so why have the 27 non-extreme insurance loss or category **C** windstorms, which rank highly in LF, not been identified as having caused extreme insured loss by the re/insurance industry? Is there a particular characteristic of category **C** storms that distinguishes them from the 23 extreme insurance loss storms?

Three of the storms in category C are documented in Lamb and Frydendahl (1991), dated 18/01/83, 01/02/83 and 09/02/88, and can therefore be explored in more detail. Lamb and Frydendahl (1991) state that all 3 of these windstorms caused damage and loss of life, so why are they not included in the list of 23 extreme insurance loss storms? All 3 windstorms occurred mainly over the North sea and the British Isles (Lamb and Frydendahl, 1991) meaning insured loss was localised and did not occur over mainland Europe. They occured early in the period of interest when less people may have been insured and/or overall population might have been lower. These factors may go some way in explaining why these 3 storms were not included in the list of extreme insurance loss storms.

Characteristics of the 50 selected windstorm footprints were compared, differentiating between categories **A**, **B** and **C**. Many characteristics were compared and population density affected and the date of occurrence were found to be of most

interest. Define:

$$P_{a_i} = \sum_{j=1}^{J} H(x_i(s_j) - 25)p(s_j)$$

$$P_{r_i} = \sum_{j=1}^{J} H(x_i(s_j) - c(s_j))p(s_j)$$

representing the total population affected by windstorm i for the absolute and relative exceedance thresholds repsectively.

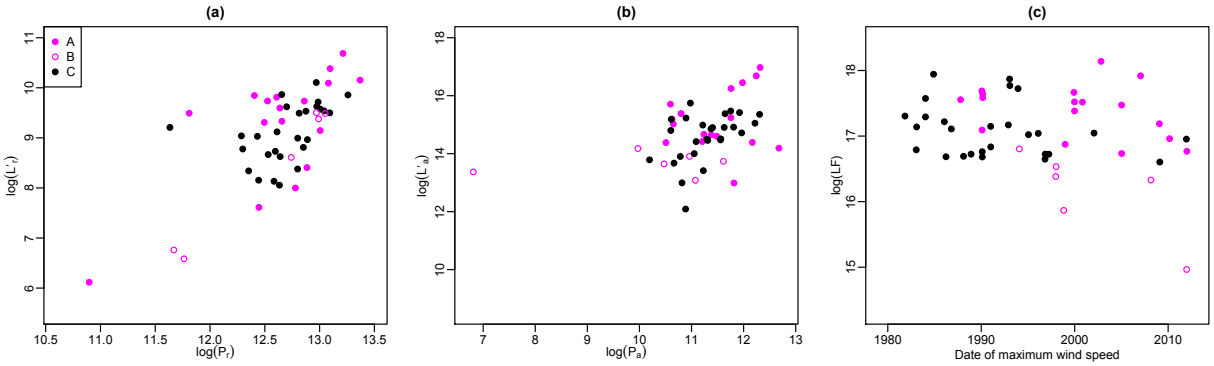

Figure 3.15. The relationship between the population affected by damaging winds and loss functions using (a) the relative 98$^{\text{th}}$ percentile threshold $c(s_j)$ and (b) the absolute threshold of 25ms^{-1}, and (c) the relationship between the date of occurrence of U_{max} and LF. Category **A**, **B** and **C** storms are identified.

Figure 3.15 (a) and (b) reveals that all types of windstorm (category **A**, **B** and **C**) affect a similar overall population. This is particularly true when the absolute 25ms^{-1} threshold is used. Other thresholds of wind gust speed (e.g. 15ms^{-1}, 30ms^{-1}) were investigated but the same conclusion was made in all cases. Figure 3.15 (a) does, however, indicate that the category **C** storms tend to have slightly lower values of L'_r and P_r, suggesting that winds that exceed the relative 98th percentile threshold occur in lesser populated regions for category **C** storms. This may reduce the insurance loss incurred and might explain why category **C** storms are not considered to have caused as much insured losses as the 23 in Table 3.1. The population density as at 2005 is used to calculate L'_r and P_r (from Haylock (2011)). Using population density corresponding to the correct year for each windstorm would improve the analysis and might give a greater distinction between the categories of windstorms. In addition, Figure 3.15 (c) shows that, in general, the category **C** storms occur earlier in the period when less of the population may have been insured, potentially explaining why category **C** storms have not been identified as causing extreme insurance loss.

3.7. Conclusions

Within this chapter windstorm footprints are related to insured loss by exploring which storm severity measures are best at representing insured loss, based on historical footprints. Due to the incompleteness of historical European wide insured loss data, this was quantified in terms of their ability to classify 23 extreme insurance loss windstorms. This methodology was also used to select 50 extreme windstorms for the XWS catalogue.

Nine storm severity measures were introduced, based on existing measures reviewed in detail in Section 2.2.2. These nine measures ranged in complexity and form. Many other storm severity measures could have been included in this investigation, using different damage thresholds and exposure variables, however those included allowed for a number of insights into the relationship between intensity and exposure characteristics.

A strong positive relationship was found to exist between the footprint damage area ($> 25\text{ms}^{-1}$), A, and the absolute threshold loss function, L_a, suggesting that including the intensity of the wind gust speeds in the exceedance area gives little extra information about the severity of the windstorm. Strong dependency was also shown to exist between loss functions with and without population density included, consistent in the extremes. Similarly, suggesting that including population density in the measure adds very little extra information about the severity. This is found to be because of the difference in spatial scale of the wind gust speed excesses and population densities. The composite measure LF has equal dependence on A and U_{max} in the extremes, suggesting that the peak intensity along the track and the footprint damage area are equally important within the composite measure for extreme windstorms.

The area of damaging winds, A, was the best footprint based classifier of insured loss. This measure was shown to outperform the spatial 90^{th} percentile wind gust speed, a measure of the footprint peak intensity, and the loss function severity measures, most commonly used in previous studies. Including the magnitude of the exceedance of the damage threshold and population density, as a proxy for exposure, within the severity measure was shown to be detrimental to the relationship between the measure and insured loss. In addition, using a relative or absolute threshold made very little difference when relating the loss function severity measures to insured loss. The composite severity measure, LF, the product of the footprint damage area and the cube maximum 925hPa wind speed along the track, was shown to give the best overall representation of insured loss, both using raw footprints and disaggregated

footprints. This composite measure was therefore used to select the extreme storms for the XWS catalogue. This suggests that windstorms that cause large insured losses have either a large footprint damage area or a high maximum 925hPa wind intensity, or both.

As explained by Bonazzi et al. (2012), during an extreme windstorm event, insured losses occur at multiple locations, therefore exploring the spatial structure of individual storms is also relevant to understanding how windstorm footprints relate to total insured loss. Characterising a footprint into a scalar severity measure therefore does not capture this possible variation in spatial structure and local intensity. This motivated Bonazzi et al. (2012) to develop a model for the bivariate distribution of footprint wind gust speeds at pairs of locations, in terms of the local intensity and spatial dependence, to explore which of these statistical characteristics has the most influence on insured loss, approximated by a storm severity measure.

In a similar way, in the remainder of this thesis a statistical model for windstorm footprints will be developed and used in a sensitivity study to explore which statistical footprint characteristic, local intensity or spatial dependence structure, has more influence on insured loss. Insured loss will be approximated by the most related footprint based storm severity measure A. Furthermore, this model based sensitivity study will help to explain which statistical footprint characteristic has changed in the last 20 years, causing the observed decline in insured loss in the mid 1990's (Mark, 2013), reflected in the trend in annual means of A.

4. Bivariate modelling of windstorm footprints

4.1. Introduction

This chapter presents the development of a bivariate model for windstorm footprint wind gust speeds at pairs of locations. This is the first stage in developing a spatial model for windstorm footprints over the whole European domain which will be used to explore which of the statistical footprint characteristics, the local intensity or the spatial dependence structure, have the most influence on insured loss. The insured loss will be approximated by the area of the footprint that exceeds 25ms^{-1} over land, found to be the best footprint based storm severity measure at representing insured loss.

The model developed within this chapter will be based on all 5730 windstorm events from extended winters (October - March) 1979 - 2012 and will represent pairs of locations throughout the European domain. The bivariate model must be extendable to higher dimensions to allow for simulation of synthetic windstorm footprints, and must be able to realistically represent joint losses.

Bonazzi et al. (2012) recently developed a bivariate model for footprint wind gust speeds for the same purpose. This model is, however, based on North-West Europe, only 135 extreme windstorm footprints and wind gust speeds at pairs of locations are assumed to be asymptotically dependent, an assumption which is not validated. Bonazzi et al. (2012) primarily focused on the very extreme wind gust speeds, while the model developed here aims to realistically model the relationship between damaging wind gust speeds, which occur lower in the tail of the joint distribution. The Bonazzi et al. (2012) bivariate model will be applied alongside the bivariate model developed in this chapter as a comparison.

This chapter aims to address three questions:

- How should the marginal distribution of damaging wind gust speeds at each

location be modelled?

- How should the dependence in extreme wind gust speeds at pairs of locations be modelled, i.e. is there extremal dependence between locations?
- Can a bivariate model for damaging wind gust speeds be developed that has the potential to be extended to represent the entire footprint and realistically represent joint losses?

The general modelling approach will first be introduced. The model specification, fitting method and validation will then be presented. The ability of the bivariate model to realistically represent joint losses, approximated by footprint area, will then be explored.

4.2. General approach

Extreme value theory is a well established topic in statistics with many domains of application (Coles 2001; Davison et al. 2012). These concepts are particularly useful when modelling natural hazards that by their nature are rare events in the tail of the distribution (e.g., Coles and Tawn 2005, Jagger and Elsner 2005, Blanchet et al. 2009). However, there has been little use of extreme value theory to represent extreme windstorms or extratropical cyclones (Economou et al., 2014). Lionello et al. (2008) used the Generalised Extreme Value (GEV) distribution to model pressure minima over Europe to investigate changes in future cyclone climatology; Della-Marta et al. (2009), Della-Marta and Pinto (2009) and Sienz et al. (2010) used a Generalised Pareto Distribution (GPD) model to analyse future changes in extreme wind intensity; and Economou et al. (2014) used a Bayesian hierarchical model (BHM) to spatially and temporally model extreme cyclones. These papers focus on extremes of cyclone track variables rather than windstorm footprints. Della-Marta et al. (2009) used the GPD to model the exceedences of scalar storm severity measures above a suitably chosen threshold, applied to peak gust wind fields, to investigate the return period of windstorms over Europe; and Bonazzi et al. (2012) used the GPD to model the marginal distribution of footprint wind gust speeds at meteorological stations.

As discussed in Section 2.2.3, a number of approaches for modelling multivariate extremes have been used in previous studies. These include Bayesian Hierarchical Models (BHMs), also known as latent variable models, max-stable random fields and copulas. As explained in Section 2.2.3, each method has advantages and disadvantages in terms of flexibility, dimensionality and complexity of implementation. When modelling bivariate extremes, it is of particular importance to correctly rep-

resent the extremal dependence, which characterises the joint behaviour of variables as they become large. Coles et al. (1999) explained how the extremal dependence between variables can take one of two forms: asymptotic dependence, where large values of the variables tend to occur simultaneously, or asymptotic independence, where large values rarely occur together. Most available statistical models for multivariate extremes are suitable for either one case or the other, but not both. Bonazzi et al. (2012) used an extreme value copula to model the dependence in windstorm footprint wind gust speeds at pairs of locations. This approach assumes asymptotic dependence, an assumption that is not validated. Within this chapter the extreme dependence property will be explored to validate the choice of model used, ensuring the correct representation of the joint distribution of footprint wind gusts.

The bivariate model for footprint wind gust speeds developed within this chapter must be easily extendable to larger dimensions, to enable modelling of the footprint over the whole European domain, but can be in finite dimensions since simulation from the model will be on a grid of a given resolution. The modelling approach must allow for the selection of an appropriate dependence model depending on what extremal dependence behaviour is identified in the data. From the discussions in Section 2.2.3, these requirements suggest the use of a copula approach, which can represent asymptotically dependent variables using extreme value copulas, or asymptotically independent variables using the Gaussian copula. If a mixture of asymptotic dependence and independence is identified, more complex methods such as the method proposed by Heffernan and Tawn (2004) (see Section 2.2.3) could be explored. Using the copula approach has the added advantage of being directly comparable to the Bonazzi et al. (2012) model. In addition, since the copula approach represents the joint distribution of random variables in terms of their marginal distributions and their mutual dependence separately, the local intensity and spatial dependence structure will be represented separately within the model parameters, allowing for their influence on insured loss to be easily explored within a sensitivity study.

The copula is a function that is used to model the joint distribution of the two variables transformed to uniform margins using their marginal distributions. Mikosch (2006) suggested that this transformation to uniform margins is mathematically trivial and may obscure important features of the data that are visible on their original scale. However, Davison et al. (2012) argued that the implicit separation of the marginal distributions of variables from their dependence structure provides a unifying framework for modelling multivariate data. The copula methodology is well developed in comparison to max-stable theory (Blanchet and Davison, 2011), which can only represent asymptotic dependence or complete dependence (see Sec-

tion 2.2.3), and a large number of copula functions are available to choose from giving a great deal of flexibility in model specification. Therefore, despite possible drawbacks, the use of copulas to model multivariate extremes has become very popular with one of the driving forces for the popularity being their application in the context of financial risk management (Mikosch, 2006). The continued popularity of the copula method stems from their relative ease of use, applicability to a wide range of situations (Renard and Lang, 2007) and, since the copula often has only one unknown parameter, the model fitting and interpretation of the dependence is simplified. It could be argued, however, that using only one parameter to explain dependence could limit the flexibility of the model.

Sklar's Theorem (Nelson (2006), pages 17-24) states that the 2-dimensional joint distribution of random variables X_1 and X_2, $F(x_1, x_2) = \Pr(X_1 \leq x_1, X_2 \leq x_2)$, may be written as:

$$F(x_1, x_2) = C\{F_1(x_1), F_2(x_2)\}$$

where $F_1(x_1) = \Pr(X_1 \leq x_1)$ and $F_2(x_2) = \Pr(X_2 \leq x_2)$ are the univariate marginal distributions of X_1 and X_2, and C, known as the copula, is a 2-dimensional distribution on $[0,1]^2$ that contains complete information about the dependence between X_1 and X_2. The function C is uniquely determined for distributions F_1 and F_2 with absolutely continuous margins (Davison et al., 2012).

Firstly, the form of the marginal distributions F_1 and F_2 will be explored, followed by the selection of an appropriate copula model. These elements of the model, followed by the overall model, will then be validated to ensure realistic representation of joint losses for pairs of locations throughout the domain, approximated by footprint area.

Bonazzi et al. (2012) also used a copula approach in their bivariate footprint model. Their model has different specifications of both the marginal distributions and copula function from those of the model developed here and will therefore be directly compared throughout as a validation of the modelling choices made.

4.3. Modelling the marginal distributions

4.3.1. Marginal distribution specification

Let X be a vector of maxima, each over δ independent random variables with a common distribution function. Then by the Extremal Types Theorem, first discovered by Fisher and Tippett (1928), and later proved in full generality by Gnedenko

(1943), there exists sequences of constants $\{a_\delta > 0\}$ and $\{b_\delta\}$ such that

$$\Pr\{(X - b_\delta)/a_\delta \leq z\} \to G(z) \quad \text{as } \delta \to \infty$$

where appropriate choices of $\{a_\delta > 0\}$ and $\{b_\delta\}$ stabilise the location and scale of $(X - b_\delta)/a_\delta$ as δ increases (Coles, 2001), and G is a non-degenerate distribution function. Then G belongs to one of three classes of extreme value distributions known as the Gumbel, Fréchet and Weibull families. These three families can be combined into a single family of models known as the Generalised Extreme Value family, having the distribution function of the form:

$$G(z) = \exp\left\{-\left[1 + \xi\left(\frac{z-\mu}{\sigma}\right)\right]^{1/\xi}\right\}$$

defined on the set $\{z : 1 + \xi(z-\mu)/\sigma > 0\}$ and the parameters satisfy $-\infty < \mu < \infty$, $\sigma > 0$ and $-\infty < \xi < \infty$ where μ is the location, σ is the scale and ξ is the shape parameter.

For a pair of locations, (s_1, s_2), the wind gust speeds taken from the 5730 windstorm footprints in the data set (Section 2.3.1), denoted X_1 and X_2, are a vector of maxima, since they are the maximum wind gust speed to occur in each location over the 72 hour period of each windstorm event. The variables X_1 and X_2 should therefore, by the Extremal Types Theorem, be modelled using the GEV distribution. Since the aim is to develop a model that realistically represents extreme, damaging footprint wind gust speeds, it is important that the marginal distributions fit very well in the upper tail. To improve the fit in the upper tail the GEV marginal distribution is based only on wind gust speeds that exceed a high threshold. This will be referred to as the left-truncated GEV distribution. In doing this, the lower wind gust speeds do not penalise the fit of the distribution in the upper tail and wind gust speeds are always distributed as being non-negative, consistent with their physical nature.

The vectors of maxima X_1 and X_2 are distributed according to the GEV distribution; therefore $X_1 > v_1$ and $X_2 > v_2$ must be distributed according to the left-truncated GEV distribution, where v_1 and v_2 are left truncation thresholds.

Let G and g denote the GEV distribution and density functions respectively and G^* and g^* the GEV distribution and density functions truncated at threshold t respec-

tively. Then for random variable $Z = (z_1, .., z_n)$ which follows a GEV distribution:

$$\begin{aligned} \Pr(Z \leq z) &= G(z; \mu, \sigma, \xi) \\ \Rightarrow \Pr(Z \leq z | Z > t) &= G^*(z; \mu, \sigma, \xi) = \frac{G(z; \mu, \sigma, \xi) - G(t; \mu, \sigma, \xi)}{1 - G(t; \mu, \sigma, \xi)}, \end{aligned}$$

and

$$g^*(z; \mu, \sigma, \xi) = \frac{g(z; \mu, \sigma, \xi)}{1 - G(t; \mu, \sigma, \xi)},$$

which is a valid probability density function because it integrates to 1 above the truncation threshold, t.

Here, the random variables X_1 and X_2 are thought to be left-truncated GEV distributed above thresholds v_1 and v_2 respectively, however the distribution for each variable is needed for the whole support of the variable, to allow for the validation of the fit of the model against the data e.g. quantile-quantile plots. Below the truncation threshold the distribution is approximated by the empirical distribution function of the variable, and above the truncation threshold the left-truncated GEV distribution is scaled to ensure the density integrates to 1 over the whole support of the variable. That is,

$$\Pr(X_j \leq x) = F_j(x; \mu_j, \sigma_j, \xi_j) = \begin{cases} \hat{F}_j(x) & x \leq v_j \\ (1 - \lambda_j) + \lambda_j G_j^*(x; \mu_j, \sigma_j, \xi_j) & x > v_j \end{cases}$$

for $j = 1, 2$, where $\hat{F}_j$ is the empirical distribution of X_j and $\lambda_j = 1 - \hat{F}_j(v_j)$.

In comparison, Bonazzi et al. (2012) used the Generalised Pareto Distribution (GPD) to model the margins. The GPD is a broadly applicable model for the tail of a sequence of independent and identically distributed random variables, above a high threshold (Coles, 2001). Let X be an arbitrary term in a sequence of independent and identically distributed random variables then, for large enough γ and $X > \gamma$:

$$\Pr(X \leq x) \approx H(x; \alpha, \beta) = 1 - \psi \left(1 + \frac{\alpha(x - \gamma)}{\beta}\right)^{-1/\alpha}$$

where $\psi = H(\gamma; \alpha, \beta)$, $\alpha \neq 0$ is the shape parameter, equivalent to ξ of the GEV distribution, and β is the scale parameter.

4.3.2. Marginal distribution fitting

Maximum Likelihood (ML) estimation, popularised by Fisher (1912), is a well known and commonly used parametric method for inference in statistics. This method finds the set of parameters that maximises the probability of observing the data given the model. Under mild regularity conditions (Cox and Hinkley, 1974) the ML estimator is asymptotically normally distributed, unbiased and fully efficient (Diggle and Ribeiro, 2007).

Let $X = x_1, ..., x_n$ denote the vector of maximum wind gust speeds associated with the $n = 5730$ windstorm footprints, at a given location. For the ML estimation approach to be applicable $x_1, ..., x_n$ should be independent. This may not be the case for all x_i because, as discussed in Sections 2.1.3 and 2.3.1, there is a possibility that more than one windstorm passes through the European region in a given 72 hour period. However, in previous studies, the 72 hour event duration is most commonly used to achieve the greatest independence between windstorm events. Haylock (2011) discussed how choosing a shorter period would increase the risk of one windstorm generating more than one footprint, violating the independence between events required to carry out statistical analysis, while Donat et al. (2011b) explained how increasing the event duration to 5 days results in some consecutive storms being considered as one event, so some events would be lost from the data set.

Assuming independence between windstorm events, the parameters of the left-truncated GEV distribution fitted to X are estimated by maximising the likelihood function:

$$L((\mu, \sigma, \xi); x_1, ..., x_n) = \prod_{x_i : x > v} \frac{g(x_i; \mu, \sigma, \xi)}{1 - G(v; \mu, \sigma, \xi)}$$

where g and G are the density and cumulative distribution function for the standard GEV respectively. This maximisation of the likelihood function is done using the Nelder and Mead (1965) optimisation method.

In order to fit this left-truncated GEV marginal distribution, an appropriate left truncation threshold, v, must be selected for each location. This threshold is specified as a quantile of the wind gust speed at each location to ensure that the GEV is fit to a large enough sample of the data. The quantile threshold must be high enough that the lower wind gust speeds do not negatively impact on the fit of the distribution in the upper tail, but low enough that all damaging wind gust speeds are modelled. In Chapter 3, the area of the footprint exceeding 25ms^{-1} over land was found to be the best severity measure, based on the footprint only, at representing insured loss. Wind gust speeds above this threshold must therefore be modelled. As discussed in Section 2.2.2, a number of studies (e.g., Klawa and Ulbrich 2003,

Leckebusch et al. 2007, Pinto et al. 2007) argue that a local threshold should be used, to account for local adaptation to the wind climatology. They use the local 98$^{\text{th}}$ percentile of the climatology wind gust speed as a loss threshold. The 60% quantile is found to be the highest quantile such that all wind gust speeds greater than or equal to 25ms^{-1} and the local 98$^{\text{th}}$ percentile climatology wind gust speed, calculated for the reference period October 1979 to March 2012 (shown in Figure 3.3), are modelled.

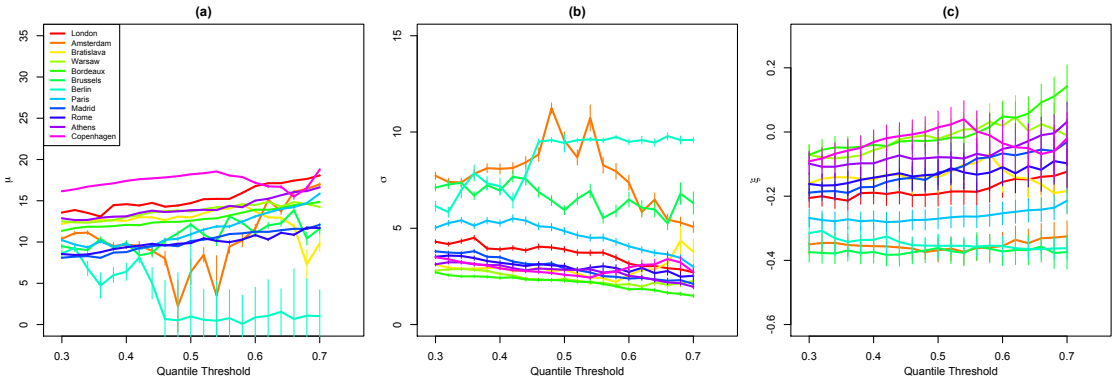

Figure 4.1. The sensitivity of the left-truncated GEV location, (a), scale, (b), and shape, (c), parameters to a change in left-truncation threshold. The vertical lines show the 95% confidence intervals based on the assumption of asymptotic normality of ML estimates.

Figure 4.1 shows the sensitivity of the left-truncated GEV location, scale and shape parameters, μ, σ and ξ to a change in the quantile left truncation threshold for 12 locations in Europe. These 12 locations are chosen because they are located throughout the European domain so representing the general fit in different regions. On the whole, the estimated parameters are not very sensitive to the threshold choice, although some variation exists in μ and σ for Amsterdam and Bordeaux. These two parameters, however, compensate one another in the left-truncated GEV distribution function. The 60% quantile threshold gives stable values for the parameters in most cases, ensures damaging wind gust speeds are modelled and is high enough in the tail to ensure a good fit to extreme winds.

The GPD margins must also be fitted above an appropriately high threshold, now to satisfy extreme value theory (Coles, 2001). Similarly, threshold selection involves fitting the GPD at a range of thresholds and observing the stability of the parameter estimates (Coles, 2001). The argument for this methodology is that, if a GPD is a reasonable model for excesses of a threshold ν_0, then excesses of a higher threshold ν should also follow a GPD (Davison and Smith, 1990). The shape parameter, α and the modified scale, defined as $\beta^* = \beta_\nu - \alpha\nu$, where β_ν is the scale parameter estimate when using threshold ν, should be consistent above ν_0 if ν_0 is a valid threshold for

excesses to follow the GPD.

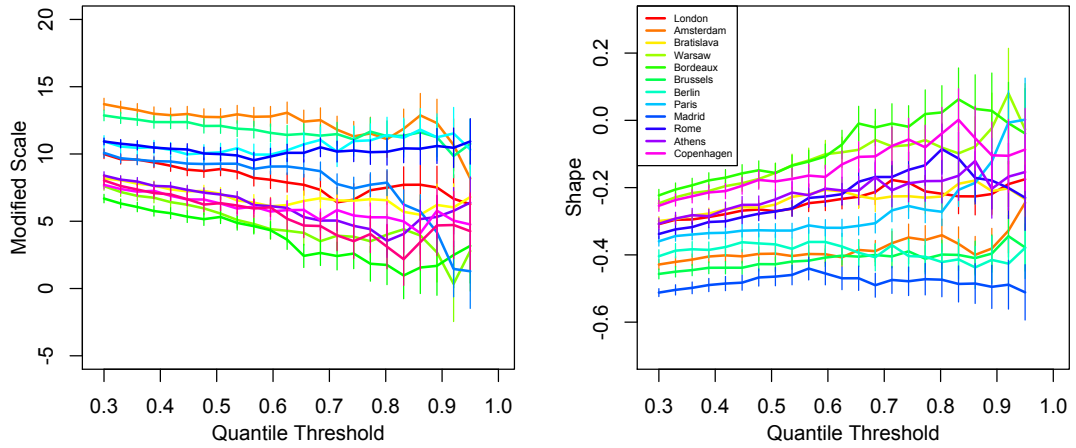

Figure 4.2. The sensitivity of GPD (a) modified scale and (b) shape parameters to a change in GPD threshold. The vertical lines show the 95% confidence intervals based on the assumption of asymptotic normality of ML estimates.

Again, a quantile threshold is preferable to ensure enough data is used to fit the distribution. Figure 4.2 shows the sensitivity of the modified scale, β^* and shape, α, parameters of the GPD distribution to a change in threshold for the same 12 locations. A quantile threshold of 0.7 is chosen since both the modified scale and the shape parameter vary systematically for quantile thresholds below this, but remain relatively constant above. In some locations, however, the 70% quantile is greater than the 25ms^{-1} damage threshold. The minimum of the 70% quantile and 25ms^{-1} is therefore used to ensure that all damaging winds are modelled.

4.3.3. Marginal distribution validation

The fit of the GEV, left-truncated GEV and GPD marginal distributions are compared for four locations, London, Amsterdam, Berlin and Paris. These particular locations are shown because they are places that generally experience high wind gust speeds and must therefore be well modelled to ensure accurate loss estimates. Figure 4.3 shows that all three distributions appear to fit well with the left-truncated GEV and GPD fits looking very similar. It is hard to tell from these plots how well the distributions are fitting in the very upper tail of the distribution which is most important for modelling realistic losses. Quantile-Quantile (Q-Q), and in particular return level plots represent the upper tail of the distribution better. These are shown in Figures 4.4 and 4.5.

Figures 4.4 and 4.5 show that both the left-truncated GEV adjusted marginal distribution ($F(x), x > v$) and the GPD distribution ($H(x), x > \gamma$) are a very good fit to the data at the 4 locations presented, with the left-truncated GEV performing much better in the upper tail of the distribution compared to the standard GEV ($G(x)$). This is found to be consistent throughout the European domain.

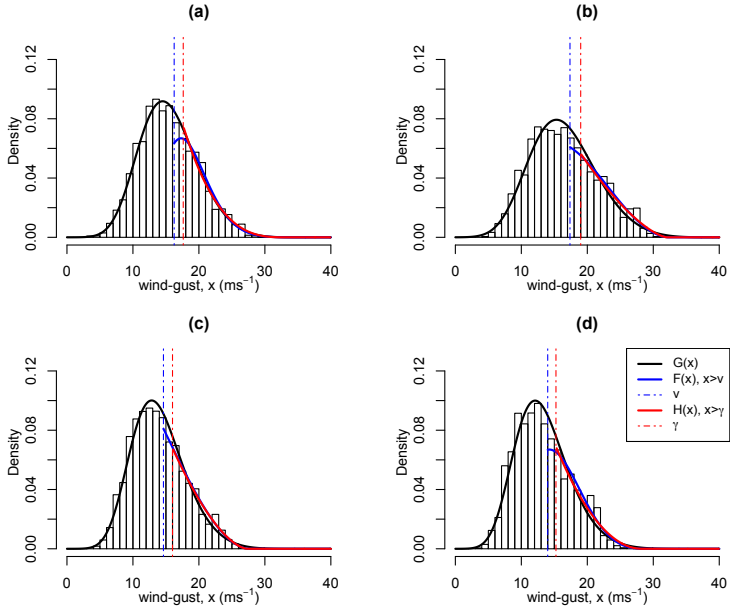

Figure 4.3. Histogram for footprint wind gusts at (a) London, (b) Amsterdam, (c) Berlin and (d) Paris with the best fit GEV, $G(x)$, left-truncated GEV, $F(x)$ for $x > v$ and GPD, $H(x)$ for $x > \gamma$ distributions. Vertical lines show the thresholds v and γ, above which the left-truncated GEV and GPD are fit.

Figure 4.6 (a)-(d) shows the resulting location μ, scale σ and shape ξ parameters and the truncation threshold v of the left truncated GEV for each land location in the domain. Figure 4.6 (e) shows the upper limit of the distribution calculated as $\mu - \sigma/\xi$ valid for $\xi < 0$. When $\xi \geq 0$ the upper tail of the distribution is infinite.

For most locations the shape parameter is negative meaning there is a finite upper limit to the GEV distribution. An upper limit to the wind gust speed at a given location is physically reasonable, so a negative shape parameter is preferable. The Alps region of Italy is an exception to this, where a positive shape parameter leading to an infinite upper tail is found to be because the majority of the distribution of wind gust speeds are very low with a small number of higher wind gusts extending the upper tail of the distribution. The location and truncation threshold are highest in the UK and the north coast of mainland Europe as would be expected since the

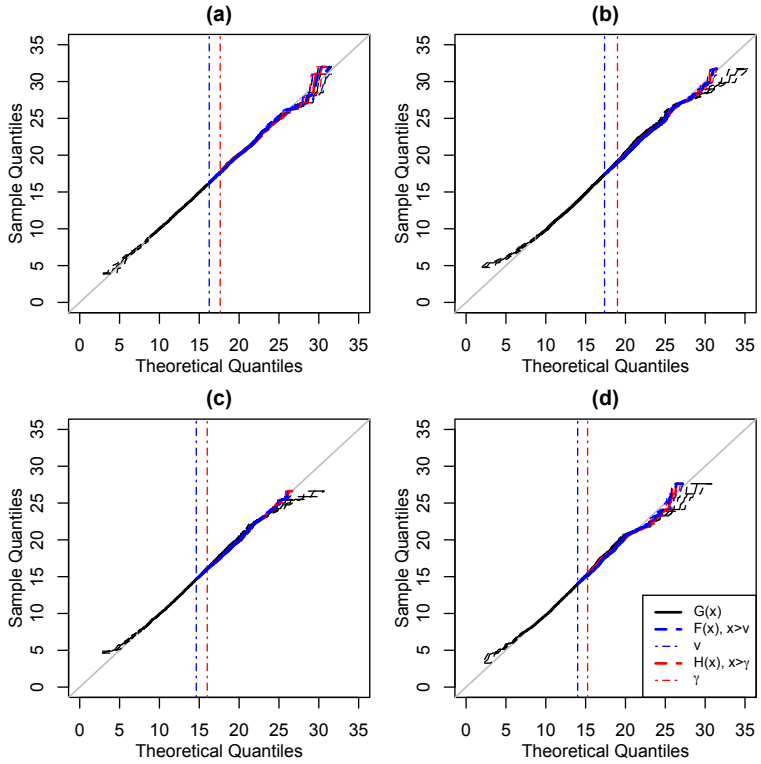

Figure 4.4. Q-Q plots for the GEV, $G(x)$, left-truncated GEV, $F(x)$, and GPD, $H(x)$, marginal distributions fitted to 4 cities in Europe. Dashed lines represent bootstrap 95% confidence intervals.

majority of windstorms cross this region and wind gusts are known to be generally stronger in the north (see Figure 3.1). The scale parameter is relatively consistent throughout the domain, with the exception of south Germany where this parameter is much higher, compensated by a lower location parameter. The quantile-quantile and return level plots for these unusual locations still show that the left-truncated GEV is a good fit to the data.

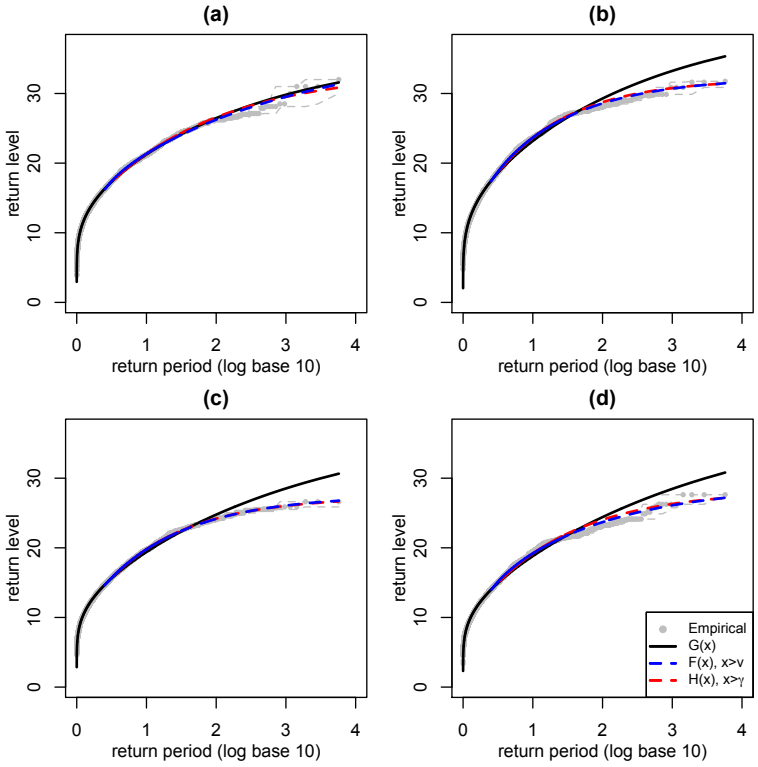

Figure 4.5. Return level plots for the GEV, $G(x)$, left-truncated GEV, $F(x)$, and GPD, $H(x)$, marginal distributions fitted to 4 cities in Europe. Dashed lines represent bootstrap 95% confidence intervals of the empirical return levels.

4.4. Modelling bivariate dependency

4.4.1. Bivariate dependence model specification

The bivariate dependence between footprint wind gusts at pairs of locations will be modelled using a copula function. Mikosch (2006) describes how the most popular copula functions are mostly chosen because they are mathematically convenient and the rationale for their applications is often not well justified. Bonazzi et al. (2012) is an example of this, where the justification for using the Gumbel copula is that it is "in line with many examples of this model found in the literature" (Bonazzi et al. 2012, page 1773). Here the copula selection will be thoroughly validated, based on the extremal dependence properties identified within the data. Initially some example copula functions will be introduced, the extremal dependence properties of the data will then be assessed and an appropriate copula selected to model the bivariate dependency in footprint wind gust speeds at pairs of locations.

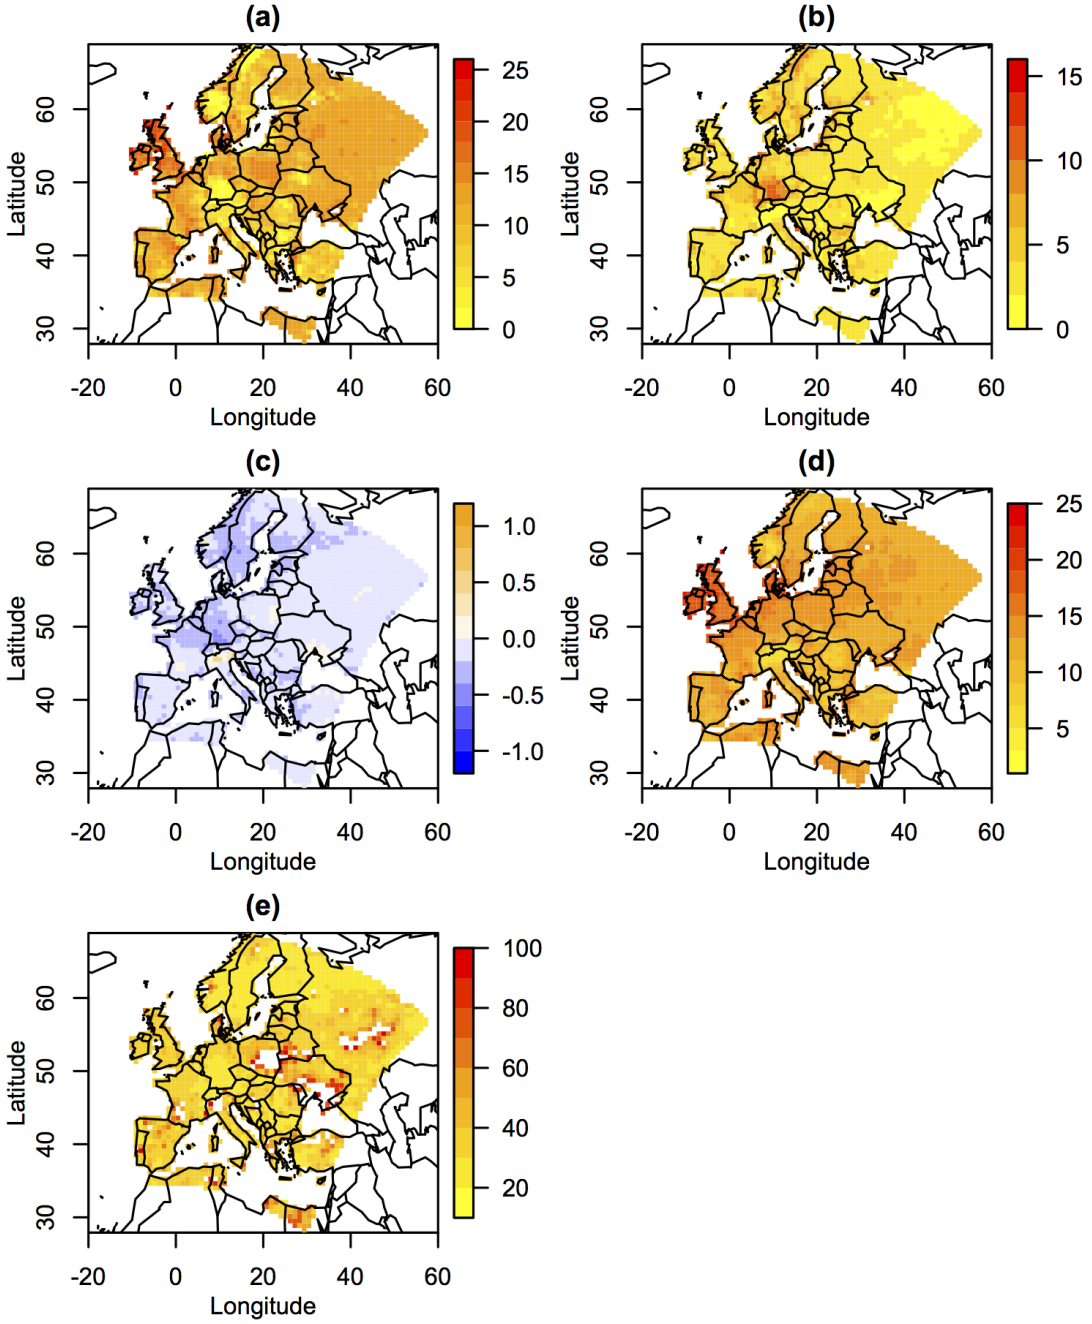

Figure 4.6. left-truncated GEV (a) location, (b) scale, (c) shape, (d) truncation threshold and (e) upper limit, for each location in the domain over land. Locations with a positive shape parameter will have no finite upper limit and are therefore left white in (e).

Let

$$U = F_1(X_1)$$
$$V = F_2(X_2)$$
$$C(u, v) = \Pr(U \leq u, V \leq v)$$

There are infinitely many parametric copula functions that can represent C, four of which are defined below, shown graphically in Figure 4.7. These four copula functions are compared here because they represent a range of dependence and extremal dependence behaviour.

1. **Bivariate Independent Copula**:
   $C(u, v) = uv$
   The joint distribution of (U, V) is bivariate uniform since $f(u, v) = \frac{\partial^2 C}{\partial u \partial v} = 1$.

2. **Bivariate Gaussian Copula**:
   $C(u, v) = \Phi(\Phi^{-1}(u), \Phi^{-1}(v)) = \int_{-\infty}^{\Phi^{-1}(u)} \int_{-\infty}^{\Phi^{-1}(v)} \frac{1}{2\pi(1-\rho^2)^{1/2}} \exp\left\{-\frac{s^2-2\rho st+t^2}{2(1-\rho^2)}\right\} dsdt$
   Where Φ is the standard Gaussian cumulative distribution function, $(\Phi^{-1}(u), \Phi^{-1}(v))^T \sim MVN\left(\binom{0}{0}, \binom{1 \quad \rho}{\rho \quad 1}\right)$ and ρ is the correlation between $\Phi^{-1}(u)$ and $\Phi^{-1}(v)$.

3. **Bivariate Gumbel Copula**:
   $C(u, v) = \exp[-\{(-\log(u))^r + (-\log(v))^r\}^{1/r}]$
   Where $r = 1$ represents independence and $r = \infty$ represents perfect dependence. The Gumbel copula belongs to the family of multivariate extreme value copulas (Galambos, 1987).

4. **Bivariate Perfectly Dependent Copula**:
   $C(u, v) = min(u, v)$.
   Since X_1 and X_2 are perfectly dependent, the variable transformed to uniform margins as U and V, will be equal. The joint distribution of (U, V) becomes univariate, based on U or V, depending on whether u or v is smaller.

As explained in Section 2.2.3, different copula functions represent different extremal dependence behaviour. The empirical extremal dependence behaviour must therefore be explored and the most appropriate copula function selected to model the dependence. This extremal dependence behaviour can be quantified using two measures, introduced by Coles et al. (1999). These measures are defined as

$$\chi(u) = \Pr(V > u | U > u) \quad \text{for } 0 \leq u \leq 1 \tag{4.1}$$

$$\bar{\chi}(u) = \frac{2\log\Pr(U > u)}{\log\Pr(U > u, V > u)} - 1 \quad \text{for } 0 \leq u \leq 1 \tag{4.2}$$

The measure $\chi(u) \in [0, 1]$, also known as the Extremal Dependence Coefficient, represents the conditional probability of V being extreme given U is extreme. If the asymptotic limit of $\chi(u)$ is zero, as $u \to 1$, U and V are said to be asymptotically independent, while if this limit is non-zero, U and V are said to be asymptotically dependent.

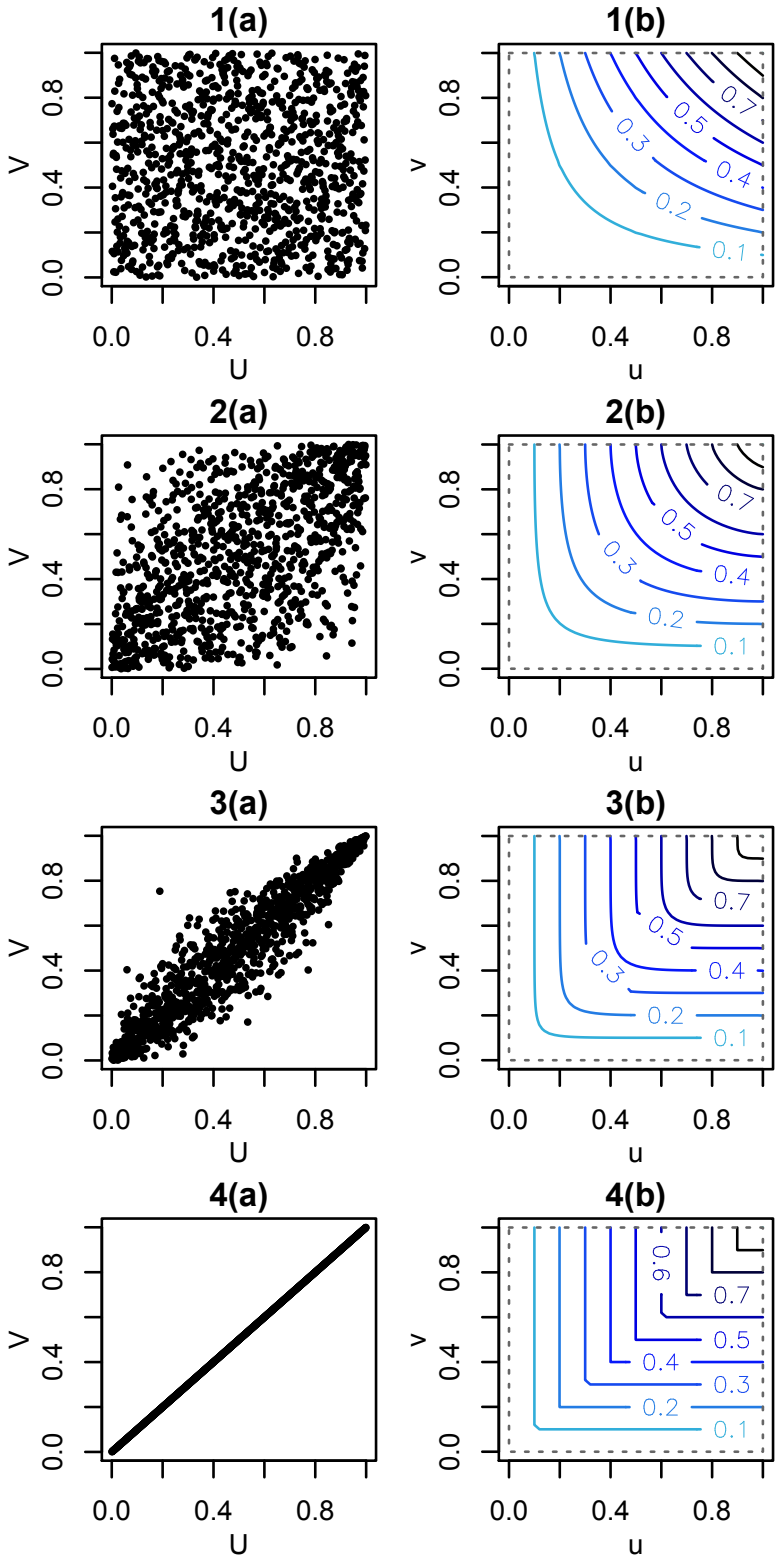

Figure 4.7. (a) Scatter plots of (U, V) simulated from bivariate copula functions 1-4 (defined above) where $\rho = 0.6$ and $r = 5$, and (b) Contour plots of the associated copula functions.

Denote $\chi = \lim_{u \to 1} \chi(u)$, then when variables are asymptotically independent $\chi = 0$ and is therefore unable to provide information on the relative strength of dependence within this class (Coles et al., 1999). To overcome this Coles et al. (1999)

introduce the non-vanishing extremal dependence measure $-1 < \bar{\chi}(u) \leq 1$. Denote $\bar{\chi} = \lim_{u \to 1} \bar{\chi}(u)$ then the complete pair $(\chi, \bar{\chi})$ is required as a summary of extremal dependence (Coles et al., 1999): $(\chi > 0, \bar{\chi} = 1)$ represents asymptotic dependence, and the value of χ measures the strength of asymptotic dependence, or $(\chi = 0, \bar{\chi} < 1)$ represents asymptotic independence, and the value of $\bar{\chi}$ represents the strength of the non-asymptotic dependence.

The extremal dependence measures, $\chi(u)$ and $\bar{\chi}(u)$, can be defined in a number of ways, aiding in calculating the asymptotic limit for the four copula functions previously introduced:

$$\begin{aligned}
\chi(u) &= \Pr(V > u | U > u) \quad \text{for } 0 \leq u \leq 1 \\
&= \frac{\Pr(V > u, U > u)}{\Pr(U > u)} \\
&= \frac{\bar{C}(u,u)}{1-u} \\
&= \frac{1 - 2u + C(u,u)}{1-u} \\
&= 2 - \frac{1 - C(u,u)}{1-u} \\
&\approx 2 - \frac{\log(C(u,u))}{\log(u)} \quad \text{as } u \to 1
\end{aligned}$$

$$\begin{aligned}
\bar{\chi}(u) &= \frac{2\log\Pr(U > u)}{\log\Pr(U > u, V > u)} - 1 \quad \text{for } 0 \leq u \leq 1 \\
&= \frac{2\log(1-u)}{\log\bar{C}(u,u)} - 1
\end{aligned} \tag{4.3}$$

where $\bar{C}(u,v) = \Pr(U > u, V > v) = 1 - u - v + C(u,v)$ is known as the survivor function. The extremal dependence property can therefore be identified for each of the four copula functions by observing the behaviour of $\chi(u)$ and $\bar{\chi}(u)$ as $u \to 1$:

1. **Bivariate Independent Copula**:
   $\bar{C}(u,u) = (1-u)^2$
   $\chi(u) = \frac{(1-u)^2}{1-u} = 1 - u \to 0$ as $u \to 1$
   $\bar{\chi}(u) = \frac{2\log(1-u)}{\log((1-u)^2)} - 1 = 0 < 1$
   $(\chi = 0, \bar{\chi} < 1) \Rightarrow$ **Asymptotic independence**

2. **Bivariate Gaussian Copula**:
   $\bar{C}(u,u) \sim c_\rho(-\log(1-u))^{\rho/(1+\rho)}(1-u)^{2/(1+\rho)}$, (Coles et al., 1999)
   where $c_\rho = (1+\rho)^{3/2}(1-\rho)^{1/2}(4\pi)^{-\rho/(1+\rho)}$
   $\chi(u) = \frac{\bar{C}(u,u)}{1-u} \to 0$ as $u \to 1$
   $\bar{\chi}(u) = \frac{2\log(1-u)}{\log\bar{C}(u,u)} - 1 = \rho < 1$ as $u \to 1$, giving a useful interpretation of $\bar{\chi}$

$(\chi = 0, \bar{\chi} < 1) \Rightarrow$ **Asymptotic independence**

3. **Bivariate Gumbel Copula**:
   $\chi(u) \approx 2 - \frac{(2\log(u)^r)^{1/r}}{\log(u)} = 2 - 2^{1/r} > 0$
   $C(u,u) = \exp(((-2\log(u))^r)^{1/r}) = u^{-2}$
   $\bar{\chi}(u) = \frac{2\log(1-u)}{\log(1-2u+u^{-2})} - 1 = 2\frac{\log(1-u)}{\log((1-u)2u^2)} - 1 \to 1$ as $u \to 1$
   $(\chi > 0, \bar{\chi} = 1) \Rightarrow$ **Asymptotic dependence**

4. **Bivariate Perfectly Dependent Copula**:
   $\bar{C}(u,u) = 1 - 2u + u = 1 - u$
   $\chi(u) = \frac{1-u}{1-u} = 1 > 0$
   $\bar{\chi} = 2\frac{\log(1-u)}{\log(1-u)} - 1 = 1$
   $(\chi > 0, \bar{\chi} = 1) \Rightarrow$ **Asymptotic dependence**

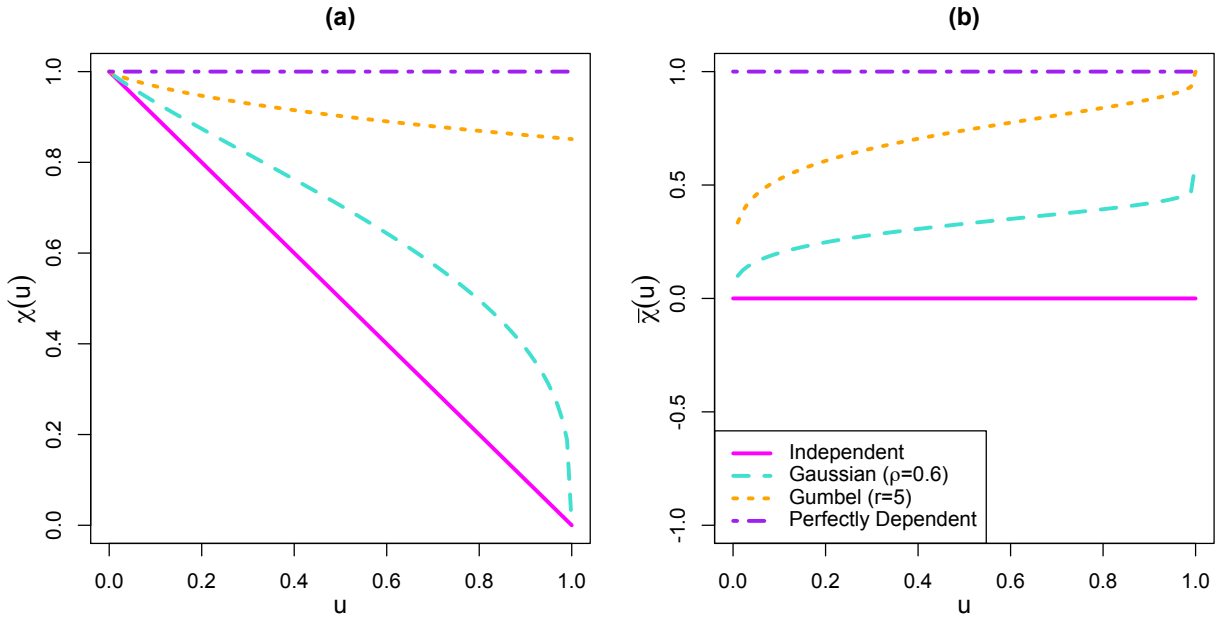

Figure 4.8. Extremal dependence measures (a) $\chi(u)$, (b) $\bar{\chi}(u)$ for $u \in [0,1]$ associated with copula functions 1-4.

The extremal dependence measures, $\chi(u)$ and $\bar{\chi}(u)$ for $u \in [0,1]$, for these four copula functions are displayed graphically in Figure 4.8. The differing behaviour in the asymptotic limit of the measures is apparent, with the independent and Gaussian copulas characterising asymptotic independence $(\chi = 0, \bar{\chi} < 1)$ and the Gumbel and perfectly dependent copulas characterising asymptotic dependence $(\chi > 0, \bar{\chi} = 1)$.

These extremal dependence measures can be calculated empirically for pairs of locations to investigate the asymptotic dependence in footprint wind gust speeds, allowing for the selection of a suitable copula function that represents the same behaviour.

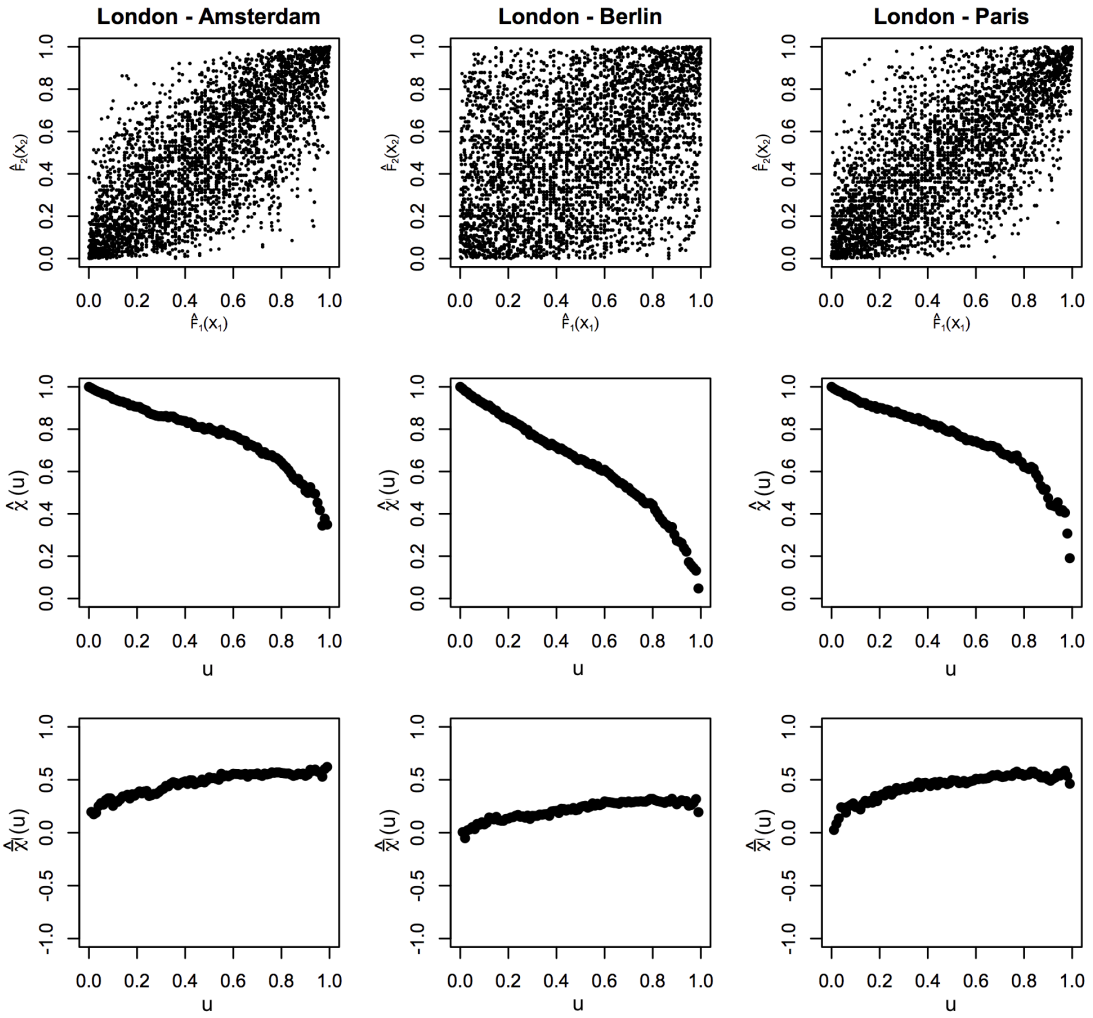

Figure 4.9. Empirical copula and empirical dependence measures $\chi(u)$ and $\bar{\chi}(u)$ for $u \in [0, 1]$, for London paired with 3 other European cities.

Figure 4.9 shows the empirical copula ($\hat{F}_1(X_1)$ against $\hat{F}_2(X_2)$, where $\hat{F}$ is the empirical cumulative distribution function) and empirical estimations of $\chi(u)$ and $\bar{\chi}(u)$, for wind gusts at London (X_1) paired with three other European cities: Amsterdam, Berlin and Paris (X_2). Referring back to Figure 4.8, it is clear from the shape of $\chi(u)$ and $\bar{\chi}(u)$ that wind gusts at these pairs of locations are neither independent nor perfectly dependent. There is an indication of asymptotic independence since in all cases $\bar{\chi} < 1$ and the curve of $\chi(u) \to 0$ as $u \to 1$, resembling the Gaussian copula in Figure 4.8. These results are representative of pairs of locations throughout the domain. However, the asymptotic limit of these measures is needed to identify asymptotic independence or dependence and cannot be found empirically due to the restricted sample size and the rarity of extreme events in the data set.

These limits can be estimated using a parametric bivariate model derived by Ledford and Tawn (1996), introduced in Section 2.2.3. This is a bivariate tail model that is

able to represent both extremal dependence classes. Let Z_1 and Z_2 denote X_1 and X_2 transformed to unit Fréchet margins, $\Pr(Z_1 \leq z) = \Pr(Z_2 \leq z) = \exp(-1/z)$. The joint survivor function for Z_1 and Z_2 takes the form

$$\Pr(Z_1 > r, Z_2 > r) \sim \mathcal{L}(r)p^{1/\eta} \tag{4.4}$$

for large r (Coles et al., 1999), where $p = \Pr(Z_1 > r) = \Pr(Z_2 > r)$, $\frac{1}{2} \leq \eta \leq 1$ is a constant and $\mathcal{L}(r)$ is slowly varying function (i.e. $\mathcal{L}(tr)/\mathcal{L}(r) \to 1$ as $r \to \infty$ for all fixed $t > 0$).

The parameter η, named the coefficient of tail dependence by Ledford and Tawn (1996), characterises the nature of the asymptotic dependence and $\mathcal{L}$ represents the relative strength of dependence for a given η. Exact independence occurs when $\eta = \frac{1}{2}$ and $\mathcal{L}(r) = 1$, and perfect dependence when $\eta = 1$ and $\mathcal{L}(r) = 1$.

By Eqn (4.4) and (4.12) and following Coles et al. (1999):

$$\begin{aligned} \chi(u) &\sim \mathcal{L}(u)(1-u)^{1/\eta - 1} \\ \bar{\chi}(u) &= \frac{2\log(1-u)}{\log\mathcal{L}\{(1-u)^{-1}\} + \frac{1}{\eta}\log(1-u)} - 1 \\ &\to 2\eta - 1 \quad \text{as } u \to 1 \end{aligned} \tag{4.5}$$

and so $\bar{\chi} = 2\eta - 1$. If $\eta = 1$ and $\mathcal{L}(r) \to c$ as $r \to \infty$, with $0 < c \leq 1$, $(\chi = c, \bar{\chi} = 1)$, the variables are asymptotically dependent of degree c (Coles et al., 1999). Alternatively, if $\eta < 1$, $(\chi = 0, \bar{\chi} = 2\eta - 1)$, the variables are asymptotically independent with non-asymptotic dependence of degree $2\eta - 1$.

As an aside, this model is not suitable for modelling the dependence in the footprint wind gust data because the theory is restricted to the bivariate case and therefore cannot be extended to model the whole domain. However, since the model is able to represent both asymptotic dependence and independence, the η parameter estimated using this model can be used to identify which class of asymptotic behaviour is present in the data.

Ledford and Tawn (1996) derive a method for estimating η using the joint survivor function along the diagonal, equivalent to the univariate distribution of variable $T = \min\{Z_1, Z_2\}$, known as the structure variable. Then $\Pr(T > r) = \Pr(Z_1 > r, Z_2 > r)$ and the distribution of the exceedance of T above some a high threshold w satisfies

$$\Pr(T > w + t | T > w) \sim \frac{\mathcal{L}(w+t)}{\mathcal{L}(w)}(1 + t/w)^{-1/\eta} \sim (1 + t/w)^{-1/\eta}$$

since $\mathcal{L}$ is a slowly varying function. This can be expressed in terms of the GPD, giving

$$\Pr(T > w + t | T > w) \sim (1 + \xi t/\sigma)^{-1/\xi}$$

where $\xi = \eta$ and $\sigma = \eta w$, and η is therefore estimated as the shape parameter of this univariate GPD by maximum likelihood.

Alternatively, here η is estimated as in Ferro (2007), who uses the Ledford and Tawn (1996) model in the application of forecast verification of rare, extreme weather events. The original variables can be transformed to exponential margins: $Z_j = -\log(1 - \hat{F}_j(X_j))$ and $\Pr(Z_1 \leq z) = \Pr(Z_2 \leq z) = 1 - \exp(-z)$. Realisations $T_i = \min\{Z_{1i}, Z_{2i}\}$ are constructed from the data $\{(Z_{1i}, Z_{2i}) : i = 1, ..., n\}$, where here $n = 5730$, the number of windstorm footprint wind gust speeds at each location, and the exceedance of T_i above a high threshold w is well approximated by a Poisson process. The probability of T_i exceeding threshold $t \geq w$ therefore has the form

$$\Pr(T_i > t) = \frac{1}{n}\exp\left[-\left(\frac{t-\alpha}{\eta}\right)\right] \quad \text{for } t \geq w \tag{4.6}$$

where α is a location parameter and $0 < \eta \leq 1$ is a scale parameter which, as before, is the coefficient of tail dependence with the same properties and interpretation as in Ledford and Tawn (1997). These parameters can be estimated by maximising the likelihood

$$L(\eta, \alpha) = \exp\left\{-\exp\left[-\left(\frac{w-\alpha}{\eta}\right)\right]\right\} \prod_{i:T_i>w} \frac{1}{\eta}\exp\left[-\left(\frac{T_i-\alpha}{\eta}\right)\right] \tag{4.7}$$

This maximisation can be performed analytically to give an estimate of η and α.

$$\hat{\eta} = \min\left\{1, \frac{1}{m}\sum_{i:T_i>w}(T_i - w)\right\} \tag{4.8}$$

$$\hat{\alpha} = w + \hat{\eta}\log(m) \tag{4.9}$$

where m is the number of T_i exceeding w.

As in Ledford and Tawn (1997) the asymptotic dependence behaviour between variables is examined by observing maximum likelihood estimates and profile likelihood based confidence intervals for the coefficient of tail dependence, η, over a range of values of structure variable threshold w. Taking the same approach as Ledford and Tawn (1996), two variables are said to be asymptotically dependent if the line $\eta = 1$ is contained within these confidence intervals for a majority of the range w, other-

wise the variables are asymptotically independent.

For a given threshold w, the $100(1-a)\%$ profile likelihood confidence interval for η is the set of values η_0 such that the two-sided test of the null hypothesis $H_0 : \eta = \eta_0$ would not be rejected at the a significance level. The value of η_0 is allowed to exceed 1 to give a non-truncated confidence interval. The associated likelihood ratio test statistic, calculated using Eqn (4.7) is $2(\log L(\hat{\eta}, \hat{\alpha}) - \log L(\eta_0, \hat{\alpha}_0))$, where $\hat{\eta}$ and $\hat{\alpha}$ are the maximum likelihood estimates from the full model and $\hat{\alpha}_0$ is the maximum likelihood estimate of α when $\eta = \eta_0$. The likelihood ratio test statistic is asymptotically chi-square distributed and the null hypothesis is true if and only if

$$\begin{aligned} &2(\log L(\hat{\eta}, \hat{\alpha}) - \log L(\eta_0, \hat{\alpha}_0)) < \chi^2_{1-a}(1) \\ \Leftrightarrow\ &\log L(\eta_0, \hat{\alpha}_0) > \log L(\hat{\eta}, \hat{\alpha}) - \chi^2_{1-a}(1)/2 \end{aligned}$$

where $\chi^2_{1-a}(1)$ is the $1-a$ quantile of the χ^2 distribution with 1 degree of freedom.

Towe et al. (2013) use a similar, likelihood ratio test approach to identify the class of extremal dependence between storm peak significant wave height, wind speed and surface level pressure in the Northern North sea, based directly on the Ledford and Tawn (1996) model, testing $\bar{\chi} = 1$ v $\bar{\chi} < 1$. They note that the maximum likelihood estimates of $\bar{\chi}$ will depend on the choice of threshold w. They exercise careful consideration for their choice of threshold and present the results for when the 80% quantile is used. Here, however, the results will be presented for a range of w, giving a more complete picture of the test for extremal dependence.

Figure 4.10 shows the maximum likelihood estimates for η for the pairs of locations in Figure 4.9 (top panel) and for London, Amsterdam and Paris paired with a neighbouring location with separation distance $\sim$ 25km (bottom panel). The neighbouring extremal dependence behaviour is presented because the closest locations will have the strongest dependence. The non-neighbouring locations are asymptotically independent since over the range of structure variable thresholds used, the line $\eta = 1$, corresponding to asymptotic dependence, is not within the 95% profile likelihood confidence intervals for estimated η. This conclusion is consistent throughout the domain for all pairs of non-neighbouring locations tested. There is a stronger dependence between neighbouring locations as would be expected; however, over most of the range of structure variable thresholds used, the line $\eta = 1$ is still not within the 95% profile likelihood confidence intervals for estimated η. Figure 4.10 (f), does however suggest Paris and the neighbouring location may be asymptotically dependent, suggesting the assumption of asymptotic independence throughout the domain may be incorrect in a small number of cases.

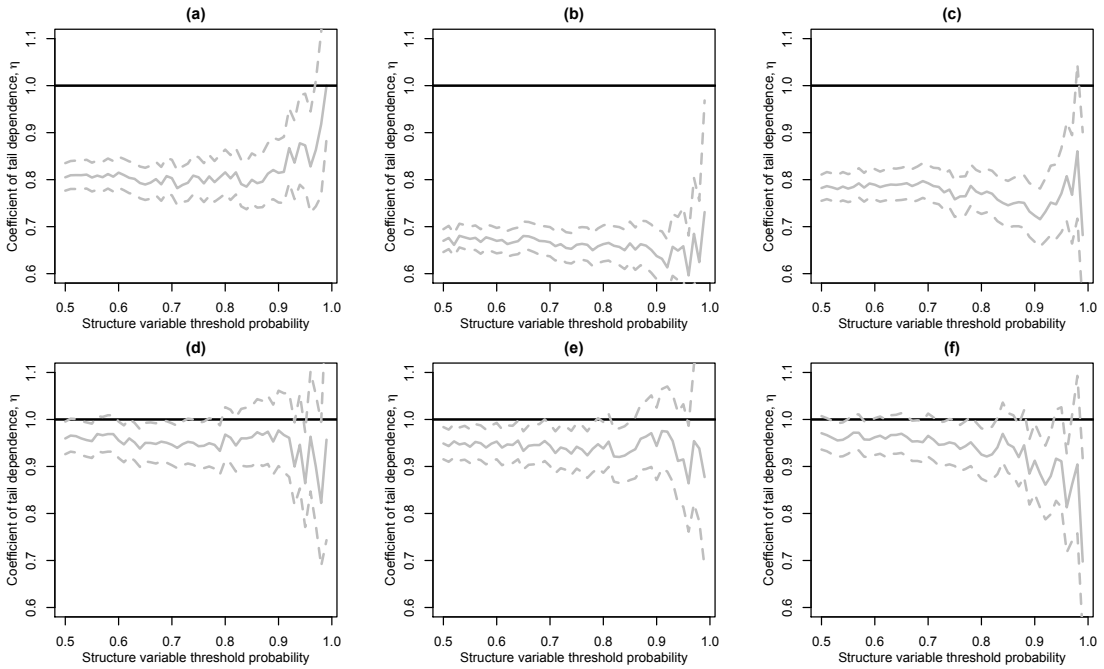

Figure 4.10. Maximum likelihood estimates (solid) and 95% profile likelihood confidence intervals (dashed) of η for structure variable threshold w in the range of the $0.5 - 1$ quantile of T for London paired with (a) Amsterdam, (b) Berlin and (c) Paris, and neighbouring locations paired with (d) London, (e) Amsterdam and (f) Paris (separation distance $\sim$ 25km).

So which type of dependence model should be used to model footprint wind gust speeds? When modelling extremes, such as windstorm footprint wind gusts, it is preferable to use extreme value copulas because they satisfy properties of max-stability, allowing for extrapolation beyond the observational range. This class of models, however, assumes asymptotic dependence or complete independence between variables and is therefore not suitable for modelling the dependence in wind gust speeds which have been found to be dominantly asymptotically independent. The Ledford and Tawn (1996) model is able to represent both asymptotic dependence classes and satisfies max-stability, however, as previously mentioned is restricted to the bivariate case and only applicable when both variables are large. The conditional model developed by Heffernan and Tawn (2004) overcomes these issues and allows for the modelling of both extremal dependence classes, in a multivariate setting, when just one of the variables is large. However, in order to extend the Heffernan and Tawn (2004) model to model the whole European domain, containing $J = 14872$ locations, 4 parameters must be estimated for each pair of locations, which would be extremely computationally expensive. For simplicity and based on the test for extremal dependence presented in Figure 4.10, which identified that asymptotic independence dominates throughout the region, an asymptotically independent copula approach will be used to model the bivariate dependence in windstorm footprint wind gust speeds.

The Gaussian copula imposes asymptotic independence. This copula does not satisfy max-stability, however, as discussed in Section 2.2.3, it has been used to model asymptotically independent extremes in the literature (Renard and Lang, 2007; Bortot et al., 2000). Bortot et al. (2000) demonstrated that for modelling extremes that are asymptotically independent the multivariate Gaussian copula model is robust, has simple diagnostics, easily interpretable parameters and extends straightforwardly to higher dimensions. In addition, Bortot et al. (2000) carry out a simulation study to examine the performance of the bivariate Gaussian copula model, fit to the upper tail of the data, relative to other tail models including the Gumbel model. They conclude that for asymptotically independent parent populations the Gaussian tail model is sufficiently flexible to provide accurate inferences for tail probability estimates, relative to the sampling uncertainty in such estimates and, even for asymptotically dependent parent populations, the estimation error of the Gaussian tail model is acceptably small. This suggests that, even if some pairs of neighbouring locations are asymptotically dependent, the Gaussian copula should provide a reasonable fit.

The data are transformed to standard Gaussian margins using the marginal left-truncated GEV distribution fitted to each location, a method known as Gaussian anamorphosis, and treated as multivariate Gaussian. The dependence between each pair of locations is then modelled as the correlation between the left-truncated standard Gaussian variables. This is similar to the novel multivariate Gaussian tail model developed by Bortot et al. (2000), derived from the joint tail of a multivariate Gaussian distribution with marginal transformation based on the GPD.

Since the dependence between locations is modelled as the correlation between standard Gaussian variables, it is directly relatable to the correlation function used in geostatistics which models the dependence in the spatial field as a function of separation distance and direction. Extending the model to cover the whole domain is greatly simplified if the dependence can be represented in this geostatistical way.

The resulting GEV-Gaussian model is defined on the region $(v_1, \infty) \times (v_2, \infty)$, where

both variables are large, and is specified as

$$
\begin{aligned}
(X_1, X_2)^T &= (F_1^{-1}(\Phi(Y_1)), F_2^{-1}(\Phi(Y_2)))^T \\
(Y_1, Y_2)^T &\sim MVN(0, \Sigma) \\
\Sigma &= \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix} \\
F_j(x) &= (1 - \lambda_j) + \lambda_j \left(\frac{G_j(x) - G_j(v_j)}{1 - G_j(v_j)} \right) \quad \text{for } x > v_j \\
\lambda_j &= 1 - \hat{F}_j(v_j) \\
G_j &= GEV(\mu_j, \sigma_j, \xi_j), \text{ i.e. } G_j(x) = \exp\left\{ - \left[1 + \xi_j \left(\frac{x - \mu_j}{\sigma_j} \right) \right]^{1/\xi_j} \right\}
\end{aligned}
$$

Bonazzi et al. (2012) used the Gumbel copula to model the dependence in footprint wind gust speeds at pairs of locations. The Gumbel copula belongs to the family of multivariate extreme value copulas (Davison et al., 2012) and therefore satisfies the max-stability property. The Gumbel copula, however, assumes asymptotic dependence between variables which Figure 4.10 suggests is an incorrect assumption. The resulting GPD-Gumbel model is applicable when one varibale is large and is specified as

$$
\begin{aligned}
(X_1, X_2)^T &= (H_1^{-1}(U_1), H_2^{-1}(U_2))^T \\
C(u_1, u_2) &= \Pr(U_1 \leq u_1, U_2 \leq u_2) = \exp\{-((-\log u_1)^r + (-\log u_2)^r)^{1/r}\} \\
H_j &= GPD(\alpha_j, \beta_j) \text{ i.e. } H_j(x) = 1 - \left(1 + \frac{\alpha_j (x - \gamma_j)}{\beta_j} \right)^{-1/\alpha_j}
\end{aligned}
$$

4.4.2. Bivariate dependence model fitting

Let Y_1 and Y_2 denote the Gaussian transformed wind gust speeds at locations s_1 and s_2 respectively, that is

$$
\begin{aligned}
Y_1 &= \Phi^{-1}(F_1(X_1)) \quad \text{for } X_1 > v_1 \\
Y_2 &= \Phi^{-1}(F_2(X_2)) \quad \text{for } X_2 > v_2
\end{aligned}
$$

where Φ is the standard Gaussian cumulative distribution function and F_1 and F_2 are the marginal left-truncated GEV cumulative distribution functions fitted to wind gust speeds at locations s_1 and s_2, X_1 and X_2 respectively.

The Gaussian copula has one parameter, ρ, the correlation between Y_1 and Y_2. Since these variables are transformed from the data above a threshold, $X_j > v_j$, they too must be greater than a threshold, $Y_j > \nu_j$ where $\nu_j = \Phi^{-1}(F_j(v_j))$. In-

ference for this model is made complicated by the fact that the bivariate pair may exceed the specified threshold in just one of its components. To overcome this a truncated likelihood method is used and ρ is estimated by maximising

$$L(\rho; (y_{11}, y_{21}), ..., (y_{1n}, y_{2n})) = \prod_{Y_1>\nu_1, Y_2>\nu_2} \frac{\phi(y_{1i}, y_{2i})}{p} \tag{4.10}$$

where ϕ is the standard bivariate Gaussian density function with correlation ρ and

$$p = Pr\{X_1 > v_1, X_2 > v_2\} = Pr\{Y_1 > \nu_1, Y_2 > \nu_2\} = \int_{\nu_2}^{\infty}\int_{\nu_1}^{\infty} \phi(y_{1i}, y_{2i}) dy_1 dy_2$$
$$= 1 - \Phi(\nu_1) - \Phi(\nu_2) + \Phi(\nu_1, \nu_2)$$

where $\Phi(\nu_1, \nu_2)$ is calculated using a numerical integration algorithm (Genz, 1992). This method therefore only uses data that are above the left truncation threshold for both locations, with the justification being that these are the only data that have been accurately transformed to standard Gaussian.

The truncated likelihood method for estimating ρ imposes the assumption of a bivariate Gaussian distribution on the data. Alternatively, ρ can be estimated using the model developed by Ledford and Tawn (1996) which does not impose any distributional assumptions. By Eqn 4.5, $\bar{\chi} = 2\eta - 1$ and it was shown that $\bar{\chi} \sim \rho$ for bivariate Gaussian random variables, therefore ρ can be estimated as $\rho = 2\eta - 1$, where η is determined from the Ledford and Tawn (1996) model.

This method, however, results in a biased estimate of η and therefore ρ, as shown by Ledford and Tawn (1996) and results from a simulation study shown in Table 4.1. The simulations involved 1000 repetitions of samples of 2500 points from the bivariate Gaussian distribution and the parameter $\hat{\rho}_{Led} = 2\hat{\eta} - 1$ is estimated using the Ledford and Tawn (1996) methodology with structure variable threshold, w, equal to the 95% quantile of the simulated variable, as in Ledford and Tawn (1996). Table 4.1 shows the median and standard deviation (in brackets) of the bias in the estimated $\hat{\rho}_{Led}$, for the bivariate normal dependence structure for a range of correlations, ρ. The truncated likelihood method for estimating $\hat{\rho}_{Trunc}$, Eqn (4.10), is also applied to the same simulated data to compare the biases in the two methods. In addition the simulation study is repeated using a bivariate Gumbel dependence structure for a range of dependence parameters $r = 1/\alpha$ for comparison, as in Ledford and Tawn (1996). Since $\rho - \hat{\rho} = 2(\eta - \hat{\eta})$, the bias in ρ is calculated based on the bias in η for the two models where the true $\eta = 1$ because the Gumbel dependence structure imposes asymptotic dependence.

When the simulated data has a Gaussian dependence structure, similarly to Led-

ρ	0.00	0.25	0.50	0.75	0.9
$\rho - \hat{\rho}_{Led}$	0.006 (0.041)	0.058 (0.047)	0.098 (0.050)	0.140 (0.049)	0.148 (0.040)
$\rho - \hat{\rho}_{Trunc}$	-0.003 (0.061)	0.009 (0.069)	-0.003 (0.050)	0.000 (0.019)	0.000 (0.006)
α	0.99	0.95	0.85	0.6	0.4
$\rho - \hat{\rho}_{Led}$	0.946 (0.046)	0.738 (0.053)	0.410 (0.060)	0.134 (0.047)	0.100 (0.035)
$\rho - \hat{\rho}_{Trunc}$	0.698 (0.018)	0.658 (0.019)	0.573 (0.019)	0.393 (0.016)	0.232 (0.011)

Table 4.1.. Median and standard deviation (in brackets) of the bias in the estimated correlation parameter $\hat{\rho}$, for the bivariate Gaussian dependence structure with correlation ρ and then Gumbel dependence structure with $r = 1/\alpha$. The first rows of the two dependence structure studies relate to the method of estimating ρ from η in the Ledford and Tawn (1996) model and the second row relates to the truncated likelihood method for estimating ρ.

ford and Tawn (1996), a positive bias is identified in $\rho - \hat{\rho}_{Led}$. This is because the model is based on asymptotic theory in the limit as $w \to \infty$ and using a lower threshold, necessary to ensure the model is being fit to enough data, results in the higher order terms neglected in Eqn (4.4) being non-negligible (Ledford and Tawn, 1996). Increasing the threshold would reduce the bias but increase the standard deviations of the estimated parameter. The standard deviations associated with $\hat{\rho}_{Led}$ are similar for all values of ρ. The bias associated with the truncated likelihood method, $\rho - \hat{\rho}_{Trunc}$, is much lower and decreases to approximately 0 as ρ increases. The largest standard deviations associated with $\hat{\rho}_{Trunc}$ are similar to those for $\hat{\rho}_{Led}$ but become much smaller as ρ increases.

When the simulated data has a Gumbel dependence structure, again as was seen by Ledford and Tawn (1996), the bias associated with $\hat{\rho}_{Led}$ is small for lower values of α, when the dependence between variables is strong. However, as α increases, corresponding to a weaker degree of asymptotic dependence, the estimates of ρ are more biased. Again, Ledford and Tawn (1996) note that this effect is due to the structure variable threshold being too low. The bias associated with the truncated likelihood method, $\rho - \hat{\rho}_{Trunc}$, is lower than that for the $\hat{\rho}_{Led}$ method when α is large. The bias decreases as α decreases because the correlation is getting stronger, however this decrease in the bias smaller than for $\hat{\rho}_{Led}$ because the asymptotic dependence is not accounted for.

This simulation study shows that the truncated likelihood method of estimated ρ is preferable when the underlying process is asymptotically independent or very weakly asymptotically dependent and will therefore be used here.

Fitting the Gumbel copula for the Bonazzi et al. (2012) copula model, requires the estimation of the dependence parameter, r. Again, the marginal GPDs are fit

above a threshold, complicating the inference of the model. For consistency with the methodology used by Bonazzi et al. (2012), a censored likelihood method, as defined in chapter 8 of Coles (2001) is used to estimate r at the same time as the GPD marginal parameters. The overall bivariate likelihood is a product of four partial likelihoods that account for both components being above or below thresholds, or one component above and the other below. For example if $X_1 > \gamma_1$ and $X_2 \leq \gamma_2$, there is information in the data concerning the marginal X_1 component, but not the X_2 component so the likelihood contribution for such a point is:

$$Pr\{X_1 = x_1, X_2 < \gamma_2\} = \frac{\partial H}{\partial x_1}\bigg|_{(x_1,\gamma_2)}$$

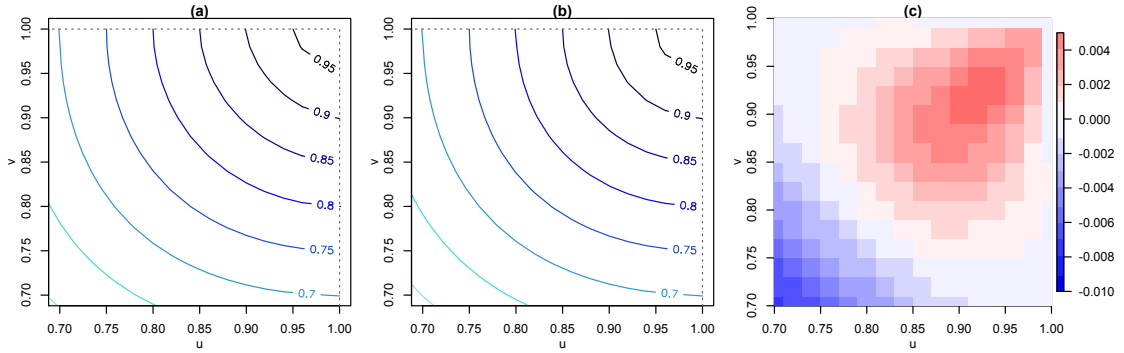

Figure 4.11. The best fit (a) Gaussian and (b) Gumbel bivariate copula functions for London paired with Amsterdam and (c) shows the difference between these functions. The functions are displayed for u and v greater than 0.7 because of the truncated nature of the models.

Figure 4.11 shows the best fit Gaussian and Gumbel bivariate copula models for wind gust speeds in London and Amsterdam. The Gaussian dependence parameter is estimated to be $\rho = 0.788$ and the Gumbel dependence parameter $r = 1.942$. The resulting functions appear to be very similar in (a) and (b), however the difference between the functions (Gumbel copula surface minus the Gaussian copula surface), shown in (c) displays how the Gumbel copula exhibits a stronger dependence in the extremes of the variables, due to the assumption of asymptotic dependence.

4.4.3. Bivariate dependence model validation

The challenge of applying goodness-of-fit (GoF) tests for copula models relates to the general difficulty of testing multivariate probability densities or distribution functions (Schoelzel and Friederichs, 2008). While the evaluation of univariate distributions is well documented, the study of GoF tests for copula models has emerged only recently as a challenging inferential problem (Berg, 2009). Some examples of

goodness-of-fit (GoF) tests for copula models in the literature include the multidimensional χ^2-tests, as in Fermanian (2005), which are simple to apply but require arbitrary probability binning and an adequate sample size (Schoelzel and Friederichs, 2008). More refined GoF tests for copula models are based on the Probability Integral Transform that project the multivariate problem into a univariate distribution allowing standard univariate tests to be applied (Schoelzel and Friederichs, 2008). Since the field is still in its infancy, guidelines and recommendations are sparse and Mikosch (2006) goes as far as to say that the best GoF tests for copulas are still unresolved.

Rather than apply these generic copula GoF tests, the models are tested to ensure they are fit for purpose, which here means having the correct extremal dependence relationship between pairs of locations. Comparing the Gaussian and Gumbel copula models, which have different extremal dependence representations, will give a further indication of which is most appropriate for representing the dependence in the data.

Firstly, the assumed extremal dependence relationships of the Gaussian and Gumbel copula models will be validated by comparing the extremal dependence measures $\chi(u)$ and $\bar{\chi}(u)$ evaluated empirically and using each of the models. These measures are evaluated using a simulation approach for the Gaussian model, and directly from the copula function for the Gumbel model (see Appendix A.1).

The Ledford and Tawn (1996) model (Eqn (4.4)) will be included in the comparison for interest. This model can represent the diagonal of a copula, i.e. $C(u, v)$ for $u = v$, and will be referred to as the power law copula. Suppose $\Pr(X_1 \leq x_1) = \Pr(X_2 \leq x_2) = u$. Let Z_1 and Z_2 denote X_1 and X_2 transformed to Exponential margins, as in Ferro (2007), and $T = \min\{Z_1, Z_2\}$, then, by Eqn (4.6):

$$\begin{aligned}
\chi(u) &= \Pr(X_2 > x_2 | X_1 > x_1) \\
&= \Pr(X_1 > x_1, X_2 > x_2)/\Pr(X_1 > x_1) \\
&= \Pr(T > -\log(1-u))/(1-u) \\
&= \frac{1}{n}\exp\left(\frac{\alpha}{\eta}\right)\exp\left(\frac{\log(1-u)}{\eta}\right)/(1-u) \\
&= \frac{1}{n}\exp\left(\frac{\alpha}{\eta}\right)(1-u)^{1/\eta-1}
\end{aligned} \tag{4.11}$$

$$\bar{\chi}(u) = \frac{2\log\Pr(X_1 > x_1)}{\log\Pr(X_1 > x_1, X_2 > x_2)} - 1$$
$$= \frac{2\log(1-u)}{\log\left(\frac{1}{n}\exp\left(\frac{\alpha}{\eta}\right)\right) + \log((1-u)^{1/\eta})} - 1 \tag{4.12}$$

The parameters η and α are estimated using a structure variable threshold probability of 0.85. Ferro (2007) explains how the chosen threshold should be high enough that the asymptotic model is valid but low enough that enough data are used to estimate the parameters. The good fit of the power law copula in Figure 4.12 indicates that this threshold is appropriate

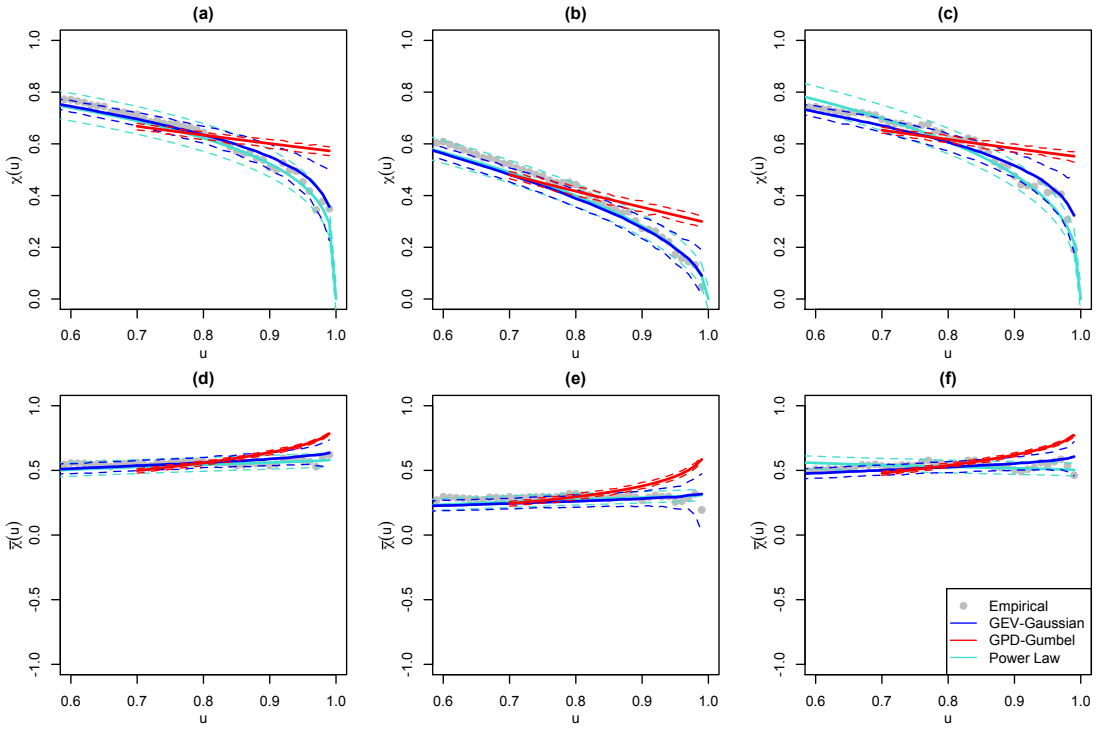

Figure 4.12. Extremal dependence measures evaluated empirically and using the Gaussian, Gumbel and power law copulas for London paired with (a)/(d) Amsterdam, (b)/(e) Berlin and (c)/(f) Paris. The measure $\chi(u)$ is shown in the top row and $\bar{\chi}(u)$ in the bottom row. The 95% confidence intervals are based on parametric bootstrapping for the Gaussian copula (see Appendix A.1) and percentile bootstrapping for the Gumbel copula and the profile likelihood for the power law copula.

From Figure 4.12 it is clear that the Gaussian copula is more successful than the Gumbel copula at capturing the empirical bivariate dependence structure, particularly in the extremes. The Gumbel copula, which assumes asymptotic dependence, is overestimating the conditional probability of joint extremes while the Gaussian copula, which assumes asymptotic independence, captures the empirical relationship well. The Gaussian and power law copulas are almost identical, further indicating that the Gaussian copula is performing well and extremal independence is a valid as-

sumption for the model. The power law copula also shows that the $lim_{u \to 1} \chi(u) = 0$ for all three pairs, again indicating extremal independence.

Figure 4.12 validates the fit of the copula models along the diagonal of the copula, but what if damage occurs in one location at a lower quantile threshold than another? Berlin and Paris, for example, experiences lower winds than London, on average, because most storms follow a track further to the North-West of Europe. Suppose that Berlin and Paris are therefore less well adapted to strong winds so the same hazard could exist in Berlin and Paris at a local quantile wind gust speed 5% and 10% lower than in London respectively. Conversely, in general stronger wind gust speeds are experienced in Amsterdam than London and the same hazard may exist there when the quantile is 10% higher. Equivalent plots to those in Figure 4.12 with the Berlin quantile threshold 5% lower, and the Amsterdam quantile threshold 10% higher than London, are presented in Figure 4.13. The power law copula only represents the diagonal of the copula and therefore cannot be included in this plot.

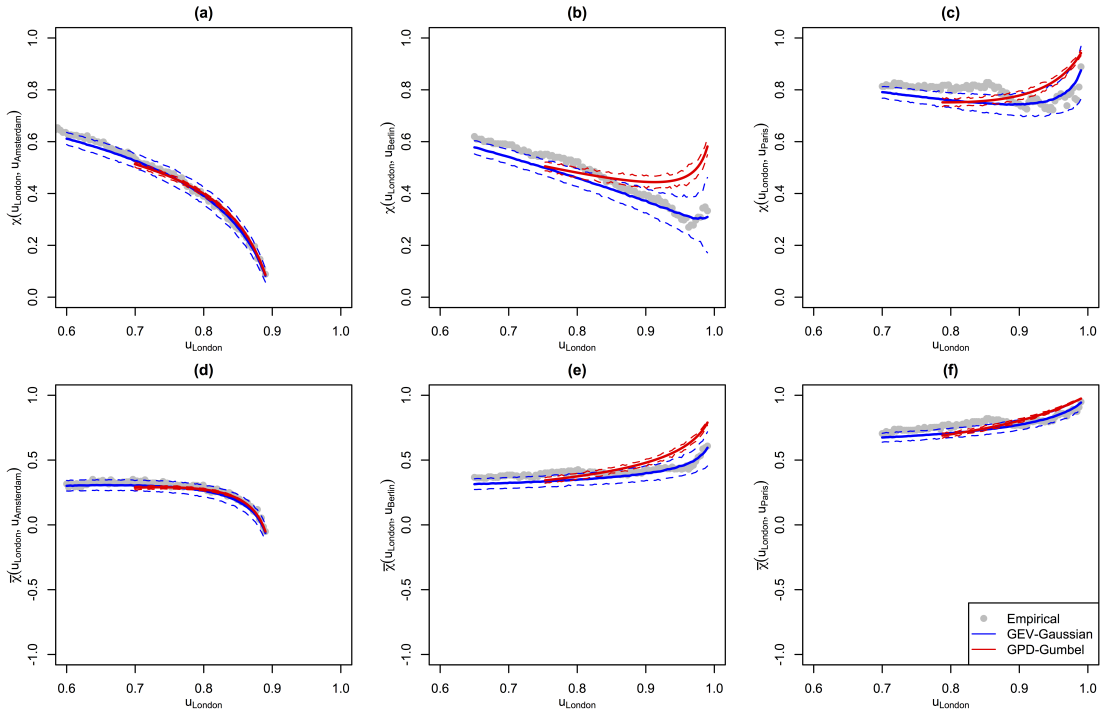

Figure 4.13. Extremal dependence measures evaluated empirically and using the Gaussian and Gumbel copulas for London paired with Amsterdam, Berlin and Paris. The measure $\chi(u_{London}, u_{Other})$ is shown in the top row (a)-(b) and $\bar{\chi}(u_{London}, u_{Other})$ in the bottom row (c)-(d), where $u_{Amsterdam} = u_{London} + 0.1$, $u_{Berlin} = u_{London} - 0.05$ and $u_{Paris} = u_{London} - 0.1$. The 95% confidence intervals are based on parametric bootstrap for the Gaussian copula (see Appendix A.1) and percentile bootstrapping for the Gumbel copula.

From Figure 4.13 it can be seen that the Gaussian copula is performing equally

well on the off diagonal of the copula. The Gumbel copula still over estimates the dependence between London and Berlin, but it is performing better for London and Paris and particularly well for London and Amsterdam. This is because the London and Amsterdam and London and Paris quantile thresholds are 10% apart; therefore extreme thresholds are not modelled in both locations at the same time and the inappropriate assumption of asymptotic dependence does not have as much of an affect on the fit of the model. When the difference in quantile threshold is only 5%, in the case of London-Berlin, the model is high enough in the tail for both locations that the inappropriate assumption of asymptotic dependence is still causing the Gumbel copula to over estimate the dependence between locations.

The extremal dependence measures are a function of the copula and are not affected by the fit of the marginal distribution (shown in Appendix A.2). This means that the results shown above are not sensitive to the accuracy of the fitting of the marginal distributions.

4.5. Bivariate model validation

Both the left-truncated GEV marginal distributions and the Gaussian copula that make up the GEV-Gaussian model have been shown to be a good fit to the data. The model has the potential to be extended to the whole domain; therefore all that remains is to validate the model as a whole to ensure it is able to realistically represent joint losses. In order to do this, two conceptual loss functions will be introduced as a proxy for re/insurance loss. Again, the same validation will be applied to the GPD-Gumbel model as a comparison.

The conceptual loss functions are defined as having the form:

$$L(X_1, X_2) = H(X_1 - t_1) + H(X_2 - t_2)$$

where t_1 and t_2 are high loss thresholds at each location and $H(n)$ as a Heaviside or indicator function $H(n) = 1$ if $n > 0$ and $H(n) = 0$ otherwise. For the first function the loss threshold t_j is the 98^{th} percentile of climatology wind gust speed at location s_j, calculated from October-March 1979-2012, (Figure 3.3). For the second function $t_1 = t_2 = 25\text{ms}^{-1}$. The two conceptual loss functions will therefore be denoted L_{98} and L_{25} respectively.

These particular functions are used to represent loss based on the conclusions of Chapter 3 of this thesis. A number of storm severity measures were compared for

how well they represented insured loss. It was found that severity measures that included the magnitude of the exceedance above the damage threshold (25ms^{-1}) and population density were less representative of insured loss than the severity measure that characterised exceedance area only. The exceedance of the 98^{th} percentile of climatology wind gust speed is included because it is a loss threshold commonly used in the literature (e.g., Klawa and Ulbrich 2003, Leckebusch et al. 2007, Pinto et al. 2007).

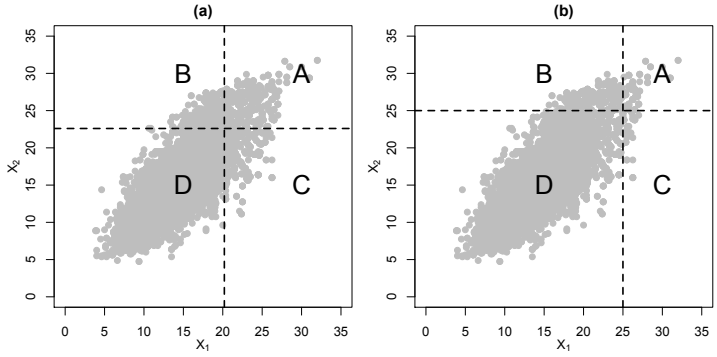

Figure 4.14. The Empirical bivariate distribution of wind gusts at London and Amsterdam with regions A, B, C and D indicating where the bivariate conceptual loss functions (a) L_{98} and (b) L_{25} take values 2, 1, 1 and 0 respectively.

The bivariate loss function takes the value of 0, 1 or 2 depending on the exceedance of the high loss threshold. This is demonstrated graphically for London and Amsterdam in Figure 4.14. Using these two different loss thresholds will validate how well the bivariate model is able to realistically represent the data at different extremities since the 25ms^{-1} threshold is generally further into the tail of the joint distribution compared to the climatology percentile threshold. The conceptual loss distributions are calculated empirically and using each of the bivariate models where, again, a simulation method is used for the GEV-Gaussian model (see Appendix A.1), shown in Figure 4.15

Figure 4.15 clearly shows that the GEV-Gaussian model is better at realistically representing joint losses for these pairs of locations. This further indicates that the dependence in extremes should be modelled using an asymptotically independent model. The GPD-Gumbel model overestimates the probability of loss equalling 0 for L_{98} and L_{25} and 2 for L_{25} and consistently underestimates the probability that loss equals 1. This indicates that within this model the dependence is too strong in the extremes, due to the assumption of asymptotic dependence. This analysis was repeated using a number of other pairs of locations, with the same resulting conclusions.

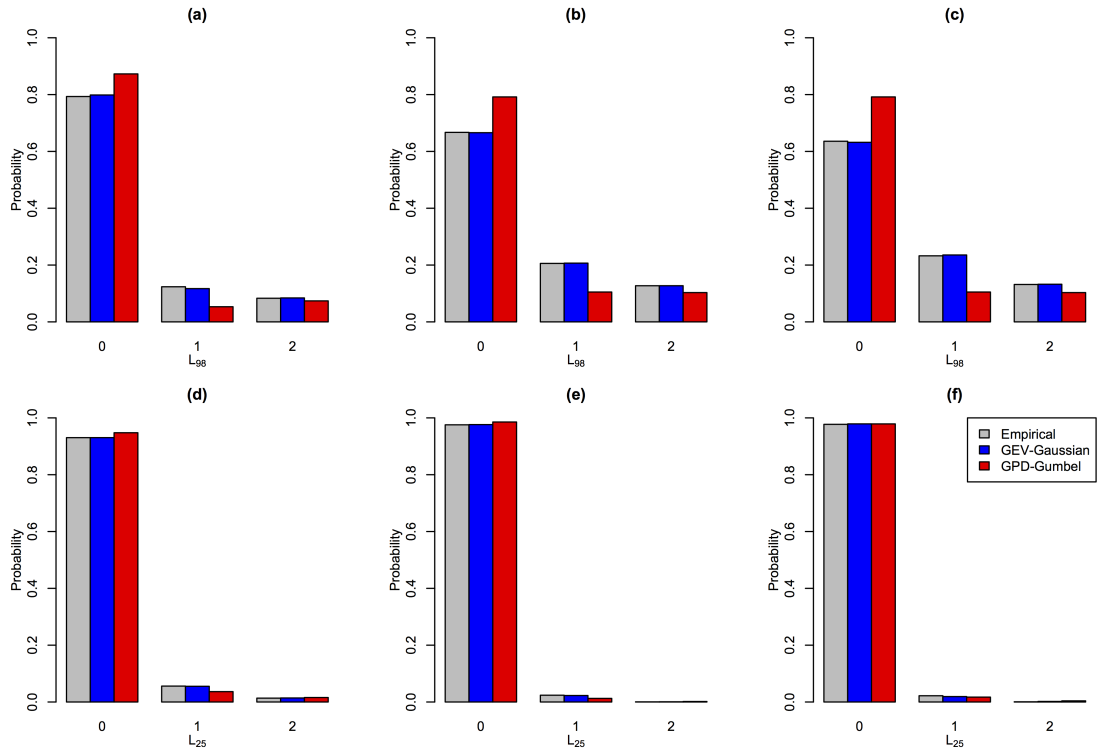

Figure 4.15. Empirical, GEV-Gaussian model and GPD-Gumbel model loss distributions for (a)-(c) conceptual loss function L_{98} and (d)-(f) conceptual loss function L_{25}, for London paired with (a)/(d) Amsterdam, (b)/(e) Berlin and (c)/(f) Paris.

To gain a greater understanding of how well these models realistically represent joint losses throughout the domain the loss distributions associated with the two conceptual loss functions are presented for London paired with all other locations in the domain. This is done in terms of the measures

$$\chi_t = \Pr(X_2 > t_2 | X_1 > t_1) = a/(a+c)$$
$$\bar{\chi}_t = \frac{2\log\Pr(X_1 > t_1)}{\log\Pr(X_1 > t_1, X_2 > t_2)} - 1 = \frac{\log((a+c)/n)}{\log(a/n)} - 1$$

where a and c are the number of points in sections A and C of the diagram in Figure 4.14 (a) and (b) for L_{98} and L_{25} respectively and n is the total number of points.

These measures are used because it can be shown that the loss distribution can be represented in terms of χ_t (see Appendix A.3) and the non-vanishing extremal dependence measure $\bar{\chi}_t$ will give information about the dependence between locations when χ_t falls to zero, therefore aiding in the validation of the models.

The resulting measures are denoted χ_{98} and $\bar{\chi}_{98}$ for the L_{98} conceptual loss function and χ_{25} and $\bar{\chi}_{25}$ for the L_{25} conceptual loss function. As before, these measures are estimated by a simulation method for the GEV-Gaussian model (see Appendix

A.1). The presented results are representative of pairs of locations throughout the domain.

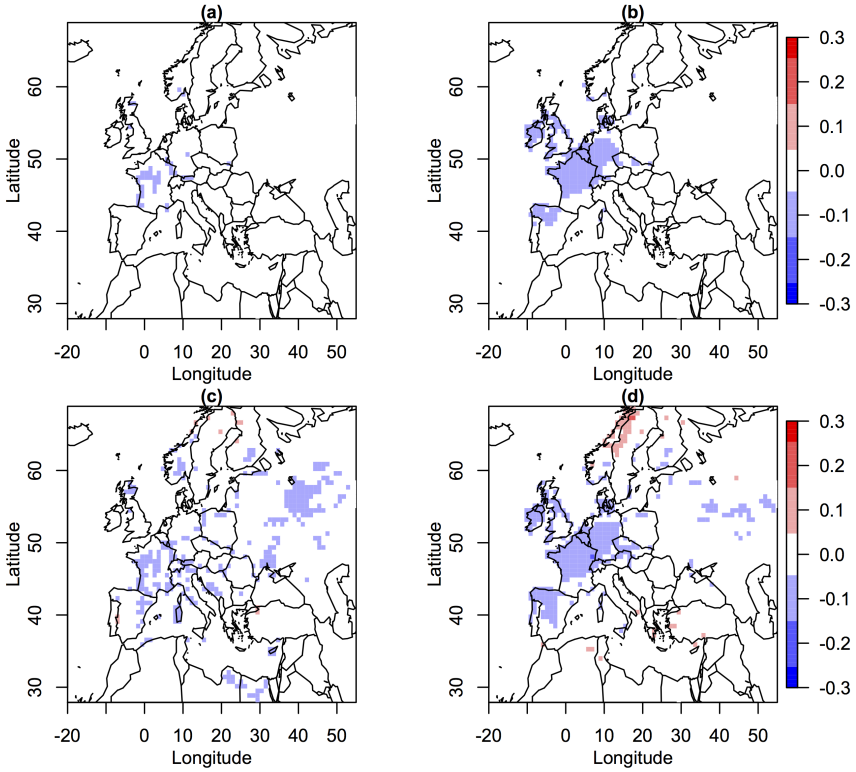

Figure 4.16. The difference between empirical and modelled χ_{98} for (a) the GEV-Gaussian model and (b) the GPD-Gumbel model and the difference between empirical and modelled $\bar{\chi}_{98}$ for (c) the GEV-Gaussian model and (d) the GPD-Gumbel model, for London paired with all other locations over land.

Figures 4.16 and 4.17 show that, for both conceptual loss functions, the GEV-Gaussian model out performs the GPD-Gumbel model. The GEV-Gaussian model slightly over estimates the conditional dependence in losses represented by χ_t and $\bar{\chi}_t$ for the higher threshold loss function L_{25} for nearby locations with some small areas are underestimated. The lower threshold loss function L_{98} is very well represented by the GEV-Gaussian model with only a slight underestimation of $\bar{\chi}_{98}$ in some places. This indicates that the GEV-Gaussian model is, in general, able to realistically represent joint losses and the conditional dependence in losses at pairs of locations throughout the domain.

The GPD-Gumbel model underestimates the conditional dependence in losses represented by χ_t and $\bar{\chi}_t$ for the lower threshold loss function, L_{98}, but overestimates this dependence for the higher threshold loss function, L_{25}. This is consistent with

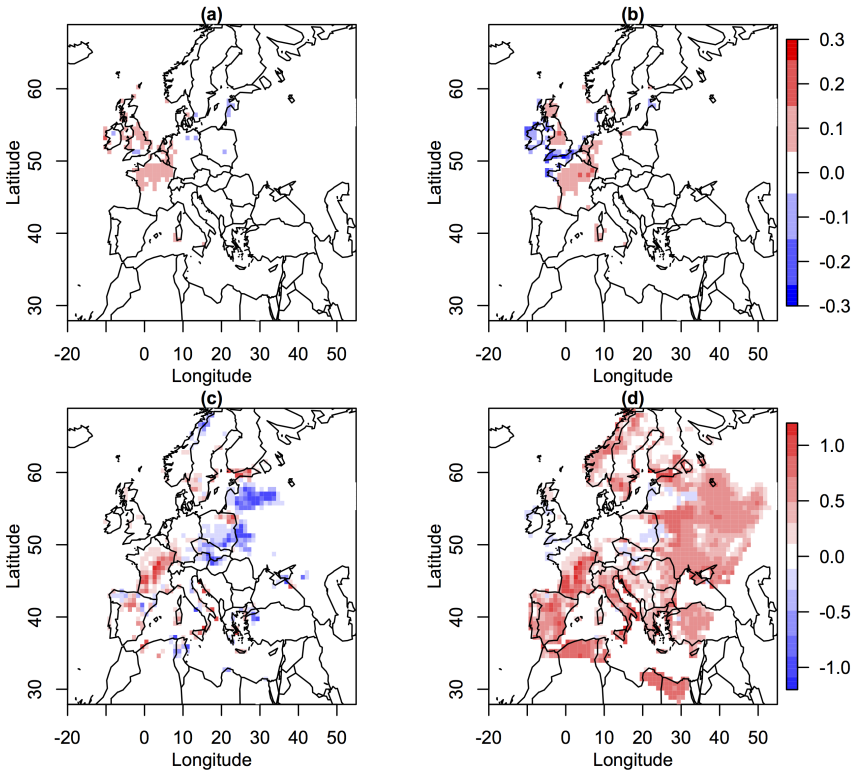

Figure 4.17. As in Figure 4.16 but for χ_{25} and $\bar{\chi}_{25}$

the over and under estimation of these measures in Figure 4.12. This is particularly evident in Figure 4.17 (d), where $\bar{\chi}_{25}$ is greatly over estimated for London paired with locations at a medium to large separation distance apart. Again, this is showing how the asymptotic dependence assumption of Gumbel copula results in overestimation of the dependence in the joint extremes.

4.6. Conclusions

This chapter presents the development and validation of a bivariate model for extremal windstorm footprint wind gust speeds, referred to as the GEV-Gaussian model. Various methods for modelling multivariate extremes were discussed and compared here and in Section 2.2.3. The copula approach was chosen due to its simplicity, flexibility and the potential for it to be extended to more dimensions. This copula methodology allows the joint distribution of two variables to be represented in terms of the marginal distributions of each variable and a copula function which models the dependence between the variables when the margins are transformed to the same scale.

Since the footprint wind gust speeds at a given location are a set of maxima, the GEV distribution is an appropriate model for the marginal distributions. A left-truncated GEV was shown to fit the data better in the upper tail of the distribution and also ensured wind gust speeds were modelled as non-negative.

Infinitely many copula functions exist for modelling the dependence between variables. A copula selection investigation was presented, based on extremal dependence behaviour. Using the method for identifying asymptotic independence or dependence explained by Ledford and Tawn (1996), footprint wind gust speeds were found to be asymptotically independent throughout the domain, with the possible exception of some neighbouring locations. The Gaussian copula, which assumes asymptotic independence, was therefore used to model the dependence within the bivariate model. This copula method was chosen over the conditional Heffernan (2000) model due to simplicity. This Gaussian dependence model was validated and shown to be a good fit to the data, confirming that more complex, flexible models do not need to be used to accurately model the dependence between footprint wind gust speeds.

The ability to realistically represent joint losses was validated by comparing the empirical joint probability of loss with those evaluated using the model. This was done for two conceptual loss functions with differing loss thresholds for comparison. The GEV-Gaussian model was shown to relatively realistically represent the joint loss distributions for London paired with three other locations. The ability for the model to represent the conditional dependence in losses for London paired with all other locations over land was also investigated, again showing that the model is able to realistically represent the relationship between losses at pairs of locations.

Throughout the GEV-Gaussian model was compared to the GPD-Gumbel model formulation used by Bonazzi et al. (2012). This model assumes asymptotic dependence and is therefore shown to overestimate the dependence in footprint wind gust speeds and hence the probability of joint losses. This indicates that using a Gaussian copula to model the dependence is preferred.

The Gaussian copula has other advantages in that it is easily extended to more than two dimensions (as will be discussed in Sections 5.2) and, since the dependence between locations is modelled as the correlation between standard Gaussian variables, it is directly relatable to the correlation functions used in geostatistical models.

5. Geostatistical modelling of windstorm footprints

5.1. Introduction

In the previous chapter, a bivariate model for damaging windstorm footprint wind gust speeds was presented as an initial step in footprint modelling. The bivariate model was rigorously validated and was shown to represent realistic joint losses. The bivariate model will now be extended to all land locations in the European domain, allowing for the simulation of synthetic windstorm footprints over land with specified characteristics, determined by the parameters of the statistical model. This footprint simulation model will be used in Chapter 6, in a sensitivity analysis experiment to investigate which windstorm footprint characteristics are most damaging or important for determining insured loss.

This footprint simulation model must correctly model dependence between locations. The model must be computationally inexpensive for simulation to ensure that it can be used in the sensitivity analysis experiment which will require a large number of synthetic footprints. As in the previous chapter, it is also essential that the model is able to simulate synthetic footprints that realistically represent joint losses. This chapter aims to address three main questions:

- How should the spatial dependence in wind gust speeds be modelled?
- Can a function of distance and direction capture the correlation structure of wind gusts?
- Can the windstorm footprint model represent realistic losses across the domain?

The general modelling approach will first be introduced in Section 5.2. The model specification, fitting and validation will then be presented, answering the questions above and ensuring that the footprint simulation model is fit for purpose in the ways specified above.

5.2. General approach

The intended aim is to extend the bivariate GEV-Gaussian copula model for extreme footprint wind gusts developed and validated in the previous chapter. This copula model uses the left-truncated GEV distribution to model the marginal distribution of wind gusts at each location and the bivariate Gaussian copula to model the dependence between pairs of locations. A simple extension of this bivariate model formulation would be to use a multivariate Gaussian copula to model the dependence between a finite set of locations in the domain. The joint distribution of wind gust speeds at J locations is then

$$\begin{aligned}\Pr(X_1 \leq x_1, X_2 \leq x_2, ..., X_J \leq x_J) &= F(x_1, x_2, ..., x_J) \\ &= \Phi\{\Phi^{-1}(F_1(x_1)), \Phi^{-1}(F_2(x_2)), ..., \Phi^{-1}(F_J(x_J)); \Sigma\} \\ &= \Phi(y_1, y_2, ..., y_J)\end{aligned}$$

where Φ is the standard Normal cumulative distribution function, Φ_J is the J-dimensional multivariate standard Normal cumulative distribution function, F_j is the left-truncated GEV cumulative distribution function, fit to wind gust speeds at location s_j above the truncation threshold v_j and $y_j = \Phi^{-1}(F_j(x_j))$ are standard Gaussian above the threshold $\nu_j = \Phi^{-1}(F_j(v_j)$. The dependence within the whole domain is then determined by the covariance matrix for Gaussian transformed wind gust speeds Y, with the $(i,j)^{\text{th}}$ entry

$$\begin{aligned}\Sigma_{ij} &= \text{Cov}(Y_i, Y_j) \\ &= \sqrt{\text{Var}(Y_i)\text{Var}(Y_j)}\text{Cor}(Y_i, Y_j) \\ &= \text{Cor}(Y_i, Y_j)\end{aligned}$$

since, by definition, $\text{Var}(Y_j) = 1$ for all j.

A few recent studies have adopted a multivariate copula approach including Davison et al. (2012), who fit a variety of models to annual maximum rainfall data for 51 locations in Switzerland to investigate which are the most appropriate for modelling extremal data. Another example is Bortot et al. (2000) who use a multivariate copula model with GPD margins and a Gaussian copula to model extreme sea-levels at 3 locations to investigate the probability of flooding. A multivariate copula model is feasible in these examples because few locations are being considered. However, the windstorm footprint data to be modelled here consists of 14872 locations (see Section 2.3.1), resulting in a very large 14872×14872 covariance matrix, Σ, which would be poorly estimated and difficult to simulate from.

The computational cost can be greatly reduced by modelling the dependence between locations using a covariance function, $C(\cdot)$, a function of distance between locations. This distance can be a measure of separation distance only or of both the distance and direction of separation. The most common form for the covariance function to take is for the dependence (i.e. correlation/covariance) between locations s_i and s_j to decrease as the distance increases. For example, the correlation may be thought to decay exponentially:

$$C(d) = \exp\left(-\frac{d}{\phi}\right)$$

where $1/\phi$ is the rate of exponential decay of the covariance with distance, d. The parameters of covariance functions therefore have simple interpretations and can be easily varied, simplifying the intended sensitivity study. This approach is a key element of geostatistics.

Geostatistics is a large and rapidly developing domain of statistics, with standard texts on the topic including Cressie (1993), Stein (1999), Wackernagel (2003), Banerjee et al. (2004), Schabenberger and Gotway (2005) and Diggle and Ribeiro (2007). This field of statistics originated in the mining industry in the 1960s with initial developments by Daniel G. Krige, a mining engineer in South Africa and Georges Matheron, a French mathematician and geologist who created the Paris School of Mines. The original aim was to predict ore grade for blocks of ore based on samples from neighbouring blocks. Geostatistics is now accepted as the main method for spatial data analysis in many disciplines (Haskard, 2007).

Davison et al. (2012) identify three common data types used in geostatistics: spatial point processes, where the observation sites are treated as random; areal data, where interpolation between observation sites may not be interpretable; and point-referenced or geostatistical data, where the observations are at fixed sites, between which interpolation makes sense. Footprint wind gust speeds are assumed to take the form of geostatistical data with the centre of the 25km grid-cell used as the location of wind gust speed associated with that grid-cell.

The basic format for univariate goestatistical data, as defined by Diggle and Ribeiro (2007), is

$$(s_j, y_j), \quad j = 1, ..., J$$

where s_j identifies a spatial location, typically in two-dimensional space, and y_j is a scalar value of the measurement or response variable associated with location s_j. If n measurements of the variable exist at each location, y_j becomes a vector of length

n,

$$y_j = (y_{1j}, y_{2j}, ..., y_{nj}).$$

The response variable is defined throughout a continuous study region, $s_j \in D$, where the sampling design of the location s_j is either deterministic (e.g. on a grid) or independent of the process which generates y_j (Diggle and Ribeiro, 2007). Each $y_{.j}$ is a realisation of a random variable Y_j whose distribution is dependent on the value of the underlying continuous stochastic process $Z(s)$ at location s_j, which is not directly observable (Diggle and Ribeiro, 2007).

The resulting geostatistical model has two elements: a real valued stochastic process $\{Z(s) : s \in D\}$, known as the signal; and a multivariate distribution for random variable $Y = (Y_1, ..., Y_J)$ conditional on $Z(\cdot)$, known as the response (Diggle and Ribeiro, 2007).

The best explored approach to modelling point-referenced geostatistical data are to suppose that $Z(s)$ follows a Gaussian process (Davison et al., 2012). A Gaussian spatial process $\{Z(s) : s \in D\}$ is a stochastic process with the property that for any collection of locations $s_1, ..., s_j \in D$ the joint distribution of $Z = \{Z(s_1), ..., Z(s_j)\}$ is multivariate normal (MVN). The process is made up of two parts; the mean term, $\mu(s) = E(Z(s))$, which corresponds to the global or first-order behaviour and the error term $\epsilon(s)$ which captures the local, or second-order, behaviour through the positive definite covariance function, $C(s, s') = \text{Cov}\{Z(s), Z(s')\}$ (Diggle and Ribeiro, 2007).

Let μ_Z be the J element mean vector with elements $\mu(s_j)$ and Σ_Z be the $J \times J$ matrix with elements $\Sigma_Z(s_i, s_j) = C(s_i, s_j)$. Then $Z(s) = \mu(s) + \epsilon(s)$ and $\epsilon \sim MVN(0, \Sigma_Z)$ so $Z \sim MVN(\mu_Z, \Sigma_Z)$. The covariance function is often called the covariogram when in graphical form.

Figure 5.1 shows the empirical spatial covariogram for Gaussian transformed wind gust speeds centred on London, Amsterdam and Berlin, i.e. $C(s_0, s')$ where s_0 is the nearest grid point to London, Amsterdam or Berlin. The dependence structure is generally similar for the three locations suggesting that representing the dependence in the domain using a covariance function of separation distance may be feasible. There appears to exist a larger correlation length in the west-east directions compared to the north-south direction suggesting that the dependence may be a function of separation direction as well as distance. This is thought to be because most windstorms move across the European region from west to east (Hanley and Caballero, 2012). It is interesting to note that the covariance does not decay to

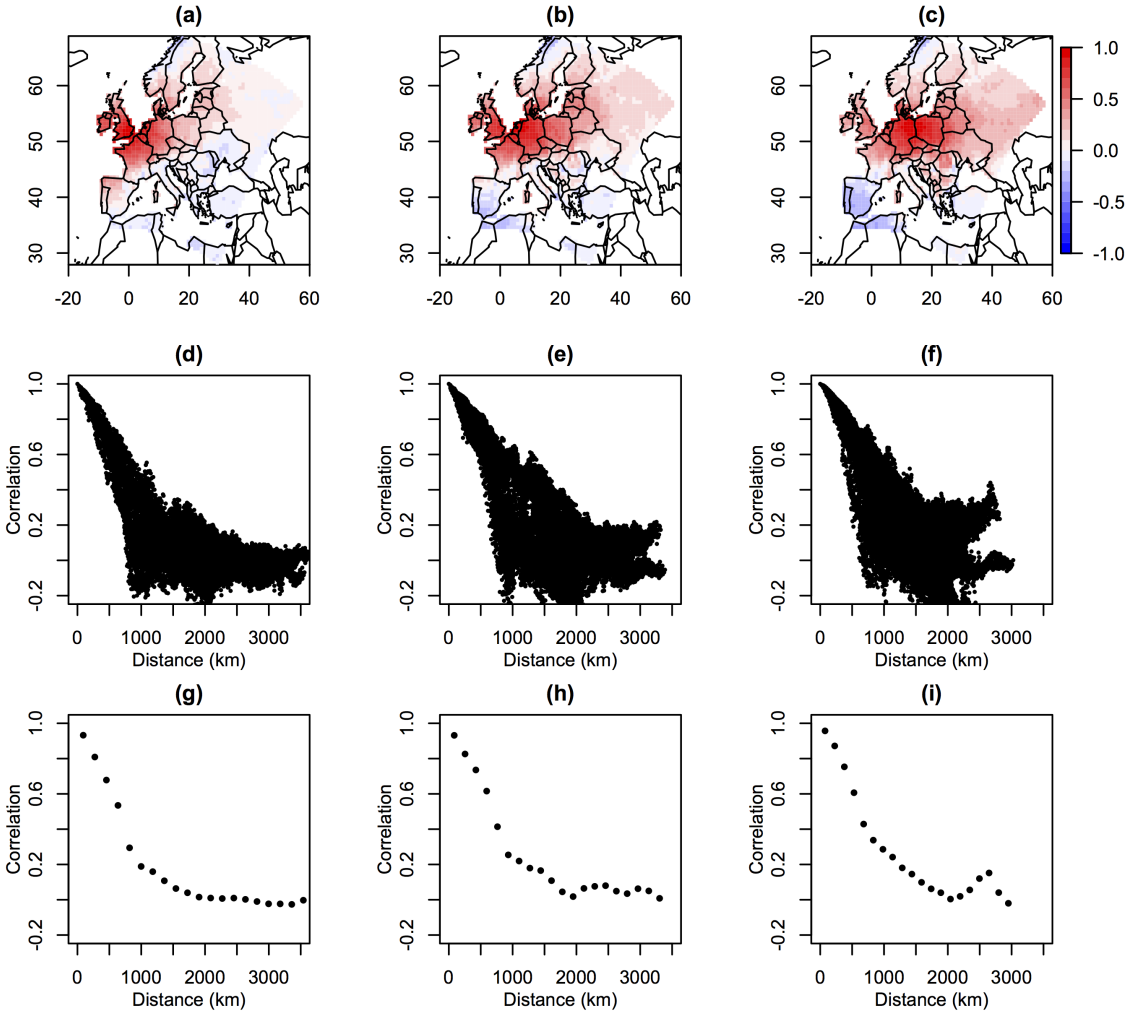

Figure 5.1. Empirical correlation between (a) London, (b) Amsterdam and (c) Berlin and all other land locations over land, plotted against distance in (d), (e) and (f) respectively and for distance binned average correlation in (g), (h), (i) respectively.

zero but becomes negative for a number of pairs of locations with large separation distances. This negative relationship can be interpreted physically as representing how an area of atmospheric low pressure creating strong cyclonic winds will be balanced by a nearby area of high pressure, having weaker wind speeds. The distance binned empirical correlations (Figure 5.1 (g)-(i)) do, however, decay to approximately zero for all three cases. When observing the same empirical correlation plots for other locations in the domain, the spatial dependence structure is observed to vary, particularly in the south of Europe, with locations having either larger or smaller correlation ranges relative to those seen in Figure 5.1. However, it can be argued that, since regions which experience higher wind gust speeds (see Figure 3.1), which are therefore more influential in ensuring the joint losses are modelled correctly, i.e. north-west Europe, have a more consistent dependence structures, there is potential for the use of a geostatistical covariance function for modelling spatial dependence.

The general approach will therefore be to model the spatial dependence in the Gaussian transformed wind gust speeds using a geostatistical covariance function. This methodology is preferred over the max-stable approach of modelling multivariate spatial extremes because, as explained in the previous chapter, all max-stable models assume asymptotic extremal dependence between locations, while using this Gaussian dependence methodology to model the spatial dependence will correctly impose asymptotic extremal independence between locations.

5.3. Marginal distributions

Marginal distribution specification, fitting and goodness-of-fit are as presented in Section 4.3. Fitting the geostatistical spatial covariance function is complicated because of missing values in the data. Therefore a slight variation to the approach in Section 4.3 is used to produce the Gaussian transformed wind gust speeds to which the covariance function is fitted. The left truncated GEV, which is used for the margins, is only valid above a left truncation threshold, v, selected as being the 60% quantile wind gust speed at each location. To ensure no missing values are present in the data, wind gust speeds below the truncation threshold must also be transformed to standard Gaussian. Gaussian transformed wind gust speeds at location, s_j, are therefore now created for wind gust speeds both above and below the truncation threshold:

$$Y_j = \Phi^{-1}(F_j(X_j))$$

for F_j defined as in Section 4.3:

$$F_j(x) = \begin{cases} \hat{F}(x) & x \leq v_j \\ (1-\lambda_j) + \lambda_j \left(\frac{G_j(x) - G_j(v_j)}{1 - G_j(v_j)} \right) & x > v_j \end{cases}$$

where $\lambda_j = 1 - \hat{F}(v_j)$, $\hat{F}$ is the empirical distribution function of X_j and $G_j(x)$ is the GEV distribution function with parameters μ, σ and ξ.

5.4. Modelling spatial dependence

5.4.1. Assumptions

Geostatistical approaches are based on the assumption of spatial second-order stationarity. A process is said to be second-order stationary if the mean is constant in

space; $\mu(s) = \mu$ for all s, and $\text{Cov}\{Z(s), Z(s')\}$ does not depend on the locations s and s', only the distance between them, $d = s - s'$, i.e. $\text{Cov}\{Z(s), Z(s')\} = C(d)$. The variance of a second-order stationary process is also constant, $\sigma^2 = C(0)$.

The Gaussian transformed footprint wind gust speeds Y_j at each location s_j are, by definition, univariate standard Gaussian distributed if the left truncated GEV is well specified, that is

$$Y_j \sim N(0, 1) \quad \text{for } j = 1, ..., J.$$

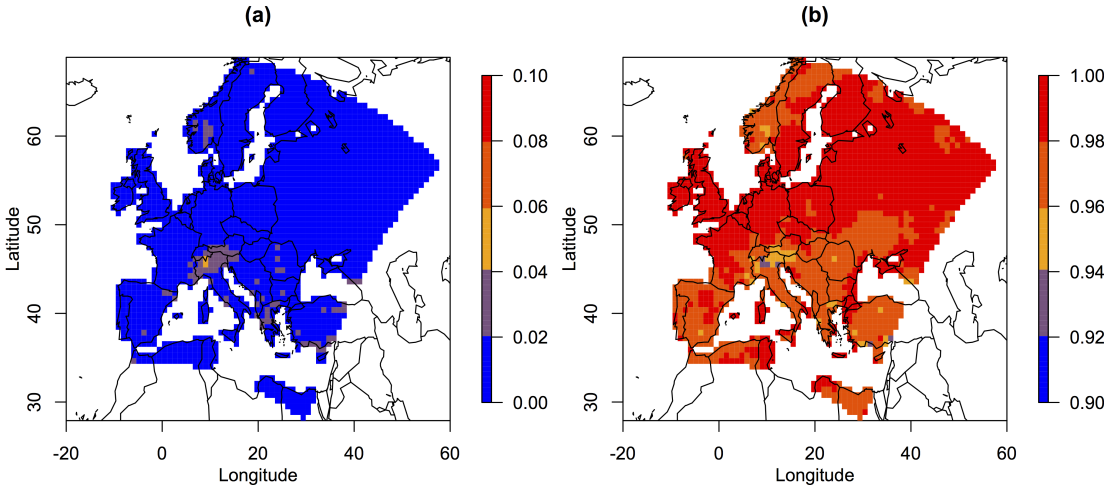

Figure 5.2. The empirical (a) mean and (b) the variance of Gaussian transformed footprint wind gusts y_j at each location s_j over land within the model domain.

Figure 5.2 shows that the empirical mean and variance of the Gaussian transformed footprint wind gusts are approximately equal to 0 and 1 respectively for most locations in the domain, with the mean in the range (0.003, 0.086) and the variance in the range (0.895, 0.993). The greatest deviation from standard Gaussian is in the north of Italy. Referring back to Figure 4.6, this is the region in which the left-truncated GEV marginal distributions have a positive shape parameter, meaning there is no upper limit to the distribution, because the empirical distribution is made up of mainly low values and a small number of much greater extreme values causing the fitted GEV tail to extend infinitely. Figure 5.1 suggests that the covariance could be a stationary function of separation distance and does not depend strongly on the location. It can therefore be concluded that the mean, variance and covariance structure are relatively constant throughout the domain, suggesting second-order stationarity and standard normal behaviour. This also implies that the mean trend surface part of the Gaussian process model can be ignored and since the

variance throughout the domain is unity, the correlation and covariance functions are synonymous.

As an aside, a weaker assumption, relied upon in classical geostatistical theory (Matheron, 1974), is that of intrinsic stationarity which asserts that the increment process $I_d = \{Z(s) - Z(s+d)\}$ is stationary for all difference vectors $d \in \mathbb{R}^2$. Then $E(Z(s) - Z(s+d))$ and $E(Z(s) - Z(s+d))^2$ do not depend on s and the variogram, rather than the covariance function, is used to model the dependence in the Gaussian field, defined as

$$\gamma(d) = \frac{1}{2}E(Z(s) - Z(s+d))^2, \qquad d \in \mathbb{R}^2.$$

Any stationary process is intrinsically stationary, but the converse is not true (Banerjee et al., 2004). Gneiting et al. (2001) explains how the variogram is more general than the covariance function but less well understood theoretically, with many important properties, characterisations and decomposition theorems established for covariance functions only. The covariance function will be used here but similar spatial dependence model fitting could be carried out using the variogram.

5.4.2. Dependence model specification

A stationary spatial process is said to be isotropic if the covariance depends only on the Euclidean distance between locations, i.e. the length of $s - s'$, not on the orientation, then $C(d) = C(\|d\|)$. Here, initially an isotropic model will be fit to the data. The anisotropy identified in Figure 5.1 will then be explored to achieve the most appropriate model. As noted by Banerjee et al. (2004), isotropy is a convenient assumption because of its simplicity, interpretability and there are a number of relatively simple parametric functions available as candidates for the covariance function. Three such isotropic covariance functions are the Exponential, Gaussian and Matérn covariance functions, which have parametric equations:

$$\text{Exponential: } C(t) = \begin{cases} \sigma^2 \exp\left(-\frac{t}{\phi}\right) & \text{if } t > 0 \\ \tau^2 + \sigma^2 & \text{otherwise} \end{cases}$$

$$\text{Gaussian: } C(t) = \begin{cases} \sigma^2 \exp\left(-\frac{t^2}{\phi^2}\right) & \text{if } t > 0 \\ \tau^2 + \sigma^2 & \text{otherwise} \end{cases}$$

$$\text{Matérn: } C(t) = \begin{cases} \frac{\sigma^2}{2^{\nu-1}\Gamma(\nu)} \left(\frac{2\sqrt{\nu}t}{\phi}\right)^{\nu} K_\nu \left(\frac{2\sqrt{\nu}t}{\phi}\right) & \text{if } t > 0 \\ \tau^2 + \sigma^2 & \text{otherwise} \end{cases}$$

where $t = \|d\|$, $\Gamma(\cdot)$ is the gamma function and K_a is the modified Bessel function of order a. The parameters of the covariance functions are the spatial effect variance σ^2; the covariance length scale ϕ; the measurement error variance or non-spatial effect variance τ^2, known as the nugget effect; and, unique to the Matérn covariance function, the shape parameter ν, often referred to as the smoothness parameter (Banerjee et al., 2004). Each parameter will now be introduced in more detail.

The spatial effect variance σ^2, also known as the sill in the context of variogram fitting, represents the global variance. This parameter would therefore equal 1 if the underlying data are standard Gaussian.

The covariance length scale, ϕ, is a measure of the separation distance at which locations are considered to be independent. Banerjee et al. (2004) noted that for the Exponential covariance function, zero covariance is reached asymptotically and ϕ is therefore infinite. They introduced the notion of the 'effective range', t_0, the distance at which the covariance drops to 0.05. By setting $\exp(-t_0/\phi) = 0.05$, it can be derived that $t_0 \approx 3\phi$, since $\log(0.05) \approx -3$, giving ϕ a more intuitive interpretation. A similar calculation can be done for other covariance functions, for example for the Gaussian covariance function $t_0 = \sqrt{3}\phi$.

As explained by Diggle and Ribeiro (2007), the nugget effect, τ^2, has a duel interpretation as either measurement error or spatial variation on a spatial scale smaller than the smallest distance between any two points, or a combination of these two effects.

These three parameters are common to all three covariance functions introduced above, however the Matérn model has another parameter, the shape ν, which adds extra flexibility. This parameter determines the analytic smoothness of the underlying process (Diggle and Ribeiro, 2007), meaning the number of times it is mean-square differentiable (see Appendix B). Specifically a process is $[\nu] - 1$ times mean-square differentiable, where $[\nu]$ is the smallest integer greater than or equal to ν. For $\nu < 1$ the process is non mean-square differentiable and is said to be mean-square continuous. When $\nu = 0.5$ the Matérn reduces to the exponential covariance function and as $\nu \to \infty$ the Matérn approaches the Gaussian covariance function. When the value of ν is a half integer, $m + \frac{1}{2}$ where m is an integer, the modified Bessel function $K_\nu(\cdot)$ takes an explicit form and the Matérn covariance function is a product of $\exp(-2\sqrt{\nu}d/\phi)$ and a polynomial of degree m (Haskard, 2007). The Matérn covariance function gives a broad range of shapes of covariance

functions, allowing any degree of smoothness at the origin, from continuous but non-differentiable when $\nu < 1$ to infinitely differentiable as $\nu \to \infty$. This flexibility means the Matérn covariance function is often preferred in practise, however the shape parameter, ν, can be problematic to estimate (Diggle and Ribeiro (2007), Stein (1999)). More detail is given on this in Section 5.4.3.

A number of parameterisations for the Matérn covariance function exist with the parameterisation proposed by Handcock and Wallis (1994) presented above. An alternative parameterisation, named the Whittle covariance model, has scale parameter $\varsigma = \phi/2\sqrt{\nu}$. Diggle and Ribeiro (2007) explained how within the Whittle parameterisation ν and ς are non-orthogonal. Suppose the true correlation structure of a process is Whittle with parameters ν and ς then the best fitting approximation with order $\nu^* \neq \nu$ will also have $\varsigma^* \neq \varsigma$ so the scale parameters corresponding to different orders of Matérn covariance are not directly comparable. The Matérn parameterisation is preferred over the Whittle parameterisation because the shape, ν, and scale, ϕ, are almost orthogonal in their effect on the induced covariance structure (Diggle and Ribeiro, 2007). Haskard (2007) explains how in the Matérn covariance function ν primarily determines the behaviour of the covariance function at the origin and ϕ broadly determines how quickly the covariance decreases with distance, roughly independently of ν. Haskard (2007) goes on to suggest that using the Matérn parameterisation should therefore improve the stability of covariance model parameter estimation.

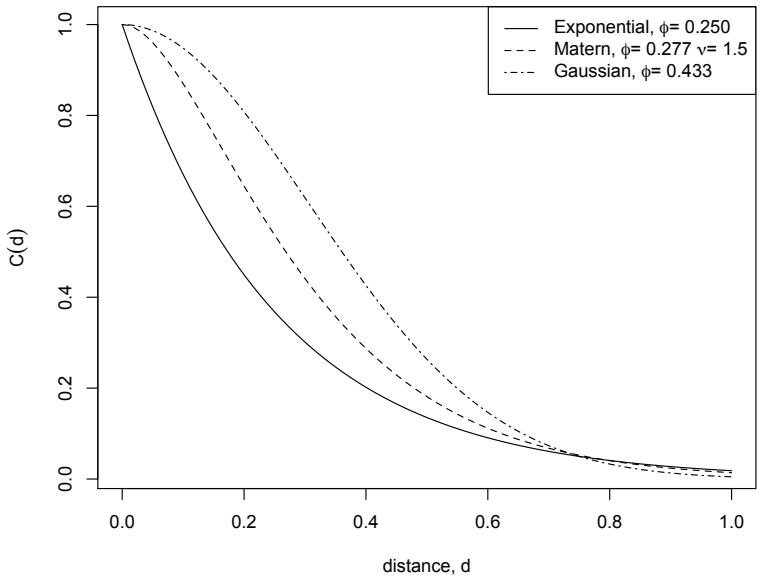

Figure 5.3. Exponential, Matérn ($\nu = 1.5$) and Gaussian covariance functions with scale parameter ϕ adjusted to give the same practical range.

Figure 5.3 shows the covariance functions for the exponential model (equivalent to the Matérn with $\nu = 0.5$), the Matérn model with $\nu = 1.5$ and the Gaussian model (equivalent to the Matérn with $\nu = \infty$). For all functions $\sigma^2 = 1$ and $\tau^2 = 0$. The values of ϕ have been adjusted so all three functions have the same practical range, $t_0 = 0.75$, calculated by setting the covariance to 0.05 and $d = 0.75$. For the Matérn function the appropriate value of ϕ is ascertained by plotting the covariance function for $d = 0.75$ and $\phi \in (0, 1)$ and finding which value of ϕ corresponds to a covariance of 0.05. The covariance initially decays quickest for the exponential model followed by the Matérn ($\nu = 1.5$) model meaning locations that are close together are less correlated for the exponential model, with the Gaussian model having the greatest correlation at small distances making the process smoother. This behaviour of the covariance functions near the origin is based on the increasing value of ν in the equivalent Matérn functions.

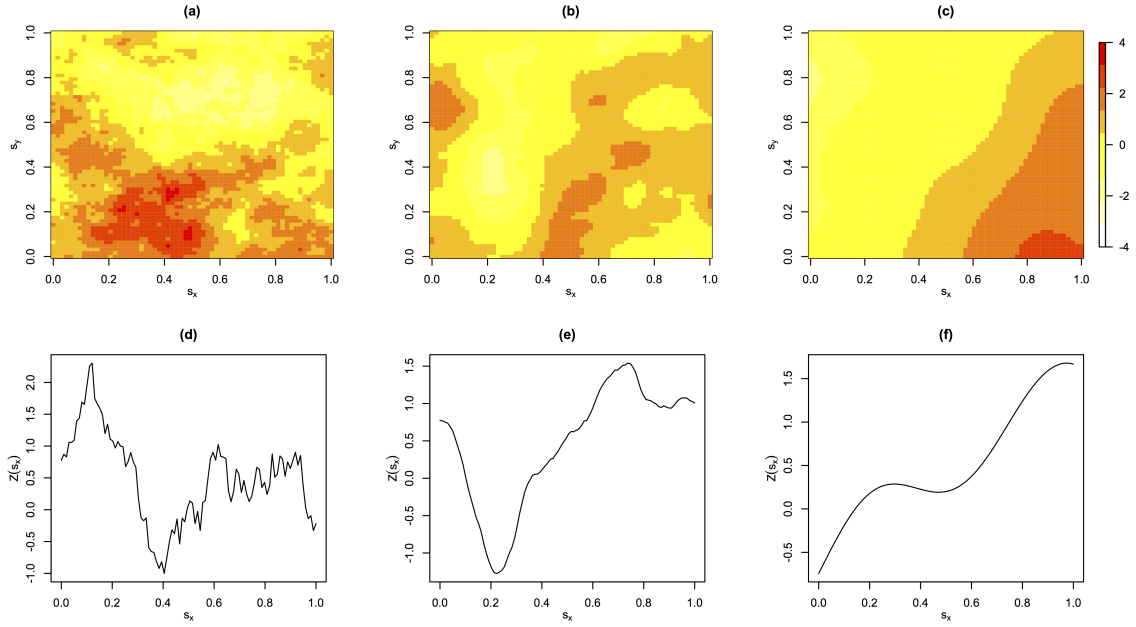

Figure 5.4. (a)-(c) two-dimensional simulations from Gaussian processes with Exponential, Matérn ($\nu = 1.5$) and Gaussian covariance functions respectively. (d)-(f) one-dimensional representations of the same Gaussian processes respectively, a cross-section of the two-dimensional simulation at $s_y = 0.5$.

This increasing analytic smoothness of the processes as ν increases has a large effect on the resulting covariance structure of the Gaussian process. This can be seen in Figure 5.4 (a)-(c) which show simulated two-dimensional realisations of Gaussian processes and (d)-(f) which show a one-dimensional cross-section for each of the three covariance functions in Figure 5.3. The most noticeable difference is between the non-differentiable exponential covariance structure, which gives a much rougher process compared to the differentiable Matérn ($\nu = 1.5$) and Gaussian covariance structures.

Rather than imposing a mean-square differentiability property by fitting the exponential or Gaussian covariance functions, the Matérn covariance function will be used here, allowing the data to define this property by estimating the shape parameter ν.

Spatial processes are often not isotropic, that is the covariance depends on the separation direction as well as the separation distance. Anisotropy can be accounted for within a geostatistical model, while still allowing for the use of a simple isotropic covariance function, by transforming the domain via a stretching and rotation of the coordinate axes. Therefore anisotropy is defined by additional parameters, θ, the anisotropy angle, and spatial scale parameters in the longitudinal, ϕ_1, and latitudinal, ϕ_2, directions. Alternatively, as in Diggle and Ribeiro (2007), an anisotropic ratio can be defined which is simply ϕ_1/ϕ_2. Here, ϕ_1 and ϕ_2 will be represented separately so they can be varied separately, allowing for greater flexibility in spatial dependence structures within the sensitivity study in the next chapter of this thesis. A process with anisotropy in spatial coordinates $s = (s_1, s_2)$ can be converted to an isotropic process in coordinates $\tilde{s} = (\tilde{s}_1, \tilde{s}_2)$ by the transformation

$$(\tilde{s}_1, \tilde{s}_2) = (s_1, s_2) \begin{pmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{pmatrix} \begin{pmatrix} 1/\phi_1 & 0 \\ 0 & 1/\phi_2 \end{pmatrix}. \tag{5.1}$$

Figure 5.5 illustrates this transformation. The directional effect is clearly visible with longitudinal scale 4 times as large as latitudinal scale and an axes rotation of $-\pi/3$, resulting in the longitude pointing along the diagonal. Anisotropy will be identified by fitting both an isotropic and anisotropic Matérn covariance function to the spatial data and observing which is the better fit to the data.

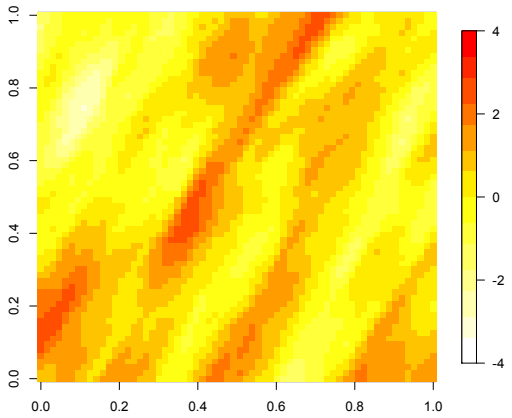

Figure 5.5. A simulation from an anisotropic Gaussian spatial process with a Matérn ($\nu = 1.5$) covariance structure with anisotropy parameters $\phi_1 = 0.277, \phi_2 = \phi_1/4 = 0.069$ and $\theta = -\pi/3$.

5.4.3. Dependence model fitting

In the general case, fitting a spatial Gaussian process to candidate data requires the estimation of the mean and covariance function parameters. This means making inference about the underlying spatial process, $Z(s)$, based upon the partial realisation $Y_j : j = 1, ..., J$, at spatial locations $s_j : j = 1, ..., J$.

The mean function $\mu(s)$ can most simply be represented by a linear regression model on K spatial covariates, $\{w_1(s), ..., w_K(s)\}$, such as longitude, latitude and altitude,

$$\mu(s) = \beta_0 + \sum_{k=1}^{K} \beta_k w_k(s),$$

where $\beta = (\beta_0, \beta_1, ..., \beta_K)$ must be estimated, based on the data Y. The number of covariance function parameters to be estimated depends on whether the data are isotropic or anisotropic. When isotropy is assumed

$$Y \sim MVN(W\beta, \Sigma(\tau^2, \sigma^2, \phi, \nu)),$$

where the shape parameter ν is only included for the Matérn covariance function. Alternatively, if anisotropy is assumed the covariance function includes the parameters to stretch and rotate the spatial domain, therefore

$$Y \sim MVN(W\beta, \Sigma(\tau^2, \sigma^2, \theta, \phi_1, \phi_2, \nu)).$$

Parameters can be estimated using maximum likelihood (ML). Suppose that at each location the response variable Y_j is observed n times. Then define y_{ij} as the i^{th} observation of the response variable at the j^{th} location for $i = 1, ..., n$ and $j = 1, ..., J$. Then define $y_{i\cdot}$ as the i^{th} observation of the response variable for all J locations and $y_{\cdot j}$, all n observations of the response variable at location j. Then the log-likelihood $l(\beta, \psi; Y)$, where $\psi = (\tau^2, \sigma^2, \phi, \nu)$ in the isotropic case and $\psi = (\tau^2, \sigma^2, \theta, \phi_1, \phi_2, \nu)$ in the anisotropic case, is derived from the multivariate Gaussian density function as,

$$\begin{aligned}
f_Y(y_{\cdot 1}, ..., y_{\cdot J}) &= L(\beta, \psi; Y) && (5.2)\\
&= \prod_{i=1}^{n}((2\pi)^J |\Sigma(\psi)|)^{-\frac{1}{2}} \exp\left(-\frac{1}{2}\sum_{i=1}^{n}(y_{i\cdot} - W\beta)^T \Sigma^{-1}(\psi)(y_{i\cdot} - W\beta)\right) && (5.3)\\
l(\beta, \psi; Y) &= \log(L(\beta, \psi; Y)) && (5.4)\\
&= -\frac{nJ}{2}\log(2\pi) - \frac{n}{2}\log|\Sigma(\psi)| - \frac{1}{2}\sum_{i=1}^{n}(y_{i\cdot} - W\beta)^T \Sigma^{-1}(\psi)(y_{i\cdot} - W\beta)
\end{aligned}$$

maximisation of which yields the ML estimates of the model parameters. This maximisation is done here numerically using the quasi-Newton optimisation algorithm

introduced by Byrd et al. (1995).

An alternative approach for covariance model fitting, known as the Least Squares (LS) approach, is also commonly discussed (e.g. Diggle and Ribeiro 2007, Stein 1999 and Banerjee et al. 2004). The LS approach involves estimating β, removing the trend surface from the data and then fitting the covariance function, via a curve fitting algorithm, to the empirical binned covariogram, defined as

$$C_{ij}^{*} = \frac{1}{N_{B_{ij}}} \sum_{(k,l):\|s_k - s_l\| \in B_{ij}} \mathrm{Cov}(Y_k, Y_l) \tag{5.5}$$

where $N_{B_{ij}}$ is the number of pairs of locations that have separation distance in bin B_{ij}. The LS approach is equivalent to a non-linear regression analysis, treating the empirical covariogram as the response variable and the separation distance-bin midpoint as the corresponding explanatory variable (Diggle and Ribeiro, 2007). The ML approach has a number of advantages over the Least Squares (LS) approach. The ML approach allows for the mean trend parameter vector β to be estimated at the same time as the covariance function parameters, ψ while this must be done in two separate steps for the LS approach. The LS parameter estimates for the covariance function are very sensitive to the choice of empirical covariogram bins so a significant advantage of the ML approach is that it does not require this arbitrary bin selection, using all separation distances within the data.

One of the main objections to using the ML approach is the assumption of multivariate normality of the data, explicit in the definition of the log-likelihood function in Eqn (5.3). The data at each location are univariate normally distributed, which is only a necessary and not sufficient condition for multivariate normality. Pardo-Igúzquiza (1998) and Lark (2000) noted that the actual multivariate distribution that underlies the data can, in fact, never be verified because the data are only a realisation from this distribution. Pardo-Igúzquiza (1998) argued that the multivariate normal distribution is therefore a natural assumption because it is the distribution of maximum entropy when all that is known is the mean and covariance matrix, which fully explain a multivariate Gaussian process. Pardo-Igúzquiza (1998) identified how the log-likelihood is made up of two terms that depend on the covariance parameters:

$$(1) \quad \log|\Sigma(\psi)|$$

$$(2) \quad \sum_{i=1}^{N} (y_i - W\beta)^T \Sigma^{-1}(\psi)(y_i - W\beta).$$

Term (1) is a general measure of the spatial uncertainty which decreases as the degree of spatial correlation increases. Term (2) is a weighted least-squares criterion for the fit of the model which increases as the degree of spatial correlation increases. The parameter selection therefore finds the model with the strongest spatial dependence structure while not incurring an excessive penalty for the squared error term. This suggests that the set of parameters that maximises Eqn (5.3) is the rational choice even when the distributional assumption does not hold.

The ML approach is more computationally intensive than the LS approach. The $J \times J$ covariance matrix in Eqn (5.3) must be inverted at each iteration of the optimisation algorithm used to maximise the log-likelihood. This can be infeasible if J is very large. The computation can be made quicker by using the Cholesky decomposition of the covariance matrix,

$$\Sigma = LL^T,$$

where L is a lower triangular matrix. It is quicker to compute the inverse of the triangular matrix L and L can be used to calculate the inverse of Σ as

$$\Sigma^{-1} = (L^{-1})^T(L^{-1}).$$

It is also quicker to calculate the determinant of the triangular matrix L and

$$|\Sigma| = |L|^2.$$

To reduce computational expense further, Stein (1999) identified the possibility of using spectral methods to approximate the likelihood. Since any joint density can be written as a product of conditional densities based on some ordering of the observations, the computation is lessened by conditioning on only some of the nearest observations when computing the conditional densities. The smaller the number of neighbouring locations used the quicker the computation but the worse the approximation of the true joint density. Stein (1999) also suggested dividing the spatial domain into a number of subregions and calculating the likelihood for each subregion separately and multiplying them together. Again, the smaller these subregions the quicker the computation but the worse the approximation of the likelihood and a trade off must be made. It may also be that the correlation scale is larger than these subregions and will therefore be poorly approximated. For both of these methods arbitrary neighbourhood choices must be made so the sensitivity to these choices should also be investigated.

Stein (1999) went on to discuss further computational issues of the ML approach, noting that the method will not necessarily converge to the global maximum if more

than one local maxima exists, which may be the case if the likelihood function is not twice differentiable. Warnes and Ripley (1987) also showed an example of multiple local maxima when using the exponential covariance function, although these results are contested by Stein (1999) who claimed to be unaware of any examples of multiple local maxima in the likelihood for the Matérn model. He did note, however, that the global maximum may not be within the parameter space and may occur as the shape parameter $\nu \to \infty$. Stein (1999) and Warnes and Ripley (1987) agree that there can exist long ridges within the parameter space along which the likelihood is nearly constant. This can lead to numerical problems when using iterative procedures for finding the maximum. Stein (1999) explained how the presence of these ridges is not a sign of a problem with the likelihood methods, rather an entirely correct indication that the data provides no information for distinguishing between parameter values along the ridge. Stein (1999) suggested plotting the likelihood surface in order to detect these ridges.

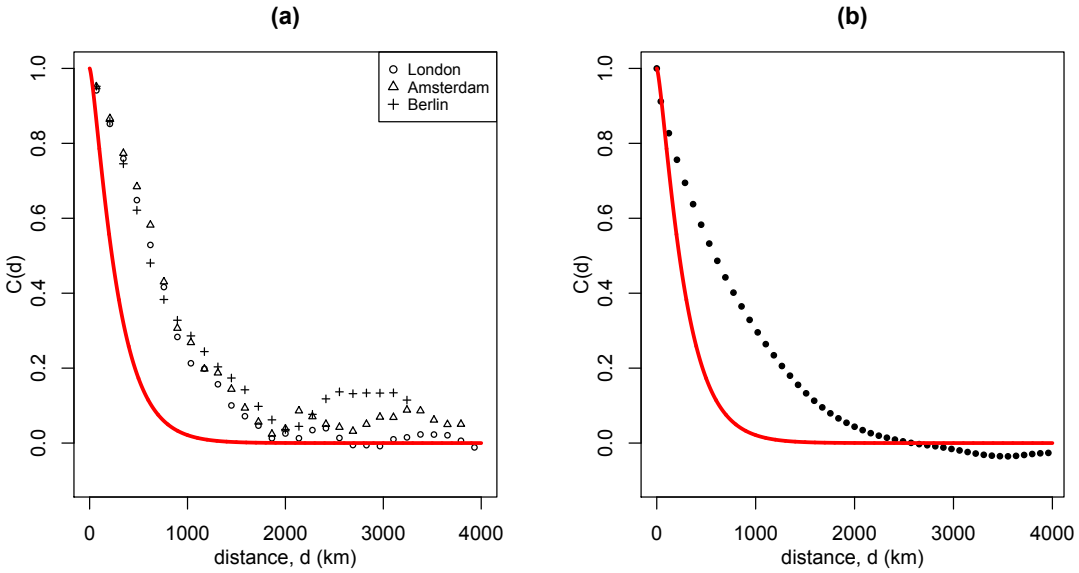

Figure 5.6. The best fit isotropic Matérn model (red line) compared with (a) the empirical binned covariograms for London, Amsterdam and Berlin paired with all other locations and (b) the empirical binned covariogram for all pairs of locations together.

It was shown in Figure 5.2 that, throughout the domain, the empirical mean of the Gaussian transformed wind gust speeds is approximately constant and equal to zero. It is therefore concluded that the mean function can be ignored when modelling this data as a Gaussian process and only the covariance function parameters need to be estimated. The isotropic Matérn covariance function was fit to the data using the ML approach where $\sigma^2 = 1$, a reasonable assumption based on Figure 5.2. The set of parameters to estimate is therefore $\psi = (\tau^2, \phi, \nu)$ in equation (5.3), estimated as $\phi = 260.031, \nu = 0.776, \tau^2 = 0$. Figure 5.6 shows the fit of this Matérn model

to the empirical binned covariogram (Eqn (5.5)). It is clear from Figure 5.6 that the ML estimate of the best fit Matérn covariance structure fails to represent the dependence structure of the data, but why is this?

Diggle and Ribeiro (2007) and Stein (1999) explained how, in their experience, the quadratic approximation of the log-likelihood surface, i.e. estimating ϕ and ν together, is often difficult because the parameters are poorly identified, leading to ridges and plateaus in the likelihood surface. Diggle and Ribeiro (2007) advised choosing a fixed value of ν from a discrete set of values that cover different degrees of mean-square differentiability, for example $\{0.5, 1.5, 2.5\}$, optimising individually for ϕ. The best fit is then found by observing which value of ν maximises the likelihood. This method is known as a profile likelihood approach and avoids the optimisation algorithm getting trapped at a local maximum rather than the global maximum in the likelihood surface.

The profile likelihood approach is carried out for $\nu = \{0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0, 1.2, 1.5, 2.5\}$ with the resulting models for a selection of these values of ν shown in Figure 5.7 (a) and the corresponding maximum log-likelihood value for each ν shown in Figure 5.7 (b).

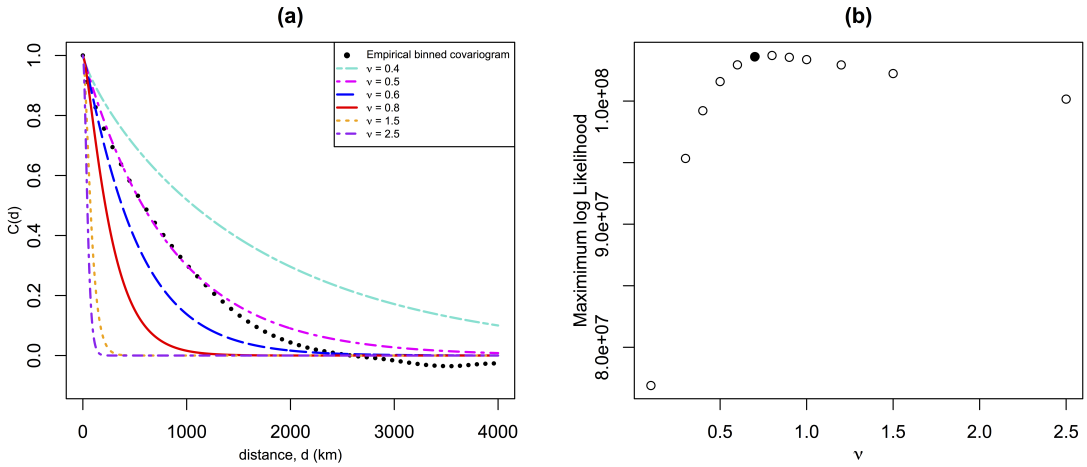

Figure 5.7. (a) The best fit isotropic Matérn model using the profile likelihood approach of fixing the shape parameters ν with the empirical binned covariogram and (b) the maximum log-likelihood for a range of fixed values of ν, with the maximum at $\nu = 0.8$ indicated by a solid point.

Figure 5.7 (b) shows that the global maximum of the likelihood surface is at $\nu \approx 0.8$ indicating that the original ML optimisation did indeed find the global maximum and not a local maximum, as may have been the issue, suggesting that the likelihood surface does not have multiple maxima. The maximum log-likelihood is, however, relatively similar for $\nu \in (0.5, 2.5)$ suggesting that a plateau or ridge may exist in

the likelihood surface.

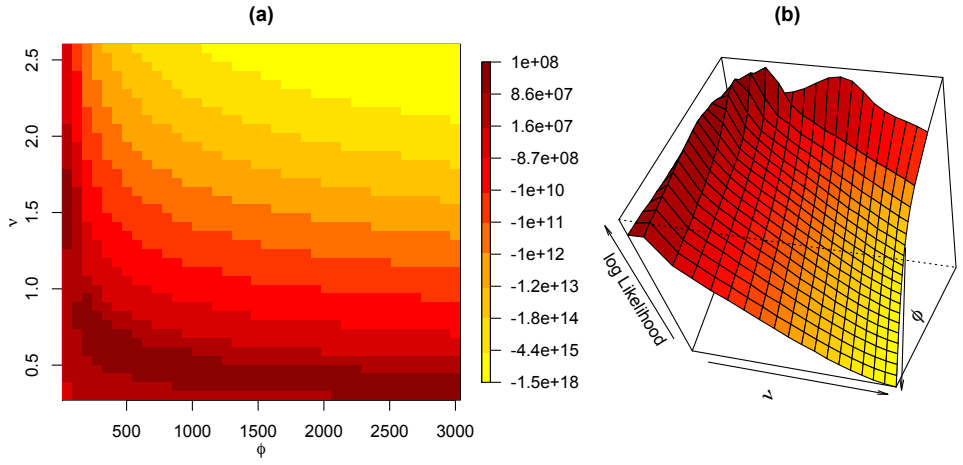

Figure 5.8. The log-likelihood surface based on the data for isotropic Matérn covariance parameters ν and ϕ in (a) 2 dimensions and (b) 3 dimensions. The colour scale is partitioned into quantiles of the log-likelihood at 0.1 intervals.

As suggested by Stein (1999), any ridges can be identified by plotting the likelihood surface over a range of values of ν and ϕ, as shown in Figure 5.8. It is clear that there is a significant ridge in the log-likelihood surface, along which the data provides no or very little information for distinguishing between best fit parameter values.

Referring to Figure 5.7 (a) it is clear that most of the Matérn models that result from the parameter estimates along this ridge in the likelihood surface do not fit the empirical dependence structure for the whole range of separation distances. An explanation of this is hinted at by Diggle and Ribeiro (2007), in an example application of the ML approach for fitting the Matérn covariance function to surface elevation data. Diggle and Ribeiro (2007) pointed out that the likelihood criterion automatically takes account of the fact that sample covariances become less precise as the separation distance increases, meaning that more weight is given to pairs of locations with smaller separation distances within the ML estimate of the covariance parameters. This explains the bad fit for larger separation distances for most of the models in Figure 5.7 (a), rather the ML approach is focused on fitting the covariance structure for locations with smaller separation distances.

The exact weighting for different separation distances can be investigated by observing the construction of the likelihood function. The part of the function that

depends on the data

$$\sum_{i=1}^{n}(y_i - W\beta)^T\Sigma^{-1}(\psi)(y_i - W\beta)$$

is weighted by the inverse of the covariance matrix, $\Sigma^{-1}(\psi)$, also known as the precision matrix. The precision matrix is a sparse matrix and has an interpretation in terms of partial correlations and partial variances.

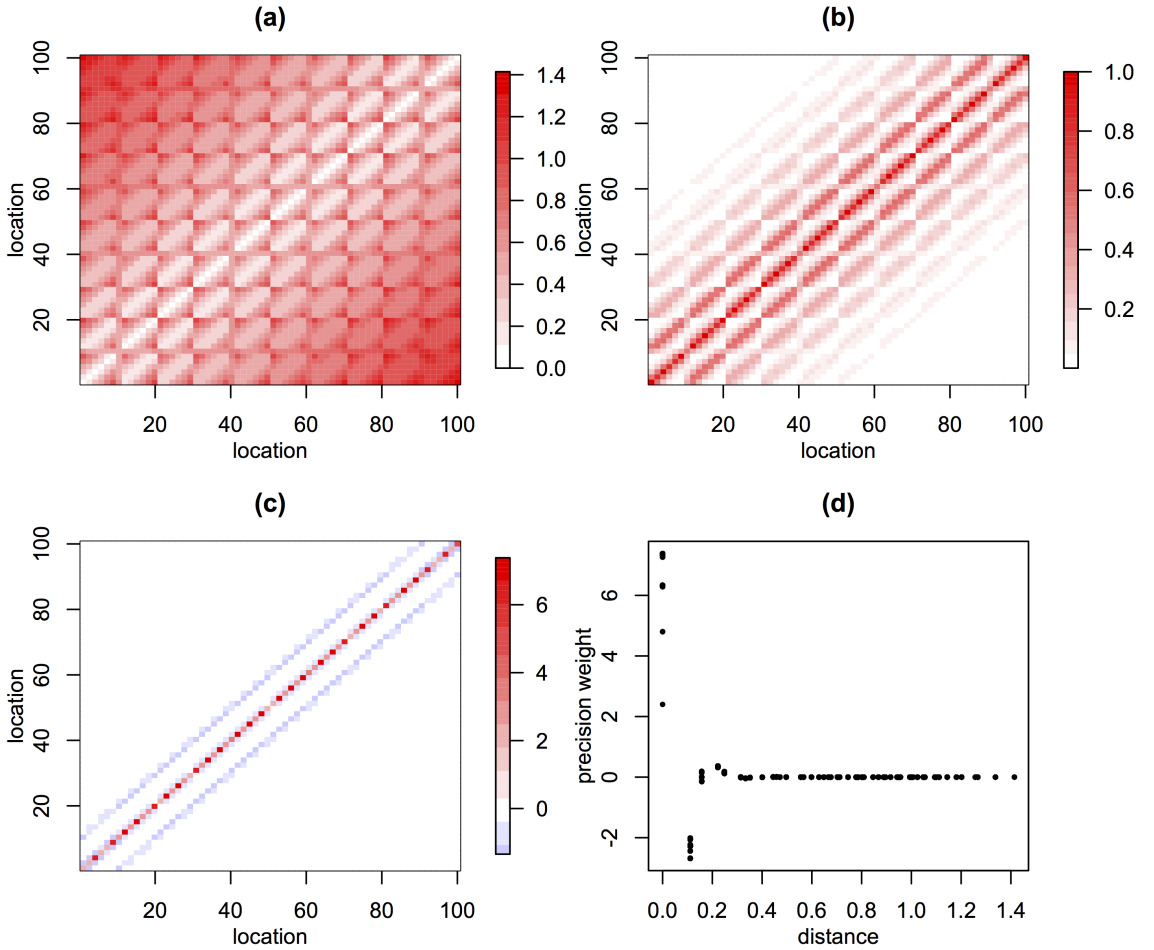

Figure 5.9. (a) The distance matrix for a set of 100 locations on a 10×10 grid in the domain $(0, 1)^2$, (b) the associated Matérn covariance matrix where $\tau^2 = 0, \nu = 0.8$ and $\phi = 0.2$ (chosen arbitrarily), (c) the associated inverse covariance matrix or precision matrix, and (d) the precision matrix plotted against the distance matrix representing the weight given to locations separated by a given distance within the likelihood function.

A toy example of the precision matrix weights given to a set of locations on a 10×10 grid in the domain $(0, 1)^2$ for a simulated Gaussian process with a Matérn dependence structure is shown in Figure 5.9. The inverse covariance matrix in Figure 5.9 (c) is very sparse with most entries equalling approximately zero. The highest weights are on the diagonal of the precision matrix coinciding with a separation

distance of zero and negative weighting is given to the next smallest separation distance of 0.1. This is reflected in Figure 5.9 (d) which shows that within the likelihood function almost no weight is given to locations separated by more than 0.3. Therefore pairs of locations separated by more than 0.3 have effectively no influence in the likelihood function and the resulting fit of the covariance function, while pairs of locations separated by less than 0.1 have a very large influence.

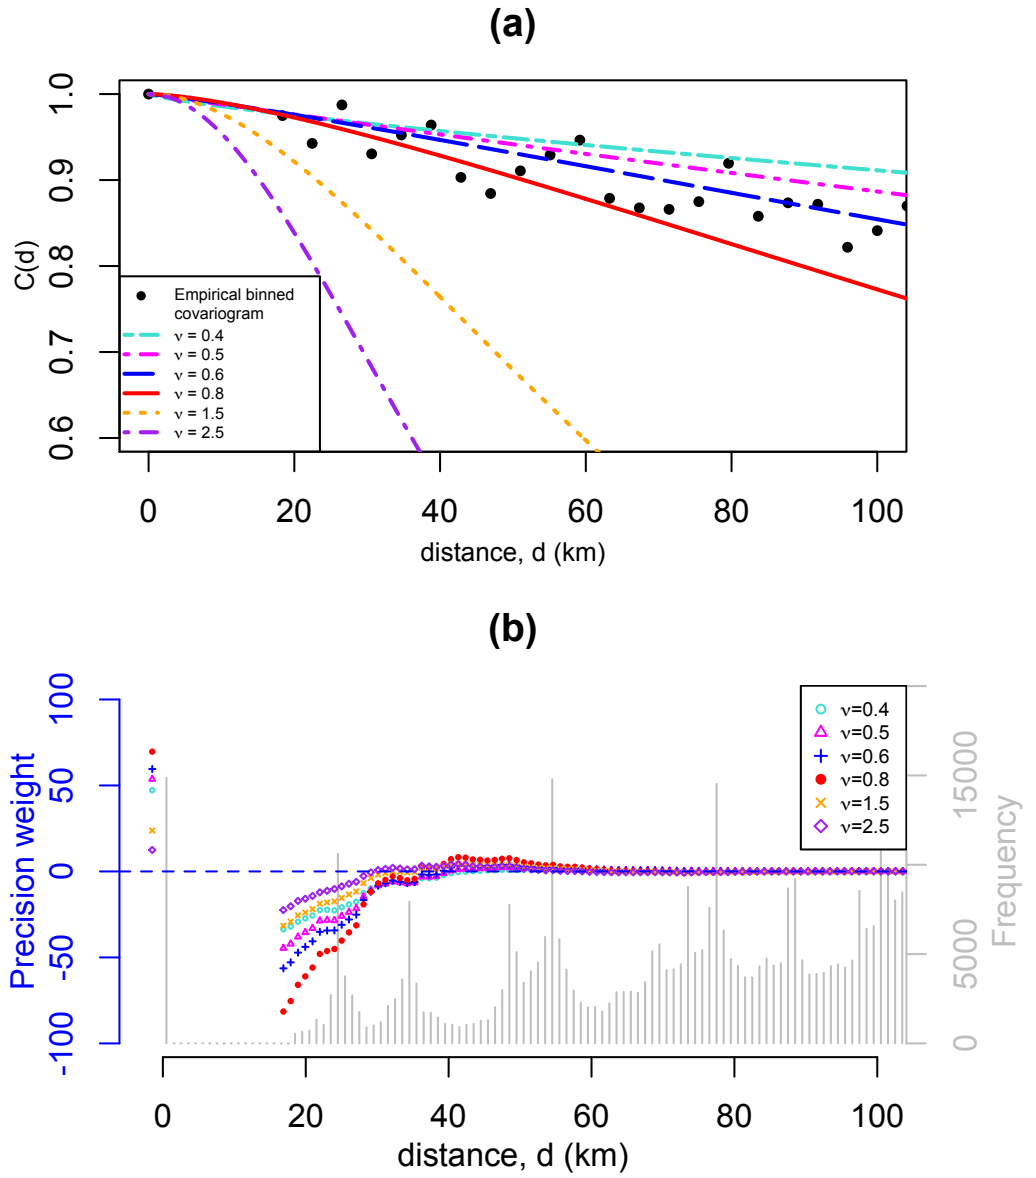

Figure 5.10. (a) As in Figure 5.7 (a) but for separation distance $d \in (0, 100)$km and (b) the precision matrix (inverse of the covariance matrix) weights for binned separation distances $d \in (0, 100)$km with histogram showing the frequency of each separation distance bin within the data.

For each of the profile likelihood fits shown in Figure 5.7, the precision matrix (inverse covariance matrix) weight can be calculated for each pair of locations and averaged over separation distance bins. These weights are shown in Figure 5.10 (b)

and the associated fits for separation distance $d \in (0, 100)$ km shown in Figure 5.10 (a). For all of the fits the precision weights decay to approximately zero at a separation distance of 50km which is very low considering locations within the domain are separated by up to 4000km. A histogram of separation distances is added to Figure 5.10 since separation distances with a higher frequency will also have more influence on the likelihood. It can be concluded that the likelihood function is weighting the fit of the covariance model to a very small proportion of the dependence structure, explaining why the maximum likelihood is large for models that fit badly to the data for higher separation distances (Figure 5.7). For example the model associated with the ML estimate of $\nu = 0.8$ does indeed fit the empirical binned covariogram very well for separation distances less than 50km, seen in Figure 5.10.

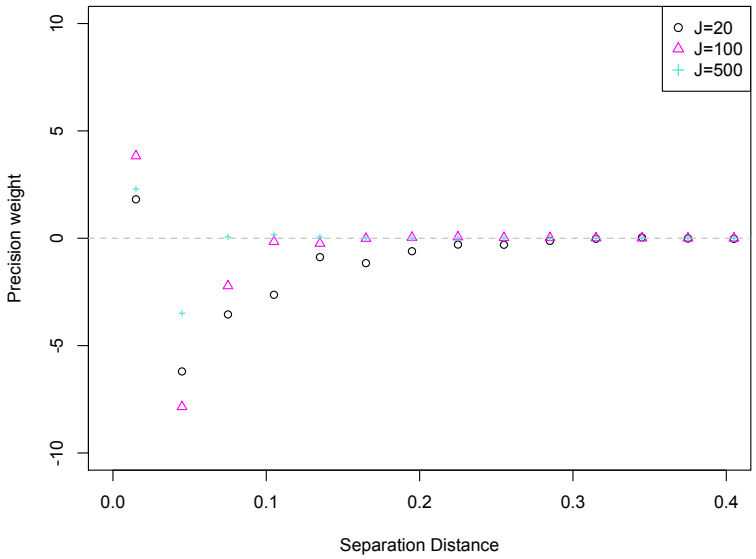

Figure 5.11. The precision matrix (inverse covariance matrix) weights for different separation distances based on the Matérn covariance function with $\tau^2 = 0, \nu = 0.8$ and $\phi = 0.2$ for different numbers of locations in the domain.

This large weighting of small separation distances in the likelihood function is due to the large number of locations within the domain for the footprint data set. Figure 5.11 shows how the precision matrix weights vary for different numbers of locations, J, sampled at random from the domain $(0, 1)^2$. The larger the value of J, the quicker the precision matrix weights decay to zero with separation distance and therefore the more relative weight is placed on very small separation distances. In conclusion, the data provides no information for distinguishing between parameter values along the ridge in the likelihood because a relatively large weight is placed on small

separation distances in the likelihood function due to the large number of locations ($J = 14872$) in the data set.

One way to overcome this issue may be to fit the model to a number of smaller sets of locations sampled from the original locations. This would increase the weighting towards larger separation distances, however it would require the arbitrary selection of an appropriate sample size. Alternatively the model that best fits the dependence structure for a majority of the separation distance range could be chosen from the profile likelihood models. From Figure 5.7 (a) it is clear that this is the Matérn model with shape parameter $\nu = 0.5$. This model fits the data well at small separation distances as was seen in Figure 5.10 (a), however the fit is not as good for separation distances greater than 1500km where the covariance in the model does not decay to zero quickly enough and cannot capture the negative empirical covariances that exist for very large separation distances. Using this shape parameter in the Matérn covariance model has the additional simplification of being equivalent to the exponential covariance model which has one fewer parameter than the Matérn covariance model. The parameters of this isotropic exponential model are estimated as $\tau^2 = 0$ and $\phi = 831.703$. The physical interpretation of this best fit model is that wind gusts are a rough, non-differentiable spatial process, with covariance between Gaussian transformed wind gust speeds reducing exponentially with separation distance at a rate of $1/\phi$ per km.

If the decay in covariance with separation distance was exponential then the relationship between the natural logarithm of the bin averaged empirical covariance and separation distance should be linear. This relationship is presented in Figure 5.12. The relationship is linear for separation distances $0 - 1500$km but the decay is quicker than exponential for separation distances greater than 1500km, as was shown in Figure 5.7 (a). In conclusion, this exponential model fits very well for small and medium separation distances and less well for large separation distances. The variation in covariances at larger separation distances is greater, as can be seen in Figure 4.9 (d)-(f), therefore the fit of the model here is less precise and the fit of the model for smaller separation distances is more representative of the dependence structure in the domain.

The mean-squared differentiability of the spatial process will remain the same, irrespective of whether the process is isotropic or anisotropic. This property was explored in the isotropic case for simplicity and clear visualisation of the likelihood surface. Since the spatial dependence within the data is best modelled when $\nu = 0.5$, the existence of anisotropy is now investigated by fitting the anisotropic exponential

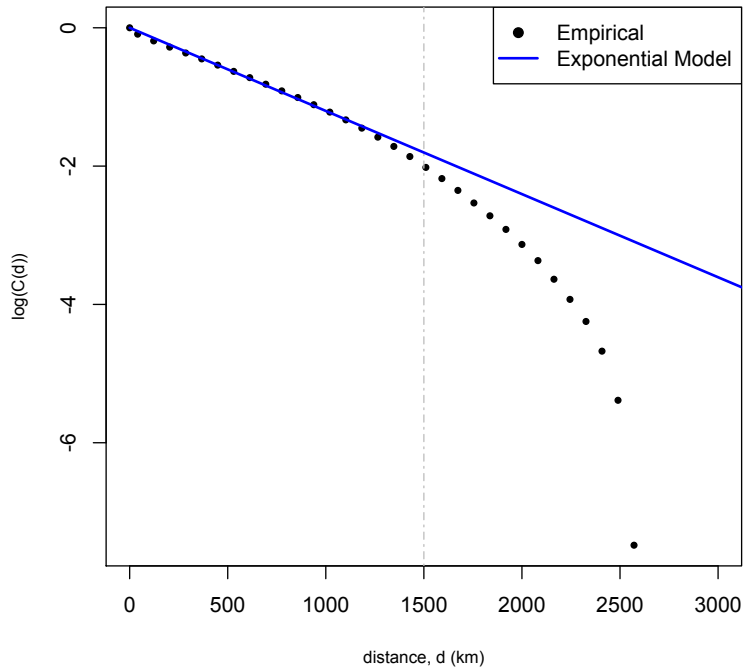

Figure 5.12. The relationship between separation distance and the natural logarithm of the empirical covariance. The line $x = 1500$ indicates where the relationship becomes non-linear.

covariance model to the data and identifying whether this results in a better fit than the isotropic model.

The set of parameters in equation (5.3) becomes $\psi = (\tau^2, \phi_1, \phi_2, \theta)$. These are estimated as $\tau^2 = 0$, $\phi_1 = 869.543$, $\phi_2 = 797.802$ and $\theta = -0.0179 = -\pi/175.5$ degrees. These parameter estimates are inevitably uncertain, however, since the purpose of the model is to explore the sensitivity of insured loss to changes in the parameters, their uncertainty range is not important, rather that they produce synthetic footprints that give realistic losses.

The best fit anisotropic exponential covariance model is shown in Figure 5.13. The separation distances are represented in relation to the transformed spatial coordinates $\tilde{s}_j$ as in equation 5.4.2, therefore the distances in Figure 5.13 are $\tilde{d}_{ij} = \|\tilde{s}_i - \tilde{s}_j\|$. The exponential model with unit covariance length scale is then plotted alongside to assess the fit of this anisotropic exponential model. Similarly to the isotropic model, the model fits very well for small and medium separation distances, however the covariance in the data decays quicker than the exponential model for larger separation distances.

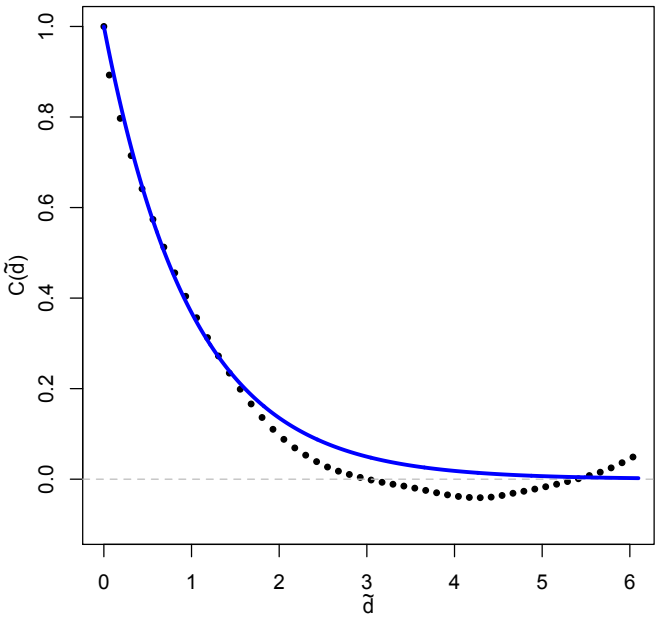

Figure 5.13. The empirical binned covariogram for spatial coordinates transformed using the ML estimated anisotropic covariance parameters (black dots), with the unit scale exponential covariance model added to assess the fit of the anisotropic exponential dependence model (blue line).

This anisotropic model has a greater likelihood than the isotropic model. This is to be expected because it has more parameters, however, since the longitudinal covariance/correlation length scale ϕ_1 is approximately 70km longer than the latitudinal covariance/correlation length scale, ϕ_2, this anisotropic model better represents how windstorms pass over the European region from west to east.

A comparison of the covariance/correlation functions for the isotropic and anisotropic models is presented in Figure 5.14, illustrating the scale of the spatial dependence within the domain and how the anisotropic dependence structure better represents the expected movement of windstorms across Europe. The anisotropic model will therefore be used to represent the spatial dependence within the windstorm footprint model. Including anisotropy parameters within the windstorm footprint model will also allow for greater variation in the spatial dependence structure of simulated footprints in the sensitivity study in the next chapter.

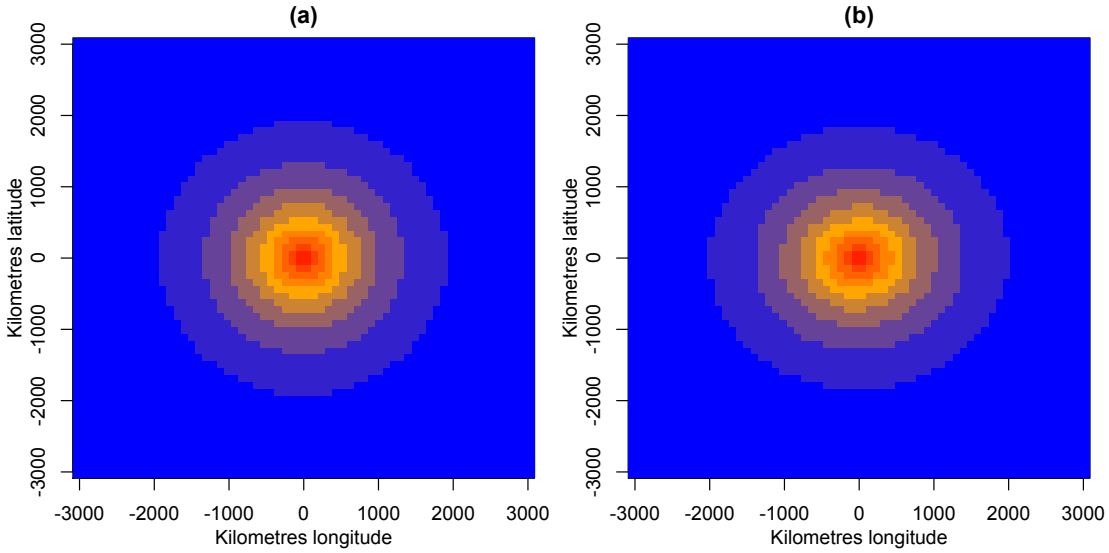

Figure 5.14. Best fit (a) isotropic and (b) anisotropic exponential covariance functions.

Footprint simulation

The windstorm footprint simulation model is specified as

$$X_j = F_j^{-1}(\Phi(Y_j))$$

$$F_j(x) = \begin{cases} \hat{F}(x) & x \leq v_j \\ (1-\lambda_j) + \lambda_j \left(\frac{G_j(x) - G_j(v_j)}{1 - G_j(v_j)} \right) & x > v_j \end{cases}$$

$$\begin{aligned} \lambda_j &= 1 - \hat{F}_j(v_j) \\ G_j &= GEV(\mu_j, \sigma_j, \xi_j), \text{ i.e. } G_j(x) = \exp\left\{ -\left[1 + \xi_j \left(\frac{x - \mu_j}{\sigma_j} \right) \right]^{1/\xi_j} \right\} \\ Y &= (Y_1, ..., Y_J) \sim MVN(0, \Sigma(\psi)), \quad \psi = (\phi_1, \phi_2, \theta) \\ \Sigma(\psi) &= \exp(-\tilde{d}) \\ \tilde{d} &= \|\tilde{s}_j - \tilde{s}_k\| \\ \tilde{s}_j &= (\tilde{s}_{j1}, \tilde{s}_{j2}) = (s_{j1}, s_{j2}) \begin{pmatrix} cos(\theta) & -sin(\theta) \\ sin(\theta) & cos(\theta) \end{pmatrix} \begin{pmatrix} 1/\phi_1 & 0 \\ 0 & 1/\phi_2 \end{pmatrix} \end{aligned}$$

A synthetic footprint is therefore produced by unconditional simulation of a Gaussian random field at all locations in the domain based on the exponential anisotropic covariance model with parameters ϕ_1, ϕ_2 and θ. At each location, this Gaussian random field is then transformed from standard Gaussian to left truncated GEV using the marginal distribution fitted at each location, s_j, based on parameters μ_j, σ_j and ξ_j and assuming $\hat{F}_j(v_j) = 0.6$ because the 60% quantile of wind gust speeds at each

location is used as the left truncation threshold when fitting the marginals (see Section 4.3).

Using this statistical framework allows for fast simulation of synthetic footprints. One thousand footprints are simulated in 3035 seconds, i.e. each footprint takes approximately 3 seconds to simulate. As discussed in Section 2.2.1, catastrophe modellers often use numerical weather and climate models to simulate artificial windstorm footprints. Simulating footprints using numerical models is very computationally expensive, requiring the use of a super computer. The statistical footprint simulation model presented above is therefore a much quicker and computationally cheaper way of producing synthetic footprints.

5.4.4. Dependence model validation

Figure 5.15 shows a comparison of the best fit anisotropic exponential covariance function and the empirical covariance/correlation for London, Berlin and Bratislava paired with all other locations over land. Bratislava is included rather than Amsterdam to demonstrate the variation in the ability of the covariance function in representing the empirical dependence.

Figure 5.15 shows that for these particular locations the dependence structure is, in general, well represented by the anisotropic exponential model, however in the model, the dependence decays too quickly as separation distance increases, particularly in the longitudinal direction, for London and Berlin, and too slowly for the more south-eastern location, Bratislava. Similar comparisons for other locations in the domain show a variety of consistencies with the exponential covariance model. This is a very simplistic way of modelling the spatial dependence structure but, in general the visual comparison is good, particularly in the north-west of Europe, where the most extreme wind gust speeds and damage occur.

To further validate how well this covariance function, which is the same throughout the domain, can capture the empirical dependence between locations, the extremal dependence measures $\chi(u)$ and $\bar{\chi}(u)$, for quantile thresholds $u \in (0.6, 1)$ (Eqns (4.12)) are calculated for London paired with Amsterdam, Berlin and Paris based on 6000 simulated footprints and compared to the equivalent using the bivariate Gaussian copula model from Chapter 4 and the empirically calculated values, as in Figure 4.12. Figure 5.16 shows this comparison and indicates that, for these 3 pairs of locations, the covariance function slightly underestimates the dependence

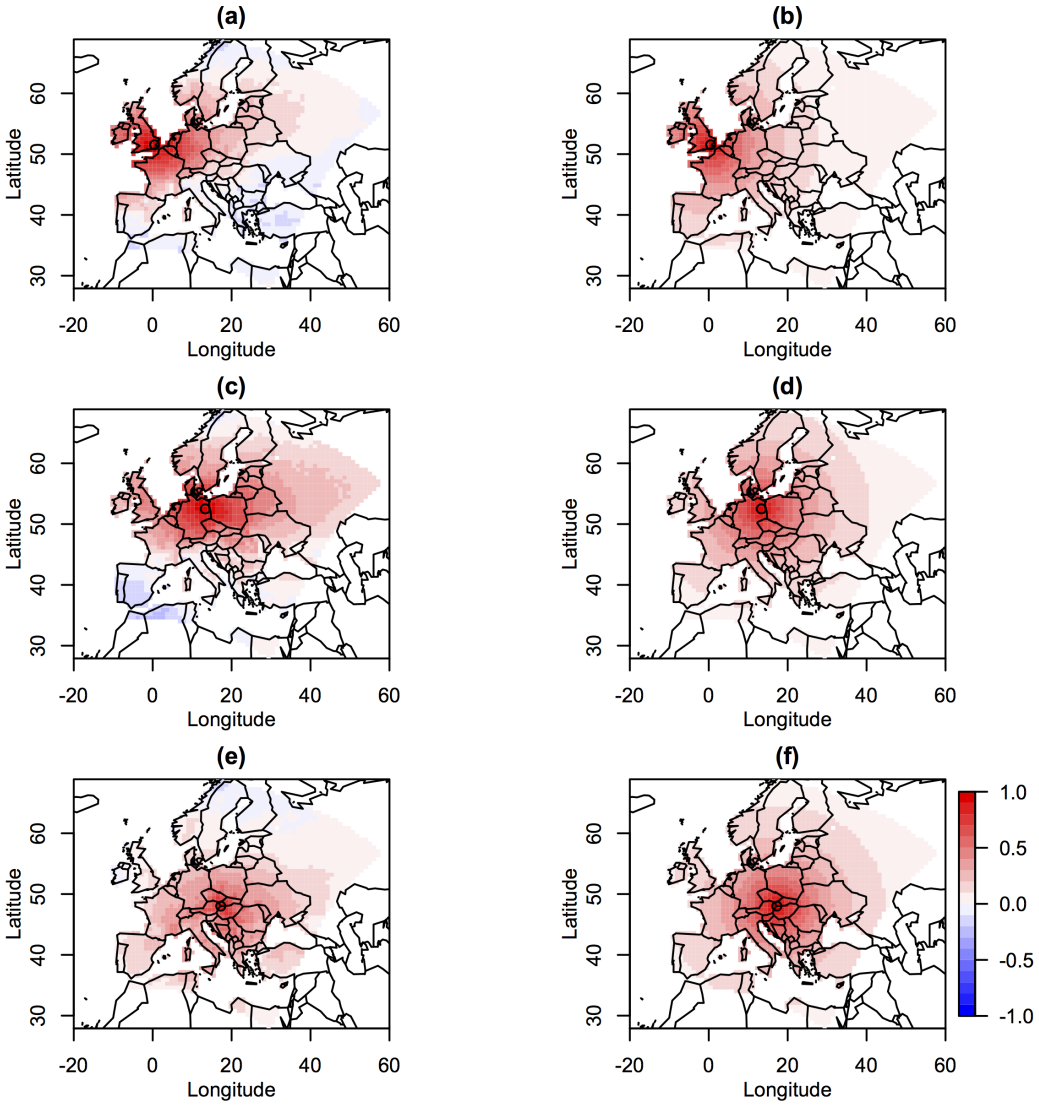

Figure 5.15. A comparison of the empirical covariance (first column) and best fit anisotropic exponential covariance function (second column) centred on (a)/(b) London, (c)/(d) Berlin and (e)/(f) Bratislava, included rather than Amsterdam to demonstrate the variation in the ability of the covariance function to represent the empirical dependence.

between wind gust speeds in all 3 cases, reflecting the conclusions of Figure 5.15. The general fit to the empirical extremal dependence is, however, good considering that the covariance model has been fit to the whole domain rather than a particular pair of locations as was the case for the bivariate Gaussian copula model.

The anisotropic exponential dependence structure can also be tested using a cross validation method. The locations at which the footprint data occur are divided into a training set and a validation set. The model is fit to the footprint data at the training locations and the resulting model is used to predict the wind gust speed at the validation locations for individual footprints by an optimal prediction method called simple kriging (Diggle and Ribeiro, 2007). Simple kriging aims to

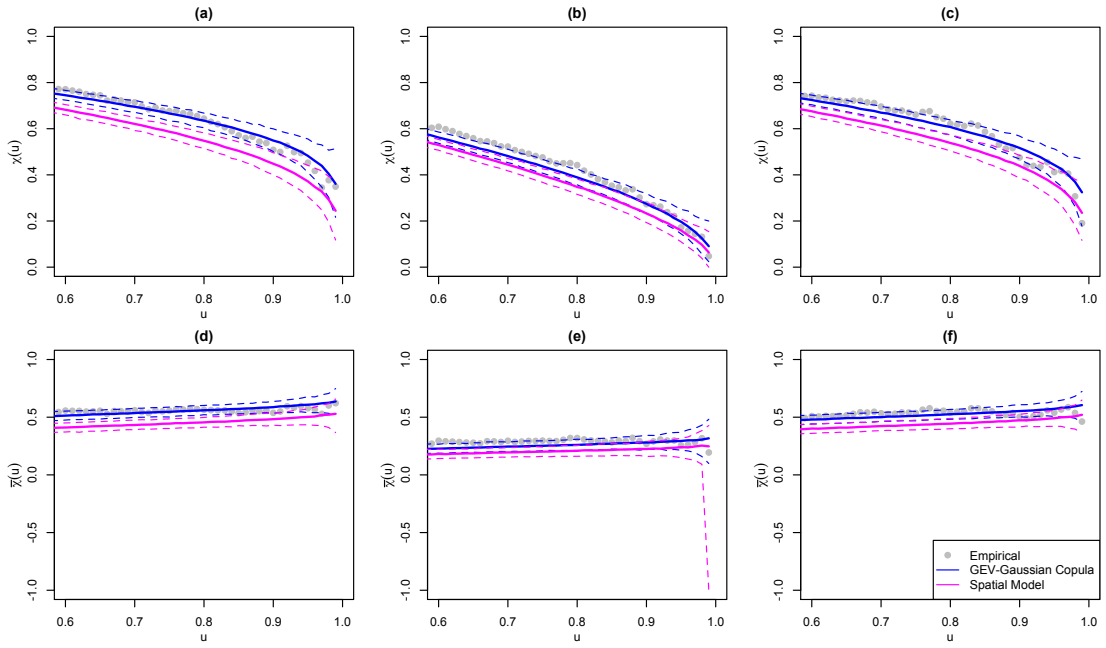

Figure 5.16. Extremal dependence measures evaluated empirically, using the Gaussian bivariate copula model and the anisotropic exponential covariance model. The measure $\chi(u)$ is shown in the top row (a)-(c) and $\bar{\chi}(u)$ in the bottom row (d)-(f), with 95% parametric bootstrap confidence intervals (see Appendix A.1).

find the value of the random field at unobserved location s_0, $Z(s_0)$, given the observations of the random field and the covariance model. The minimum mean squared error predictor at s_0, $\hat{Z}(s_0)$ is the function of the data at k training locations, $Y = (Y_1, ..., Y_k)$ which minimises the quantity $E[\{\hat{Z}(s_0) - Z(s_0)\}^2]$ and therefore $\hat{Z}(s_0) = E[Z(s_0)|Y] = \sum_{j=1}^{k} w_j(s_0)(Y_j)$, a linear function of the observations, Y_j, and the covariances between each training location and the unobserved location s_0, $w_j(s_0)$ (Diggle and Ribeiro, 2007). The predicted and true wind gusts speeds are compared to investigate how well the model dependence structure represents the data.

This comparison is shown in Figure 5.17. The wind gust speeds at the validation locations are, on the whole, very well predicted by the model, particularly for extreme wind gust speeds. The error in the prediction is greatest around Longitude 0 and 20, where most of the validation locations are either near to a coast or the end of the domain and the prediction will therefore be based on fewer surrounding locations. The error is greater for locations in lower latitudes. This is because the dependence structure in the south of Europe is more variable and less consistent with the best fit covariance model which better represents the north-west of Europe.

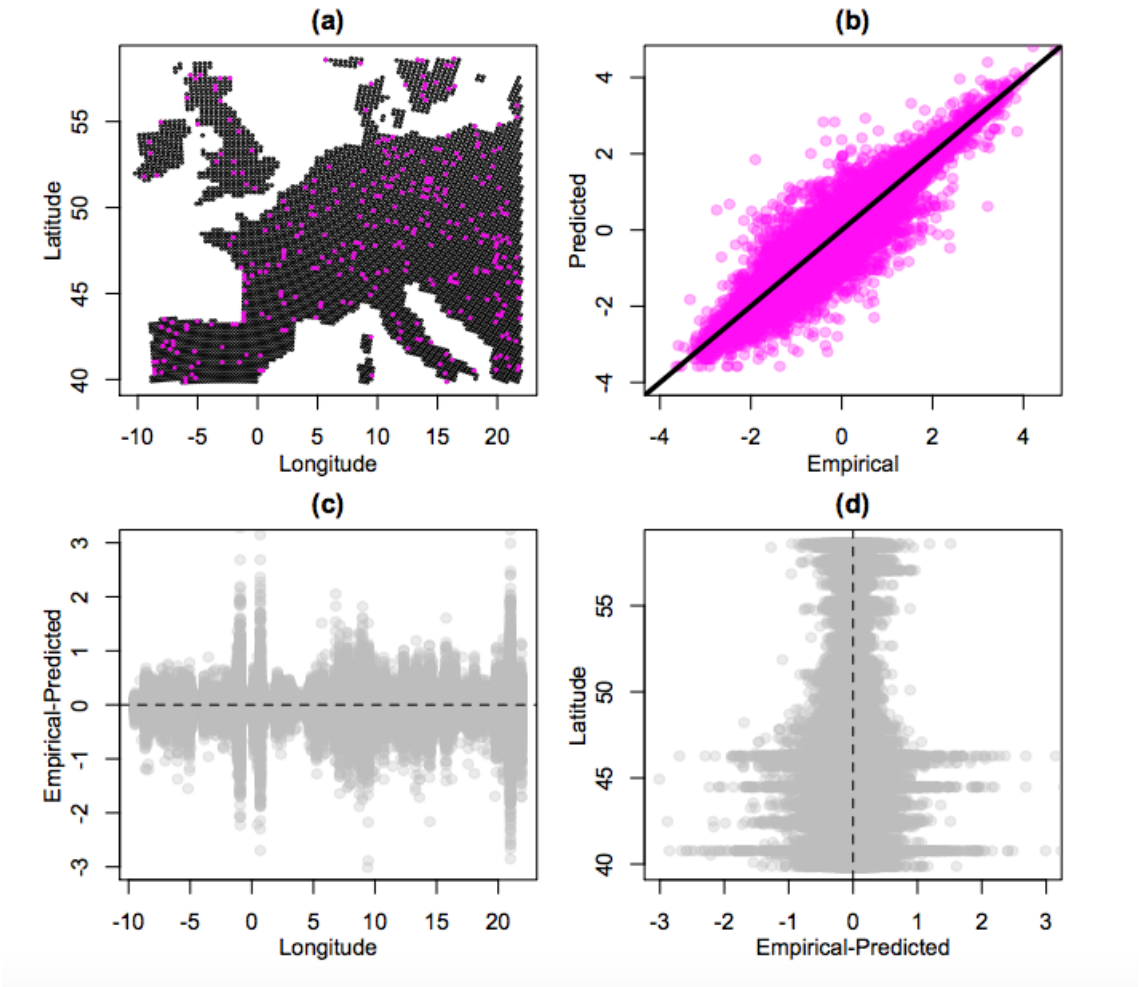

Figure 5.17. (a) The position of the validation locations, (b) scatter plot of the predicted and true Gaussian transformed wind gust speeds at the validation locations for 500 randomly selected footprints, (c) scatter plot of (Empirical-Predicted) against longitude and (d) against latitude.

5.5. Geostatistical windstorm footprint model validation

As an initial visual validation of the ability for the model to produce realistic footprints, footprints simulated from the spatial model are compared to randomly selected historical footprints from the data set of 5730, presented in Figures 5.18 and 5.19. Visual observation of these two sets of footprints shows that the synthetic footprints are hard to differentiate from the historical. They have similar smoothness, peak intensities, spatial dependence scales, and variation in location and intensity.

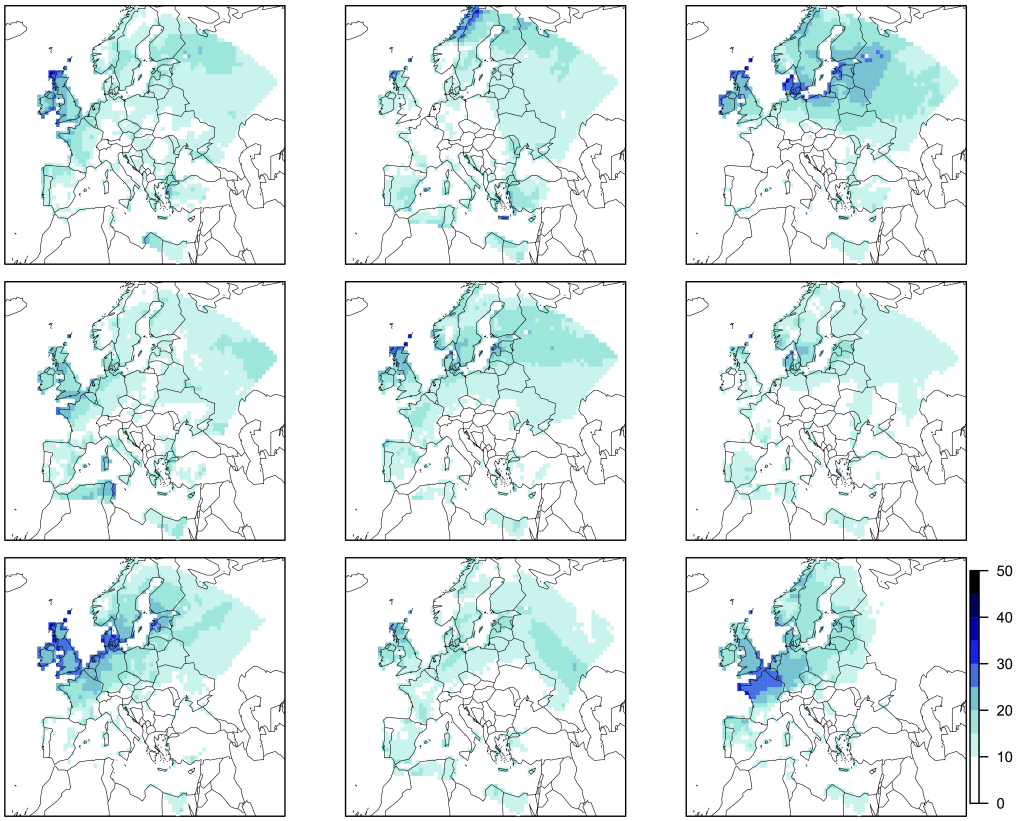

Figure 5.18. Nine randomly selected historical footprints from the data set of 5730.

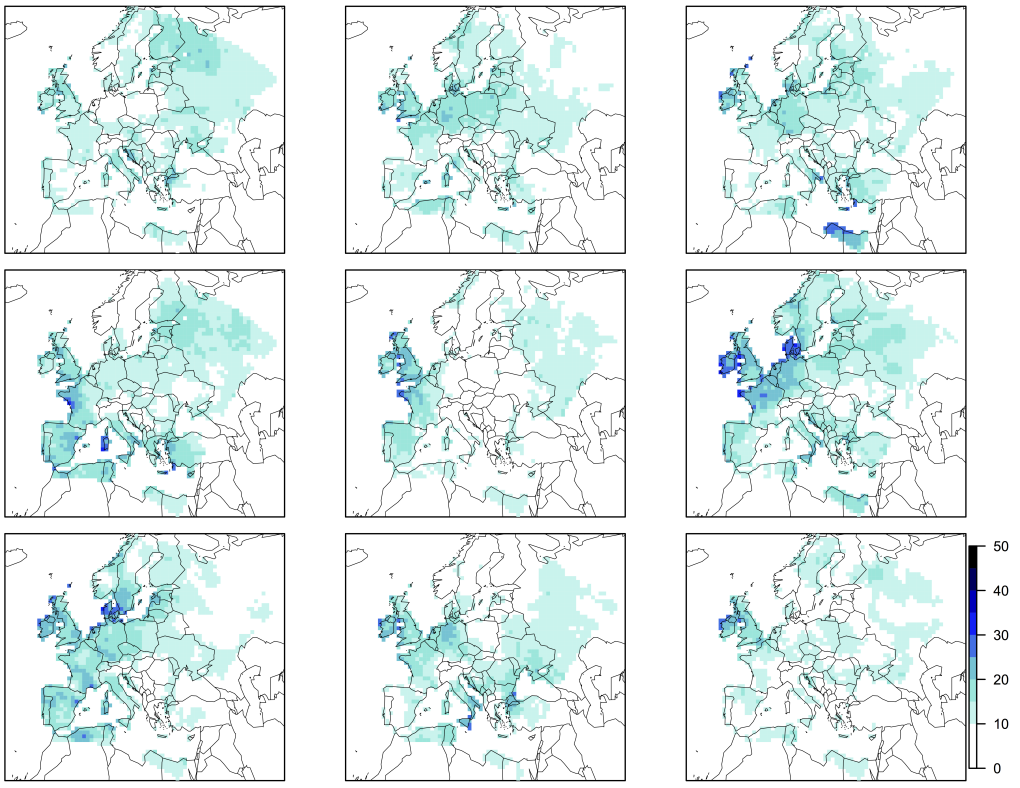

Figure 5.19. Nine synthetic footprints simulated using the statistical spatial model with an anisotropic exponential covariance structure throughout the domain and left truncated GEV margins at each location.

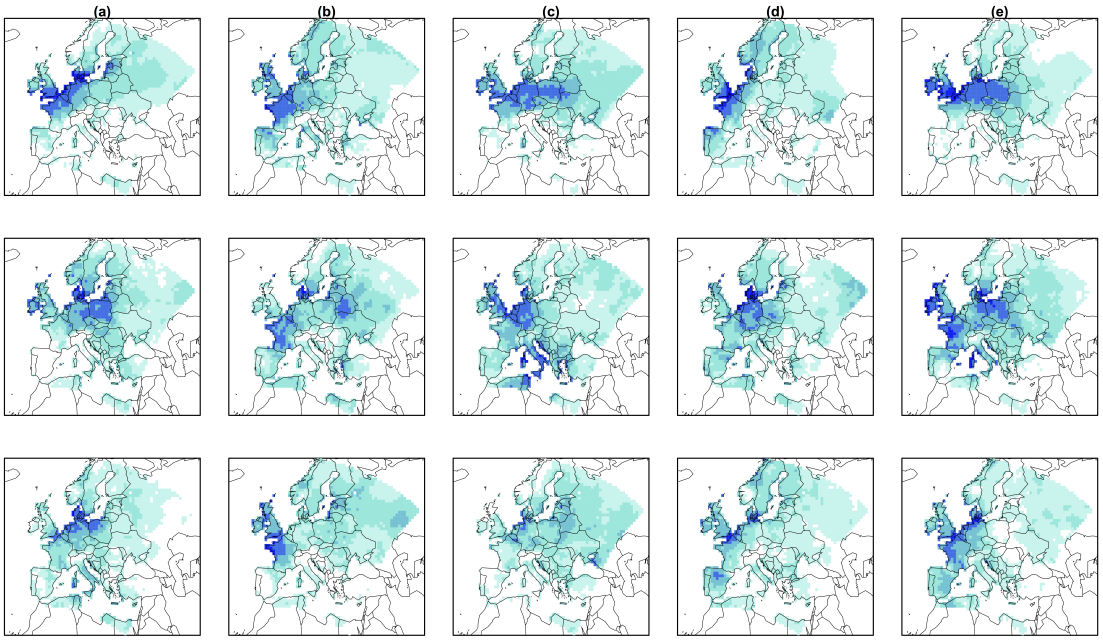

Figure 5.20. Historical footprints (top row) and synthetic footprints with the same area over land exceeding 25ms^{-1} (middle row) and the 98th percentile of the climatology wind gust speed (bottom row) for windstorms (a) Daria, (b) Lothar, (c) Kyrill, (d) the Great Storm of '87 and (e) Jeanette. Colour legend as in Figure 5.18.

Another interesting check is to ensure that the statistical model is able to simulate windstorm footprints as extreme as those experienced in the historical data set. Figure 5.20 shows the footprints for five of the most extreme historical windstorms; Daria, Lothar, Kyrill, the Great Storm of '87 and Jeanette, and synthetic footprints taken from a set of 500,000 simulated footprints which have the same area over land exceeding the loss thresholds of 25ms^{-1} and the 98$^{\text{th}}$ percentile of the climatology (October-March 1979-2012) wind gust speed. The simulated footprints in Figure 5.20 resemble the historical extreme footprints.

These visual validation results are very promising but the simulated and historical footprints can be compared in more detail by calculating particular characteristics. In Chapter 3 of this thesis, it was shown that the area of the footprint over land that exceeds 25ms^{-1} and the maximum 925hPa wind speed over land were the best individual predictors of insured loss. Therefore the area, A, and peak intensity, represented by the 90$^{\text{th}}$ percentile of the wind gusts speed for a given footprint, x_{90} (because the simulation model does not produce 925hPa wind speeds) are compared for historical and synthetic footprints, shown in Figure 5.21. Figure 5.21 shows that the spatial model has very similar relationship between these two storm severity measures compared to when calculated empirically from the historical footprints. In particular, the strong extremal dependence between the measures is captured by

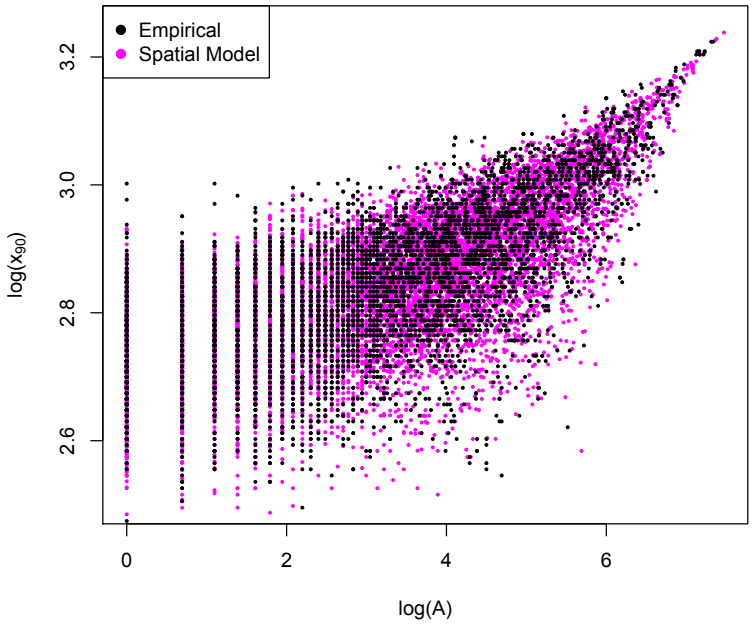

Figure 5.21. Scatter plot of the logarithm of the area of the footprint exceeding 25ms^{-1} and the logarithm of the 90$^{\text{th}}$ percentile of the footprint wind gusts speeds for simulated footprints (purple) and historical footprints (black).

the simulated footprints.

As in Chapter 4, the model must also be validated to ensure it can realistically represent joint losses, and is therefore fit for purpose for the sensitivity analysis experiment in the final chapter. The same two conceptual loss functions, L_{25} and L_{98}, defined in Section 4.5, are used to validate this. These conceptual loss functions characterise the area of the footprint that exceed the corresponding damage threshold. Therefore, L_{25} is equivalent to A in Figure 5.21. Rather than being a bivariate loss distribution where loss equals either $0, 1$ or 2, depending on the exceedance of a loss threshold, the loss associated with a windstorm footprint is now multivariate and can range from 0 to 14872 (the number of locations/grid cells over land) that is

$$L(X_1, ..., X_J) = H(X_1 - t_1) + \cdots + H(X_J - t_J)$$

where t_1,..., t_J are high loss thresholds at each location and $H(n)$ as a Heaviside or indicator function $H(n) = 1$ if $n > 0$ and $H(n) = 0$ otherwise, as in Section 4.5.

Figure 5.22 shows a comparison of the loss distributions for the 5730 historical footprints and 500,000 synthetic footprints, where the losses, L_{25} and L_{98}, are shown on a logarithmic scale to reduce skewness. For both conceptual loss functions the

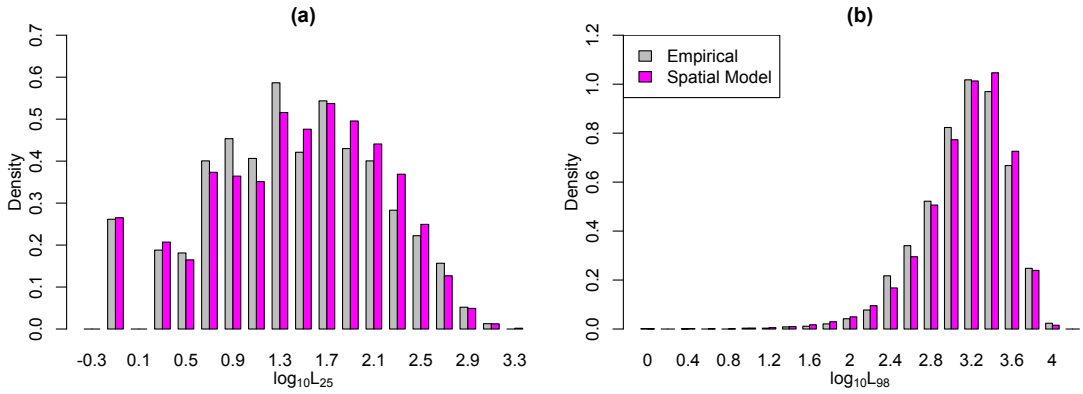

Figure 5.22. The distribution of log base 10 losses according to multivariate conceptual loss functions L_{25} and L_{98}.

synthetic footprints have a slightly higher proportion of footprints causing medium and higher losses, however the general shapes of the loss distributions are very similar to the empirical, showing that the model is able to simulate footprints with realistic joint losses.

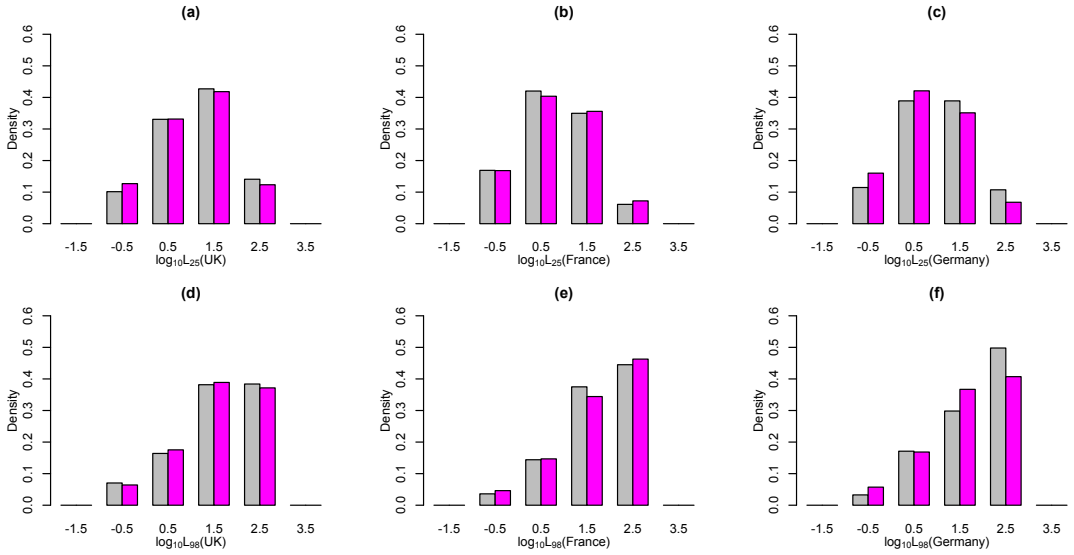

Figure 5.23. The distribution of log base 10 losses according to multivariate conceptual loss functions L_{25} and L_{98} for (a)/(d) the UK, (b)/(e) France and (c)/(f) Germany.

It is also important that the losses and the relationship between losses in different regions of Europe are well represented by the model. This is because re/insurers can diversify their portfolio of clients, ensuring that they insure locations that experience losses at different times.

As in Figure 5.22, Figure 5.23 shows a comparison of the loss distribution for the 5730 historical footprints and 500,000 synthetic footprints for the UK, France and

Germany, where again the losses, L_{25} and L_{98}, are shown on a base 10 logarithmic scale. Some small variations between the historical and simulated footprints exist, for example, in Germany a smaller proportion of higher L_{98} losses occur in the synthetic footprints compared to historical, while the opposite is true for France. However, again, in general the distributions are very similar.

5.6. Conclusions

This chapter has presented the specification, fitting and validation of a spatial model for the simulation of synthetic windstorm footprints. A geostatistical approach is taken, which uses a function of the separation vector between locations to model the spatial dependence throughout the domain. This approach is used for simplicity and to reduce the computational expense of simulation compared to a multivariate copula approach due to the large number of locations in the data set. Using this approach also ensures asymptotic extremal independence between locations, as was identified in Chapter 4.

The Matérn covariance model was initially fitted to the spatial data because the extra shape parameter allows the data to specify the mean-squared differentiability of the process rather than being predetermined by the covariance model choice. Problems arose initially, when fitting the isotropic Matérn covariance function using the ML approach. A ridge in the likelihood surface identified a number of possible parameter combinations with very similar likelihoods. It was found that the data provided no information for distinguishing between parameter combinations along this ridge in the likelihood because a relatively large weight was placed on small separation distances in the likelihood function due to the large number of locations in the data set.

The $\nu = 0.5$ Matérn covariance model, equivalent to the exponential covariance model, gave the best fit to the binned empirical covariogram over a large range of separation distances and was therefore selected over the other parameter combinations along the ridge. The anisotropic exponential covariance function was shown to be a better representation of the eastward transit of windstorms over Europe and will therefore be used rather than the isotopic model. Including anisotropy parameters within the windstorm footprint model will also allow for greater variation in the spatial dependence structure of simulated footprints in the sensitivity study in the next chapter.

The best fit anisotropic exponential covariance model was not able to capture the

dependence perfectly. However, the match with the empirical dependence structure was generally very good considering the model was fit to the whole domain and not individual pairs of locations.

This spatial model allows fast simulation of synthetic windstorm footprints. Firstly, a Gaussian random field with the specified anisotropic exponential covariance model is simulated. This random field is then transformed to left truncated GEV individually at each location using the marginal distributions and the empirical distribution of wind gust speeds at that location. Using this methodology, a windstorm footprint can be simulated in $\approx$3 seconds.

The ability for this statistical model to produce realistic footprints was validated, based on the 5730 historical footprints. The synthetic and historical footprints are visually extremely similar and the relationship between the storm severity measures, A and x_{90}, found in the historical footprints are very well represented by the simulation model. The multivariate loss distributions associated with the conceptual loss functions introduced in Chapter 4, which represent the area that exceeds a damage threshold, show that the synthetic footprints realistically represent joint losses throughout the domain. The same validation was done for different regions within Europe and showed that the model can also realistically represent losses and the relationship between losses in these different regions.

6. Sensitivity analysis of windstorm footprint characteristics and insured losses

6.1. Introduction

The geostatistical windstorm footprint model, presented in Chapter 5, will now be used in a sensitivity study to explore which of the model parameters, characterising local intensity and spatial dependence characteristics of a footprint, are most influential on insured loss. In Chapter 3, it was shown that area of the footprint that exceeds 25ms^{-1} over land was the best footprint storm severity measure at representing insured loss. This measure will therefore be used to approximate insured loss within this sensitivity study.

Bonazzi et al. (2012) is one of the very few studies that have investigated windstorm loss sensitivity. They estimated the loss associated with 135 historical events using the storm severity measure in Eqn (2.4), for the regions South-UK and West-Denmark. For this pair of regions, they fitted two bivariate Gumbel extreme value copula models to the estimated losses; one for events that occurred in a positive North Atlantic Oscillation (NAO) phase, and one for events that occurred in a negative NAO phase. They used combinations of the storm frequency, intensity and spatial dependence parameters from these two bivariate models to simulate losses, and investigated how these parameter variations influenced the loss return period curves. They found that the storm frequency parameter had most influence for shorter return periods, while the spatial dependence parameter had most influence for longer return periods.

This chapter will present a more thorough Sensitivity Analysis (SA) experiment, using a formal methodology recently developed as part of the CREDIBLE project, named PAWN (derived from authors names) (Pianosi and Wagener, 2015). Saltelli

et al. (2004) defined SA as a method for exploring how the variability in the output of a model can be apportioned to different sources of variability in the model input parameters, i.e. the relative influence of each input on the output. SA techniques are usually applied to deterministic models, where the output is completely determined by the input parameters. In other words, the same input produces the same output when the model is rerun. However, the statistical windstorm footprint model used in this investigation is stochastic: the same input will produce different outputs because of the inherent randomness due to the use of pseudo-random numbers (Saltelli et al., 2008). This extra source of intrinsic variation will therefore also be explored. This application of SA to a stochastic natural hazard model is novel with no examples found in previously published studies.

This chapter addresses the following questions:

- How does the stochasticity of the windstorm footprint model affect the SA experiment?
- Are spatial dependence structure characteristics or the local intensity characteristics of a footprint most influential on insured loss?

Firstly, the methodology used within this SA experiment will be introduced. The above questions will then be addressed and the most influential footprint characteristics identified.

6.2. Methodology

The parameters of the windstorm footprint model, presented in Chapter 5, can be thought of in terms of those explaining the local distribution of wind gust speeds at each location: the left truncated GEV location, μ, scale, σ and shape, ξ, and those explaining the spatial dependence structure: the angle of rotation, θ, the longitudinal correlation scale, ϕ_1, and the latitudinal correlation scale, ϕ_2. The windstorm footprint model $f(\cdot)$ uses a set of these input parameters to simulate a synthetic footprint X:

$$f(\mu, \sigma, \xi, \theta, \phi_1, \phi_2) \text{“} = \text{”} X = (X_1, X_2, ..., X_J)$$

where $\mu = (\mu_1, ..., \mu_J)$, $\sigma = (\sigma_1, ..., \sigma_J)$, $\xi = (\xi_1, ..., \xi_J)$ and J is the number of simulation locations. Here, the same $J = 14872$ land locations as in the historical footprint data set are used for simulation (see Section 2.3.1).

This chapter aims to investigate which input parameters have most relative in-

fluence on insured loss. This loss will be approximated by the area of the simulated footprint exceeding 25ms^{-1} over land, shown to be the best footprint based storm severity measure at classifying extreme insurance loss windstorms (denoted A in Chapter 3), equivalent to the conceptual loss function L_{25}.

$$loss(X) \approx L_{25}(X) = \sum_{j=1}^{J} H(X_j - 25)$$

where $H(n)$ is a Heaviside function, i.e. $H(n) = 1$ if $n > 0$ and $H(n) = 0$ if $n \leq 0$.

The input parameters are varied and their relative influence on the loss is then quantified. As previously discussed, $f(\cdot)$ is a stochastic model, meaning that the same set of input parameters will produce different simulated footprints and resulting losses when the model is rerun. Therefore, the loss associated with each set of input parameters used in the SA experiment is calculated as a summary statistic of the losses from a sample of simulated footprints. The insurance industry is interested in loss summary statistics such as Average Annual Loss (AAL), i.e. the expected loss for a given year, and Maximum Annual Loss (MAL), i.e. the most extreme loss for a given year:

$$AAL(\mu, \sigma, \xi, \theta, \phi_1, \phi_2) = \frac{1}{M} \sum_{m=1}^{M} loss(X_m)$$
$$MAL(\mu, \sigma, \xi, \theta, \phi_1, \phi_2) = \max_{m=1,\ldots,M} loss(X_m)$$

where M is the number of windstorm events per year. The relative influence of the input parameters on insured loss will therefore be explored using AAL and MAL. Including both AAL and MAL will provide an interesting comparison of the influence of the parameters on different loss summaries. The historical footprint data set used in this thesis consists of 5730 events which occurred over 33 winters, therefore $M = 174 \approx 5730/33$ simulated footprints are used to represent a simulated year. The values of outputs AAL and MAL may be sensitive to the sample of 174 footprints, or event set, used to represent a year. This source of variation will be investigated and further discussed in Section 6.3.

Figure 6.1 illustrates the concept of SA by comparing how varying μ and ξ in the footprint simulation model influences AAL. This scatter plot shows how μ has more influence on AAL than ξ because AAL increases with μ but varies in magnitude over the range of ξ.

The application of SA in the environmental modeling community has increased

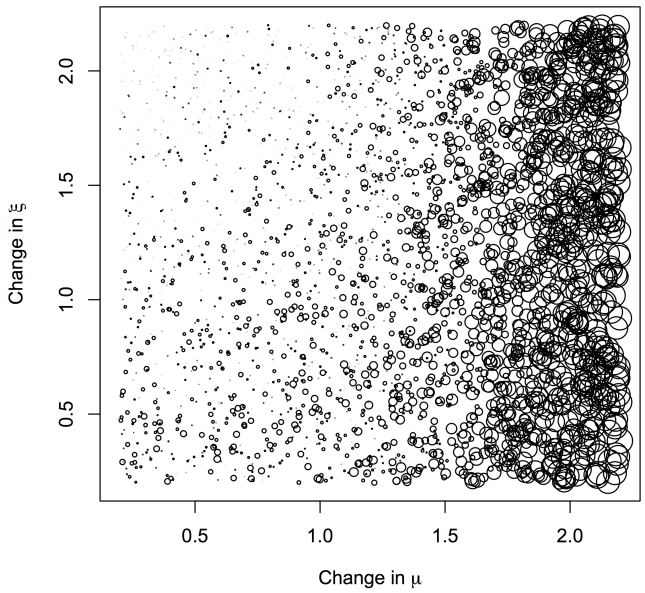

Figure 6.1. Scatter plot of the change in μ and ξ and the resulting value of AAL, represented by the diameter of the circle symbols. A change of 0.5, for example, means the parameter is reduced by 50%.

significantly in recent years due to an increasing awareness of the importance of uncertainty assessment in model-based decision-making, and the accessibility of sufficient computing resources required for computationally demanding SA methods (Pianosi et al., 2014).

The type of approach, level of complexity and purpose of SA can be very different depending on the specific application. Systematic reviews of existing SA methodologies can be found in Pianosi et al. (2014), Norton (2015) and Saltelli et al. (2008). SA methods can be classified as either "local" or "global". Local methods focus on a region of the input parameters around a point of interest, and hence assess the model output sensitivity to small variations in the input parameters. Global methods aim at assessing the overall output sensitivity across the whole credible range of the input parameters (Norton, 2015). As explained in Pianosi et al. (2014), global SA is used for a range of very diverse purposes, including: to support model calibration, verification, diagnostic evaluation or simplification e.g. Wagener et al. (2001); to prioritise efforts for uncertainty reduction e.g. Hamm et al. (2006); to analyse the dominant controls of a system e.g. Pastres et al. (1999); and to support robust decision-making e.g. Singh et al. (2014).

SA methods also use different sampling approaches to estimate the sensitivity indices: either One-At-the-Time (OAT) or All-At-the-Time (AAT). In OAT, input parameters are varied one at a time while keeping all others fixed; in AAT, all in-

put factors are let vary simultaneously. AAT methods are therefore able to capture interactions between input parameters. Anderson et al. (2014) demonstrated the usefulness of global AAT SA methods in the area of integrated assessment modeling for climate change economics. They noted that global AAT SA methods allow for significant interactions to be identified explicitly, rather than simply acknowledged or speculated upon, and the direction of the interaction effect can be observed. Pianosi et al. (2014), however, identified that AAT methods require more extensive sampling and hence a higher number of model evaluations to provide robust results compared to OAT methods.

Different SA methods are better suited to address different purposes. Saltelli et al. (2008) specified 3 key purposes for SA:

- Ranking (or Factor Prioritisation) which aims to rank the model inputs according to their relative contribution to the output variability;
- Screening (or Factor Fixing) which aims to identify which (if any) of the inputs have no influence on the output;
- Mapping which aims to determine the region of the input space that produces "significant", for example extreme, model output values.

The primary purpose here is to rank the input parameters, in terms of their relative influence on the output AAL and MAL, while accounting for possible interactions between the input parameters. The global AAT SA method known as PAWN (Pianosi and Wagener, 2015) is used within this investigation.

6.2.1. PAWN

PAWN is a density-based SA methodology. This means that the sensitivity of the model output is investigated in terms of the whole output sample Probability Density Function (PDF) created when evaluating the model at various input parameter settings, as appose to variance based approaches which only use moments of the output sample. This therefore also allows for the selection of a particular part of the output sample to be analysed, e.g. the upper tail. Density based SA methods are not commonly used in the literature because they require the estimation of a large number of empirical PDFs if multiple input values are to be explored, requiring a great deal of computational expense and the arbitrary selection of kernel density estimation bandwidths or histogram bin widths (Pianosi and Wagener, 2015).

The recently developed PAWN methodology characterises the model output sample by its Cumulative Distribution Function (CDF) rather than its PDF. The advantage

of this is that the empirical CDF (i.e. sorted output values) comes at little computational cost and does not require any density smoothing parameters (Pianosi and Wagener, 2015). The methodology is therefore very easy to implement and Pianosi and Wagener (2015) also showed how the intermediate results generated in the PAWN implementation procedure can be effectively visualised to gain further insights into the output behaviour.

The PAWN methodology characterises the sensitivity of the model output y to a given input factor x_i, as the distance between the unconditional empirical CDF of output y across all input values, $\widehat{F}_y(y)$, and the conditional empirical CDF of y when x_i is fixed and all other input parameters are allowed to vary, $\widehat{F}_{y|x_i}(y)$.

The unconditional empirical CDF, $\widehat{F}_y(y)$ is calculated by evaluating the model at a large number, N_u, of sets of input parameters. The conditional empirical CDF, $\widehat{F}_{y|x_i}(y)$, is calculated by initially sampling n conditioning values of x_i. Then for each conditioning value of x_i, N_c samples of the remaining input parameters $x_{\sim i}$ are used to evaluate the model. The method used to sample the input parameter space is discussed in Section 6.2.2.

This method of calculating the conditional empirical CDF requires $n \times N_c$ model evaluations in addition to the N_u model evaluations used to calculate the unconditional empirical CDF. To save computational expense, Pianosi and Wagener (2015) discuss an alternative approach for calculating the conditional empirical CDF which uses the original sample of N_u input and output values used to calculate the unconditional empirical CDF. For this alternative method, which will be used in this investigation, the conditional empirical CDF for a given input parameter x_i is calculated based on conditioning on 'similar' values of x_i rather than a fixed value of x_i. That is, $\widehat{F}_{y|x_i}(y|x_i \in \mathscr{I}_k)$, where $\mathscr{I}_k$ $(k = 1, ..., n)$ are n (e.g. 10) equally spaced intervals over the range of variation of x_i, equivalent to conditioning on n values of x_i.

Pianosi and Wagener (2015) noted that applying PAWN using this methodology for calculating the conditional empirical CDFs rather than resampling and conditioning on fixed values of x_i is sub-optimal, and discussed their intention to carry out further research to investigate the loss in accuracy.

For a given input parameter, x_i, and conditioning interval $\mathscr{I}_k$, the distance between the unconditional and conditional empirical CDF is quantified in terms of the Kolmogorov-Smirnov (KS) statistic:

$$\widehat{KS}(x_i \in \mathscr{I}_k) = \max_y \left| \widehat{F}_y(y) - \widehat{F}_{y|x_i}(y|x_i \in \mathscr{I}_k) \right|$$

and the resulting relative influence of input parameter x_i on the output y, known as the sensitivity index T_i, is a summary statistic of $\widehat{KS}(x_i \in \mathscr{I}_k)$, calculated over the $k = 1, 2, ..., n$ conditioning intervals. The median was used by Pianosi and Wagener (2015) and will be used in this investigation.

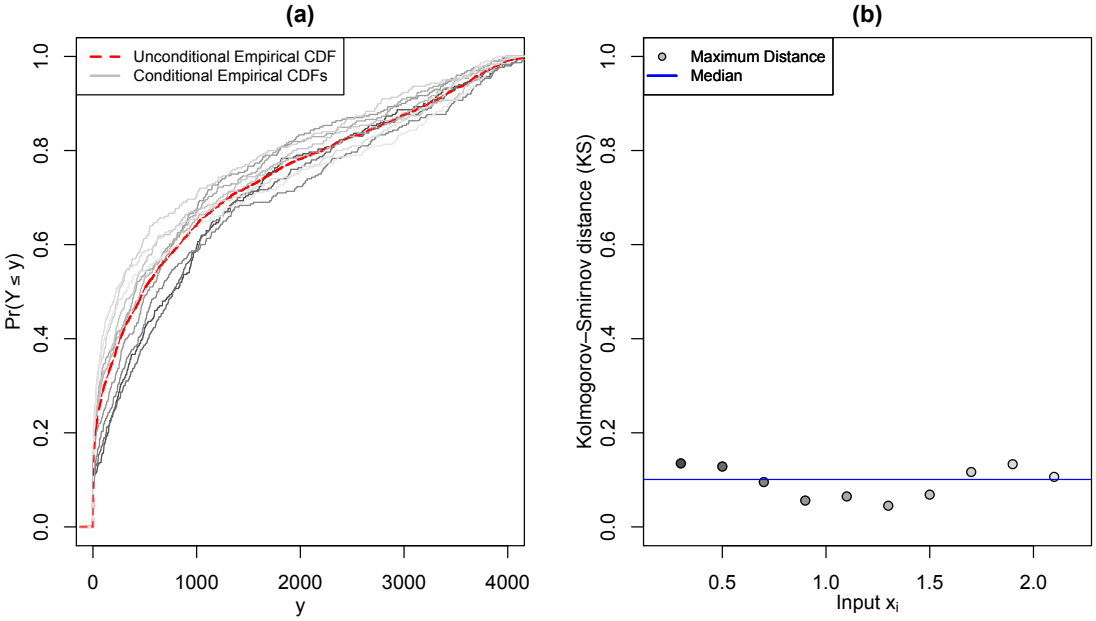

Figure 6.2. (a) The unconditional empirical CDF (red) and conditional empirical CDFs for $n = 10$ conditioning intervals, based on the input parameter x_i=σ and output AAL for input parameter range $0.2 - 2.2$ relative change (see Section 6.2.2), (b) the Kolmogorov-Smirnoff (KS) distance between the unconditional empirical CDF and each conditional empirical CDF (the grey scale corresponds to the conditional empirical CDFs in (a)). The blue line shows the median of these 10 KS distances, representing the sensitivity index, T_i, for that input parameter.

Figure 6.2 (a) presents how these CDFs appear graphically for one input parameter of the footprint model, $x_i = \sigma$, based on the output AAL. The different conditional empirical CDF curves each represent one of the $n = 10$ conditioning intervals of x_i. The KS distances between the unconditional empirical CDF and each conditional empirical CDF are presented in Figure 6.2 (b), and the median of these 10 KS distances (blue line) represents the sensitivity index T_i, quantifying the relative influence of σ on AAL.

Sensitivity indices T_i, $i = 1, ..., r$ do not depend on the numerical value of the input parameters. They vary between 0 and 1, with a lower value indicating less relative influence, allowing for a direct comparison between input parameters.

This PAWN analysis is repeated multiple times for N_b bootstrap samples of the

N_u input and output samples. The mean and 95% bootstrap confidence intervals are used to quantify the input parameter space sampling uncertainty of the sensitivity index.

A schematic for this workflow is shown in Figure 6.3.

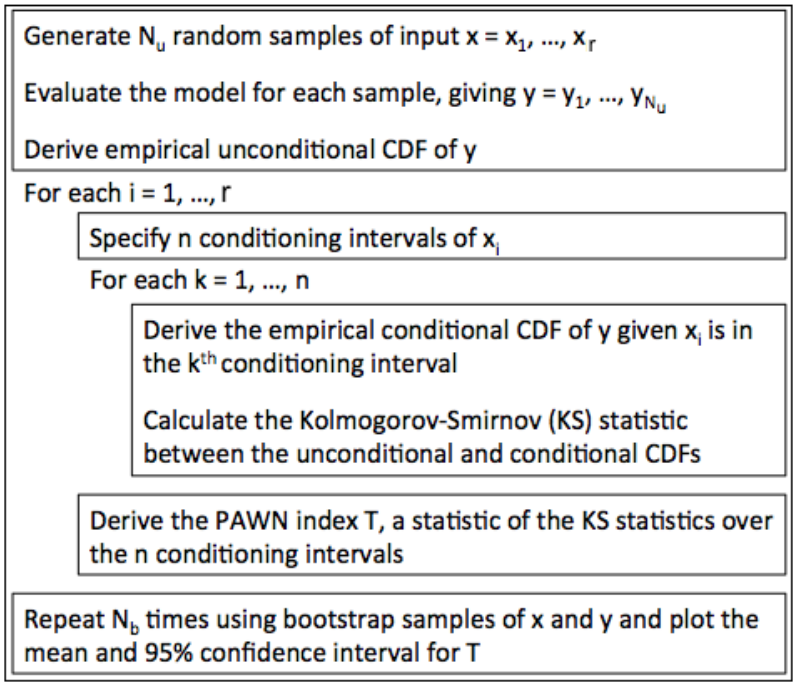

Figure 6.3. The steps in the numerical implementation of the PAWN analysis.

The PAWN indices can also be calculated over a specific sub-range of the output sample, e.g. the upper tail $y > t$, where t is a predefined threshold:

$$\widehat{KS}(x_i \in \mathscr{I}_k) = \max_{y>t} \left| \widehat{F}_y(y) - \widehat{F}_{y|x_i}(y|x_i \in \mathscr{I}_k) \right| \tag{6.1}$$

This calculation requires no new sampling or CDF calculation and is therefore extremely easy and computationally cheap to apply, a particular benefit of this SA method. Within this investigation the relative influence of the input parameters on the upper tail of the output sample of AAL and MAL will also be explored because the most extreme values of these measures are of particular interest from a damage point of view.

6.2.2. Sampling the input parameter space

A Latin Hypercube (LHC) sampling scheme is used to create the N_u sets of input parameters used in the PAWN SA experiment. LHC sampling is a stratified random sampling technique which ensures all portions of each input parameter is represented in the sample. As in the LHC sampling method explained by McKay et al. (2000), the range of each input parameter is divided into N_u strata of equal marginal probability, $1/N_u$, and one value is sampled from each stratum. The N_u sampled values of each input parameter are then matched at random. The advantage of using the LHC method is that each input parameter range is represented in a fully stratified manner, giving more reliable SA results. An example of a LHC sample from a two-dimensional input parameter space, where $N_u = 3000$, is shown in Figure 6.4.

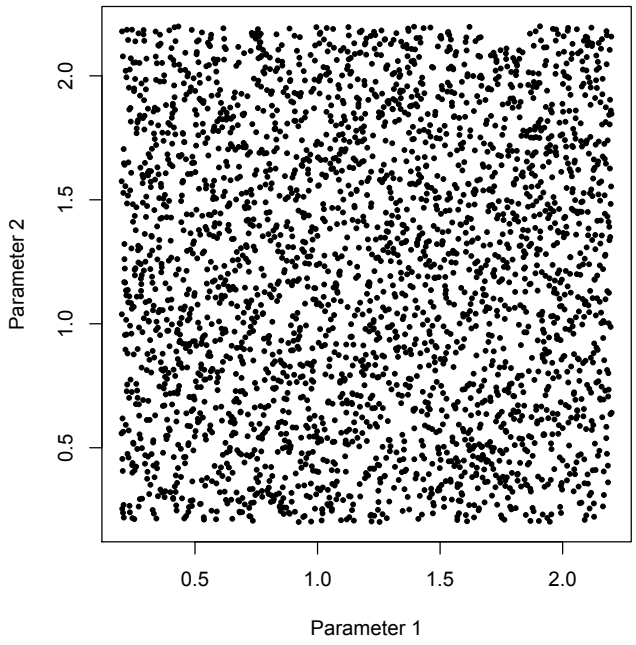

Figure 6.4. An example of a Latin Hypercube (LHC) sample for a two-dimensional input parameter space where $N_u = 3000$.

For this investigation, the input parameter space will be sampled by generating values that represent a relative change from the best fit parameters estimated in Chapter 5, denoted $p_\mu, p_\sigma, p_\xi, p_\theta, p_{\phi_1}$ and p_{ϕ_2}. Relative changes are used because they have a straightforward interpretation and are directly comparable between input parameters e.g. at different locations. The left truncated GEV parameters are different for each grid cell, however the same relative change will be applied at all grid cells, for example if $p_\mu = 0.9$ then all left truncated GEV location parame-

Table 6.1.. Input parameter ranges used to explore the sensitivity of input parameter ranking to this specification.

Input Parameter	p_μ	p_σ	p_ξ	p_θ	p_{ϕ_1}	p_{ϕ_2}
Parameter range 1	(0.2, 2.2)	(0.2, 2.2)	(0.2, 2.2)	(0.2, 2.2)	(0.2, 2.2)	(0.2, 2.2)
Parameter range 2	(0.5,1.5)	(0.5,1.5)	(0.5,1.5)	(0.5,1.5)	(0.5,1.5)	(0.5,1.5)
Parameter range 3	(0.7, 1.3)	(0.7, 1.3)	(0.7, 1.3)	(0.7, 1.3)	(0.7, 1.3)	(0.7, 1.3)
Parameter range 4	(0.5, 2)	(0.5, 2)	(-2.2, 2.2)	(-10, 10)	(0.2, 2.2)	(0.2, 2.2)

ters are decreased by 10%, throughout the domain. This means that only 6 input parameter values need to be generated for each of the N_u samples from the input parameter space.

The results of the SA experiments will depend on the range of the input parameters used because the PAWN sensitivity index, T, is calculated as a statistic over the range of the input parameter, the median (see Figure 6.2). For this reason, four separate PAWN SA experiments will be carried out, each with a different set of input parameter ranges, to explore the sensitivity of the results to this parameter range choice. Each experiment requires a new sample of N_u sets of input parameters and therefore N_u model evaluations to calculate the output AAL and MAL. The input parameter ranges used to test this sensitivity are shown in Table 6.1. Parameter range 1 represents relative ranges from an 80% decrease and 120% increase. This large interval is used to ensure that a wide range of possible model behaviour is captured. The variation in the model parameters based on the relative changes in Parameter range 1 is presented in Figure 6.5, showing that this range in the input parameters results in a large variation in both the local wind gust distribution and the footprint spatial dependence structure.

Parameter ranges 2 and 3 (Table 6.1) explore how the SA results change when the input parameter ranges are reduced to a 50% and 30% increase or decrease respectively. Parameter range 4 increases the range of input factors p_ξ and p_θ. The best fit values of ξ and θ are relatively small compared to the other parameters of the model, therefore this set of relative change parameter ranges is used to explore whether these two parameters have a greater relative influence on the model output when varied more.

The variation in ξ and θ based on Parameter range 4 is presented in Figure 6.6. The GEV shape parameter ξ varies between high positive values which will produce very heavy tailed GEV distributions with no upper limit, to low negative values which

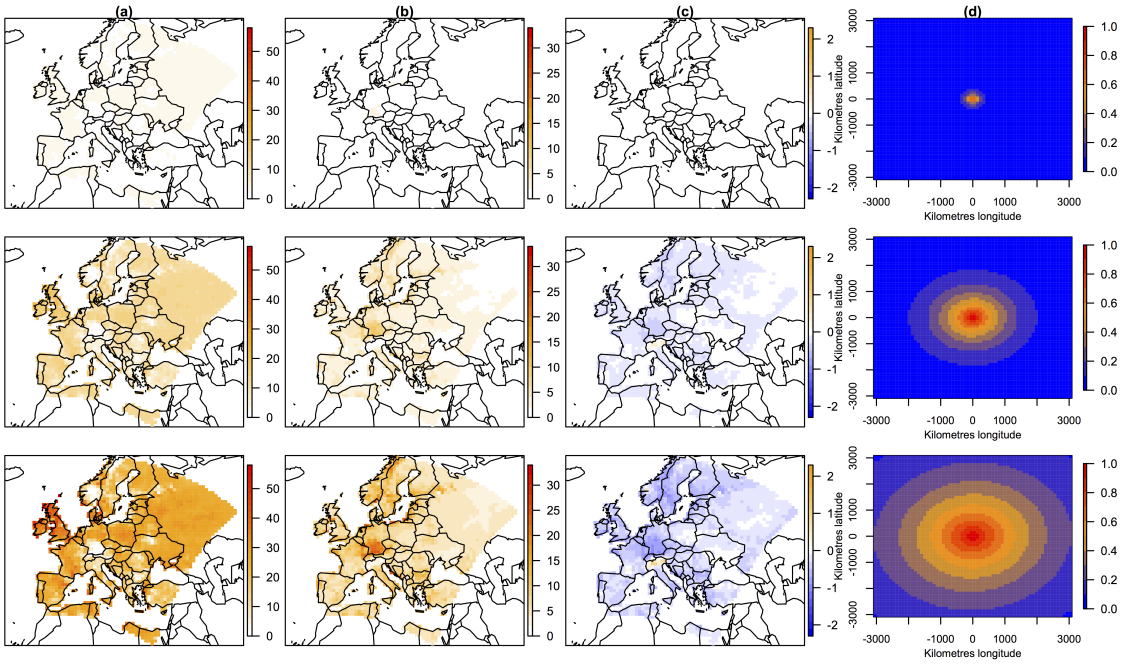

Figure 6.5. Variation in footprint model parameters when input factors (a) p_μ, (b) p_σ, (c) p_ξ and (d) p_θ, p_{ϕ_1} and p_{ϕ_2} apply a decrease of 80% (top row), apply no change (middle row) and apply an increase of 120% (bottom row). The colour scale in (d) represents the correlation between locations.

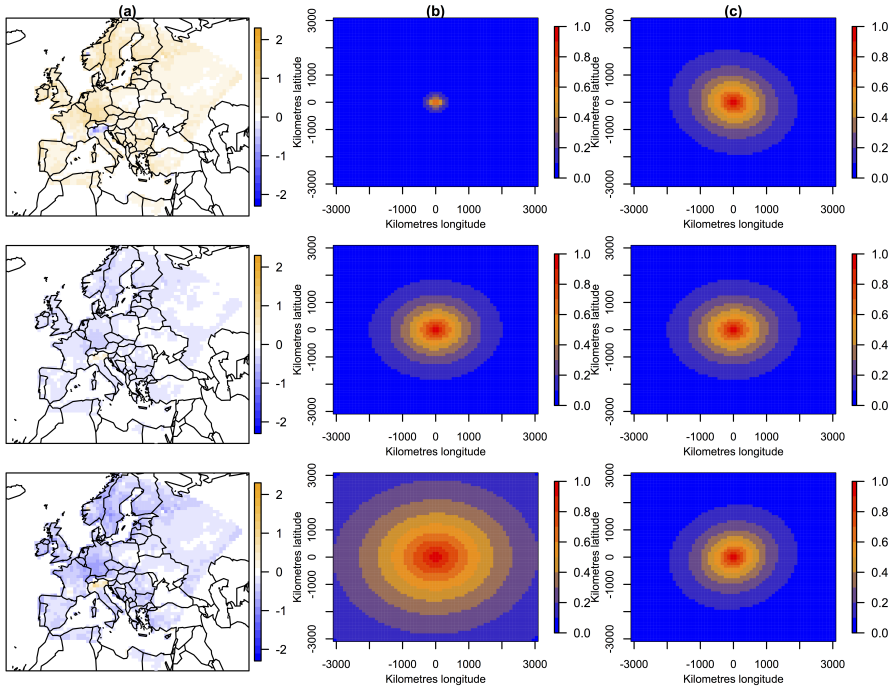

Figure 6.6. Variation in footprint model parameters when input parameters (a) p_ξ applies a change of -120% (top row), applies no change (middle row) and applies an increase of 90% (bottom row), (b) p_{ϕ_1} and p_{ϕ_2} apply a decrease of 120% (top row), apply no change (middle row) and apply an increase of 120% (bottom row), and (c) p_θ applies a decrease of 1000% (top row), applies no change (middle row) and applies an increase of 1000% (bottom row).

will produce a GEV distribution with a low upper limit. The spatial dependence rotation angle, θ, varies from the longitude axis pointing north-west to south-east to pointing south-west to north-east, therefore encompassing a much larger range in spatial dependence angle.

PAWN SA requires the specification of n, N_u and N_b, which were selected by trial and error in Pianosi and Wagener (2015). The values used in this investigation are based on those used by Pianosi and Wagener (2015). They used $n = 10$ conditioning intervals, $N_b = 200$ bootstrap samples and, for an example application, showed that the sensitivity indices start to converge when $N_u = 3000$.

6.3. Exploring the effect of stochasticity

The sensitivity of insured loss to changes in the footprint model input parameters will be explored. The insured loss associated with a given set of input parameters will be represented by the Average Annual Loss (AAL) and Maximum Annual Loss (MAL), calculated based on a sample of 174 simulated footprints (see Section 6.2). AAL and MAL may be sensitive to the particular sample of 174 footprints, or event set, used. This source of variation will depend on the realisation used, which is determined by the pseudo-random number generator seed. In a global sensitivity analysis experiment, Marrel et al. (2012) identified how a slight variation in the random seed can lead to very different sensitivity results. To directly quantify the relative influence of the event set used on the output of the model they treated the pseudo-random number generator as an additional input parameter of the SA experiment. Following Marrel et al. (2012), the same will be done here to quantify the sensitivity of the measures AAL and MAL to the event set used. Since the random seed can be any positive integer, with each seed setting creating a different sequence of pseudo-random numbers, increasing the seed by a particular increment does not have the same meaning as for the other quantitative input parameters of the model. The seed will therefore be set as either 1 or 2 within the SA experiment. The six footprint model input parameters are varied according to parameter range 1 (Table 6.1).

Figure 6.7 shows the resulting PAWN indices, $T_1, ..., T_7$ for each of the model pa-

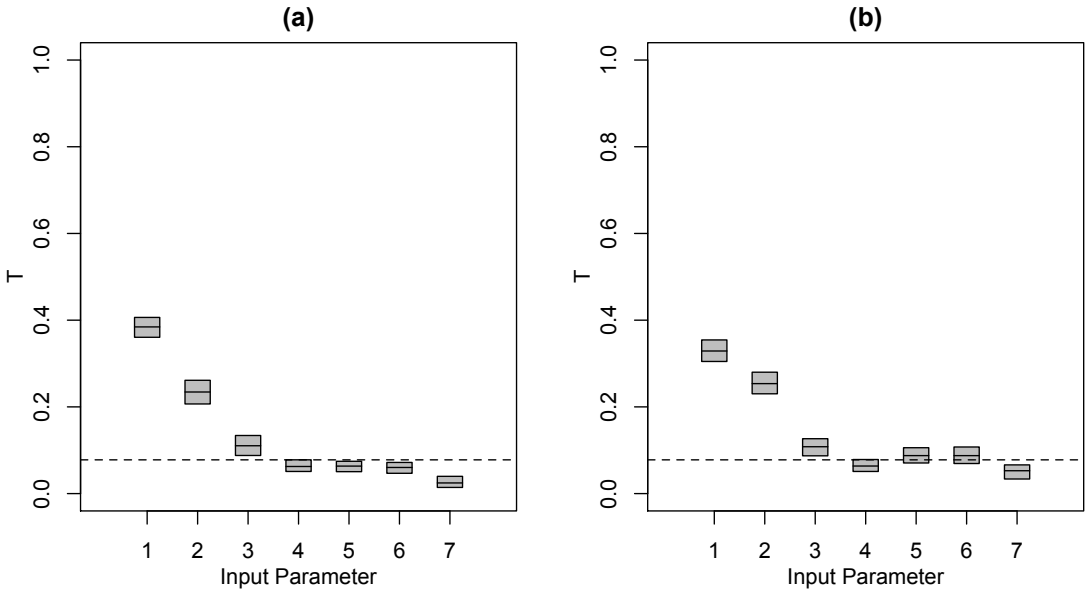

Figure 6.7. PAWN indices for each of the model parameters $p_\mu, p_\sigma, p_\xi, p_\theta, p_{\phi_1}$ and p_{ϕ_2} (input parameters 1-6) and the random seed (input parameter 7) for output (a) AAL and (b) MAL. The horizontal dashed line represents the Kolmogorov-Smirnov critical value for a 95% confidence level.

rameters (input parameters 1-6) and the random seed (input parameter 7). The PAWN index for the random seed is below the Kolmogorov-Smirnov critical value at a 95% confidence level, meaning that the distance between the unconditional and conditional empirical CDFs is not statistically significant. The choice of the random seed therefore has negligible influence on AAL and MAL, relative to the other input parameters. Using a different set of 174 simulated footprints would therefore give similar SA results. The relative influence of the six model parameters will be discussed in Section 6.4.

The results of this experiment indicate that only one event set needs to be used to calculate AAL and MAL, for each input parameter setting, to give representative SA results. The sensitivity to the event set may, however, increase if a smaller spatial domain were used due to the reduced affect of the law of large numbers when spatially summing over fewer locations.

6.4. Exploring the relative influence of footprint parameters on Average Annual Loss (AAL) and Maximum Annual Loss (MAL)

The intermediate results generated in the PAWN implementation procedure for AAL and MAL using parameter range 1 are presented in Figures 6.8 and 6.9 respectively.

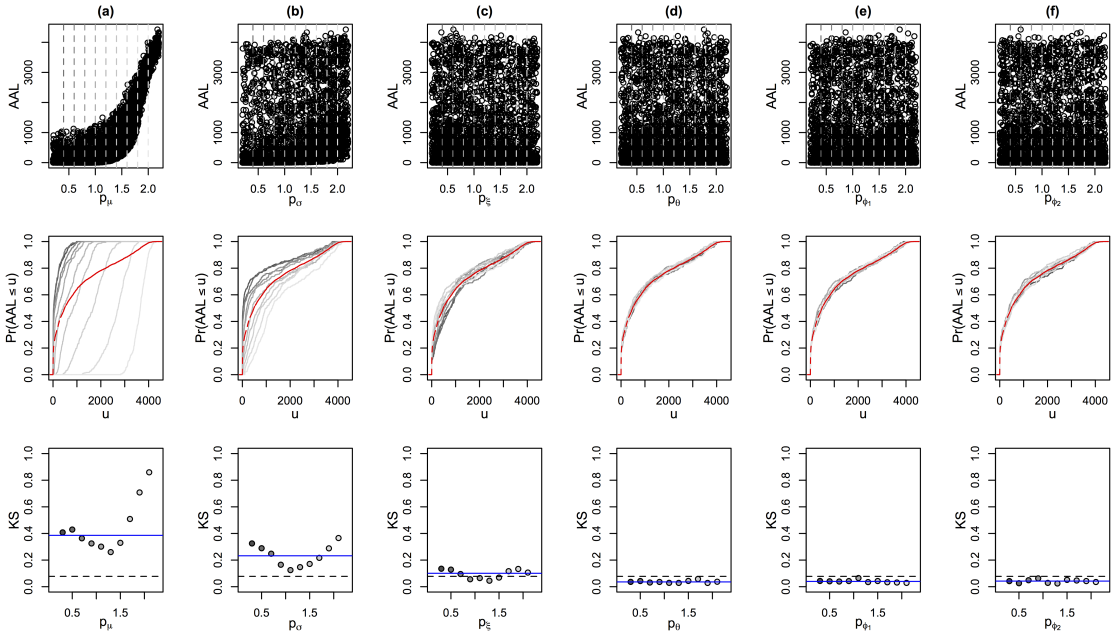

Figure 6.8. (a)-(f) for input parameters $p_\mu, p_\sigma, p_\xi, p_\theta, p_{\phi_1}, p_{\phi_2}$ respectively, top row: scatter plot of input parameter against output AAL, with dashed lines showing the partitions between the $n = 10$ conditioning intervals, middle row: the unconditional empirical CDF (red) and conditional empirical CDFs for the n conditioning intervals, bottom row: the Kolmogorov-Smirnov (KS) distance between the unconditional empirical CDF and each conditional empirical CDF for the n conditioning intervals. The blue line represents the median of KS over the n intervals and the dashed line represents the KS critical value at a 95% confidence level. The grey scale in the conditioning cluster partitions corresponds to the equivalent conditional empirical CDF and KS distance.

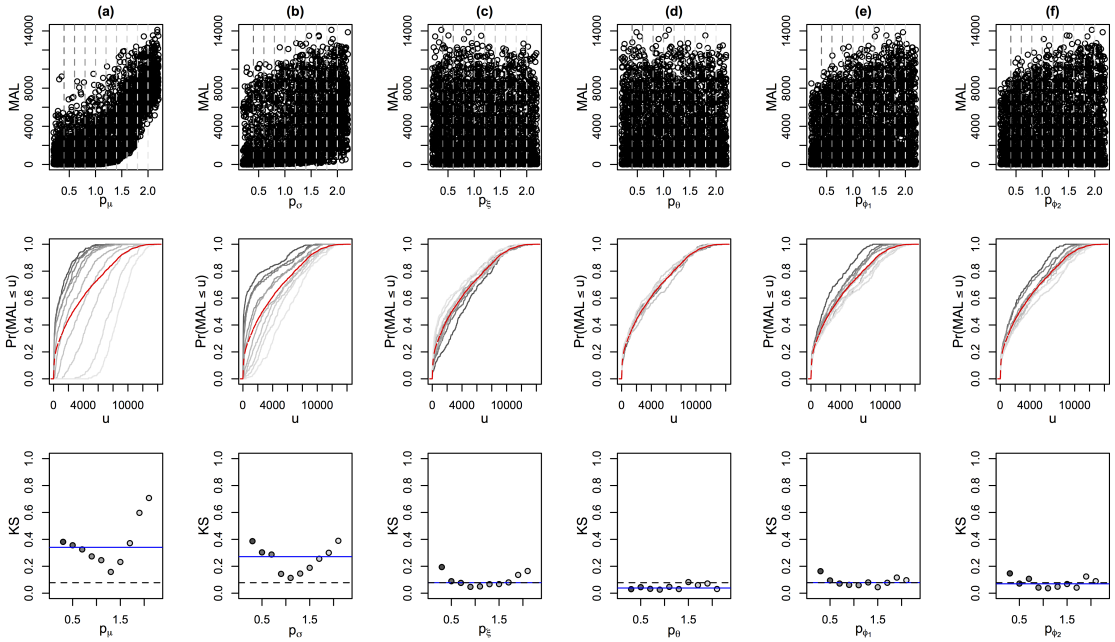

Figure 6.9. As Figure 6.8 but for output MAL.

The intermediate PAWN results are similar for the two model outputs, AAL and MAL. The top rows of Figures 6.8 and 6.9 (a) show that both AAL and MAL have a strong dependence on p_μ, the relative change in the left truncated GEV location parameter. This relationship is reflected in the large variation in conditional empirical CDFs (middle row) and resulting KS statistics (bottom row), particularly for large increases, $p_\mu \geq 1.5$. The input parameter p_σ also has some relative influence on both AAL and MAL, although this relative influence is less than p_μ in both cases. The input parameter p_ξ has a median KS statistic (blue line) just above/on the KS critical value (dashed line) for AAL and MAL respectively showing that it has a relatively small influence on the model outputs. However, these intermediate results allow for the interpretation of the influence of input parameters for different conditioning intervals. Figures 6.8 (c) and 6.9 (c) show that, for output AAL, the individual KS statistics for each of the conditioning intervals indicate that p_ξ has an influence on the output when less than 0.7 or greater than 1.5, representing a 30% decrease and 50% increase in the left truncated GEV shape parameter. Similarly for MAL, p_ξ has an influence on the output when less that 0.5 or greater then 1.6. The input parameters p_θ, p_{ϕ_1} and p_{ϕ_2} have negligible relative influence on the output AAL, even for very large values, with median and all individual KS statistics below the KS critical value. The relative influence of p_{ϕ_1} and p_{ϕ_2} is slightly higher for MAL, although still almost negligible relative to p_μ and p_σ. The resulting PAWN indices $T_1, .., T_6$ for AAL and MAL are shown in Figure 6.10 (a) and (e) respectively, reflecting the summary of Figures 6.8 and 6.9.

Figure 6.10 shows that, for all four parameter ranges, the input parameters that determine the local intensity of footprint wind gusts, p_μ, p_σ and p_ξ (input parameters 1-3) have more relative influence on both AAL and MAL than the input parameters that determine the spatial structure of the footprints, p_θ, p_{ϕ_1} and p_{ϕ_2} (input parameters 4-6).

For both AAL and MAL the relative influence of p_σ increases as the parameter ranges decrease, exceeding the relative influence of p_μ for parameter ranges 2 and 3, for MAL and parameter range 3, for AAL. This indicates that both the left-truncated GEV location and scale have the greatest relative influence on insured loss. The relative influence of p_ξ remains low for parameter ranges 1-3 but slightly increases for parameter range 4, for both AAL and MAL. This suggests that the left-truncated GEV shape parameter, ξ, has more influence on loss when allowed to vary within a larger range.

The relative influence of the spatial dependence parameters are consistently neg-

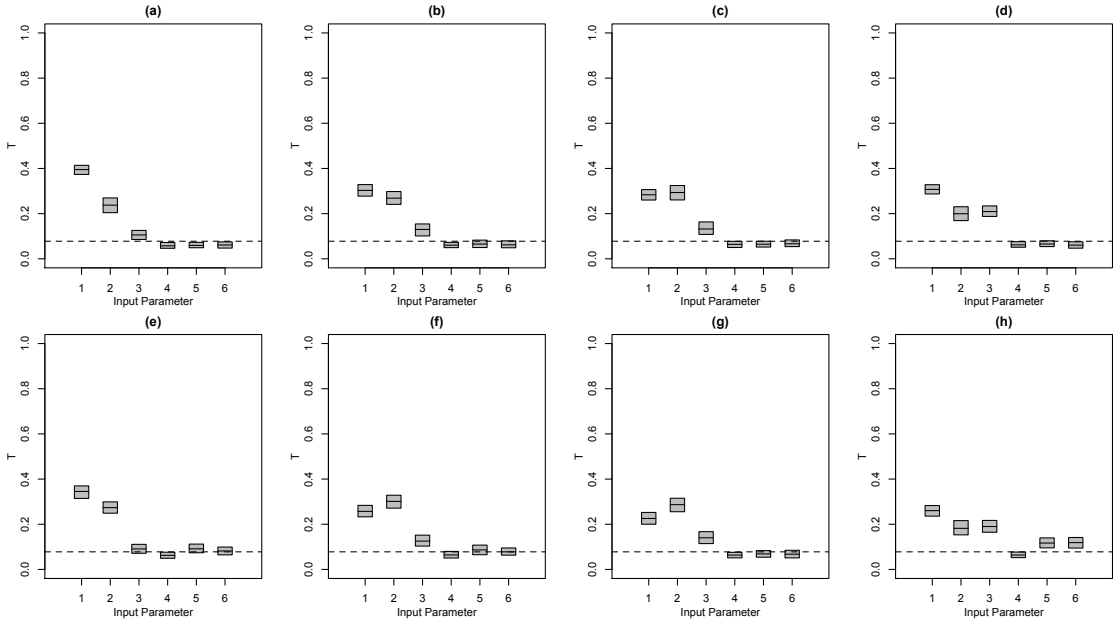

Figure 6.10. PAWN sensitivity indices for AAL (top row) and MAL (bottom row) for input parameter ranges 1, 2, 3 and 4 in each column (a)/(e), (b)/(f), (c)/(g) and (d)/(h) respectively. Parameter ranges are defined in Table 6.1. Input parameters $p_\mu, p_\sigma, p_\xi, p_\theta, p_{\phi_1}, p_{\phi_2}$ are labelled 1-6 respectively. Boxes represent bootstrap 95% confidence intervals and black lines represent the bootstrap mean. The dashed line is the KS critical value.

ligible for AAL. This agrees with derivation in Appendix A.3, which shows that for pairs of locations, irrespective of the loss function used, the expected loss (equivalent to AAL) does not depend on the relationship between the locations. The spatial dependence parameters have slightly higher relative influence on MAL and this influence is significant for parameter range 4, although still smaller than the relative influence of the local intensity parameters.

Increasing the left-truncated GEV location or scale parameter shifts or extends the local wind gust speed distribution upwards, increasing the probability of wind gust speeds exceeding 25ms^{-1}. These parameters therefore directly influence AAL and MAL. The spatial dependence parameters have a more indirect effect on the exceedance of the loss threshold, possibly explaining their relatively low influence.

Figure 6.11 shows the PAWN indices associated with each of the 6 input parameters varied according to parameter range 1, when considering AAL and MAL greater than predefined thresholds, as in Eqn 6.1: $t = 200$, the 90% quantile of historical footprint areas, $t = 670$, the 99% quantile of historical footprint areas and $t = 1533$, the maximum historical footprint area. The resulting PAWN indices show that as

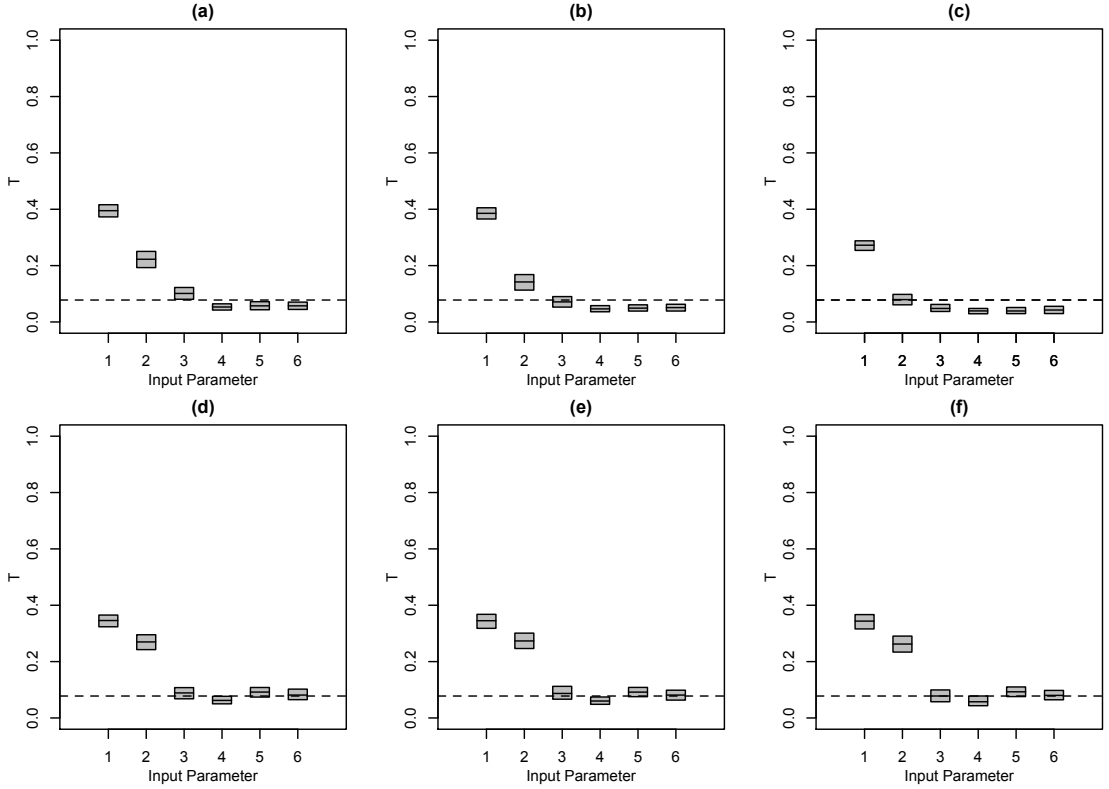

Figure 6.11. PAWN indices associated with each of the 6 input parameters when considering AAL (top row) and MAL (bottom row) greater than (a)/(d) 200, (b)/(e) 670 and (c)/(f) 1533. Boxes represent bootstrap 95% confidence intervals and black lines represent the bootstrap mean.

the threshold increases and the analysis moves further into the upper tail of the output sample of AAL, the greater the relative influence of p_μ and the lesser the relative influence of p_σ and p_ξ. For MAL, the relative influence of p_σ remains high as well as p_μ. The left-truncated GEV location parameter is the most important for determining extreme AALs and the left-truncated GEV scale parameter is also important for determining extreme MALs.

Figure 3.12 (b) in Chapter 3 showed that the area of the footprint exceeding 25ms^{-1} over land sharply declined in the mid 1990s. The results of this SA experiment imply that this decline was due to a change in the local intensity of footprint wind gust speeds, rather than the spatial dependence structure. This can be validated by fitting the windstorm footprint model to historical footprints in the two periods before and after the observed drop, 1979-1996 and 1997-2012.

Figure 6.12 shows the difference between the left-truncated GEV parameters at each location over land, when fit to all 5730 historical footprints and when fit to those events that occurred in the period 1979-1996 (top row), and 1997-2012 (bottom row). The parameters for the period before the 1997 are approximately equal to

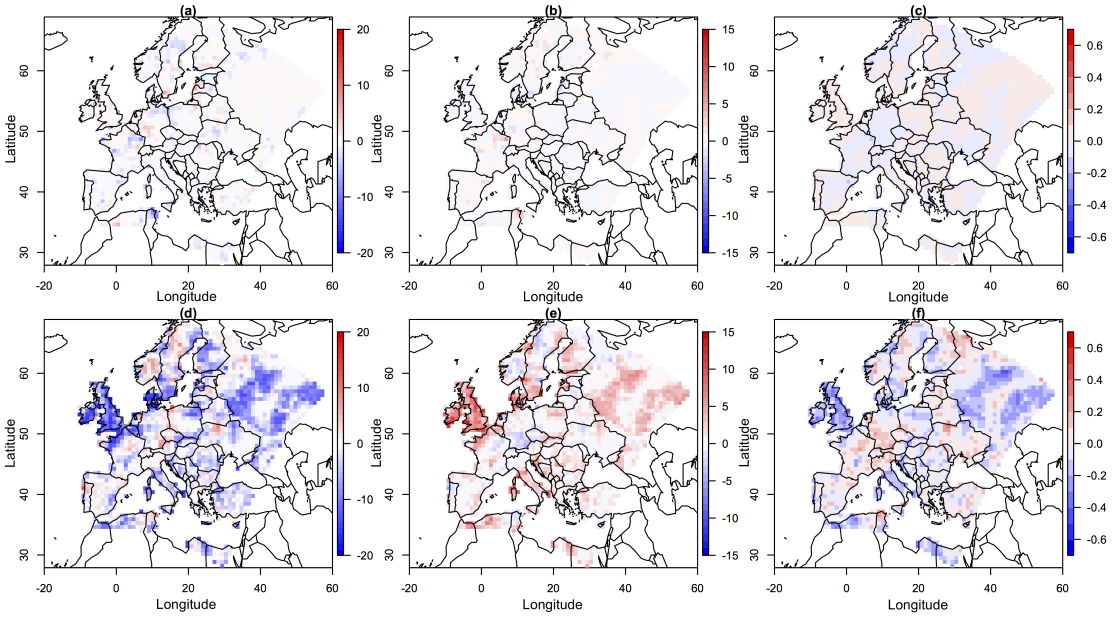

Figure 6.12. The difference between the left-truncated GEV parameters μ, σ and ξ in columns (a)/(d), (b)/(e) and (c)/(f) respectively, when fit to all land locations in the European domain using all 5730 historical footprints and those in 1979-1996 (top row), and 1997-2012 (bottom row).

those fit to the whole data set for all locations over Europe. In the later period, however, the location parameter is much smaller than when fitted to the whole data set, particularly in the UK and north-west Europe, the region in which most extreme wind gust speeds occur. This decrease in the location parameter is, however, slightly compensated for by an increase in the scale parameter. The shape parameter is also lower in this region for the later period, implying that the upper limit of the local intensity distributions are lower in this period.

Figure 6.13 shows the empirical binned covariogram when the spatial domain is transformed using the spatial dependence parameters fit to all the 5730 historical footprints (black circles), those in 1979-1996 (turquoise squares), and 1997-2012 (pink triangles), which should be well modelled by the unit range exponential covariance function. The transformed empirical binned covariances are consistent for the three periods, showing that the spatial dependence parameters have not changed in these historical periods.

This comparison of historical time periods has shown that the lull in windstorm footprint area and insured loss (Mark, 2013) is indeed due to a change in the local intensity of footprint wind gust speeds and not a change in the spatial dependence structure of windstorm footprints.

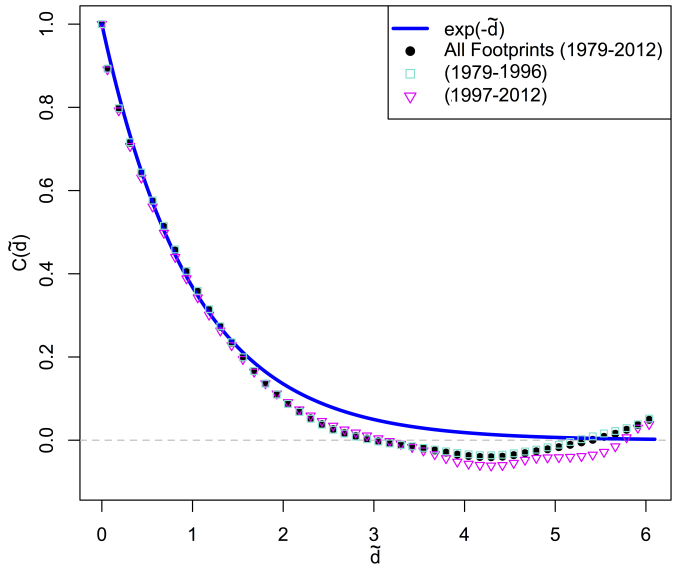

Figure 6.13. A comparison of the empirical binned covariogram when the spatial domain is transformed using the spatial dependence parameters, ϕ_1, ϕ_2 and θ, fit to all of the 5730 historical footprints (black circles), those in 1979-1996 (turquoise squares), and those in 1997-2012 (pink triangles).

6.5. Conclusions

A formal SA experiment has been carried out to explore which windstorm footprint model parameters have the greatest influence on insured loss, represented by the average annual and maximum annual footprint area exceeding 25ms^{-1}, AAL and MAL. PAWN, a global All-At-the-Time density-based SA methodology, was used for this investigation.

Both model outputs AAL and MAL, used to represent loss due to the stochasticity of the footprint model, were found to be insensitive to the sample of footprints used to represent a given year, determined by the pseudo-random number generator value. This meant that only one event set was needed to calculate AAL and MAL, for each set of input parameters within the SA experiment to give representative results.

For all four parameter ranges presented, the input parameters that determine the local intensity of footprint wind gusts, p_μ, p_σ and p_ξ had more relative influence on both AAL and MAL than the input parameters that determine the spatial structure of the footprints, p_θ, p_{ϕ_1} and p_{ϕ_2}. This result was consistent for the upper tail of the output samples of AAL and MAL, suggesting that the left-truncated GEV parameters are also most important for determining very extreme losses.

The windstorm footprint model was fitted to events before and after the observed sharp decline in the area of the footprint exceeding 25ms^{-1} over land in the mid 1990s. The spatial dependence structure remains constant throughout the period, while the left-truncated GEV location and shape parameters are lower in the later period, reducing the local intensity of wind gust speeds, particularly in the UK and North-west Europe. The recent lull in insured losses can therefore be attributed to a change in the local intensity of footprint wind-gust speeds, rather than a change in the spatial dependence structure.

7. Conclusions

This chapter summarises the main results of this thesis and suggests possible directions for further development.

7.1. Summary

In Chapter 3, it was shown that the storm severity measure quantifying the area of damaging wind gust speeds over land ($> 25\text{ms}^{-1}$) was the best footprint classifier of insured loss. This measure was shown to outperform the spatial 90^{th} percentile wind gust speed, a measure of the footprint peak intensity, and the loss function severity measures, most commonly used in previous studies. Including the magnitude of the exceedance of the damage threshold and population density, as a proxy for exposure, within the severity measure was shown to be detrimental to the relationship between the measure and insured loss. In addition, using a relative or absolute threshold made very little difference when relating the loss function severity measures to insured loss. A composite severity measure, the product of the footprint area and the cube maximum 925hPa wind speed along the track, was shown to give the best overall representation of insured loss, both using raw footprints and disaggregated footprints. This composite measure was therefore used to select the extreme storms for the XWS catalogue. Furthermore, it was shown that the composite severity measure and footprint area (conceptual loss) experienced a sharp decline in the mid 1990's, consistent with the decline experienced in windstorm related insured losses (Mark, 2013). This exploration was novel because the conclusions represent the whole of Europe and were based on a large data set of 5730 high resolution windstorm footprints.

In Chapter 4, the left-truncated GEV was shown to be a reasonable model for the marginal distribution of footprint wind gust speeds exceeding the 60^{th} percentile at each location over land in the European domain. Wind gust speeds at pairs of locations throughout the domain were found to be asymptotically independent. The Gaussian copula was found to provide a good model for the bivariate dependence between locations and outperformed the Gumbel copula proposed by Bonazzi et al.

(2012). The GEV-Gaussian bivariate copula model was able to realistically represent bivariate conceptual losses for pairs of locations throughout the domain.

In Chapter 5, the dependence between all Gaussian transformed wind gust speeds were shown to be well captured using a geostatistical covariance function of separation distance and direction. Problems arose when fitting the Matérn covariance function. A ridge in the likelihood surface identified a number of possible parameter combinations with very similar likelihoods. It was shown that this ridge was caused by a relatively large weight being placed on small separation distances in the likelihood function due to the large number of locations in the footprint data set. The best fit was shown to be achieved when the Matérn covariance function shape parameter was 0.5 (equivalent to an exponential covariance function). Windstorm footprints were therefore modelled as non-differentiable, rough spatial processes. An anisotropic covariance function was found to better represent the most common eastward transit of windstorms over Europe. The geostatistical windstorm footprint model was shown to realistically represent joint conceptual losses for all locations and individual countries in Europe.

In Chapter 6, it was shown that the left-truncated GEV parameters of the geostatistical model, rather than the geostatistical spatial dependence parameters, have a greater influence on loss, represented by Average Annual Loss (AAL) and Maximum Annual Loss (MAL) due to the stochasticity of the footprint. This was true for all loss values and the most extreme losses. This conclusion was found to be robust to variations in the input parameter ranges used. As a consequence, when the windstorm footprint model was fit to historical footprints in the two periods, before and after the observed decline in conceptual loss in the mid 1990's, the spatial dependence structure of footprints was found to remain constant throughout the period, while the left-truncated GEV location and shape parameters were lower in the later period, particularly in the UK and north-west Europe, leading to the observed decrease in storm severity. This was a novel application of sensitivity analysis to a stochastic natural hazards model.

7.2. Directions for further development

The exploration of the relationship between storm severity measures and insured loss in this thesis was based on how successful each severity measure was at classifying notable storms that caused large insured losses. This investigation could possibly be improved by using quantitative insured loss values. However, insured loss data is hard to acquire, for example Klawa and Ulbrich (2003) were only able to obtain

losses for Germany for 11 extreme windstorms. A fully rigorous validation would necessitate the availability of a European wide insured loss values for each of the 5730 windstorm events in the data set, however any available loss data could be used to further validate the conclusions of the investigation.

As discussed in Section 2.3.1, the model generated windstorm footprint data used throughout this thesis underestimate wind gust speeds in excess of 25ms^{-1} and in locations above 500m altitude. These biases occur because the footprints are created by dynamically downscaling reanalysis data which is created by combining atmospheric observations with a climate model which represents sub grid cell scale wind gust speeds using a gust parameterisation scheme. This bias in the footprint data used could influence the results of this investigation. For example, if the wind gust speeds exceeding 25ms^{-1} were correctly modelled, increasing the footprint intensity in places, the spatial variation in the footprint wind gusts would increase and the excess gust speed and population density could therefore be more important within the storm severity measures in Chapter 3. The biases in model generated wind gust speed data are also often largest along the coasts (Della-Marta et al., 2009), where a large proportion of the population live i.e. along the south coast of the UK and north coast of mainland Europe (Figure 2.3). Again, this suggests that more realistic wind gust speeds could increase the influence of population density. This could potentially be explored in a sensitivity analysis study, by increasing the local marginal distribution parameters around the coast and observing the change in the population density affected by simulated footprints. Alternatively, the footprints could be bias corrected using the method developed in Roberts et al. (2014), and the storm severity measure exploration repeated based on these corrected footprints.

Further research is also required to understand how the analysis is effected by the uncertainty in the model generated footprints. These uncertainties arise from the various models and algorithms used to produce the footprints. The footprints may be different if an alternative reanalysis data set were used; if the reanalysis were initialised using slightly different initial conditions; if a different downscaling technique were used; if a different storm tracking algorithm were used; or if a different windstorm duration were specified, for example. The sensitivity of the analysis to these sources of uncertainty could be explored by varying these factors and observing the change in the results of the investigation. This exploration would require collaboration with the scientists that produced the windstorm footprint data but could provide vital information about the uncertainty in the conclusions of the investigation.

In some of the storm severity measures explored in Chapter 3, a spatially varying

damage threshold was used to represent the adaptation to local wind gust speed climatology. It could also be thought that this damage threshold may vary with time, depending on how the wind gust speed climatology varies throughout the period. Incorporating this within the storm severity measures, may, however be challenging since it requires the identification of a time window for which it takes a location to adapt to the local conditions. In addition, this period may even be longer than the 35 year period used here since buildings are not often updated. A time varying damage threshold may therefore be more appropriate when a much longer period is explored.

Within the windstorm footprint model, the marginal left truncated GEV distributions are fit above a spatially varying threshold. Again, a temporally varying threshold could be useful here to represent the non-stationarity in wind gust speeds at each location over the period, evident in the decline in storm severity since 1996. This would ensure that all temporally local extremes are modelled, rather than just the overall extremes from the period, which may be biased towards the earlier half of the period. This would add an additional level of complexity to the model, however, since the GEV parameters would vary in time as well as space. Alternatively a preprocessing method, such as in Eastoe (2009), could be used to transform the non-stationary wind gust speed time series at each location to stationary, allowing them to be modelled as such.

Footprint wind gust speeds were found to be asymptotically independent throughout the domain. The Gaussian copula was therefore used to model the dependence between pairs of locations based on its simplicity and extendability to a spatial Gaussian process. An alternative approach would be to test the more complex conditional dependence model proposed by Heffernan and Tawn (2004), which incorporates both classes of extremal dependence and is extendable to multiple dimensions. Extending the Heffernan and Tawn (2004) model to represent all 14872 locations in the data set would be a challenge since a very large number of parameters would need to be estimated (four parameters for each pair of locations). This would make the model more computationally expensive for simulation. Max-stable processes, much like Gaussian processes, represent the multivariate dependence using a spatial process with only a few parameters for the whole domain. Currently, however, max-stable models are too coarse to accurately describe tails of multivariate distributions with asymptotic independence and can therefore only be used to model asymptotic dependence or complete independence. Ongoing research aims to develop max-stable models that are able to represent both classes of asymptotic dependence (Wadsworth et al., 2015). Once these types of models have been developed, they could form a preferable and more flexible basis for modelling windstorm footprints.

The results of the sensitivity study could be further validated using additional sensitivity analysis methods, i.e. methods discussed in Pianosi et al. (2014). The local wind gust speed marginal parameters were shown to have more influence on loss than the spatial dependence parameters. This conclusion may be sensitive to the way in which loss is represented. Here the footprint damage area is used because it is shown to be the best footprint based representation of insured loss in Chapter 3, however it would be of interest to explore the sensitivity using other storm severity measures e.g. x_{90} or the Klawa and Ulbrich (2003) loss function.

There are therefore a number of ways that the investigation presented in this thesis could be developed in the future. In particular, there are still many interesting opportunities for the statistical modelling of windstorm footprints.

Appendices

A. Extremal dependence measure calculations and properties

A.1. Calculating the extremal dependence measures using the GEV-Gaussian and GPD-Gumbel bivariate copula models

A simulation method is used to calculate the extremal dependence measures and their confidence intervals for the GEV-Gaussian bivariate model. For a given pair of locations, with estimated dependence parameter ρ, 6000 data points are simulated from the corresponding bivariate Gaussian distribution 1000 times. One such simulation will produce a set of bivariate Gaussian data:

$$(Y_1^*, Y_2^*) \sim MVN\left(\begin{pmatrix}0\\0\end{pmatrix}, \begin{pmatrix}1 & \rho\\ \rho & 1\end{pmatrix}\right)$$

where $Y_j^* = Y_{j1}^*, Y_{j2}^*, ..., Y_{j6000}^*$

Then,

$$\begin{aligned}\chi(u) &= \Pr(F_2(X_2) > u | F_1(X_1) > u)\\ &= \frac{\Pr(F_2(X_2) > u, F_1(X_1) > u)}{\Pr(F_1(X_1) > u)}\\ &\approx \frac{\frac{1}{m}\sum_{i=1}^{m} I(Y_{1i}^* > \Phi^{-1}(u), Y_{2i}^* > \Phi^{-1}(u))}{\frac{1}{m}\sum_{i=1}^{m} I(Y_{1i}^* > \Phi^{-1}(u))}\end{aligned}$$

and

$$\bar{\chi}(u) \approx \frac{2\log(\frac{1}{m}\sum_{i=1}^{m} I(Y_{1i}^* > \Phi^{-1}(u)))}{\log(\frac{1}{m}\sum_{i=1}^{m} I(Y_{1i}^* > \Phi^{-1}(u), Y_{2i}^* > \Phi^{-1}(u)))}$$

Where $I(a)$ is an indicator function that takes the value 1 if a is true and 0 otherwise. For each $u \in (0.6, 1)$, the 95% parametric bootstrap interval is evaluated as the 2.5% and 97.5% quantiles of the extremal dependence measures calculated from the 1000 simulations from the bivariate Gaussian distribution.

For the Gumbel copula, $\chi(u)$ and $\bar{\chi}(u)$ are evaluated using the joint survivor function for the Gumbel copula for $u > H_j^{-1}(\gamma_j)$

$$\begin{aligned}\chi(u) &= \Pr(H_2(X_2) > u | H_1(X_1) > u) \\ &= \frac{\Pr(H_2(X_u 2) > u, H_1(X_1) > u)}{\Pr(H_1(X_1) > u)} \\ &= \frac{1 - 2u + C(u,u)}{1-u} \\ \bar{\chi}(u) &= \frac{2\log(1-u)}{\log(1-2u+C(u,u))} - 1\end{aligned}$$

Where C is the Gumbel copula, defined in section 4.4.1.

A.2. Showing that the extremal dependence measures are independent of the marginal distributions

Let $\chi_X(u)$ and $\bar{\chi}_X(u)$ be the measures of extremal dependence in the raw wind gusts:

$$\begin{aligned}\chi_X(u) &= \frac{\Pr(X_1 > x_u, X_2 > x_u)}{\Pr(X_1 > x_u)} \\ \bar{\chi}_X(u) &= \frac{2\log\Pr(X_1 > x_u)}{\log\Pr(X_1 > x_u, X_2 > x_u)} - 1\end{aligned}$$

then these measures are equal to the extremal dependence measures for Y_1 and Y_2, $\chi(u)$ and $\bar{\chi}(u)$:

$$\begin{aligned}\chi(u) &= \frac{\Pr(Y_1 > y_u, Y_2 > y_u)}{\Pr(Y_1 > y_u)} \\ &= \frac{\Pr(\Phi^{-1}(F_1(X_1)) > \Phi^{-1}(F_1(x_u)), \Phi^{-1}(F_2(X_2)) > \Phi^{-1}(F_2(x_u)))}{\Pr(\Phi^{-1}(F_1(X_1)) > \Phi^{-1}(F_1(x_u)))} \\ &= \frac{\Pr(X_1 > x_u, X_2 > x_u)}{\Pr(X_1 > x_u)} = \chi_X(u)\end{aligned}$$

and,

$$\begin{aligned}\bar{\chi}(u) &= \frac{2\log\Pr(Y_1 > y_u)}{\log\Pr(Y_1 > y_u, Y_2 > y_u)} - 1 \\ &= \frac{2\log\Pr(\Phi^{-1}(F_1(X_1)) > \Phi^{-1}(F_1(x_u)))}{\log\Pr(\Phi^{-1}(F_1(X_1)) > \Phi^{-1}(F_1(x_u)), \Phi^{-1}(F_2(X_2)) > \Phi^{-1}(F_2(x_u)))} - 1 \\ &= \frac{2\log\Pr(X_1 > x_u)}{\log\Pr(X_1 > x_u, X_2 > x_u)} - 1 = \bar{\chi}_X(u)\end{aligned}$$

A.3. Showing how the bivariate expected loss does not depend on the relationship between locations

Interestingly, it can be shown that for two locations s_1 and s_2 the expected loss, irrespective of the conceptual loss function used, does not depend on the relationship between the locations. Let

$$\begin{aligned} p_1 = Pr(X_1 > t_1) &= \frac{a+c}{n} \\ p_2 = Pr(X_2 > t_1) &= \frac{a+b}{n} \\ \chi = Pr(X_2 > t_2 | X_1 > t_1) &= \frac{a}{a+c} \end{aligned}$$

where a, b, c are defined as in Figure 3.9 and n is the total number of events. Then, shortening $L(s_1, s_2)$ to L, for any loss function,

$$\begin{aligned} Pr(L = 1) &= \frac{c+b}{n} = p_1 + p_2 - 2\chi p_1 \\ Pr(L = 2) &= \frac{a}{n} = \chi p_1 \end{aligned}$$

Then,

$$\begin{aligned} \mathbb{E}(L) &= 1\mathrm{Pr}(L = 1) + 2\mathrm{Pr}(L = 2) \\ &= p_1 + p_2 - 2\chi p_1 + 2\chi p_1 = p_1 + p_2 \\ &= Pr(X_1 > t_1) + Pr(X_2 > t_2), \end{aligned}$$

so the expected joint loss only depends on the probability of a loss occurring in each location and does not depend on the dependence relationship between the locations.

B. Covariance model properties

B.1. Mean-square Properties

Definition taken from Diggle and Ribeiro (2007): a stochastic process $Z(s)$ is mean-square continuous if $E[(Z(s+d)-Z(s))^2] \to 0$ as $d \to 0$. Also, $Z(s)$ is mean-square differential, with mean-square derivative $Z'(s)$, if

$$E\left[\left(\frac{Z(s+d)-Z(s)}{d} - Z'(s)\right)^2\right] \to 0$$

as $d \to 0$. Higher-order mean-square differentiability is then defined sequentially, i.e. $Z(s)$ is twice mean-square differentiable is $Z'(s)$ is mean-square differentiable, and so on.

Bibliography

Anderson, B., E. Borgonovo, M. Galeotti, and R. Roson, 2014: Uncertainty in climate change modeling: Can global sensitivity analysis be of help? *Risk Analysis*, **34 (2)**, 271–293.

Banerjee, S., B. P. Carlin, and A. E. Gelfand, 2004: *Hierarchical Modeling and Analysis for Spatial Data*. Chapman and Hall/CRC, New York.

Berg, D., 2009: Copula goodness-of-fit testing: an overview and power comparison. *The European Journal of Finance*, **15**, 675–701.

Beven, K. J., et al., 2016: Epistemic uncertainties and natural hazard risk assessment – part 2: Different natural hazard areas, in Review.

Blanchet, J. and A. C. Davison, 2011: Spatial modeling of extreme snow depth. *The Annals of Applied Statistics*, **5 (3)**, 1699–1725.

Blanchet, J., C. Marty, and M. Lehning, 2009: Extreme value statistics of snowfall in the Swiss Alpine region. *Water Resources Research*, **45**, doi:10.1029/2009WR007916.

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

Born, K., P. Ludwig, and J. G. Pinto, 2012: Wind gust estimation for mid-European winter storms: towards a probabilistic view. *Tellus A*, **4**, doi:10.3402/tellusa.v64i0.17471.

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

Bortot, P. and J. A. Tawn, 1998: Models for the extremes of markov chains. *Biometrika*, **85**, 851–867.

Brasseur, O., 2001: Development and application of a physical approach to estimating wind gusts. *Monthly Weather Review*, **129**, 5–25.

Byrd, R. H., P. Lu, J. Nocedal, and C. Zhu, 1995: A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, **16 (5)**, 1190–1208.

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

Clark, K. M., 2002: The use of computer modeling in estimating and managing future catastrophe losses. *The Geneva Papers on Risk and Insurance*, **27**, 181–195.

Coles, S., 2001: *An Introduction to Statistical Modeling of Extreme Values*. Springer.

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

Coles, S. and J. Tawn, 2005: Bayesian modelling of extreme surge on the UK east coast. *Philosophical Transactions of the Royal Society*, **363**, 1387–1406.

Coles, S. G. and J. A. Tawn, 1996: Modelling extremes of the areal rainfall process. *Journal of the Royal Statistical Society: Series B (Methodology)*, **58 (2)**, 329–347.

Coles, S. G. and D. Walshaw, 1994: Directional modelling of extreme wind speeds. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, **43 (1)**, 139–157.

Cox, D. R. and D. V. Hinkley, 1974: *Theoretical Physics*. Chapman and Hall, London.

Cressie, N. A. C., 1993: *Statistics for Spatial Data*. Wiley, New York.

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

Davies, T., M. J. P. Cullen, A. J. Malcolm, M. H. Mawson, A. Staniforth, A. A. White, and N. Wood, 2005: A new dynamical core for the Met Office's global and regional modelling of the atmosphere. *Quart. J. Roy. Meteorol. Soc.*, **131**, 1759–1782.

Davison, A. and R. Smith, 1990: Models for exceedances over high thresholds. *Journal of the Royal Statistical Society. Series B (Methodology)*, **52 (3)**, 393–442.

Davison, A. C., S. A. Padoan, and M. Ribatet, 2012: Spatial modeling of spatial extremes. *Statistical Science*, **27 (2)**, 161–186.

de Hann, L., 1984: A spectral representation for max-stable processes. *Ann. Probab.*, **12**, 1194–1204.

Dee, D., et al., 2011: The ERA interim reanalysis: configuration and performance of the data assimilation system. *Quart. J. Roy. Meteorol. Soc.*, **137**, 553–597.

Della-Marta, P., H. Mathis, C. Frei, M. Liniger, J. Kleinn, and C. Appenzeller, 2009: The return period of wind storms over Europe. *International Journal of Climatology*, **29**, 437–459.

Della-Marta, P. M., M. Liniger, C. Appenzeller, D. N. Bresch, P. Köllner-Heck, and V. Muccione, 2010: Improved estimates of the European Winter windstorm climate and the risk of reinsurance loss using climate model data. *Journal of Applied Meteorology and Climate*, **49**, 2092 – 2120.

Della-Marta, P. M. and J. G. Pinto, 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**, doi:10.1029/2009GL038557.

Diggle, P. J. and P. J. Ribeiro, 2007: *Model-based Geostatistics*. Springer.

Donat, M. G., G. C. Leckebusch, S. Wild, and U. Ulbrich, 2010: Benefits and limitations of regional multi-model ensembles for storm loss estimations. *Climate Research*, **44**, 211–225.

Donat, M. G., G. C. Leckebusch, S. Wild, and U. Ulbrich, 2011a: Future changes in European winter storm losses and extreme wind speeds inferred from GCM and RCM multi-model simulations. *Natural Hazards and Earth System Science*, **11**, 1351–1370.

Donat, M. G., T. Pardowitz, G. C. Leckebusch, U. Ulbrich, and O. Burghoff, 2011b: High-resolution refinement of a storm loss model and estimation of return periods of loss-intensive storms over Germany. *Nat. Hazards Earth Syst. Sci.*, **11**, 2821–2833.

Dorland, C., R. S. J. Tol, and J. P. Palutikof, 1999: Vulnerability of the Netherland and Northwest Europe to storm damage under climate change. *Climate Change*, **43 (3)**, 513 – 535.

Eastoe, E. F., 2009: A hierarchical model for non-stationary multivariate extremes: a case study of surface-level ozone and NOx data in the UK. *Environmetrics*, **20**, 428–444.

Eastoe, E. F., S. Koukoulas, and P. Jonathan, 2013: Statistical measures of extremal dependence illustrated using measured sea surface elevations from a neighbourhood of coastal locations. *Ocean Engineering*, **62**, 68–77.

ECMWF, 2006: Ifs documentsation: Cy31r1 operational implementation. www.ecmwf.int/research/ifsdocs/CY31r1.

Economou, T., D. B. Stephenson, and C. A. T. Ferro, 2014: Spatio-temporal modelling of extreme storms. *The Annals of Applied Statistics*, **8 (4)**, 2223–2246.

Fermanian, J. D., 2005: Goodness-of-fit tests for copulas. *Journal of Multivariate Analysis*, **95**, 119–152.

Ferro, C. A. T., 2007: A probability model for veryfiying deterministic forecasts of extreme events. *Weather Forecasting*, **22**, 1089–1100.

Fisher, R. A., 1912: On the absolute criterion for fitting frequency curves. *Messenger of Mathematics*, **41**, 155–160.

Fisher, R. A. and L. H. C. Tippett, 1928: Limiting forms of the frequency distribution of the largest or smallest member of a sample. *Mathematical Proceedings of the Cambridge Philosophical Society*, **24 (2)**, 180–190.

Friederichs, P., M. Göber, S. Bentizien, A. Lenz, and R. Krampitz, 2009: A probabilistic analysis of wind gusts using extreme value statistics. *Meteorogische Zeitschrift*, **18 (6)**, 615–629.

Galambos, J., 1987: *The Asymptotic Theory of Extreme Order Statistics*. Melbourne: Krieger.

Genz, A., 1992: Numerical computation of multivariate normal probabilities. *Journal of Computational and Graphical Statistics*, **1**, 141–149.

Gnedenko, B., 1943: Sur la distribution limite du terme maximum d'une serie aleatoire. *Annals of Mathematics, Second Series*, **44 (3)**, 423–453.

Gneiting, T., S. Zoltán, and M. Schlather, 2001: Analogies and correspondences between variograms and covariance functions. *Advanced Applied Probability*, **33**, 617–630.

Gulev, S. K., O. Zolina, and S. Grigoriev, 2001: Extratropical cyclone variability in the Northern Hemisphere winter from the NCEP/NCAR reanalysis data. *Climate Dynamics*, **17**, 795–809.

Haas, R. and J. G. Pinto, 2012: A combined statistical and dynamical approach for downscaling large-scale footprints of European windstorms. *Geophysical Research Letters*, **39**.

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

Hamm, N., J. Hall, and M. Anderson, 2006: Variance-based sensitivity analysis of the probability of hydrologically induced slope instability. *Computers and Geosciences*, **32 (6)**, 803–817.

Handcock, M. S. and J. R. Wallis, 1994: An approach to statistical spatial-temporal modelling of meteorological fields. *Journal of the American Statistical Association*, **89 (426)**, 368–378.

Hanley, J. and R. Caballero, 2012: The role of large-scale atmospheric flow and rossby wave breakingin the evolution of extreme windstorms over europe. *Geophysical Research Letters*, **39**, doi:10.1029/2012GL053408.

Haskard, K. A., 2007: An anisotropic Matérn spatial covariance model: REML estimation and properties. Ph.D. thesis, University of Adelaide.

Haylock, M., 2011: European extra-tropical storm damage risk from a multi-model ensemble of dynamically-downscaled global climate models. *Natural Hazards and Earth System Science*, **11**, 2847–2857.

Heffernan, J. E., 2000: A directory of coeffcients of tail dependence. *Extremes*, **3 (3)**, 279–290.

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

Heneka, P., T. Hofherr, B. Ruck, and C. Kottmeier, 2006: Winter storm risk of residential structures: model development and application to the German state of Baden-Württemberg. *Natural Hazards and Earth System Science*, **6 (5)**, 721–733.

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

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

Howard, T. and P. Clark, 2007: Correction and downscaling of NWP wind speed forecasts. *Meteorological Applications*, **14**, 105–116.

Jagger, T. H. and J. B. Elsner, 2005: Climatology models of extreme hurricane winds near the United States. *Journal of Climate*, **19**, 3220 – 3236.

Justus, C. G., W. R. Hargraves, A. Mikhail, and D. Graber, 1977: Methods for estimating wind speed frequency distributions. *Journal of Applied Meteorology*, **17**, 350 – 353.

Keef, C., I. Papastathopoulos, and J. A. Tawn, 2013: Estimation of the conditional distribution of a multivariate variable given that one of its components is large: Additional constraints for the Heffernan and Tawn model. *Journal of Multivariate Analysis*, **115**, 396–404.

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

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

Lark, R. M., 2000: Estimating variograms of soil properties by the method-of-moments and maximum likelihood. *European Journal of Soil Science*, **51**, 717–728.

Leckebusch, G., D. Renggli, and U. Ulbrich, 2008a: Development and application of an objective storm severity measure for the Northeast Atlantic region. *Meteorologische Zeitschrift*, **17**, 575–587.

Leckebusch, G., U. Ulbrich, L. Frohlich, and J. Pinto, 2007: Property loss potential for European midlatitude storms in a changing climate. *Geophysical Research Letters*, **34**, doi:10.1029/2006GL027663.

Leckebusch, G. C., B. Koffi, U. Ulbrich, J. G. Pinto, T. Spangehl, and S. Zacharias, 2006: Analysis of frequency and intensity of European winter storm events from a multi-model perspective, at synoptic and regional scales. *Climate Research*, **31**, 59–74.

Leckebusch, G. C., A. Weimer, J. G. Pinto, M. Reyers, and P. Speth, 2008b: Extreme wind storms over Europe in present and future climate: a cluster analysis approach. *Meteorologische Zeitschrift*, **17 (1)**, 67–82.

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

Ledford, A. W. and J. A. Tawn, 1997: Modelling dependence within joint tail regions. *Journal of the Royal Statistical Society*, **59 (2)**, 475–499.

Lionello, P., U. Boldrin, and F. Giorgi, 2008: Future changes in cyclone climatology over Europe as inferred from a regional climate simulation. *Climate Dynamics*, **30**, 657–671.

Luo, W., M. C. Taylor, and S. R. Parker, 2008: A comparison of spatial interpolation methods to estimate continuous wind speed surfaces using irregularly distributed data from England and Wales. *International Journal of Climatology*, **28**, 947–959.

Mark, A., 2013: European windstorm: Such a peculiarly uncertain risk for Solvency ii. Accessed: 16/12/2015, `www.rms.com/blog/tag/european-windstorm`, accessed: 16/12/2015.

Marrel, A., B. Iooss, S. D. Veiga, and M. Ribatet, 2012: Global sensitivity analysis of stochastic computer models with joint metamodels. *Statistics and Computing, Springer Verlag (Germany)*, **22**, 833–847.

Martínez-Alvarado, O., S. L. Gray, J. L. Catto, and P. A. Clark, 2014: Sting jets in intense winter North-Atlantic windstorms. *Environmental Research Letters*, **7 (2)**, 1–8.

Matheron, G., 1974: The intrinsic random functions and their applications. *Advanced Applied Probability*, **5**, 439–468.

McKay, M. D., R. J. Beckman, and W. J. Conover, 2000: A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, **42**, 55–61.

Mikosch, T., 2006: Copulas: Tales and facts. *Extremes*, **9**, 3–62.

Mildenhall, S., S. Bowen, C. Kennedy, and A. Podlaha, 2014: Annual global climate and catastrophe report. Tech. rep., Aon Benfield.

MunichRe, 1993: Winterstürme in Europa, Publication of the Munich Re, Ordering Number 2041-e-d. Tech. rep., Munich Re.

MunichRe, 1999: Naturkatastrophen in Deutschland: Schadenerfahrungen und Schadenpotentiale, Publication of the Munich Re, Order Number 2798-e-d www.munichre.com. Tech. rep., Munich Re.

MunichRe, 2001: Winter storms in Europe (ii), Analysis of 1999 losses and loss potentials. Publication of Munich Reinsurance Company.

MunichRe, 2002: Winterstorms in Europe (iii), Analysis of loss potentials, Order Number at Munich Re: 302-03108 (German), 302- 03109 (English), 302-03110 (French) or www.munichre.com. Tech. rep., Munich Re.

Nelder, J. A. and R. Mead, 1965: A simplex algorithm for function minimization. *Computer Journal*, **7 (4)**, 308–313.

Nelson, R. B., 2006: *An Introduction to Copulas: 2^{nd} Edition*. Springer.

Neu, U., et al., 2013: Imilast: A community effort to intercompare extratropical cyclone detection and tracking algorithms. *Bulletin of the American Meteorological Society*, **94 (4)**, 529–547.

Norton, J., 2015: An introduction to sensitivity assessment of simulation models. *Envirnmental Modelling and Software*, **69**, 166–174.

Palutikof, J. P. and A. R. Skellern, 1991: Storm severity over Britain, a report to commercial union general insurance. Tech. rep., Climatic Research Unit, School of Environmental Sciences, University of East Anglia, Norwich (UK).

Panofsky, H. A., H. Tennekes, D. H. Lenschow, and J. C. Wyngaard, 1977: The characteristics of turbulent velocity components in the surface layer under convective conditions. *Boundary Layer Meteorology*, **11**, 355–361.

Pardo-Igúzquiza, E., 1998: Maximum likelihood estimation of spatial covariance parameters. *Mathematical Geology*, **30 (1)**, 95–108.

Pastres, R., K. Chan, C. Solidoro, and C. Dejak, 1999: Global sensitivity analysis of a shallow-water 3D eutrophication model. *Computer Physics Communications*, **117**, 62–74.

Payer, T. and H. Küchenhoff, 2004: Modelling extreme wind speeds at a German weather station as basic input for a subsequent risk analysis for high-speed trains. *Journal of Wind Engineering and Industrial Aerodynamics*, **92**, 241–261.

Perrin, O., H. Rootzén, and R. Taesler, 2006: A discussion of statistical methods used to estimate extreme wind speeds. *Theoretical Applied Climatology*, **85**, 203–215.

Pianosi, F. and T. Wagener, 2015: A simple and efficient method for global sensitivity analysis based on cumulutive distribution functions. *Envirnmental Modelling and Software*, **67**, 1–11.

Pianosi, F., T. Wagener, J. Rougier, J. Freer, and J. Hall, 2014: Sensitivity analysis of environmental models: A systematic review with practical workflow. *Vulnerability, Uncertainty, and Risk*, 290–299.

Pinto, J., E. Frohlich, G. Leckebusch, and U. Ulbrich, 2007: Changing European storm loss potentials under modified climate conditions according to an ensemble simulation of the ECHAM5/MPI-OM1 GCM. *Natural Hazards Earth Systems Science*, **7**, 165–175.

Pinto, J. G., M. K. Karremann, K. Born, P. M. Della-Marta, and M. Klawa, 2012: Loss potentials associated with European windstorms under future climate conditions. *Climate Research*, **54 (1)**, 1–20.

Pinto, J. G., C. P. Neuhaus, G. C. Leckebusch, M. Reyers, and M. Kerschgens, 2010: Estimation of wind storm impacts over Western Germany under future climate

conditions using a statistical-dynamical downscaling approach. *Tellus Series A: Dynamic Meteorology and Oceanography*, **62 (2)**, 188–201.

Renard, B. and M. Lang, 2007: Use of a Gaussian copula for multivariate extreme value analysis: some case studies in hydrology. *Advances in Water Resources*, **30**, 897–912.

Roberts, J. F., et al., 2014: The XWS open access catalogue of extreme windstorms in Europe from 1979 to 2012. *Natural Hazards and Earth System Science*, **14**, 2487–2501.

Saltelli, A., M. Ratto, T. Andres, F. Campolongo, J. Cariboni, D. Gatelli, M. Saisana, and S. Tarantola, 2008: *Global Sensitivity Analysis: A Primer*. Wiley.

Saltelli, A., S. Tarantola, R. Campolongo, and M. Ratto, 2004: *Sensitivity Analysis in Practice: A Guide to Assessing Scientific Models*. John Wiley and Sons, Ltd.

Schabenberger, O. and C. A. Gotway, 2005: *Statistical Methods for Spatial Data Analysis*. Chapman and Hall/CRC, Boca Raton, FL.

Schoelzel, C. and P. Friederichs, 2008: Multivariate non-normally distributed random variables in climate research - introduction to the copula approach. *Nonlinear Processes in Geophysics*, **15 (5)**, 761–772.

Schwierz, C., P. Köllner-Heck, E. Z. Mutter, D. N. Bresch, P. Vidale, M. Wild, and C. Schär, 2010: Modelling European winter wind storm losses in current and future climate. *Climate Change*, **101**, 485–514.

Segers, J., 2012: Max-stable model for multivariate extremes. *REVSTAT - Statistical Journal*, **10**, 61–82.

Seregina, L. S., R. Haas, K. Born, and J. G. Pinto, 2014: Development of a wind gust model to estimate gust speeds and their return periods. *Tellus A*, **66**, doi: 10.3402/tellusa.v66.22905.

Sheridan, P., 2011: Review of techniques and research for gust forecasting and parameterisation: Forecasting research technical report 570. Tech. rep., Met Office.

Shoji, T., 2006: Statistical and geostatistical analysis of wind: A case study of direction statistics. *Computers and Geosciences*, **32**, 1025 – 1039.

Sienz, F., A. Schneidereit, R. Blender, and K. Fraedrich, 2010: Extreme value statistics for North Atlantic cyclones. *Tellus A*, **62**, 347–360.

Sigma, 2004: No 1/2004: Many fatalities, comparatively moderate insured losses. Tech. rep., Swiss Reinsurance Copmany.

Sigma, 2006: No 2/2006: Natural catastrophes and man-made disasters in 2005. Tech. rep., Swiss Reinsurance Copmany.

Sigma, 2007: No 2/2007: Natural catastrophes and man-made disasters in 2006: Low insured losses. Tech. rep., Swiss Reinsurance Company.

Sigma, 2009: No 2/2009: Natural catastrophes and man-made disasters in 2008: North America and Asia suffer heavy losses. Tech. rep., Swiss Reinsurance Copmany.

Sigma, 2011: No 1/2011: Natural catastrophes and man-made disasters in 2010: A year of devastating and costly events. Tech. rep., Swiss Reinsurance Copmany.

Sigma, 2012: No 2/2012: Natural catastrophes and man-made disasters in 2011: Historic losses surface from record earthquakes and floods. Tech. rep., Swiss Reinsurance Copmany.

Sigma, 2013: No 2/2013: Natural catastrophes and man-made disasters in 2012: A year of extreme weather events in the US. Tech. rep., Swiss Reinsurance Company.

Singh, R., T. Wagener, R. Crane, M. E. Mann, and L. Ning, 2014: A vulnerability driven approach to identify adverse climate and land use change combinations for critical hydrologic indicator thresholds: Application to watershed in Pennsylvania, USA. *Water Resources Research*, **50**, 3409–3427.

Stein, M. L., 1999: *Interpolation of Spatial Data: Some Theory of Kriging*. Springer, New York.

Towe, R., E. Eastoe, J. A. Tawn, Y. Wu, and P. Jonathan, 2013: The extremal dependence of storm severity, wind speed and surface level pressure in the Northern North sea.

Wackernagel, H., 2003: *Multivariate Geostatistics: An Introduction with Applications*. Springer, New York.

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

Wadsworth, J. L., J. A. Tawn, A. C. Davison, and D. M. Elton, 2015: Modelling across extremal dependence classes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology) - Accepted/In press*.

Wagener, T., D. Boyle, M. Lees, H. Wheater, H. Gupta, and S. Sorooshian, 2001: A framework for development and application of hydrological models. *Hydrology and Earth System Sciences*, **5**, 13–26.

Warnes, J. J. and B. D. Ripley, 1987: Problems with likelihood estimation of covariance functions of spatial Gaussian processes. *Biometrika*, **74 (3)**, 640–642.

Zappa, G., L. C. Shaffrey, and K. I. Hodges, 2013: The ability of CMIP5 models to simulate North Atlantic extratropical cyclones. *Journal of Climate*, **26**, 5379–5396.

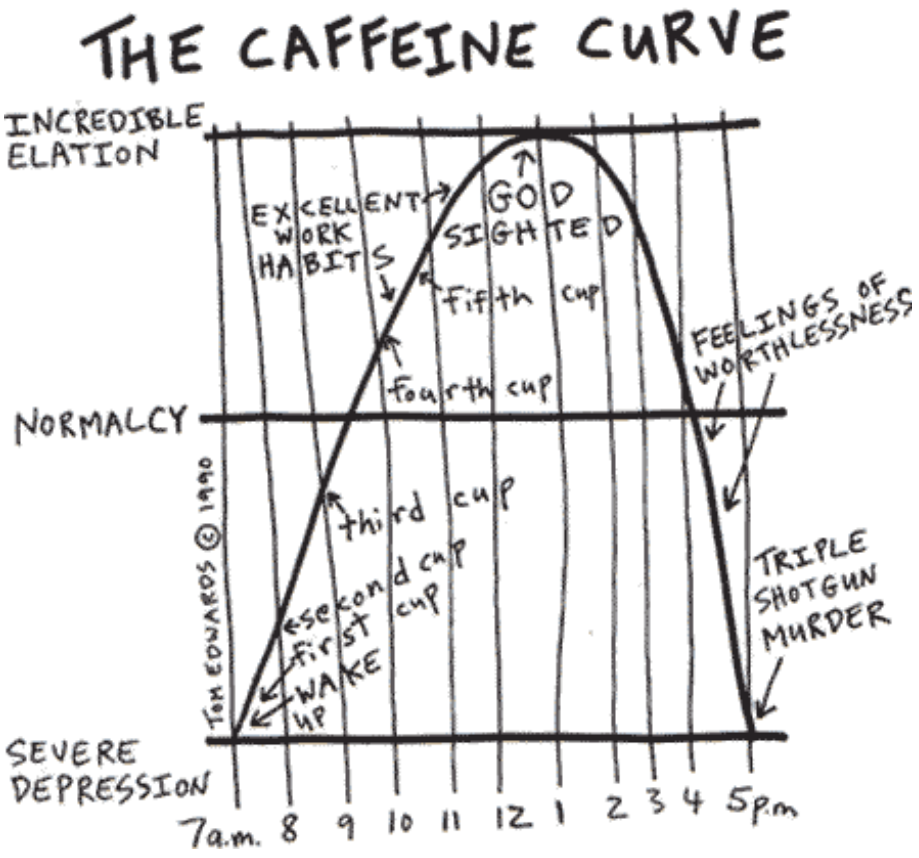
THE CAFFEINE CURVE
INCREDIBLE ELATION
NORMALCY
SEVERE DEPRESSION
EXCELLENT WORK HABITS
GOD SIGHTED
fifth cup
fourth cup
third cup
second cup
first cup
WAKE UP
FEELINGS OF WORTHLESSNESS
TRIPLE SHOTGUN MURDER
TOM EDWARDS © 1990
7a.m. 8 9 10 11 12 1 2 3 4 5p.m