## Methodology and forecast products for the optimal offering of ancillary services from wind in a market environment

Peder Bacher, Henrik Madsen

DTU Compute, Technical University of Denmark, DK-2800 Lyngby, Denmark

Pierre Pinson DTU Electro, Technical University of Denmark, DK-2800 Lyngby, Denmark

Stig Bousgaard Mortensen, Henrik Aalborg Nielsen, ENFOR A/S, DK-2970 Hørsholm, Denmark

July 15, 2014

### Contents

1	Intr	oduction	3
2	Gen	eralized extreme value distributions	5
	2.1	Non-stationary GEV distribution	7
	2.2	Maximum likelihood estimation	7
	2.3	Semi-parametric and non-parametric maximum likelihood estimation	7
3	Extr	eme value models for DK1 wind power forecast errors	9
0	3.1		9
	0.12	3.1.1 Wind power measurements	9
		3.1.2 Wind power forecasts	9
	3.2	Block maxima	13
		3.2.1 Selection of block length	13
	3.3	Single horizon models	16
		3.3.1 Model selection procedure	16
		3.3.2 Model for all observations	17
		3.3.3 Model for block maxima above a threshold of forecasted wind	
		power	22
	3.4	Semi-parametric models	27
		3.4.1 Model for all observations	28
		3.4.2 Model for block maxima above a threshold of forecasted wind	
		power	32
	3.5	Non-parametric model	35
	3.6	Comparison between the fitted models	36
	3.7	Forecasts	39
4	Extr	eme value models for DK2 wind power forecast errors	42
-	4.1	Data DK2	42
	1.1	411 Wind power measurements	42
		4.1.2 Wind power forecasts	42
	4.2	Block maxima	46
	1.2	4.2.1 Selection of block length	46
	4.3	Single horizon models	49
	1.0	4.3.1 Model for all observations	49
		4.3.2 Model for block maxima above a threshold of forecasted wind	
		power	54
	4.4	Semi-parametric models	58
		4.4.1 Model for all observations	58
	4.5	Non-parametric model	64
	4.6	Comparison between the fitted models	64
	4.7	Forecasts	66

5	Dise	cussion and conclusion	69
	5.1	Discussion	69
	5.2	Further work	70
	5.3	Conclusion	71

# Chapter 1 Introduction

In this report models for extreme negative wind power forecast errors are presented. The models can be applied to estimate levels below which the wind power rarely drops. Such levels could be call called "certain-levels" or "guaranteed levels" of wind power, well knowing that full guarantee never can be given. The estimated levels are obtained by modelling the error from already existing wind power forecasting software, this is accomplished by modelling the residuals with statistical extreme value techniques.

The forecasts can be used in the operation of power systems with significant amounts of wind power for example in the planning of ancillary power services, where the level of available wind power with a high degree of certainty is important to know.

The presented extreme value models are applied to negative forecast residuals from state-of-the-art wind power forecast software. This enables the estimation of return levels below which the extreme wind power forecast error events occur only at a specified rate, e.g. once a month or once every year. The techniques allows extrapolation beyond the available data period. In the study data from 1.5 years is used. It consists of hourly wind power production in the two regions of Denmark (DK1 and DK2) and corresponding wind power forecasts. The wind power forecasts are generated using the software WPPT<sup>1</sup>. and are based on the outcome of three numerical weather prediction models. They cover horizons from 1 to 42 hours ahead in time and are updated each hour.

In the report a range of extreme value models are suggested. They are of increasing complexity and a model selection is carried using statistical measures and test. A normal procedure when building forecasting models is to divide the data into a learning and a test set to cross-validate the results in order to avoid over-fitting the models. This is hardly ever possible for extreme value analysis, instead the model selection and evaluation sole rely on statistical techniques such as correlation measures, likelihood ratio-tests and information about uncertainty, for example in the form of confidence bands on parameter estimates and predictions.

The foundation to statistical extreme value theory was set by Fisher and Tippett in 1928 and later developed by Gumbel (1958). Since then it has been used for modelling extremes in a wide range of applications. Typical applications are for estimation of extreme weather induced phenomena, for example extreme water

pbac: More on applications, how can it be used. Ancillary reserve power services

<sup>&</sup>lt;sup>1</sup>www.enfor.dk/products/wppt.html

levels in a river, wind levels or at sea for design of dykes (de Haan and de Ronde, 1998). In insurance and finance the extreme value modelling is widespread (Embrechts et al., 1997). Extreme value statistics for energy and power applications is also widely used, for example for planning in wind power operation (Horvat et al., 2013) and peak wind prediction (Cook, 1982) and (Friederichs and Thorarinsdottir, 2012). Several books provide comprehensive introductions to extreme value theory, for example Coles (2001) and Beirlant et al. (2006). A really good overview of available extreme value analysis software is given by (Stephenson and Gilleland, 2005). In the present study the R R Core Team (2013) package extRemes Gilleland and Katz (2011) is used.

In the following chapter a very short introduction to extreme value statistics is given. This is following by a chapter in which the data, the models, model selection, result evaluations, and finally examples of applications for wind power in DK1 are presented. Next chapter is similar, just DK2. The final chapter consists of a discussion of the results, suggestions for further work and a conclusion.

### Chapter 2

# Generalized extreme value distributions

In this section the statistical extreme value theory used in the modelling are described. All the models in this study are based on the generalized extreme value (GEV) distribution, which has the cumulative distribution function (CDF)

$$F_{Y}(y;\mu,\sigma,\xi) = \exp\left\{-\left[1+\xi\left(\frac{y-\mu}{\sigma}\right)\right]^{-1/\xi}\right\}$$
(2.1)

where y is a realization of the stochastic variable Y, which is formed by taking maximum of blocks

$$Y = \max(Z_1, Z_2, \dots, Z_N) \tag{2.2}$$

where  $Z_i$  are i.i.d. stochastic variables and the three parameters of the distribution are

- the location  $\mu \in \mathbb{R}$ ,
- the scale  $\sigma > 0$ ,
- the shape  $\xi \in \mathbb{R}$ .

By the extreme value theorem the GEV distribution is the only possible limit distribution of properly normalized maxima of a sequence of independent and identically distributed random variables. The notation used for the models is

$$Y \sim \text{GEV}(\mu, \sigma, \xi) \tag{2.3}$$

indicating that the stochastic variable follows the GEV. The probability density function (PDF) of the GEV is

$$f_{Y}(y;\mu,\sigma,\xi) = \frac{1}{\sigma} \left[ 1 + \xi \left( \frac{y-\mu}{\sigma} \right) \right]^{(-1/\xi)-1} \exp\left\{ - \left[ 1 + \xi \left( \frac{y-\mu}{\sigma} \right) \right]^{-1/\xi} \right\}$$
(2.4)

Too see how the parameters influence the PDF it is plotted for different values varying the parameters one at a time in Figure 2.1. In the upper plot it is seen how the location parameter simply displace the PDF, but doesn't change its shape. In the middle plot it is seen how the scale parameter affects the PDF more narrow or wide. Finally, in the lower plot the effect of the scale parameter on the PDF. It



Figure 2.1: The PDF of the GEV for different values of the parameters. In the upper plot the location parameter is varied, in the middle plot the scale parameter is varied, and in the lower plot the shape parameter is varied.

changes the PDF rather drastically and actually this parameter makes either the PDF bounded upwards (negative shape value), not bounded (shape is zero), or bounded downwards (positive shape). Clearly, also interactions between the scale and the shape parameters will affect the shape of the PDF. For more details see (Coles, 2001) or (Beirlant et al., 2006).

#### 2.1 Non-stationary GEV distribution

In a non-stationary GEV distribution the parameter values in the GEV are modelled as a function of co-variates. Including dependence in the shape parameter led to very unstable results and therefore it was is kept constant in all the models. In the applied models the location and scale parameters

$$Y \sim GEV(f_{\mu}(m{x}) \,,\, f_{\sigma}(m{x}) \,,\, \xi)$$

are modelled as functions of co-variates *x*. Thus the PDF becomes

$$f_{Y}(y; f_{\mu}(x), f_{\sigma}(x), \xi) = \frac{1}{f_{\sigma}(x)} \left[ 1 + \xi \left( \frac{x - f_{\mu}(x)}{f_{\sigma}(x)} \right) \right]^{(-1/\xi) - 1}$$

$$\exp \left\{ - \left[ 1 + \xi \left( \frac{x - f_{\mu}(x)}{f_{\sigma}(x)} \right) \right]^{-1/\xi} \right\}$$
(2.5)

The functions in the applied models are polynomials of up to second order of a single input

$$f_{\mu}(x) = \mu_0 + \mu_1 x + \mu_2 x^2$$

$$f_{\sigma}(x) = \sigma_0 + \sigma_1 x + \sigma_2 x^2$$
(2.6)

#### 2.2 Maximum likelihood estimation

The parameters are estimated using maximum likelihood techniques. The likelihood function is formed by

$$L(\boldsymbol{\theta}; \mathcal{Y}) = \prod_{i=1}^{N} f_{Y}(y_{i}; f_{\mu}(\boldsymbol{x}_{i}), f_{\sigma}(\boldsymbol{x}_{i}), \xi)$$
(2.7)

where  $\mathcal{Y}$  denotes all the observed data, i.e. a series of block maxima  $\{y_i\}$  and corresponding co-variates  $\{x_i\}$  and  $\theta$  are all the parameters which are estimated. The maximum likelihood estimates are obtained by minimizing the negative log-likelihood, which is implemented in the used R software package extRemes (Gilleland and Katz, 2011).

#### 2.3 Semi-parametric and non-parametric maximum likelihood estimation

Semi- and non-parametric extreme value models (see (Kotz and Nadarajah, 1988) and Heffernan and Tawn (2004)) are fitted using weighted maximum likelihood estimation (Coles, 2001). The models are fitted local in one or several dimensions. In the applied semi-parametric models they are fitted locally in the dimension of the horizon. This means that for a forecast horizon  $k_f$  the observations from the all horizons k are included and that the observations are weighted down with the distance  $\Delta k_f = k_f - k$ , such that more weight is put to the observations close the

forecast horizon. Note that this means that observations from all horizons are included in the likelihood, hence the likelihood in Equation (2.7) is extended to include all horizons and furthermore a part is added to achieve the weighting. The weights are calculated using an Epanechnikov kernel function

$$w(\Delta k) = \frac{3}{4h} (1 - \frac{\Delta k^2}{h^2}) I_{\{\Delta k \le h\}}$$
(2.8)

where *h* is the bandwidth and  $I_{\{\Delta k \le h\}}$  is the indicator function.

In the non-parametric models the model is furthermore fitted locally in the covariates, by calculating the weights also as a function of, here a single co-variate, by using a multiplicative kernel

$$w(\Delta k, \Delta x) = \frac{3}{4h_{k}} (1 - \frac{\Delta k^{2}}{h_{k}^{2}}) I_{\{\Delta k \le h_{k}\}} \cdot \frac{3}{4h_{x}} (1 - \frac{\Delta x^{2}}{h_{x}^{2}}) I_{\{\Delta x \le h_{x}\}}$$
(2.9)

where  $\Delta x = x_f - x$  is the distance in the co-variate dimension. For more details, see Section 3.4 and 3.5 where the applied semi- and non-parametric models described.

## **Chapter 3**

# Extreme value models for DK1 wind power forecast errors

In this chapter extreme value models for DK1 wind power forecasts are presented, applied and evaluated.

#### 3.1 Data DK1

In this section an exploratory data analysis is carried out to give an overview of the data for DK1 used in the modelling and to point out some relevant features for the modelling carried. The data used are covering 1.5 years and consist of hourly measurements of wind power in DK1 and corresponding wind power forecasts.

#### 3.1.1 Wind power measurements

The time series of hourly wind power measurements for DK1 are denoted by

$$\{P_t: t = 1, \dots, N_h\}$$
 (3.1)

where *t* is the hours since 2011–07–01 00:00:00 UTC and  $N_{\rm h} = 14159$  and the last time point is 2013–02–09 23:00:00 UTC. The time marks the end of the hour. In Figure 3.1 the series is plotted together with the rated wind power level for DK1. It is seen that the rated power level increase slightly over the period and that a few periods are missing: in total 449 values are missing (around 2.5 week).

#### 3.1.2 Wind power forecasts

Based on numerical weather predictions (NWPs) and the wind power measurements, wind power forecasts using the software WPPT <sup>1</sup> has been generated for DK1. Three different NWPs was used as input to the WPPT . The measurements are treated as online (available with very little delay), such that the current value of measured wind power can be taken into account when generating a new forecast. This assumption was found reasonable, since in an operational scenario where auxiliary services are provided from wind turbines then a fast data connection and on-line communication will be available. The forecasts are updated every hour

pbac: Ref. to a paper??

<sup>&</sup>lt;sup>1</sup>www.enfor.dk/products/wppt.html



Figure 3.1: The time series of hourly wind power and rated wind power in DK1.

and are up to a horizon of 42 hours ahead. The wind power forecasts are denoted by

$$\{\hat{P}_{t|t-k}: t = 1, \dots, N_h\}$$
 (3.2)

Scatter plots for nine selected k hours horizons: 1,2,3,6,12,18,24,32 and 48 hours, are shown in Figure 3.2. It is seen that the scatter (the level of residuals) does seem to increase with the horizon, which is anticipated, however interestingly the most extreme event (the event with the highest negative residual) at the 18,24 and 32 hours horizons is more extreme than the most extreme event at the 42 hours horizon.

In the modelling the residuals for each horizon are used. They are

$$\epsilon_{k,t} = P_t - \hat{P}_{t|t-k} \tag{3.3}$$

for the *k*′th horizon.

In Figure 3.3 the negative residuals (i.e.  $-\epsilon_{k,t}$ ) are plotted for four selected horizons, k = 1, 3, 24 and 42 hours. It is clearly seen how the high negative residuals are correlated in time, such that an extreme event consists of more than high negative residual. Again it is seen how the largest most extreme event occurs at the 24 hours horizon.



Figure 3.2: DK1 wind power measurements versus forecasts for nine selected horizon *k* hours.



Figure 3.3: The DK1 negative residuals versus the wind power forecasts. The points are connected with lines to their neighboring points in time (i.e. the previous and next hour).

#### 3.2 Block maxima

In this section it is described how the block maxima are computed. The block maxima are taken of the negative residuals by

$$\{v_{k,i} = \max(-\epsilon_{k,(i-1)24n+1}, -\epsilon_{k,(i-1)24n+2}, \dots, -\epsilon_{k,(i-1)24n+24n}): i = 1, 2, \dots, N\}$$
(3.4)

where *n* is the block length in days and  $N = N_h/(24n)$  is the number of blocks. The time *t* at which the maximum  $v_{k,i}$  occurred within the block is denoted with  $t_i$ . The forecasted wind power for horizon *k* at time  $t_i$  is noted by  $P_{t_i|t_i-k}$ , i.e. the value at the time at which the block maximum occurred.

In Figure 3.4 the block maxima for three selected horizons and block lengths are plotted versus the forecasted wind power. Of course as the block length increases the number of block maxima decreases, hence a too long block length will decrease the information and hence increase the model variance. Also clear is it that the block maxima reach the upper bound (the one-to-one function, marked with the red dashed line on the plots) for forecasted wind power in the lower range.

#### 3.2.1 Selection of block length

It is important to select a suitable block length to balance model bias (to short block lengths) and variance (to long block lengths). The length of the blocks are found by fitting a GEV distribution model for each horizon for block lengths between 1 day to 15 days. The model which is fitted is

$$v_{k_{\rm f}} \sim GEV(\mu_0 + \mu_1 P_{t|t-k_{\rm f}}, \sigma_0 + \sigma_1 P_{t|t-k_{\rm f}}, \xi)$$
 (3.5)

where  $v_{k_f}$  is the block maximum random variable for forecast horizon  $k_f$ ,  $P_{t|t-k_f}$  is the forecasted wind power and the parameters estimated are  $\theta = (\mu_0, \mu_1, \sigma_0, \sigma_1, \xi)$ . The location of the GEV is modelled as linear effect (1st order polynomial) of forecasted wind power, the scale also as a linear of the forecasted wind power and the shape as a constant. Note that for each horizon only the only forecasted wind power for that particular horizon k used. This model is a relatively complicated model, it is used since, as described in Section 3.4.1, it is found as a suitable model for most horizons. The plots in Figure 3.5 are of the parameter estimates for each horizon k as a function of the block length. Also included in the plots as red lines are local regression estimates fitted using the R function loess() with span=0.7 (R Core Team, 2013). Considering the local regression estimates (red lines) and starting from the upper left plot of the location intercept  $\mu_0$ . It is seen that naturally  $\mu_0$ increases with the block length. The location slope  $\mu_1$  increases for block lengths below 5 days and stays around a constant level until a block length of 12 days. The scale intercept  $\sigma_0$  and slope  $\sigma_1$  decrease until a block length around 7 days from where they decrease at a slower rate. Finally, the shape  $\xi$  also reach a rather constant level around a block length of 6 to 8 days. From these considerations it is found that for block lengths below around 7 days the parameter estimates (except the location) are more biased (i.e. more dependent on the block length) than above, where the variance increase. Hence a block length of 7 days is found suitable for obtaining the best balance between bias and variance and this block length is used in the remaining of the study.



Figure 3.4: DK1 block maxima for three selected horizons *k* and block lengths *n*.



Figure 3.5: The parameters of the model in Equation (3.5) fitted for each horizon and for each block lengths n = 1 to n = 15 days. The red lines are local regression estimates.

#### 3.3 Single horizon models

In this section, first a procedure for model selection is presented and then it is applied to select models fitted to a single horizon. The objective is to identify suitable models for each horizon. Two approaches to the fitting the models are found reasonable

- with keep all block maxima
- or remove block maxima in the lower range of forecasted wind power, where they clearly reach the upper bound.

For each approach a model selection is carried out.

#### 3.3.1 Model selection procedure

Models are selected using a forward selection procedure based on likelihood-ratio tests, similar to the procedure applied in (Bacher and Madsen, 2011), for more details see Madsen and Thyregod (2010). The implementation in the R package extRemes (Gilleland and Katz, 2011) are used to carry out for the maximum likelihood estimation. First the simplest model with no dependence is fitted

$$v_{k_{\rm f}} \sim GEV(\mu_0, \sigma_0, \xi) \tag{3.6}$$

where  $v_{k_f}$  is the block maximum random variable for horizon  $k_f$  and then in each step in the model selection the order of the polynomials for the location and scale are increased, up to a second order effect for each parameter

$$v_{k_{\rm f}} \sim GEV(\mu_0 + \mu_1 P_{t|t-k_{\rm f}} + \mu_2 P_{t|t-k_{\rm f}}^2,$$

$$\sigma_0 + \sigma_1 P_{t|t-k_{\rm f}} + \sigma_2 P_{t|t-k_{\rm f}}^2,$$
(3.7)

where  $P_{t|t-k_f}$  is the forecasted wind power. Note that for the single horizon models only values for the horizon k are used, for the extended models (semi- and nonparametric) presented values for more horizons are used and the model above is exchanged with the particular second order model. Note also that a dependence of the shape parameter is not considered in any of the models, this was tried but the results obtained by adding this dependence gave very high variance and uncertainty of predictions and it is generally recommended not to include dependence of the shape parameter.

The model selection procedure is carried out with no manual interaction. First a zero order models for each parameter is fitted. Then in each iteration:

- a model where the order of the location polynomial is increased by one, and a model where the order of the scale polynomial is increased by one, are fitted.
- The model with the highest likelihood is tested against the model selected in last iteration using a log-likelihood ratio test (Madsen and Thyregod, 2010).
- If the test has a significant *p*-value (below 0.05) then the extended model is selected and a new iteration is started.
- If the test has a non-significant *p*-value, then the model selection is stopped and model is kept (i.e. the model selected in the last iteration, not the extended model).

This is carried out for each horizon *k*.

Horizon	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Location order	1	1	2	1	2	1	2	1	2	2	2	1	2	1
Scale order	0	1	1	1	1	1	1	1	2	1	1	1	2	2
Horizon cont.1	15	16	17	18	19	20	21	22	23	24	25	26	27	28
Location order cont.1	1	2	1	2	2	2	2	2	1	1	1	1	1	1
Scale order cont.1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
Horizon cont.2	29	30	31	32	33	34	35	36	37	38	39	40	41	42
Location order cont.2	2	1	1	2	2	1	2	2	2	2	2	2	1	1
Scale order cont.2	1	1	0	1	1	1	2	2	2	2	2	2	1	1

Table 3.1: The order selected for the location and scale polynomials for each horizon using the forward selection procedure.

#### 3.3.2 Model for all observations

A model is selected for each horizons for all observations. The orders selected for the location and scale polynomials for the model in Equation (3.5) for each horizon k are shown in Table 3.1. It is apparent that there is some pattern in the selected orders, however the pattern is not clearly changing from e.g. a low order for short horizons to a higher order for long horizons. This does indicate that applying the selected model for each horizon will result in a high variation between the horizons.

In Figure 3.6 model diagnostic plots are shown for four horizons, where the model was fitted using the selected polynomial orders. The QQ and density plots indicates that the model follows the observations reasonably well, only the 2 to 5 most extreme values deviate some from the model fit. They are slightly higher than the fit, which this indicates that extrapolation using the model fits (i.e. prediction of return levels longer than the current data period) could tend to be biased to the lower side. Furthermore it is seen from the return level plots for the longer horizons (for 18 and 32 hours), that most of the observations at the lower level of forecasted wind power reach the upper bound. This influences the fit. In this range the higher quantiles of the fitted GEV will follow the upper bound (the one-to-one function) and the second order allows it to "bend" away from the upper bound.

In Figure 3.7 the predicted one year return level of for each horizon as a function of the forecasted wind power is plotted as a contour plot with colors indicating the predicted level. It is clearly seen that the predicted return level varies up and down depending on the horizon, especially for the higher levels of forecasted wind power the level varies up and down as the horizon increases. Considering the horizons where the polynomial orders were selected differently, as seen in Table 3.1, for example for horizon: 33,34,41,42 versus 35 to 40, then it is clear that the model order has a high influence of the predicted return levels (especially for high levels of forecasted wind power).

A reasonable assumption would be that the return level is an increasing function of the horizon. One way of decreasing the dependence of the horizon and thus obtain less variance is to apply the same order of the polynomials for all horizons. Considering the selected polynomial orders in Table 3.1 it is found reasonable to apply

a second order polynomial for location and a first order for the scale parameter

$$v_{k_{\rm f}} \sim GEV(\mu_0 + \mu_1 P_{t|t-k_{\rm f}} + \mu_2 P_{t|t-k_{\rm f}}^2,$$

$$\sigma_0 + \sigma_1 P_{t|t-k_{\rm f}}, \xi)$$
(3.8)

which is fitted to all horizons separately. In Figure 3.8 the predicted one year return level using this model is plotted as a contour plot. Compared with the return levels predicted using the selected model for each horizon (in Figure 3.7) less variation is seen and the overall it is more clear that a increasing trend of the return level as function of the horizon.



Figure 3.6: Model diagnostic plots for four selected horizons for model for all observations. Left column are QQ-plots of the empirical distribution versus the fitted distribution. Mid column plots is an estimated empirical density (fitted using a local kernel regression) and the fitted (modeled) distribution. Right column plots are of the estimated one and five year return levels including 95% confidence bands calculated using the delta method (Coles, 2001) from the covariance of the parameters, which is derived from the Fisher information matrix.



Figure 3.7: The single horizon model for observations above a threshold, fitted with the selected orders of the polynomials. Predicted one year return level for each horizon.



Figure 3.8: Predicted one year return level using the same polynomial orders for all horizons: a second order for location and a first order for the scale parameter.

# 3.3.3 Model for block maxima above a threshold of forecasted wind power

Removing the block maxima below a threshold of forecasted wind power where they reach the upper bound could be a reasonable approach to ensure that the models are not biased for higher levels of forecasted wind power. The threshold is set for each horizon separately, simply by a visual inspection of scatter plots to find the level where the block maxima gets close to the upper bound. In Figure 3.9 the threshold used as a function of the horizon is plotted.

The forward model selection (see the description in Section 3.3.1) is carried out to select a suitable model for each horizon. The selected polynomial orders for the location and scale parameters are listed in Table 3.2. It is apparent that there is some pattern in the orders selected, however as before the pattern is not clearly changing from e.g. a low order for short horizons to a higher order for long horizons. Compared to the orders selected for all observations (see Table 3.1) the orders are lower, hence that less significant non-linear dependence of the forecasted wind power is found. This is natural since now the model doesn't need to take into account both the distribution at lower levels of forecasted wind power levels of forecasted wind power (reaching the upper bound) and doesn't need to "bend" down at higher levels of forecasted wind power. However fewer observations are included when fitting the models.

In Figure 3.10 model fit diagnostic plots for four horizons are shown. From the plots for the shorter horizons (1 and 6 hours) it is seen that the fits are quite close to the fits for all observations (compare the Figure 3.6). For the longer horizons (18 and 32 hours) it is seen that since fewer observations are used, then less information is available and the fits are more affected by the highest block maximum, especially the most extreme value at 32 hours. This leads to a less close match between the empirical and modelled density and the confidence bands becomes wider.

In Figure 3.11 the predicted one year return level as a function of the horizon and forecasted wind power using is shown. Again a rather high variance is seen caused by different selected orders as for the similar plot for all observations (plot in Figure 3.7). Now also the high variance is due to fewer observations and thus less

Horizon	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Location order	1	1	0	2	1	1	1	1	0	0	0	0	0	1
Scale order	0	1	0	1	1	1	1	1	0	1	1	2	1	1
Horizon cont.1	15	16	17	18	19	20	21	22	23	24	25	26	27	28
Location order cont.1	0	1	1	1	1	0	0	1	1	1	0	0	1	1
Scale order cont.1	1	1	1	1	1	1	1	1	1	1	0	0	0	1
Horizon cont.2	29	30	31	32	33	34	35	36	37	38	39	40	41	42
Location order cont.2	1	0	1	0	1	1	1	1	1	1	1	0	0	0
Scale order cont.2	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Table 3.2: The order selected for the location and scale polynomials for each horizon using the forward selection procedure for the single horizon model for observations above a threshold of forecasted wind power.



Figure 3.9: The thresholds, as function of the horizon, below which the block maxima are removed.

information.

Again to decrease the variance a model with the same orders is fitted to all horizons. From the orders selected for each horizon as shown in Table 3.2 it is found reasonable to use a first order polynomial for the location and a zero order polynomial for the scale parameters. The variation of the predicted one year return level is, compared to the predictions using the selected orders for each horizon decreased (compare to the similar plot in Figure 3.11) and the level is more steadily increasing with the horizon. However, especially at the longest horizons the variance is still quite high.



Figure 3.10: Model diagnostic plots for four selected horizons for the single horizon model for observations above a threshold. Left column are QQ-plots of the empirical distribution versus the fitted distribution. Mid column plots is an estimated empirical density (fitted using a local kernel regression) and the fitted (modeled) distribution. Right column plots are of the estimated one and five year return levels including 95% confidence bands calculated using the delta method (Coles, 2001) from the covariance of the parameters, which is derived from the Fisher information matrix.



Figure 3.11: Predicted one year return level for each horizon for the single horizon model for observations above a threshold of forecasted wind power and fitted with the selected polynomial orders.



Figure 3.12: Predicted one year return level for the single horizon model for observations above a threshold using the same polynomial orders for all horizons: a first order for location and a zero order for the scale parameter.

#### 3.4 Semi-parametric models

In this section semi-parametric models are described and fitted to the observations. The semi-parametric models are extensions of the models applied for the single horizons (see Equation (3.7) in the previous section). The extensions are by formed by including observations from neighboring horizons using weighted maximum likelihood (Coles, 2001). Hence the models are set up such that the horizon k is a dimension and the models are fitted locally in this dimension. This allows for the parameters to change as smooth functions of the horizon and more observations are included in the fit for each horizon. Thereby the variation is decreased. The drawbacks are that one more parameter must be set, namely the bandwidth of the kernel function used for the local weighting, and that bias is introduced especially at the boundaries, e.g. at 1 and 42 hours horizons. The weighted maximum like-lihood is implemented in the R package extRemes, where it is straight forward the fit the models.

The applied semi-parametric models are

$$v \sim GEV \Big( \mu \big( P_{t|t-k_{\mathrm{f}}}, k_{\mathrm{f}} \big)$$
 ,  $\sigma \big( P_{t|t-k_{\mathrm{f}}}, k_{\mathrm{f}} \big)$  ,  $\xi \Big)$ 

where v is the block maximum random variable,  $P_{t|t-k_f}$  is the forecasted wind power,  $k_f$  is the forecast horizon,  $\mu(P_{t|t-k_f}, k_f)$  and  $\sigma(P_{t|t-k_f}, k_f)$  are semi-parametric smooth functions for the location and the scale parameter, respectively. They are called semi-parametric because they are fitted as parametric functions in the forecasted wind power dimension and locally (non-parametric) in the horizon dimension. This is achieved by calculating the likelihood for a given horizon  $k_f$  with the density function

$$f_{v}\left(\mu_{f}\left(P_{t|t-k_{\mathrm{f}}},k_{\mathrm{f}},\boldsymbol{\mu}\right),\ \sigma_{f}\left(P_{t|t-k_{\mathrm{f}}},k_{\mathrm{f}},\boldsymbol{\sigma}\right),\ \boldsymbol{\xi}\right)$$
(3.9)

where the location function is

$$\mu_f (P_{t|t-k_f}, k_f, \mu) = \mu_0 + \mu_1 P_{t|t-k_f} + \mu_2 P_{t|t-k_f}^2 + \mu_3 \Delta k + \mu_4 \Delta k^2$$
(3.10)

and the scale function is

$$\sigma_f \left( P_{t|t-k_f}, k_f, \sigma \right) = \sigma_0 + \sigma_1 P_{t|t-k_f} + \sigma_2 P_{t|t-k_f}^2 + \sigma_3 \Delta k + \sigma_4 \Delta k^2$$
(3.11)

where  $\Delta k = k - k_f$  is the distance in the horizon dimension. The fits becomes local second order polynomial fits in the horizon dimension by down-weighting the observations with the distance  $\Delta k$  using an Epanechnikov kernel

$$w_{k}(\Delta k) = \frac{3}{4h} (1 - \frac{\Delta k^{2}}{h^{2}}) I_{\{\Delta k \le h\}}$$
(3.12)

where *h* is the bandwidth and  $I_{\{\Delta k \le h\}}$  is the indicator function. A plot of the kernel for a few values of the bandwidth 5 is shown in Figure 3.13. For the full model the parameters which are fitted with the maximum likelihood are  $\theta = (\mu_0, \ldots, \mu_4, \sigma_0, \ldots, \sigma_4, \xi)$ , note that these are fitted for each  $k_f$  which is left out in the notation for simplification. Note also that a second order polynomial is always used for the local fit in the horizon dimension (i.e.  $\mu_3, \mu_4, \sigma_3$  and  $\sigma_4$  is always included in the model), but the order of the polynomials for the forecasted wind power is selected using the forward model selection procedure as described in the following section.



Figure 3.13: Examples of the kernel function for three different bandwidths *h*.

1	2	3	4	5	6	7	8	9	10	11	12	13	14
2	2	2	2	2	2	2	2	2	1	1	1	1	1
1	1	1	1	1	1	1	2	2	2	2	2	2	2
15	16	17	18	19	20	21	22	23	24	25	26	27	28
1	1	1	2	2	2	2	2	2	2	2	1	1	1
2	2	2	1	1	1	1	1	1	1	2	2	2	2
29	30	31	32	33	34	35	36	37	38	39	40	41	42
1	1	1	1	1	1	1	1	2	2	2	2	2	2
2	2	2	2	2	2	2	2	2	2	2	2	2	2
	1 2 1 5 1 2 29 1 2	1       2         2       2         1       1         15       16         1       1         2       2         29       30         1       1         2       2         29       30         1       1         2       2	1       2       3         2       2       2         1       1       1         15       16       17         1       1       1         2       2       2         29       30       31         1       1       1         2       2       2	1       2       3       4         2       2       2       2         1       1       1       1         15       16       17       18         1       1       1       2         2       2       2       1         29       30       31       32         1       1       1       1         29       2       2       2	1       2       3       4       5         2       2       2       2       2         1       1       1       1       1         15       16       17       18       19         1       1       1       2       2         2       2       2       1       1         29       30       31       32       33         1       1       1       1       1         2       2       2       2       2	1       2       3       4       5       6         2       2       2       2       2         1       1       1       1       1         15       16       17       18       19       20         1       1       1       2       2       2         2       2       2       1       1       1         29       30       31       32       33       34         1       1       1       1       1       1         29       20       2       2       2       2       33	1       2       3       4       5       6       7         2       2       2       2       2       2       2         1       1       1       1       1       1       1         15       16       17       18       19       20       21         1       1       1       2       2       2       2         2       2       2       1       1       1       1         15       16       17       18       19       20       21         1       1       1       2       2       2       2       2         2       2       2       1       1       1       1         29       30       31       32       33       34       35         1       1       1       1       1       1       1         2       2       2       2       2       2       2       2	1       2       3       4       5       6       7       8         2       2       2       2       2       2       2       2         1       1       1       1       1       1       1       2         1       1       1       1       1       1       1       2         15       16       17       18       19       20       21       22         1       1       1       2       2       2       2       2       2         2       2       2       1       1       1       1       1         29       30       31       32       33       34       35       36         1       1       1       1       1       1       1       1         29       30       31       32       2       2       2       2       2         1       1       1       1       1       1       1       1	1       2       3       4       5       6       7       8       9         2       2       2       2       2       2       2       2       2         1       1       1       1       1       1       1       2       2         1       1       1       1       1       1       1       2       2         15       16       17       18       19       20       21       22       23         1       1       1       2       2       2       2       2       2         2       2       2       1       1       1       1       1       1         2       2       2       33       34       35       36       37         1       1       1       1       1       1       1       1       2       2         2       2       2       2       2       2       2       2       2       2	1       2       3       4       5       6       7       8       9       10         2       2       2       2       2       2       2       2       1         1       1       1       1       1       1       2       2       2       2       1         1       1       1       1       1       1       1       2       2       2         15       16       17       18       19       20       21       22       23       24         1       1       1       2 <td>1       2       3       4       5       6       7       8       9       10       11         2       2       2       2       2       2       2       2       2       1       1         1       1       1       1       1       1       1       2       2       2       1       1         1       1       1       1       1       1       1       2</td> <td>1       2       3       4       5       6       7       8       9       10       11       12         2       2       2       2       2       2       2       2       1       1       1         1       1       1       1       1       1       2       1&lt;</td> <td><math display="block"> \begin{array}{c ccccccccccccccccccccccccccccccccccc</math></td>	1       2       3       4       5       6       7       8       9       10       11         2       2       2       2       2       2       2       2       2       1       1         1       1       1       1       1       1       1       2       2       2       1       1         1       1       1       1       1       1       1       2	1       2       3       4       5       6       7       8       9       10       11       12         2       2       2       2       2       2       2       2       1       1       1         1       1       1       1       1       1       2       1<	$ \begin{array}{c ccccccccccccccccccccccccccccccccccc$

Table 3.3: The order selected for the location and scale polynomials for each horizon using the forward selection procedure for the semi-parametric model for all observations.

#### 3.4.1 Model for all observations

First the model selection described in Section 3.3.1 is carried out for the semiparametric model for each horizon. The order of the location and scale polynomials are selected. The bandwidth is set manually to a value which used for all horizons. This is carried out by predicting one year return levels and finding a bandwidth as low as possible such that the predicted return level is an increasing function of the horizon for almost all levels of forecasted wind power. It is a matter of balancing bias and variance. Using a bandwidth of 1 hour (makes exactly the same models as fitted for single horizons in the previous section) results in a too much variance. A too high bandwidth results in too much bias. From a bandwidth of 3 to 5 hours and up to 12 to 15 hours not much difference was found of predicted return levels, hence a value

$$h = 5 \tag{3.13}$$

was used.

In Table 3.3 the order of the location and scale polynomials selected for each horizon for the semi-parametric model for all observations. The orders selected are

now much more similar for all horizons compared to the orders selected for the single horizon model. This is natural since the fits for each horizon now depends on each other. Considering the orders selected it is found that using a second order polynomial for the location and a first order polynomial for the scale provides is a good compromise for all horizons, and this model is applied in the following.

Diagnostic plots for four selected horizons are shown in Figure 3.10. Note that the size of the points in the rightmost plots indicates the weight (from the kernel) put on the observation, however this is not the case for the left-most plots (QQ-plots). Considering the QQ-plots, it seems like the fits does deviate from especially the most extreme values, however since the observations are not weighted equally and are correlated since the most extreme values originate from the same extreme event. This can also be seen on points on the right column plots, where the most extreme values are located in groups with points having different size (weight) since they originate from different horizons. Again it is clearly seen how the return levels follow the upper bound for lower levels of forecasted wind power, and at high levels the return levels "bends" down from the bound.

In Figure 3.15 the predicted one year return level as a function of the forecasted wind power and horizon is plotted. The variation between the horizons has decreased compared to the plot for the similar single horizon model seen in Figure 3.8 and the overall levels are very alike, except for the very short horizons, where the semi-parametric model clearly is biased (which will also become apparent in a comparison later in the report).



Figure 3.14: Model diagnostic plots for four selected horizons for semi-parametric model for all observations. Left column are QQ-plots of the empirical distribution versus the fitted distribution. Mid column plots is an estimated empirical density (fitted using a local kernel regression) and the fitted (modeled) distribution. Right column plots are of the estimated one and five year return levels including 95% confidence bands calculated using the delta method (Coles, 2001) from the covariance of the parameters, which is derived from the Fisher information matrix. Note that the confidence bands are too narrow, since observations included from adjacent horizons are correlated, hence more points are added without adding an equal amount of information.



Figure 3.15: Predicted one year return level using the same polynomial orders for all horizons for the semi-parametric model for all observations: a second order for location and a first order for the scale parameter.

## 3.4.2 Model for block maxima above a threshold of forecasted wind power

The semi-parametric model defined in Equation (3.14) is fitted to observations where the forecasted wind power is above a selected threshold for each horizon, as described in Section 3.3.3.

First a model selection was carried out. The selected orders are listed in Table 3.4. More similar orders through out the horizons are found compared to orders selected for the similar single horizon model (listed in Table 3.2. For the short horizons, where fewer observations are removed the orders are close the orders selected for the semi-parametric model for all observations. For longer horizons, where more observations are removed, the orders decrease. From the selected orders it is found reasonable to select use a first-order polynomial for the location and a zero-order (i.e. no dependence of forecasted wind power) for the scale, however first this leads to quite unstable fits, so a zero-order for the scale was preferred..

Diagnostic plots for four selected horizons are shown in Figure 3.16. Note that the size of the points in the right most plots indicates the weight (from the kernel) put on the observation, however this is not the case for the left-most plots (QQ-plots). Considering the QQ-plots, it again seems like the fits does deviate from the observations, especially for the most extreme values, however since the observations are not weighted equally it is not possible to conclude the influence based on this.

In Figure 3.17 the predicted one year return-level for the semi-parametric model for observations above a threshold of forecasted wind power is shown. Compared to the similar plot for the single horizon model (in Figure 3.15) less variation is seen, but also a bias at the shortest and longest horizons.

Horizon	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Location order	2	2	2	2	2	2	2	2	1	1	1	1	2	2
Scale order	1	1	1	1	1	1	1	1	2	2	2	2	2	2
Horizon cont.1	15	16	17	18	19	20	21	22	23	24	25	26	27	28
Location order cont.1	2	2	1	1	1	1	2	2	2	2	1	1	1	1
Scale order cont.1	2	1	1	1	1	1	1	1	1	1	1	2	0	0
Horizon cont.2	29	30	31	32	33	34	35	36	37	38	39	40	41	42
Location order cont.2	1	1	1	1	1	2	1	1	1	1	1	1	1	1
Scale order cont.2	0	0	0	0	0	0	2	1	1	1	1	1	1	1

Table 3.4: The order selected for the location and scale polynomials for each horizon using the forward selection procedure.



Figure 3.16: Model diagnostic plots for four selected horizons for semi-parametric model for observations above a threshold. Left column are QQ-plots of the empirical distribution versus the fitted distribution. Mid column plots is an estimated empirical density (fitted using a local kernel regression) and the fitted (modeled) distribution. Right column plots are of the estimated one and five year return levels including 95% confidence bands calculated using the delta method (Coles, 2001) from the covariance of the parameters, which is derived from the Fisher information matrix. Note that the confidence bands are too narrow, since observations included from adjacent horizons are correlated, hence more points are added without adding an equal amount of information.



Figure 3.17: Predicted one year return level as a function of the forecasted wind power and the horizon for the semi-parametric model for fitted to observations above a threshold of forecasted wind power.

#### 3.5 Non-parametric model

In this section a fully non-parametric model is presented by extending the semiparametric model to be fitted also locally in the forecasted wind power dimension.

The applied non-parametric models are

$$v \sim GEV \Big( \mu \big( P_{t|t-k_{\mathrm{f}}}, k_{\mathrm{f}} \big) , \ \sigma \big( P_{t|t-k_{\mathrm{f}}}, k_{\mathrm{f}} \big) , \ \xi \Big)$$

where v is the block maximum random variable,  $P_{t|t-k_f}$  is the forecasted wind power,  $k_f$  is the forecast horizon,  $\mu(P_{t|t-k_f}, k_f)$  and  $\sigma(P_{t|t-k_f}, k_f)$  are non-parametric smooth functions for the location and the scale parameter, respectively. They are called non-parametric because they are fitted locally in the forecasted wind power dimension and the horizon dimension. This is achieved by calculating the likelihood for a given horizon  $k_f$  and wind power forecast  $P_{t|t-k_f}$  with the density function

$$f_{v}\left(\mu_{f}\left(P_{t|t-k_{\mathrm{f}}},k_{\mathrm{f}},\boldsymbol{\mu}\right),\ \sigma_{f}\left(P_{t|t-k_{\mathrm{f}}},k_{\mathrm{f}},\boldsymbol{\sigma}\right),\ \boldsymbol{\xi}\right)$$
(3.14)

where the location function is

$$\mu_f(P_{t|t-k_f}, k_f, \mu) = \mu_0 + \mu_1 \Delta P + \mu_2 \Delta P^2 + \mu_3 \Delta k + \mu_4 \Delta k^2$$
(3.15)

and the scale function is

$$\sigma_f(P_{t|t-k_f}, k_f, \sigma) = \sigma_0 + \sigma_1 \Delta P + \sigma_2 \Delta P^2 + \sigma_3 \Delta k + \sigma_4 \Delta k^2$$
(3.16)

where  $\Delta k = k - k_f$  is the distance in the horizon dimension and  $\Delta P = \hat{P}_{t_i|t_i-k} - P_{t|t-k_f}$  is the distance in the forecasted wind power dimension. The fits becomes local second order polynomial fits in both dimensions by down-weighting the observations with a multiplicative Epanechnikov kernel function

$$w(\Delta k, h_{\rm k}, \Delta \hat{P}_{t_i|t_i-k}, h_{\rm P}) = \frac{3}{4h_{\rm k}} (1 - \frac{\Delta k^2}{h_{\rm k}^2}) I_{\{\Delta k \le h_{\rm k}\}} \frac{3}{4h_{\rm P}} (1 - \frac{\Delta k^2}{h_{\rm P}^2}) I_{\{\Delta k \le h_{\rm P}\}}$$
(3.17)

where  $h_k$  and  $h_P$  are bandwidths in each dimension, and  $I_{\{\Delta k \le h\}}$  and  $I_{\{\Delta k \le h_P\}}$  are indicator functions. For the full model the parameters which are fitted with the maximum likelihood are  $\theta = (\mu_0, \ldots, \mu_4, \sigma_0, \ldots, \sigma_4, \xi)$ , note that  $\theta$  is fitted for each given  $k_f$  and  $P_{t|t-k_f}$  which is left out in the notation for simplification. Note also that a second order polynomial is always used for the local fit in the horizon dimension (i.e.  $\mu_3, \mu_4, \sigma_3$  and  $\sigma_4$  is always included in the model), but the order of the polynomials for the forecasted wind power is selected using the forward model selection procedure as described in the following section.

For the non-parametric model no procedure for model selection was carried out, the orders of the local polynomials and the bandwidths was set by manually. This was based on the selected fits for the previous simpler models in order to find an appropriate balance between bias and variance. The contour plot in Figure 3.18 is of the predicted one-year return level fitting the non-parametric model to all observations. The bandwidths used are

$$h_{\rm k} = 5 \text{ hours}, h_{\rm P} = 1500 \text{ MW}$$
 (3.18)


Figure 3.18: Predicted one year return level for each horizon with the non-parametric model fitted to all observations.

and a first order polynomial is used for the location and a zero order for the scale.

It was not found possible to decrease the model variance to obtain an increasing return level for increasing horizon and forecasted wind power, especially not at the higher levels of forecasted wind power, without introducing a high bias in other regions. Hence for the suggested fully non-parametric model it seems like too little data is available in all regions to obtain useful results. One way to develop this, which is left for further studies, could be to use a nearest neighbor approach to set the bandwidths according to the available information in different regions.

#### 3.6 Comparison between the fitted models

In this section the applied models are evaluated by comparing predicted return levels. This provides a direct view of the features of the fits, especially bias vs. variance, obtained with the suggested models.

In Figure 3.19 the predicted one year return level for the single horizon model and semi-parametric model are plotted with confidence bands, both for the fit to all observations and the fit to observations above a threshold of forecasted wind power. It is seen that for each model the single horizon and semi-parametric are not far

from each others, except for shortest (horizon 1 and 2 hours) and longest horizons (42 hours), indicating that the two semi-parametric models are biased in these regions, especially for the shortest horizons. Furthermore, from the 6 hours horizon the upper bound start to have an influence for the lower range of forecasted wind power and from above the 20 hours horizon most of the values are in this range, hence influencing heavily the fits. Finally, it is noted that the confidence bands are wider for the single horizon models compared to the semi-parametric models. The number of points included when fitting the semi-parametric models are higher, since neighboring horizons are included. However, since the forecasts are calculated from NWPs which are not updated every hour, then both the block maxima and forecasted wind power from neighboring horizons are correlated. Thus that an error in one NWP update will affect several horizons, e.g. if the NWPs are updated every six hours, then an event with a high error in one NWP update will affect six adjacent horizons. Due to this correlation the degrees of freedom are lower than the number of points resulting in too narrow confidence intervals for the semi-parametric models. In future work this should be studied to find ways to correct for this.



Figure 3.19: Predicted one year return level for horizons 1,2,3,6,12,20,28,36 and 42 hours, for the fits for the single horizons and semi-parametric models as denoted in the legend. For the two fits, where observations are removed below a threshold (blue and cyan), the points to the left side of the lower line end was not used for the fit.

#### 3.7 Forecasts

In this section a few examples of how the return-level forecasts could be visualized are given. For the two shortest horizons (1 and 2 hours) the single horizon model including all observations is used and for the longer horizons ( $k \ge 3$  hours) the semi-parametric model including all observations is used.

In Figure 3.20 forecasted return level time series for four different horizons are plotted, together with the observed and forecasted wind power, for the first half of January 2012. The return levels are for return periods of 1,2 and 6 months, and 1,2,5 and 10 years. It is seen how the return levels decrease as the horizon increase and an event where the observed wind power does drop below some return levels occur around the 5th of January.

In Figure 3.21 forecasts generated at three points in time are plotted for the horizons up to 42 hours. Return levels for same return periods as in the previous figure are plotted. Again it is seen how the return levels decrease as the horizon increase. The lower plot is of the event at the 5th of January, where clearly the observed wind power drops below the two month (0.17 years) return level and it seen how the event has five or six points below the one month return level and two points below the two month level.



Figure 3.20: Example of forecasts for four different horizons and different return levels. The return levels are for return periods of 1,2 and 6 months, and 1,2,5 and 10 years. The return levels are marked with colors in the region from the level down to the next return level.



Figure 3.21: Example of forecasts for different return periods for three different time points from 1 to 42 hours ahead.

## **Chapter 4**

# Extreme value models for DK2 wind power forecast errors

In this chapter an analysis for DK2 similar to the one for DK1 in the previous chapter is presented, however with less details and explanations than for DK1. Note that the same notation is used, e.g. the wind power is denoted by  $P_t$ , which in previous the chapter was denoting the wind power in DK1, in this chapter denotes the wind power in DK2.

#### 4.1 Data DK2

In this section an exploratory data analysis is carried out to give an overview of the data for DK2 used in the modelling and to point out some features in the data relevant for the modelling. The data used are covering 1.5 years and consists of hourly measurements of wind power in DK2 and corresponding wind power forecasts.

#### 4.1.1 Wind power measurements

The time series of hourly wind power measurements for DK2 are denoted by

$$\{P_t: t = 1, \dots, N_h\}$$
 (4.1)

where *t* is the hours since 2011–07–01 00:00:00 UTC and  $N_{\rm h} = 14159$  and the last time point is 2013–02–09 23:00:00 UTC. The time marks the end of the hour. In Figure 4.1 the series is plotted together with the rated wind power level in DK2. It is seen that the rated power level varies only slightly over the period and that a few periods are missing: in total 453 values are missing (around 2.5 week).

#### 4.1.2 Wind power forecasts

See the description of the wind power forecasts in Section 3.1.2. The wind power forecasts for DK2 are denoted by

$$\{\hat{P}_{t|t-k}: t = 1, \dots, N_h\}$$
 (4.2)

in this chapter. A scatter plot for nine selected *k* hours horizons: 1,2,3,6,12,18,24,32 and 48 hours, are shown in Figure 4.2. It is seen that the scatter (the level of resid-



Figure 4.1: The time series of hourly wind power and rated wind power in DK2.

uals) does seem to increase with the horizon, as anticipated. Different events with high negative errors are seen, the most extreme at 18 and 24 hours horizons, again, as for DK1, are more extreme than for the 24 hours horizon. Compared to the similar plots for DK1, see Figure 3.2 on page 11, it is seen that the negative error extremes are closer to the bound at zero wind power for DK2.

In the modelling the residuals for each horizon are used. They are

$$\epsilon_{k,t} = P_t - \hat{P}_{t|t-k} \tag{4.3}$$

for the *k*′th horizon.

In Figure 4.3 the negative residuals (i.e.  $-\epsilon_{k,t}$ ) are plotted for four selected horizons, k = 1, 3, 24 and 42 hours. As for DK1 it is clearly seen how the high negative residuals are correlated in time, such that an extreme event consists of more than high negative residual.



Figure 4.2: Wind power measurements versus forecasts for nine selected horizon *k* hours for DK2.



Figure 4.3: The negative residuals versus the wind power forecasts for DK2. The points are connected with lines to their neighboring points in time (i.e. the previous and next hour).

#### 4.2 Block maxima

Block maxima for DK2 are formed similarly as for DK1, see Section 3.2. In Figure 4.4 the block maxima for three selected horizons and block lengths are plotted versus the forecasted wind power.

#### 4.2.1 Selection of block length

Block maxima length is selected for DK2 similarly as for DK1, see Section 3.2.1. The plots in Figure 4.5 are of the parameter estimates for each horizon as a function of the block length. Also included in the plots as red lines are local regression estimates fitted using the R function loess() with span=0.7 (R Core Team, 2013). Compared to the similar plots for DK1 in Figure 3.5 a less clear pattern is seen for DK2. It is found that using a block length of 7 days, as for DK1, is found suitable for obtaining the best balance between bias and variance and this block length is used in the remaining of the study.



Figure 4.4: Block maxima vs. forecasted wind power for three selected horizons k and block lengths n for DK2.



Figure 4.5: The parameters of the model in Equation (3.5) fitted for DK2 for each horizon and for each block lengths n = 1 to n = 15 days. The red lines are local regression estimates.

#### 4.3 Single horizon models

In this section a model selection for models fitted to a single horizon is carried out. The model selection procedure presented in Section 3.3.1 is used.

#### 4.3.1 Model for all observations

A model is selected for each single horizon for observations above a threshold. The orders selected for the location and scale polynomials for the model in Equation (3.5) for each horizon are shown in Table 4.1. It is apparent that there is some pattern in the selected orders, however the pattern is not clearly changing from e.g. a low order for short horizons to a higher order for long horizons, indicating that applying the selected orders for each horizon will result in high variation of predictions.

In Figure 4.6 model diagnostic plots are shown for four horizons, where the model was fitted using the selected polynomial orders. The QQ and density plots indicates that the model follows the observations reasonably well, only the 2 to 5 most extreme values deviate from the model fit. They are slightly higher than the fit, which this indicates that extrapolation with the model (i.e. prediction of return levels longer than the current data period) could tend to be biased to the lower side. Furthermore it is seen from the return level plots for the longer horizons (for 18 and 32 hours), that most of the observations at the lower level of forecasted wind power reach the upper bound. This influences the fit. In this range the higher quantiles of the fitted GEV will follow the upper bound (the one-to-one function) and the second order allows it to "bend" away from the upper bound.

In Figure 4.7 the predicted one year return level of for each horizon as a function of the forecasted wind power is plotted with colors indicating the predicted level. It is clearly seen that the predicted return level varies up and down depending on the horizon, especially for the higher levels of forecasted wind power. Considering the horizons where the polynomial orders were selected differently, as seen in Table 4.1, for example for horizon: 33,34,41,42 versus 35 to 40, then it is clear that the model order has a high influence of the predicted return levels (especially for high levels of forecasted wind power).

Horizon	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Location order	1	1	1	1	1	2	2	1	2	1	1	1	2	1
Scale order	0	1	1	1	1	2	2	1	2	1	2	1	1	1
Horizon cont.1	15	16	17	18	19	20	21	22	23	24	25	26	27	28
Location order cont.1	2	2	2	2	2	2	1	2	1	1	2	2	2	2
Scale order cont.1	2	1	1	1	1	1	1	1	1	1	1	1	1	1
Horizon cont.2	29	30	31	32	33	34	35	36	37	38	39	40	41	42
Location order cont.2	1	2	2	2	2	1	2	1	2	2	2	2	2	2
Scale order cont.2	1	1	1	1	1	1	2	1	2	2	2	2	1	2

A reasonable assumption would be that the return level is an increasing function of

Table 4.1: The order selected for the location and scale polynomials for each horizon using the forward selection procedure.

the horizon. One way of decreasing the dependence of the horizon and thus obtain less variance is to apply the same order of the polynomials for all horizons. Considering the selected polynomial orders in Table 4.1 it is found reasonable to apply a second order polynomial for location and a first order for the scale parameter

$$v_k \sim GEV(\mu_0 + \mu_1 P_{t|t-k_f} + \mu_2 P_{t|t-k_f}^2, \sigma_0 + \sigma_1 P_{t|t-k_f}, \xi)$$
(4.4)

which is fitted to all horizons separately. In Figure 4.8 the predicted one year return level using this model is plotted. Compared with the return level predicted using the selected model for each horizon (in Figure 4.7) less variation and an overall increasing trend of the return level as function of the horizon are obtained.



Figure 4.6: Model diagnostic plots for four selected horizons for the single horizon model for all observations. Left column are QQ-plots of the empirical distribution versus the fitted distribution. Mid column plots is an estimated empirical density (fitted using a local kernel regression) and the fitted (modeled) distribution. Right column plots are of the estimated one and five year return levels including 95% confidence bands calculated using the delta method (Coles, 2001) from the covariance of the parameters, which is derived from the Fisher information matrix.



Figure 4.7: Predicted one year return level for each horizon for the single horizon model for all observations fitted with the selected polynomial orders.



Figure 4.8: Predicted one year return level for the single horizon model for all observations using the same polynomial orders for all horizons: a second order for location and a first order for the scale parameter.

# 4.3.2 Model for block maxima above a threshold of forecasted wind power

Removing the block maxima below a threshold of forecasted wind power where they reach the upper bound is found to be a reasonable approach to ensure that the models are not biased for higher levels of forecasted wind power. The threshold is set for each horizon separately, simply by a visual inspection of scatter plots to find the level where the block maxima gets close to the upper bound. In Figure 4.9 the threshold used as a function of the horizon is plotted.

The forward model selection (see the description in Section 3.3.1) is carried out to select a suitable model for each horizon. The selected polynomial orders for the location and scale parameters are listed in Table 4.2. It is apparent that there is some pattern in the orders selected, however as before the pattern is not clearly changing from e.g. a low order for short horizons to a higher order for long horizons. Compared to the orders selected for all observations (see Table 4.1) the orders are lower, hence that less significant non-linear dependence of the forecasted wind power is found. This is natural since now the model doesn't need to take into account both the distribution at lower levels of forecasted wind power (reaching the upper bound) and doesn't need to "bend" down at higher levels of forecasted wind power. However fewer observations are included when fitting the models.

In Figure 4.10 model fit diagnostic plots for four horizons are shown. From the plots for the shorter horizons (1 and 6 hours) it is seen that the fits are quite close to the fits for all observations (compare the Figure 4.6). For the longer horizons (18 and 32 hours) it is seen that since fewer observations are used, then less information is available and the fits are more affected by the highest block maximum, especially the most extreme value at 32 hours. This leads to a less close match between the empirical and modelled density and the confidence bands becomes wider.

It was not possible to predict return levels using the selected models for each horizon, instead to decrease the variance a single model is fitted to all horizons, as for the fit for all observations. From the orders selected for each horizon as shown in Table 4.2 it is found reasonable to use a zero order polynomial both for the location and scale parameters. Clearly, the predicted one return levels varies with the

Horizon	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Location order	1	1	1	2	0	1	0	0	0	0	0	0	0	0
Scale order	0	1	1	0	1	1	0	0	0	0	0	1	1	2
Horizon cont.1	15	16	17	18	19	20	21	22	23	24	25	26	27	28
Location order cont.1	2	0	0	2	2	0	0	0	0	0	0	0	0	0
Scale order cont.1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Horizon cont.2	29	30	31	32	33	34	35	36	37	38	39	40	41	42
Location order cont.2	0	0	0	0	0	1	1	0	0	0	0	0	0	0
Scale order cont.2	0	1	0	0	0	2	2	1	1	0	0	0	0	0

Table 4.2: The order selected for the location and scale polynomials for each horizon using the forward selection procedure for the single horizon model for observations above a threshold of forecasted wind power.



Figure 4.9: The thresholds, as function of the horizon, below which the block maxima are removed.

horizon, it simply goes up and down, this is natural, and the results is found quite reasonable.



Figure 4.10: Model diagnostic plots for four selected horizons for models for observations above a threshold. Left column are QQ-plots of the empirical distribution versus the fitted distribution. Mid column plots is an estimated empirical density (fitted using a local kernel regression) and the fitted (modeled) distribution. Right column plots are of the estimated one and five year return levels including 95% confidence bands calculated using the delta method (Coles, 2001) from the covariance of the parameters, which is derived from the Fisher information matrix.



Figure 4.11: Predicted one year return level using the same polynomial orders for all horizons for the model for observations above a threshold: a zero order for location and a zero order for the scale parameter.

#### 4.4 Semi-parametric models

The same model structure as applied for DK1 described in Section 3.4 is applied for DK2, except different orders of the parameter polynomials are selected.

#### 4.4.1 Model for all observations

A semi-parametric model is selected for DK2 as described for DK1 in Section 3.4.1. The same bandwidth

$$h = 5 \tag{4.5}$$

is used.

In Table 4.3 the order of the location and scale polynomials selected for each horizon for the semi-parametric model for all observations. The polynomial orders selected are varying for the different horizons, for shorter horizons a first or second order polynomial for the location and first order for scale seems reasonable. For the longer horizons the order for location is zero and one to two for the scale.

Since, as seen below, the focus will be on the shorter horizons, since the for longer horizons the models are not able to provide useful results, a model with a second order polynomial for the location and first order for the scale is fitted, as for DK1. The diagnostic plots for four selected horizons are shown in Figure 4.10. The predicted return levels are less varying compared to the single horizon model, however naturally a problem is that for lower levels of forecasted wind power the block maximum reach the upper bound, hence they are certainly biased in this range, and for higher levels of forecasted wind power the return levels are bending down away from the upper bound. Hence also very biased in this range, however the predictions are "stable" in a sense and this model is used in the remaining analysis.

In Figure 4.13 the predicted one year return level as a function of the forecasted wind power and horizon is plotted for horizons 1 to 12 hours. The variation between the horizons is decreased compared to the plot for the similar single horizon model seen in Figure 4.8 and the overall levels are very alike, except for the very short horizons, where the semi-parametric model clearly is biased (which will also be apparent in a comparison later in the report).

Horizon	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Location order	2	2	2	2	2	2	2	0	0	0	0	0	0	0
Scale order	1	1	1	1	1	1	1	2	2	2	1	1	1	1
Horizon cont.1	15	16	17	18	19	20	21	22	23	24	25	26	27	28
Location order cont.1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Scale order cont.1	2	2	2	2	2	2	2	2	2	1	1	1	1	1
Horizon cont.2	29	30	31	32	33	34	35	36	37	38	39	40	41	42
Location order cont.2	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Scale order cont.2	1	1	1	1	1	1	1	1	1	1	1	1	1	1

Table 4.3: The order selected for the location and scale polynomials for each horizon using the forward selection procedure for the semi-parametric model for all observations.



Figure 4.12: Model diagnostic plots for four selected horizons for semi-parametric model for all observations. Left column are QQ-plots of the empirical distribution versus the fitted distribution. Mid column plots is an estimated empirical density (fitted using a local kernel regression) and the fitted (modeled) distribution. Right column plots are of the estimated one and five year return levels including 95% confidence bands calculated using the delta method (Coles, 2001) from the covariance of the parameters, which is derived from the Fisher information matrix. Note that the confidence bands are too narrow, since observations included from adjacent horizons are correlated, hence more points are added without adding an equal amount of information.



Figure 4.13: Predicted one year return level using the same polynomial orders for all horizons for the semi-parametric model for all observations: a second order for location and a first order for the scale parameter.

Model for block maxima above a threshold of forecasted wind power The semiparametric model defined in Equation (3.14) is fitted to observations where the forecasted wind power is above a selected threshold for each horizon, as described in Section 3.3.3.

First a model selection was carried out and the selected orders are listed in Table 4.4. For the short horizons, where fewer observations are removed the orders are close the orders selected for the semi-parametric model for all observations. For longer horizons, where more observations are removed, the orders decrease, and the pattern is not clear. From the selected orders for the shorter horizons a second order for location and first order for scale seems reasonable, however fitting this results in very varying predicted return levels (not shown). It was chosen to use a first-order polynomial for the location and a zero-order (i.e. no dependence of forecasted wind power) for the scale.

Diagnostic plots for four selected horizons are shown in Figure 4.14. For shorter horizons the diagnostics seems reasonable, however for 12 hours ahead the predicted 5 years return levels are far above the upper bound, hence the fits are simply useful.

In Figure 4.15 the predicted one year return-level for the semi-parametric model for observations above a threshold of forecasted wind power is shown. Compared to the similar plot for the single horizon model (in Figure 4.13) less variation is seen, but the levels from above around 8 hours seems to increase more than anticipated.

Horizon	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Location order	2	2	2	2	2	2	2	2	0	0	0	0	0	0
Scale order	1	1	1	1	1	1	1	1	2	2	2	1	1	1
Horizon cont.1	15	16	17	18	19	20	21	22	23	24	25	26	27	28
Location order cont.1	0	0	0	0	0	0	1	0	2	2	2	2	2	2
Scale order cont.1	2	2	2	2	2	2	2	2	1	1	1	1	1	1
Horizon cont.2	29	30	31	32	33	34	35	36	37	38	39	40	41	42
Location order cont.2	2	2	2	1	1	0	1	1	2	1	1	0	1	1
Scale order cont.2	1	2	2	2	2	2	2	2	2	2	2	1	0	0

Table 4.4: The order selected for the location and scale polynomials for each horizon using the forward selection procedure.



Figure 4.14: Model diagnostic plots for four selected horizons for semi-parametric model for observations above a threshold. Left column are QQ-plots of the empirical distribution versus the fitted distribution. Mid column plots is an estimated empirical density (fitted using a local kernel regression) and the fitted (modeled) distribution. Right column plots are of the estimated one and five year return levels including 95% confidence bands calculated using the delta method (Coles, 2001) from the covariance of the parameters, which is derived from the Fisher information matrix. Note that the confidence bands are too narrow, since observations included from adjacent horizons are correlated, hence more points are added without adding an equal amount of information.



Figure 4.15: Predicted one year return level as a function of the forecasted wind power and the horizon for the semi-parametric model for fitted to observations above a threshold of forecasted wind power.

#### 4.5 Non-parametric model

No non-parametric model was fitted for DK2, it was simply not possible to obtain useful results.

#### 4.6 Comparison between the fitted models

In this section the applied models are evaluated by comparing predicted return levels. This provides a direct view of the results, especially bias vs. variance, obtained with the suggested models.

In Figure 4.16 the predicted one year return level for the single horizon model and semi-parametric model are plotted with confidence bands, both for the fit to all observations and the fit to observations above a threshold of forecasted wind power. The following is noted

- single horizon fits for all observations seems reasonable until horizons 6 hours and then starts to be varying,
- single horizon fits for observations above a threshold seems only useful for 1 and 2 hours horizons,
- semi-parametric fits for all observations seems fine for most horizons, except the shortest horizons,
- semi-parametric fits for observations above a threshold are from k=6 to increasing heavily,
- from horizon 2 to 6 hours three of the fits agree,
- above the 6 hours horizon the upper bound for lower values of forecasted wind power starts to influence heavily.



Figure 4.16: Predicted one year return level for horizons 1,2,3,6,12,20,28,36 and 42 hours, for the fits for the single horizons and semi-parametric models as denoted in the legend. For the two fits, where observations are removed below a threshold (blue and cyan), the points to the left side of the lower line end was not used for the fit.

#### 4.7 Forecasts

In this section a few examples of how the return-level forecasts can be visualized are given. As for DK1, for the two shortest horizons (1 and 2 hours) the single horizon model including all observations is used and for the longer horizons ( $k \ge 3$  hours) the semi-parametric model including all observations is used.

In Figure 4.17 forecasted return level time series for four different horizons are plotted, together with the observed and forecasted wind power, for the first half of January 2012. The return levels are for return periods of 1,2 and 6 months, and 1,2,5 and 10 years. It is seen how the return levels decrease as the horizon increase and for horizon 12 hours the predicted levels for 1 year return period and above are simply zero. As discussed above, only the shorter horizons return levels (1 to 6 hours) are reasonably accurate.

In Figure 4.18 forecasts generated at three points in time are plotted for the horizons up to 42 hours. Return levels for same return periods as in the previous figure are plotted. Again it is seen how the return levels decrease as the horizon increase.



Figure 4.17: Example of forecasts for four different horizons and different return levels. The return levels are for return periods of 1,2 and 6 months, and 1,2,5 and 10 years. The return levels are marked with colors



Figure 4.18: Example of forecasts for different return periods for three different time points from 1 hour to 42 hours ahead.

# Chapter 5

## **Discussion and conclusion**

This chapter consists of a discussion of the results, suggestions for further work and ends with the conclusions of the study.

#### 5.1 Discussion

In this section a discussion of the results is given by highlighting important points:

- The number of available values for fitting the models has a direct influence on the uncertainty of the predicted return levels. The current data spans a period of 1.5 years and from the width of the confidence bands and the variation between the applied models on the predicted return levels, it will be very useful to have data from a longer period. This is most important for longer horizons and high levels of forecasted wind power. Especially it should be possible to obtain reliable results with the non-parametric model to capturing the non-linearities between the block maxima and the forecasted wind power.
- It is apparent that the upper bound of the block maxima influences the results. Upper bounded GEV distributions are reversed Weibull or type III extreme value distributions and occur when the shape parameter is below zero  $\xi < 0$ , see also the description of GEV distributions in Section 2. This upper bound which can be estimated(Coles, 2001) and could be used for validation, since it should not be too far above the bound. The value of the shape parameter for the semi-parametric model fitted to all horizons, see Section 3.4.1, is in the range from -0.14 to -0.01 for all the horizons, hence all negative and therefore the GEVs fitted are bounded upwards.
- Block maxima for neighboring horizons are highly correlated, hence, as noted before, this will influence the widths of the confidence bands of parameters and return levels. This should be studied further to provide more reliable confidence bands should be derived. Furthermore profile-likelihood and boot-strapping confidence bands should also be calculated and compared.
- Discussion of results for DK1 versus DK2.

#### 5.2 Further work

In this section suggestions for work which fell beyond the scope of the current study are given, simply by a listing of ideas and features which could be included in further work:

- Characterization of the extreme error events, e.g. their average lengths etc. This also depend on the return period, so that events are relatively longer for short return periods. Furthermore, the error events certainly are linked to the period between NWP forecast updates, which for the NWPs used for the wind power forecasts in the present study is 6 hours.
- Different approaches to calculating the kernel bandwidths (for semi- and non-parametric models) for example using a nearest neighbor approach, or by lowering it for short horizons. Also splines could be used for the non-linear parts and knot points could by put to achieve functions with appropriate properties, e.g. a higher curvature for the lower horizons in horizon dimension.
- Plots of return levels (they are quantiles of the GEV) for the models for some specific horizon to see how actually the fitted distribution vary as a function of the forecasted wind power. The upper bounds can also be estimated for type III GEV distributions, when  $\xi < 0$ , see Coles (2001).
- An analysis of the extreme error events should be carried out to find out when they occur and if there are any systematic patterns apparent when they occur.
- As mentioned in the discussion the confidence bands should be further studied and different techniques for calculating the confidence bands should be compared.
- Generalized Pareto (GP) distributions which describe the tails of the distributions should be investigated. However, GP distributions do not have an upper bound, which can be means that they might not be as suitable as the GEV distributions, maybe truncated GP distribution (Aban et al., 2006) can be used.
- Maybe other methods for calculating the value of the forecasted wind power, instead of simply taking the value at the time point of the block maximum, could improve the fits.
- Inclusion of information of forecast uncertainty from other sources. Hence an input signal could be generated for example from a quantile of probabilistic wind power forecasts or ensemble NWP forecasts. Furthermore, for example using the wind direction is an input to the extreme value model could be an improvement, however this would require more data.
- Using meteorological knowledge the weather phenomena present when large error events occur should be investigated.

#### 5.3 Conclusion

Models for extreme negative wind power forecast errors are presented in this report. The models are applied for hourly wind power production and corresponding forecasts for the two areas in Denmark, DK1 and DK2. Generalized extreme value distributions fitted with maximum-likelihood are used. The range of studied models have an increasing complexity, from the parametric models applied to single horizon data, to semi-parametric and over to non-parametric models. A procedure providing a systematic approach to model selection is suggested and applied. Since the results cannot be cross-validated, statistics, e.g. tests and confidence bands, provide the basis for the evaluation. The predicted return level for the negative forecast errors increase as a function of the horizon and forecasted wind power. Most clearly for DK1, it is considerably lower for the 1 and 2 hours horizons, indicating that the wind power forecasts are significantly more reliable for very short horizons. From the 3 hours horizon the return level increases at a steady rate until around the 12 hours horizon from where the return level increase at a slower rate. There is clear difference between DK1 and DK2. The errors for DK1 are relatively smaller and thus the upper bound of the negative forecast errors influence on the fits. The applied semi-parametric model for all observations is found to be quite robust. For the longer horizons it is biased due to the upper bound, however it is found to be the most reliable for the present data. The non-parametric models doesn't provide useful results for the present data, they will need observations from a longer period. Finally examples of forecasts are presented together with a discussion and suggestions for further work are given.
## Bibliography

- I. B. Aban, M. M. Meerschaert, and A. K. Panorska. Parameter estimation for the truncated pareto distribution. *Journal of the American Statistical Association*, 101 (473):270–277, 2006. doi: 10.1198/016214505000000411. URL http://amstat. tandfonline.com/doi/abs/10.1198/016214505000000411.
- P. Bacher and H. Madsen. Identifying suitable models for the heat dynamics of buildings. *Energy & Buildings*, 43(7):1511–1522, 2011. ISSN 03787788. doi: 10. 1016/j.enbuild.2011.02.005.
- J. Beirlant, Y. Goegebeur, J. Segers, and J. Teugels. *Statistics of extremes: theory and applications*. Wiley. com, 2006.
- S. Coles. An introduction to statistical modeling of extreme values. Springer, 2001.
- N. Cook. Towards better estimation of extreme winds. Journal of Wind Engineering and Industrial Aerodynamics, 9(3):295 – 323, 1982. ISSN 0167-6105. doi: http://dx. doi.org/10.1016/0167-6105(82)90021-6. URL http://www.sciencedirect.com/ science/article/pii/0167610582900216.
- L. de Haan and J. de Ronde. Sea and wind: Multivariate extremes at work. *Extremes*, 1(1):7–45, 1998. ISSN 1386-1999. doi: 10.1023/A:1009909800311. URL http://dx.doi.org/10.1023/A%3A1009909800311.
- P. Embrechts, C. Klüppelberg, and T. Mikosch. Modelling Extremal Events: For Insurance and Finance. Applications of mathematics. Springer, 1997. ISBN 9783540609315. URL http://books.google.dk/books?id=BX012pICfJUC.
- P. Friederichs and T. L. Thorarinsdottir. Forecast verification for extreme value distributions with an application to probabilistic peak wind prediction. *Environmetrics*, 23(7):579–594, 2012. ISSN 1099-095X. doi: 10.1002/env.2176. URL http://dx.doi.org/10.1002/env.2176.
- E. Gilleland and R. W. Katz. New software to analyze how extremes change over time. *Eos*, 92(2):13–14, 2011.
- E. J. Gumbel. Statistics of extremes. 1958.
- J. E. Heffernan and J. A. Tawn. A conditional approach for multivariate extreme values (with discussion). *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 66(3):497–546, 2004. ISSN 1467-9868. doi: 10.1111/j.1467-9868. 2004.02050.x. URL http://dx.doi.org/10.1111/j.1467-9868.2004.02050.x.
- A. Horvat, T. Plavsic, and I. Kuzle. Application of extreme value theory for measuring risk of safe operation of power system in adriatic wind conditions. In

*EUROCON, 2013 IEEE,* pages 874–881, July 2013. doi: 10.1109/EUROCON.2013. 6625086.

- S. Kotz and S. Nadarajah. *Multivariate Extreme-Value Theory*. Wiley Online Library, 1988.
- H. Madsen and P. Thyregod. *Introduction to General and Generalized Linear Models*. CRC Press, 2010.
- R Core Team. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria, 2013. URL http://www.R-project.org/.
- A. Stephenson and E. Gilleland. Software for the analysis of extreme events: The current state and future directions. *Extremes*, 8(3):87–109, 2005.