

# TEMPORAL AND SPATIAL EVOLUTION OF EXTREME EVENTS

Jacco Groeneweg<sup>1</sup>, Sofia Caires<sup>1</sup> and Kathryn Roscoe<sup>2</sup>

In this study two statistical methods for computing time- and space-evolving extreme events are implemented and assessed. The implementation and application of the methods are carried out with a view towards the safety assessment of water defences. With this type of application in mind, the methods are used to produce time- and space-varying extreme wind fields, which are needed to force wave and hydrodynamic models. Furthermore, the results of the two methods are compared and effects of the choice of distribution, threshold and reference location on the resulting wind fields are assessed. Although the wind fields produced by the methods show more erratic spatial variations, due to the associated uncertainties, than those of the observed wind fields, they seem realistic and can in principle be used to drive wave and hydrodynamic models. However, an assessment of the results of these statistical methods in terms of the underlying physics still needs to be carried out and is part of future work.

*Keywords: extreme value analysis; extreme events; wind velocity; non-stationarity; spatial variation*

## INTRODUCTION

In compliance with the Dutch Water Act, the strength of the Dutch primary water defences must be assessed periodically for the required level of protection against hydraulic loads with return periods ranging from 250 to 10,000 years. These loads are derived probabilistically and are referred to as Hydraulic Boundary Conditions (HBC).

The quality of the derived HBC depends on the adequate performance of a number of components in the HBC chain. In this chain, statistics of waves, wind, water levels and river discharges and physical numerical model results form the input to probabilistic models. These probabilistic models in turn provide the output of the HBC chain, namely the hydraulic loads, a combination of water level and/or wave height, wave period and wave direction per location, depending on the water system (coastal region, lakes or rivers) and on the failure mechanism under consideration.

Within the HBC chain, typically only one instant during a storm is considered, e.g. the instant at which the maximum water level is achieved in the region of the dike section of interest. For the determination of the required water defence crest level this is an appropriate approach, since the maximum wave overtopping rate is typically obtained at the maximum water level. For failure mechanisms other than overtopping, the instant of the maximum water level does not necessarily lead to the critical load on the water defence. Furthermore, failure mechanisms such as dune retreat and erosion of revetments typically depend on the temporal, and therefore also spatial, variation of the storm. Currently, for most water systems, spatial variations of wind fields are schematised, often even assumed to be uniform. The latter assumption becomes even more unrealistic when temporal variations are accounted for.

In order to fill in the need for temporally and spatially varying hydraulic loads, numerical models (e.g. the wave model) should be driven by wind fields evolving in time and space and provided with fields of other relevant variables (in the case of the wave model: offshore wave conditions, water levels and currents) associated with extreme wind speeds. There are many possibilities to describe the temporal and spatial evolution of an extreme event. Two leading experts in extreme value theory, Prof. Laurens de Haan and Prof. Richard L. Smith, were consulted for advice on how to approach the problem. Each expert independently derived a semi-parametric method (Caires et al., 2011). In this study the two methods, referred to as the Method de Haan and the Method Smith, are implemented and validated.

## DESCRIPTION OF THE METHODS

In this section the implementation of the methods proposed by Prof. Laurens de Haan and Prof. Richard L. Smith is described. Each expert independently derived a semi-parametric method based on max-stable processes (see e.g. Buishand et al., 2008). The proposed methods were implemented in such a way that for a given reference location the wind fields associated with a fixed peak value at the reference location are determined. The following is a step-by-step summary of the methods as implemented.

---

<sup>1</sup> Hydraulic Engineering, Deltares, Delft, The Netherlands

<sup>2</sup> Inland Water Systems, Deltares, Delft, The Netherlands

Method de Haan	Method Smith
<p>H1. Consider the time series of (for instance) wind speed at several locations, indexed by <math>s</math>, <math>\{X(s)\}</math>. Apply the Peaks-over-Threshold (POT) method to the time series at each location <math>s</math>, determine the threshold <math>b_s</math>, and estimate the scale parameter <math>a_s</math> and the shape parameter <math>\gamma_s</math> of the Generalized Pareto Distribution (GPD).</p>	<p>S1. As in H1, and determination of the empirical distribution of the observations.</p>
<p>H2. Fix a value <math>t &gt; 1</math>. Determine thresholds <math>b_s(t) \geq b_s</math> such that <math>P\{X(s) &gt; b_s(t)\} = 1/t</math> and use the GPD fitted at each location <math>s</math> in step 1 to determine the scale parameter <math>a_s(t)</math> corresponding to the GPD describing the excesses over <math>b_s(t)</math>. This is done by using the threshold stability property of the GPD distribution: the shape parameter of the GPD is the same (i.e. <math>\gamma_s</math>) and <math>a_s(t) = a_s + \gamma_s (b_s(t) - b_s)</math>.</p>	<p>S2. As in H2 (although in principle not necessary).</p>
<p>H3. For a pre-chosen number of storms (<math>ns</math>) and hours (<math>nh</math>), select all observations that occur from <math>nh</math> hours before to <math>nh</math> hours after each <math>ns</math> highest peak at the reference location. The observations at all locations in these time windows (either above the local threshold or not) define a resulting set of successive observed storms, <math>\{X^{(1)}(s)\}</math>.</p>	<p>S3. For each location transform all observations so that (according to the POT model) they have a unit Fréchet distribution, <math>Y_s = -1/\log(F_s(X))</math>, where <math>F_s(X)</math> is the fitted location-specific GPD model augmented by the empirical distribution of observations below the threshold at the location.</p>
<p>H4. Transform the storm observations <math>\{X^{(1)}(s)\}</math> using</p> $X^{(2)}(s) := \begin{cases} \left(1 + \widehat{\gamma}_s \frac{X^{(1)}(s) - b_s(t)}{\widehat{a}_s(t)}\right)^{1/\widehat{\gamma}_s} & \text{for } \widehat{\gamma}_s \neq 0 \\ \exp\left(\frac{X^{(1)}(s) - b_s(t)}{\widehat{a}_s(t)}\right)^{1/\widehat{\gamma}_s} & \text{for } \widehat{\gamma}_s = 0 \end{cases}$	<p>S4. As in H3, but applied on transformed observations.</p>
<p>H5. Use the fitted Generalized Pareto model to estimate the <math>m</math>-year return level, <math>z_m</math> at the reference location <math>s_0</math>. Lift the transformed observations in the <math>i^{\text{th}}</math> storm from all locations by multiplying each storm value by a factor that takes the peak of the <math>i^{\text{th}}</math> storm at the reference location to <math>z_m</math>:</p> $t_{0i} = \left(\frac{z_m + \widehat{a}_{s_0}/\widehat{\gamma}_{s_0} - b_{s_0}}{X_i^{(1)}(s_0) + \widehat{a}_{s_0}/\widehat{\gamma}_{s_0} - b_{s_0}}\right)^{1/\widehat{\gamma}_{s_0}}.$	<p>S5. Use the fitted GPD model to estimate the <math>m</math>-year return level, <math>z_m</math>, at the reference location and transform <math>z_m</math> to the <math>m</math>-year level, <math>y_m</math>, on the unit Fréchet scale. Still on the unit Fréchet scale, lift the observations in the <math>i^{\text{th}}</math> storm from all locations by multiplying each storm value by a factor that takes the peak of the <math>i^{\text{th}}</math> storm at the reference location to <math>y_m</math>.</p>
<p>H6. Transform all lifted storm observations back to the original scale, using:</p> $X^{(3)}(s) := \begin{cases} \widehat{a}_s(t) \frac{\left(t_0 X^{(2)}(s)\right)^{\widehat{\gamma}_s} - 1}{\widehat{\gamma}_s} + b_s(t) & \text{for } \widehat{\gamma}_s \neq 0 \\ \widehat{a}_s(t) \log\left(t_0 X^{(2)}(s)\right) + b_s(t) & \text{for } \widehat{\gamma}_s = 0 \end{cases}$	<p>S6. Transform all lifted storm observations back from the unit Fréchet scale to the original scale. Use the inverse of the location-specific transformation of S3.</p>

As can be inferred from their description, the methods are similar and both involve the following steps:

- Extreme value analyses of the time series of the variables under consideration at each location of interest;
- Marginal transformations of the time series using the extreme value distributions fitted locally;
- Selection of storm periods;
- Uplifting of the transformed ‘storm’ data;
- Inversion of the marginal transformations.

The methods differ in the way ‘storms’ are identified, in the marginal transformations and in the lifting factors.

Except for the fits of the GPD, the methods de Haan and Smith are non-parametric. In particular, no parametric model is used to describe the temporal and spatial correlation. In fact, all the information that is contained in the selected observed storms is retained (used) in the final result. However, these methods of ‘lifting storms’ should not be confused with the procedure of simply applying a constant factor to each wind speed value in time and space; the lifting per location will depend on the local distribution of the data.

Contrary to what is the case in Method Smith, in the implementation of the Method de Haan the fitted location-specific GPD model is not augmented by the empirical distribution of the observations below the threshold at the location. Although the GPD is not strictly valid below the threshold, the choice was considered acceptable, because the method applies the max-stable processes theory “a little bit out of context”. However, given that using no augmentation can be interpreted as assuming that the GPD distribution applies right down to zero, which surely is not valid, we also considered the results of the method when the location-specific GPD model is augmented by the empirical distribution of the observations below the threshold. In that case, steps H2, H3 and H5 remain as described above and steps H1, H4 and H6 are replaced by:

- H1. plus the determination of the empirical distribution of observations, as in S1.
- H4. transform all observations so that (according to the POT model) they have a unit GPD. Base the transformation on the fitted location-specific GPD model augmented by the empirical distribution of observations below the threshold at the location.
- H6. Transform all lifted storm observations back from the unit GPD scale to the original scale. Use the inverse of the location-specific transformation of H4.

#### **APPLICATION AND ASSESSMENT OF THE METHODS**

In this section we compare uplifted 10,000-yr storms as obtained by both methods. A 10,000-yr return period is used because that is the period for which many of the water defences need to be assessed. We start by describing the data used in the analysis. Then the local extreme value analyses are described. The last subsections present the baseline results of the methods and the sensitivity of the results to the choice of distribution, threshold, and reference location. The sensitivity of the results of the Method de Haan to augmenting the fitted location-specific GPD model by the empirical distribution of the respective observations below the threshold is also presented.

##### **Description of the data**

We made use of the CFSR (Climate Forecast System Reanalysis, <http://cfs.ncep.noaa.gov/cfsr/>) dataset (Saha et al., 2010) which includes 6-hourly fields of wind velocity at 10 meters height from January 1979 until December 2009 (31 years) on a global  $0.312^\circ \times 0.312^\circ$  latitude/longitude grid. The considered CFSR 10-m wind speed and direction data were those for the region with longitudes between  $2^\circ\text{E}$  and  $8^\circ\text{E}$  and latitudes between  $50^\circ\text{N}$  and  $55^\circ\text{N}$ . This region covers the relevant Dutch water system. Figure 1 (left panel) shows the location of the CFSR grid points, in total 234, that fall in this region. Figure 1 also shows a comparison, by means of a density scatter plot, between hourly averaged 10-m wind speeds measured at the North Sea K13 platform, located at  $3.2^\circ\text{E}$  and  $53.2^\circ\text{N}$ , and the CFSR wind velocities at the nearest grid point. The figure shows that the correlation between the measured and the CFSR 10-m wind velocities is high (the correlation coefficient of the two datasets is 0.93), indicating that the CFSR wind fields can be considered realistic.

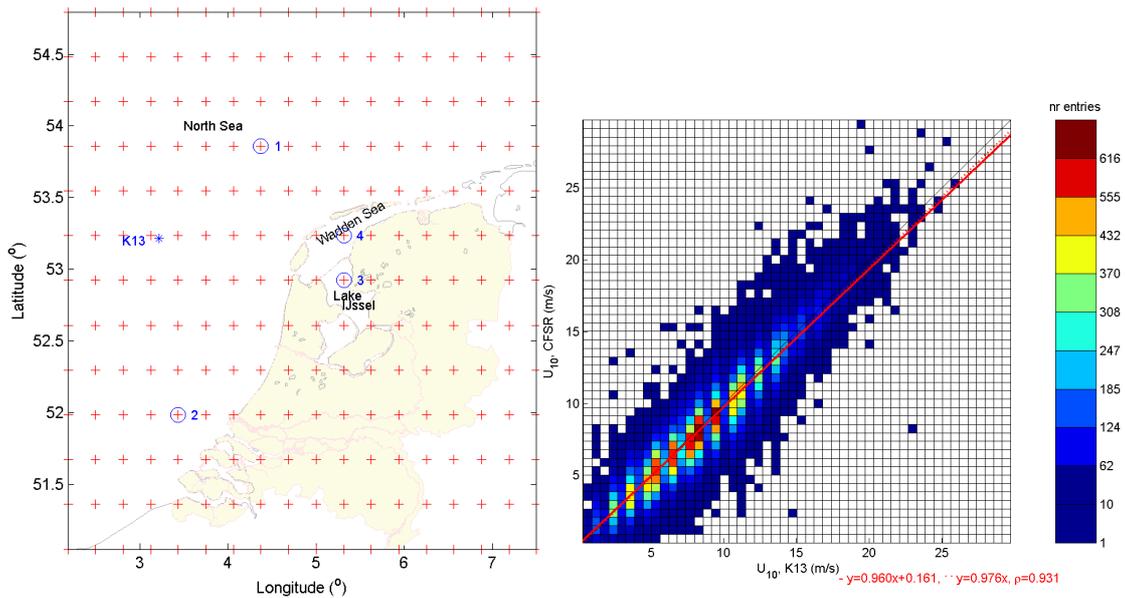


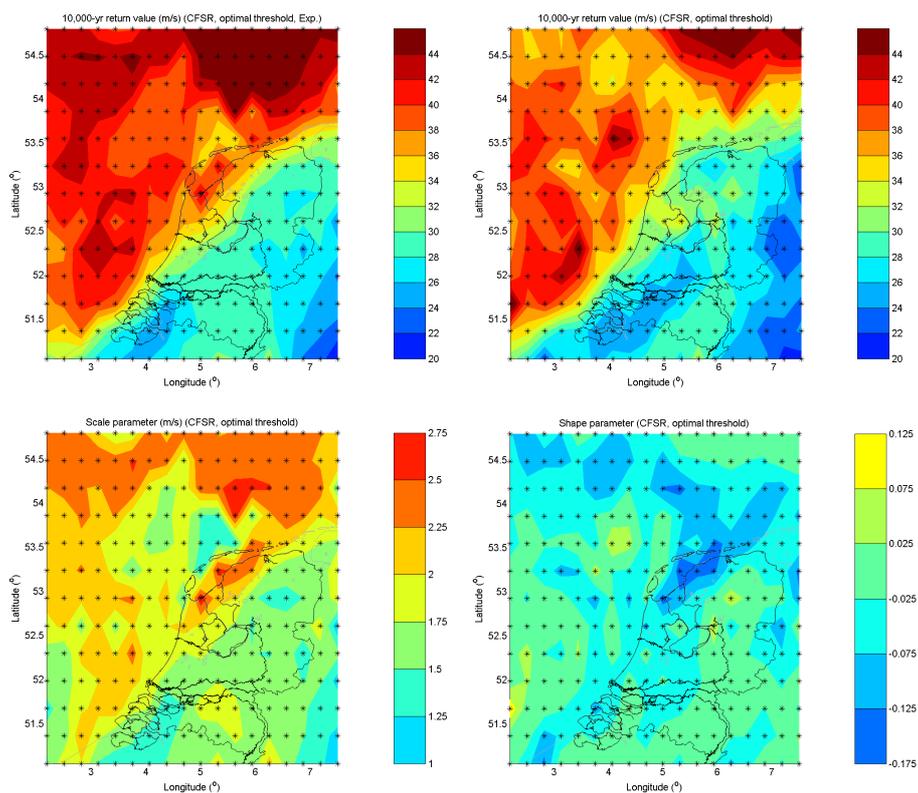
Figure 1. Left: Region covered by the considered CFSR data. Right: Comparison between CFSR and observed wind speeds.

#### Local extreme value analyses

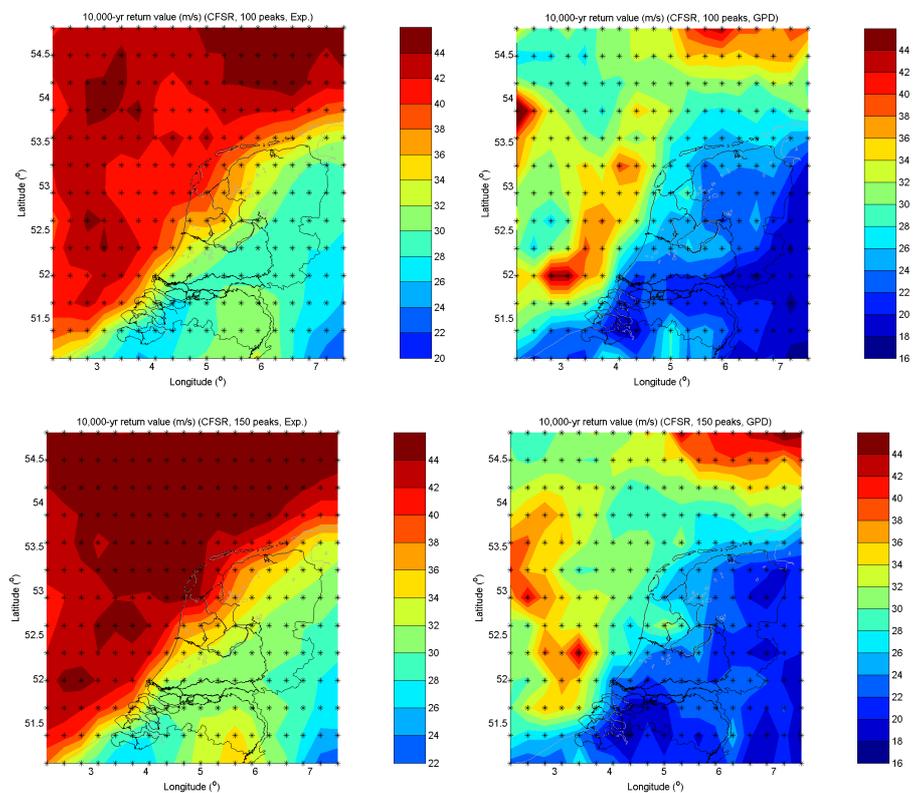
In this section we describe the extreme value analysis carried out per CFSR grid point. The analysis approach consisted of choosing a POT sample and fitting the GPD distribution with and without assuming a type I tail (i.e. fixing the shape parameter equal to zero) to the POT data. Two types of POT samples have been compiled: Using the same fixed number of peaks in each location and using an optimal threshold per location.

In order to choose the optimal threshold the stability property is used and, following the conclusions of Caires (2009), a type I tail is assumed. More precisely, we have looked for threshold values around which the estimate of the GPD shape parameter shows the least variation. Because we are also interested in choosing a threshold for which the GPD shape parameter is close to zero, we have also looked for a threshold within the stable region that would yield the shape parameter estimate closer to zero. Since 234 time series had to be analysed, we automatized the choice of the threshold. The samples used in the choice of threshold are extracted from the original time series in such a way that the data can be modelled as independent observations. This is done by a process of de-clustering in which only the peak (highest) observations in clusters of successive exceedances of a specified threshold are retained and, of these, only those which in some sense are sufficiently apart temporally (so that they belong to more or less 'independent storms') are considered as belonging to the collection of POT points. Specifically, cluster maxima less than 96 h apart are treated as belonging to the same cluster (storm).

Figure 2 shows the exponential and GPD 10,000-yr return value estimates from the automated threshold choice procedure described above. We refer to these estimates as the optimal estimates. Note, however, that the word optimal used here is not intended to suggest that the obtained estimates are as accurate as possible. It just indicates that we have aimed at obtaining a more or less justifiable threshold using the stability property. All chosen thresholds and resulting estimates are crude and have not been analysed in detail. The GPD scale and shape parameter estimates associated with these thresholds are given also in Figure 2. As can be seen in the figure, and as expected, the return values above water are generally higher than those above land. The erratic variations in adjacent grid points indicate that the estimates are spatially unsmooth. This is due to their associated uncertainties. Furthermore, as expected, the exponential estimates are smoother than the GPD estimates. However, given the associated uncertainties, the differences between the exponential and the GPD estimates are not significant. That is mostly because the automatized threshold choice procedure aims at picking a threshold for which the shape parameter estimate is (in the stable region and) close to zero. In fact, the shape parameter estimates are indeed close to zero, varying between -0.17 and 0.12, see bottom left panel of Figure 2.



**Figure 2. Estimates associated with a POT sample collected using the optimal threshold. Top panels: Local 10,000-yr return value estimates of the Exponential distribution (left) and the GPD (right). Bottom panels: Local GPD scale (left panel) and shape (right panel) parameter estimates**



**Figure 3. Top panels: Local exponential distribution (left) and GPD (right) 10,000-yr return value estimates associated with a POT sample with size 100 panel estimates. Bottom panels: Local exponential distribution (left) and GPD (right) 10,000-yr return value estimates associated with a POT sample with size 150.**

In order to check to some extent the effect of the choice of threshold in the results, we have also obtained Exponential and GPD estimates using the same fixed number of peaks in each location. Figure 3 shows the exponential and GPD 10,000-yr return value estimates using a POT sample size of 100 per location (an average of 3.3 storms per year). From the figure it can be inferred that, since the GPD return value estimates are generally lower than those in Figure 3, the shape parameters estimates are in general lower than those using the optimal threshold. The lower shape parameters associated with these POT samples result in higher return value estimates if a Type I tail were assumed. These exponential estimates are smoother than those obtained using the optimal threshold (cf. left top panels of Figure 2 and Figure 3). Note also the rather larger spatial variation of the GPD estimates in Figure 3. Such variations in the wind velocities above water are not realistic, since when a storm crosses the region under consideration wind speeds are not expected to vary so much in intensity more or less intense above certain grid locations. Such spatial variations are only expected in regions with a varying topography. Figure 3 also shows the exponential and GPD 10,000-yr return value estimates using a POT sample size of 150 per location (an average of 5 storms per year). Note that the exponential return value estimates are even higher than when considering 100 peaks, indicating that the associated GPD shape parameter estimates are negative. These estimates are in our opinion too conservative and are not considered further.

### Baseline results

In this section we present the baseline results, which are the results of the application of the methods:

- using as reference location the North Sea location with CFSR grid point coordinates of 4.4°E and 53.9 °E (location 1 in Figure 1),
- assuming that the tail of the data is of type I at every location,
- using the optimal threshold (cf. top left panel of Figure 2), and
- adjusting the distribution parameters to the threshold yielding 100 peaks at each location. Note that the adjustment is also made in the Method Smith (cf. steps H2 and S2) although that is not strictly necessary. The number of peaks is fixed at 100 because the number of peaks using the optimal threshold is generally higher than 100 and because we shall also look at the results of the methods when the 100 peaks sample is used directly to estimate extreme value distribution parameters (as in the top left panel of Figure 3).

Figure 4 shows for the storm period with the highest peak at the reference location (the highest peak of the POT sample at the reference location) three temporal snapshots of the original and uplifted wind speed fields: 6 hours before the peak, the moment of the peak and 6 hours after the peak. The sequential choice of storms is made in the transformed space in the Methods Smith and in the original space in the Method de Haan. Despite this difference in selection the resulting storm rankings (at least for the highest 100 storms) are the same in all applications we considered. For each temporal snapshot the figure contains four panels:

- The top left panel shows the original CFSR wind speed and direction data.
- The bottom left panel shows the time series of the original and uplifted wind speeds from 48h before until 48h after the peak of the storm at the reference location. The vertical line indicates the moment in time to which the surface plots correspond.
- The top right panel shows the uplifted wind speeds using the Method Smith.
- The bottom right panel shows the uplifted wind speeds using the Method de Haan.

The red asterisk in the surface plots indicates the reference location. Note that the wind directions are left unchanged in the upscaling process.

Three clear observations can be made from Figure 4. First, the uplifted fields show more spatial variation (lack of smoothness) than the original fields. Second, for the temporal snapshots shown, the fields uplifted using the Method Smith look rather similar to those using the Method de Haan. Third, the uplifted wind velocities derived with the two methods seem to coincide at the reference location for the three temporal snapshots considered. Further away from the peak of the storm the temporal change in the speeds uplifted using the Method de Haan is larger and more in line with the temporal variation in the original time series (cf. bottom left panels of the figures). Note that, as imposed, for both methods the uplifted wind speed at the peak of the storm at the reference location is equal to the 10,000-yr return value at that location.

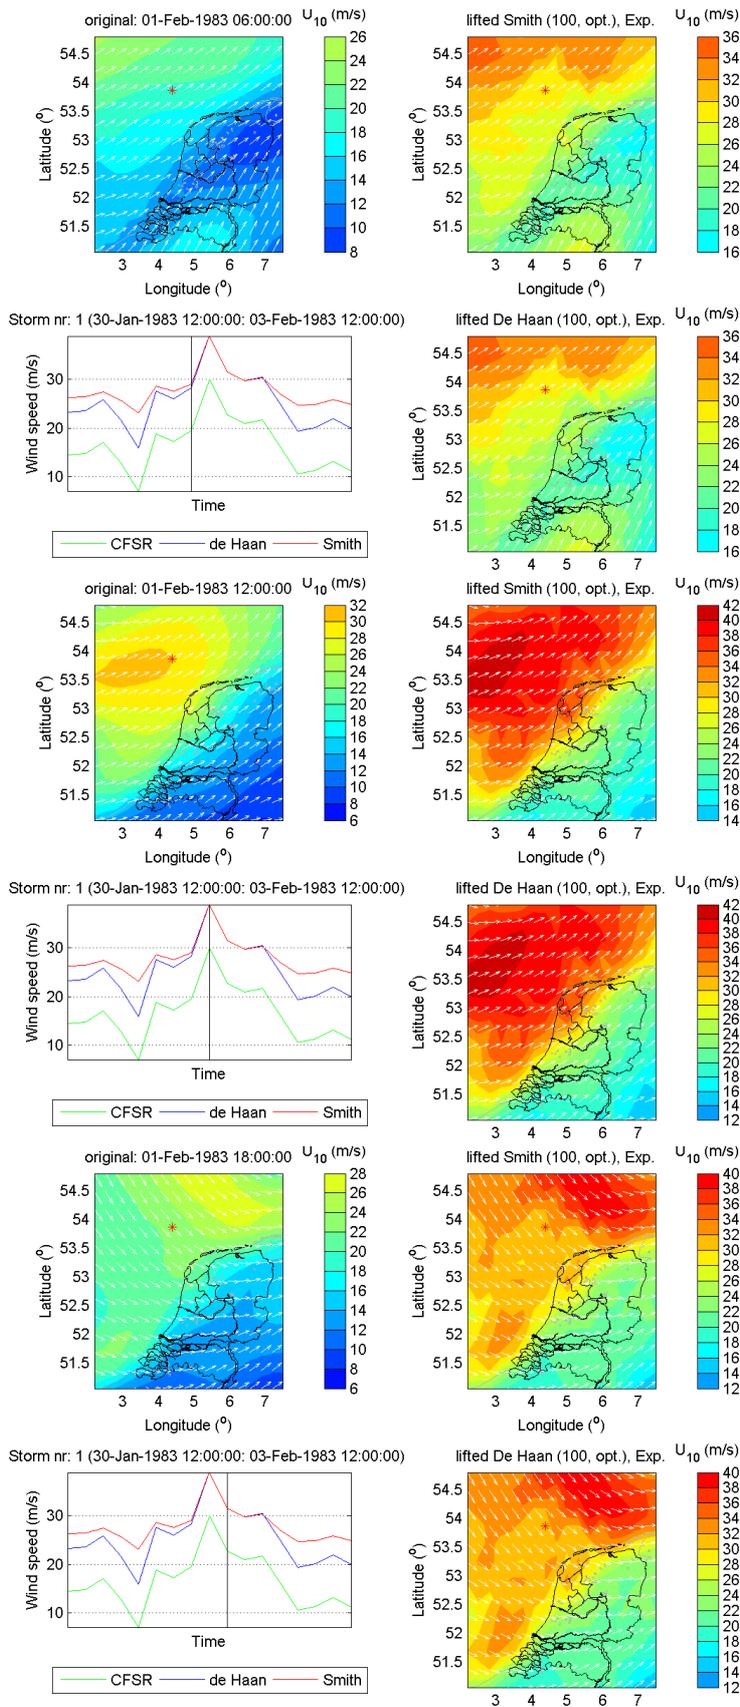


Figure 4. Original and uplifted wind fields for 01-02-1983 from 6 a.m. until 6 p.m., including wind speed time series at the reference location (4.375° E and 53.859° N).

At instants further away from the peak of the storm (figures not shown) the wind speed fields uplifted using the Method de Haan are lower than those using the Method Smith. The differences in the temporal variations of the uplifted wind speeds shown at the reference location (cf. time series in Figure 4) are also found at the other locations.

In order to assess in more detail the difference between the results of the two methods, we have computed histograms of the uplifted wind speed for the storms corresponding to the highest 30 peaks of the POT sample at the reference location and from 48h before until 48h after the peak of the storm. Given that the data is 6-hourly, the total number of values used in each histogram is  $17 \times 30$ . The uplifted wind speed histograms are presented for four locations denoted as Northern, Southern, Lake IJssel and Wadden Sea locations and indicated in Figure 1 by 1, 2, 3 and 4 respectively. The Northern location is the reference location used in the methods. The histograms are shown in Figure 5.

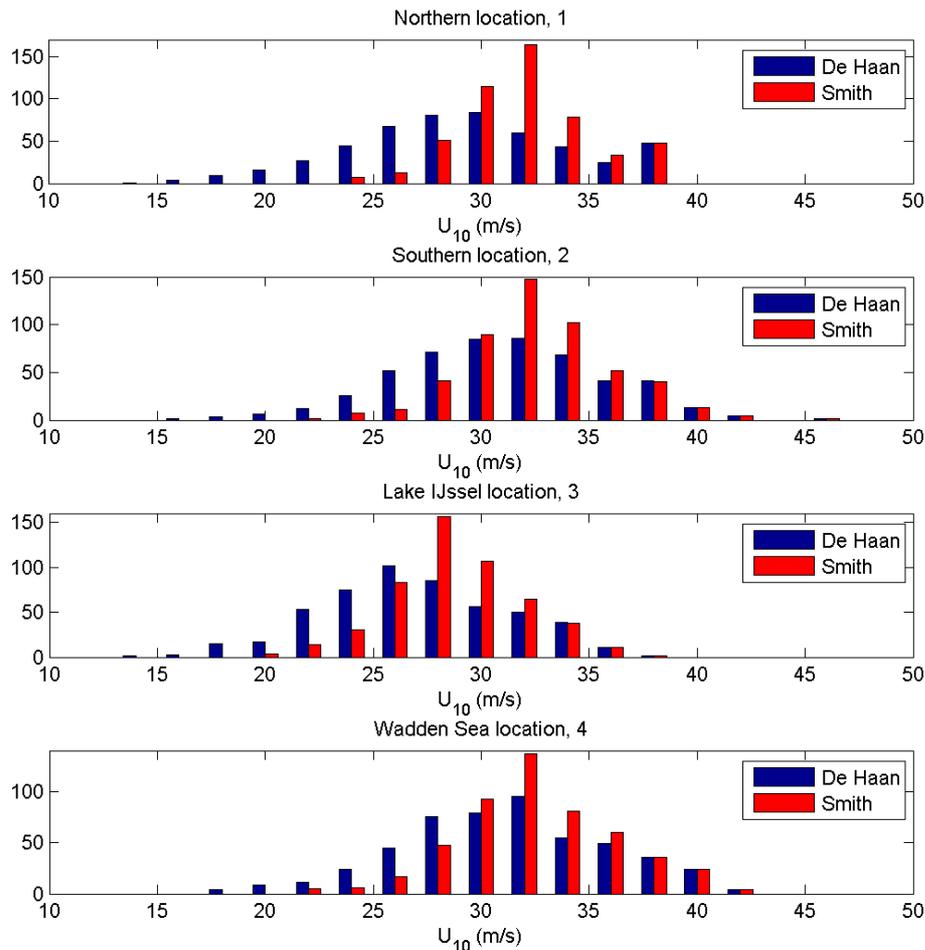


Figure 5. Histograms of the uplifted wind speeds at four locations. Optimal exponential estimates.

As was already expected from the uplifted time series comparisons (Figure 4), the spread of the data resulting from the Method de Haan is larger than that of the data resulting from the Method Smith. At the reference location the highest wind speed bin coincides, because this bin contains the highest uplifted values at this location, i.e. the 10,000-yr return value. Given that the plotted data are for 30 storms, each method produces at least 30 values falling in this bin.

Evaluating, it is difficult to say whether one method can be preferred over the other, but the larger temporal changes of the Method de Haan' uplifted wind speeds, which are close to those in observed storms, may suggest that they are more realistic. Such a conclusion can, however, only be drawn after a thorough validation of the results of the two methods. The Method Smith uplifted wind speeds in principle result in higher loads, which may prove to be more realistic.

Looking at the results in general, one can say that both method produce fields that look realistic, although one should probably try to obtain more spatially smoothly evolving fields. In principle they can be used to drive wave and hydrodynamic models. In general, it can thus be concluded that the

methods of de Haan and Smith can both be used to model the temporal and spatial evolution of multivariate extreme events.

### **Sensitivity studies**

In this section the sensitivity of the results to some of the choices made in the computation of the baseline results is analysed. We look namely at the results of fitting the GPD instead of the exponential distribution to the local wind speed data, of fitting the exponential distribution to the highest 100 POT values at each location and of using other reference locations. For the Method de Haan the sensitivity of the results for augmenting the fitted location-specific baseline model by the empirical distribution of the respective observations below the threshold is also analysed.

#### *Choice of distribution*

In this section, the results of applying the methods as in the previous section (baseline results) but without assuming that the data has a type 1 tail are presented. The presented results were thus obtained using the GPD fits to the local POT samples obtained using the optimal threshold and adjusting the distribution parameters to the threshold yielding 100 peaks at each location.

Figure 6 shows for the identified storm period with the highest peak at the reference location the original and uplifted wind speed fields 6 hours before the peak, the peak, and 6 hours after the peak. In comparison with the baseline results, the uplifted wind speeds are generally higher at locations where the shape parameter estimates are above zero and lower where they are below zero. Furthermore, the spatial variation of the shape parameters leads to larger differences between the original and uplifted fields in terms of spatial smoothness.

Figure 7 shows the histograms of the uplifted wind speeds. At the reference location, the GPD shape parameter estimate is 0.016. This value is rather close to zero and therefore the results of the two methods do not differ much from those in which the shape parameter is made equal to zero (the baseline results). However, because the 10,000 year return value at this location is now higher (40.1 m/s instead of 38.9 m/s) the spread of the histograms is slightly higher (cf. figures 5 and 7). The differences between the results of the two methods are similar to those in the baseline results: the histograms of the results of the Method Smith have a higher mean and are more peaked than those of the Method de Haan. For the Southern location the results are also rather close to the baseline results. The shape parameter estimate is -0.013 which explains the closeness of the results. For the Lake IJssel and Wadden Sea locations the shape parameter estimates are -0.106 and -0.168 respectively. These light tails lead to lower means and more peaked histograms (cf. bottom panels of Figure 7). Furthermore, from the presented results it can be concluded that the lighter the tail the closer the results of the methods.

In conclusion, these results show that the higher uncertainties of the local GPD fits in comparison to the local exponential fits lead to larger differences between the original and uplifted fields in terms of spatial smoothness. Furthermore, the lighter the GPD tails at a certain location the closer the uplifted local wind speeds of the two methods (De Haan vs. Smith).

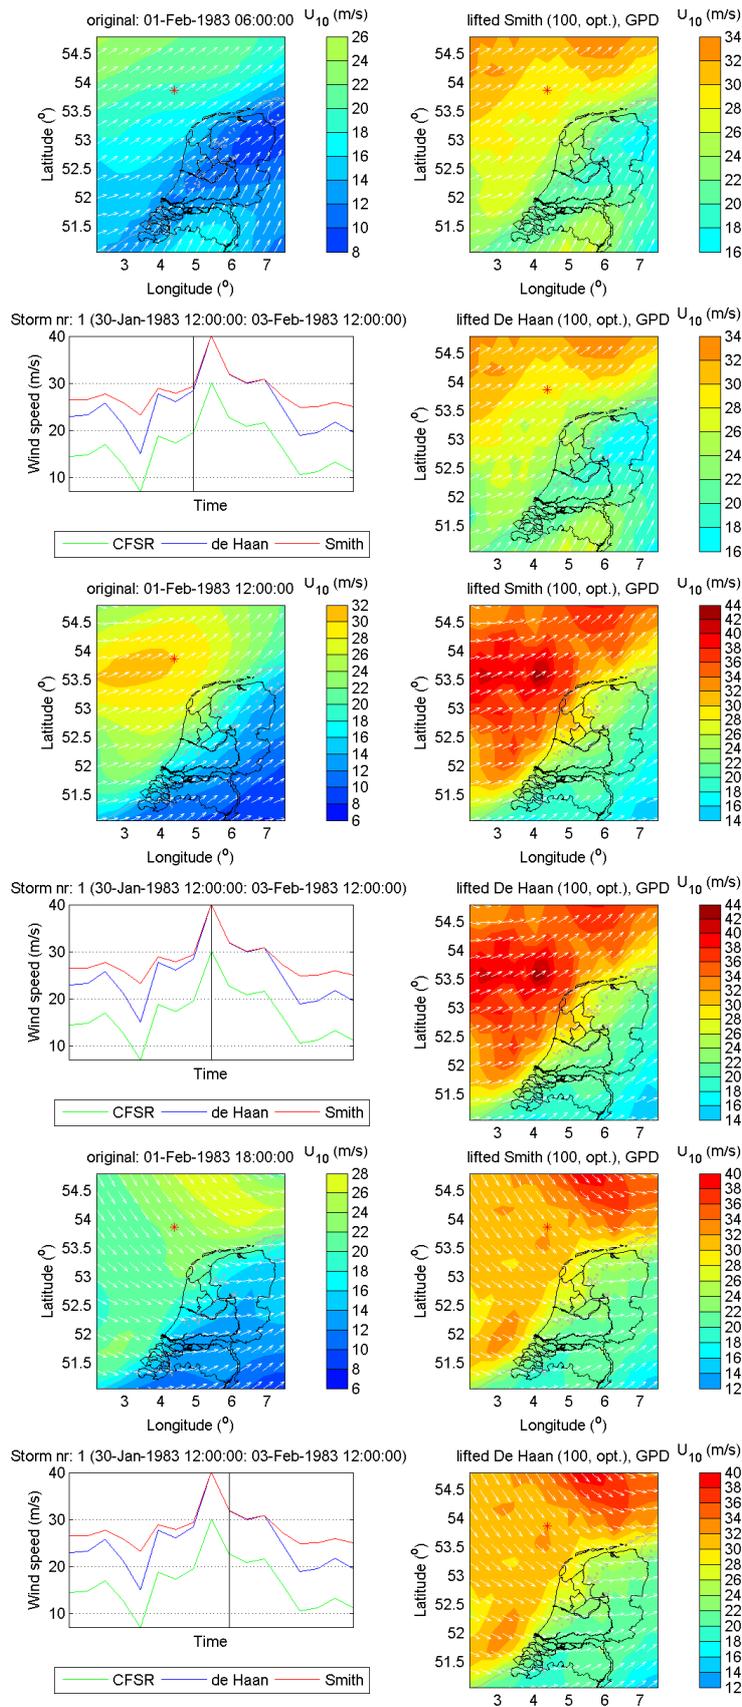


Figure 6. Original and uplifted wind fields for 01-02-1983 from 6 a.m. until 6 p.m., including wind speed time series at the reference location (4.375°E and 53.859°N). Optimal GPD estimates.

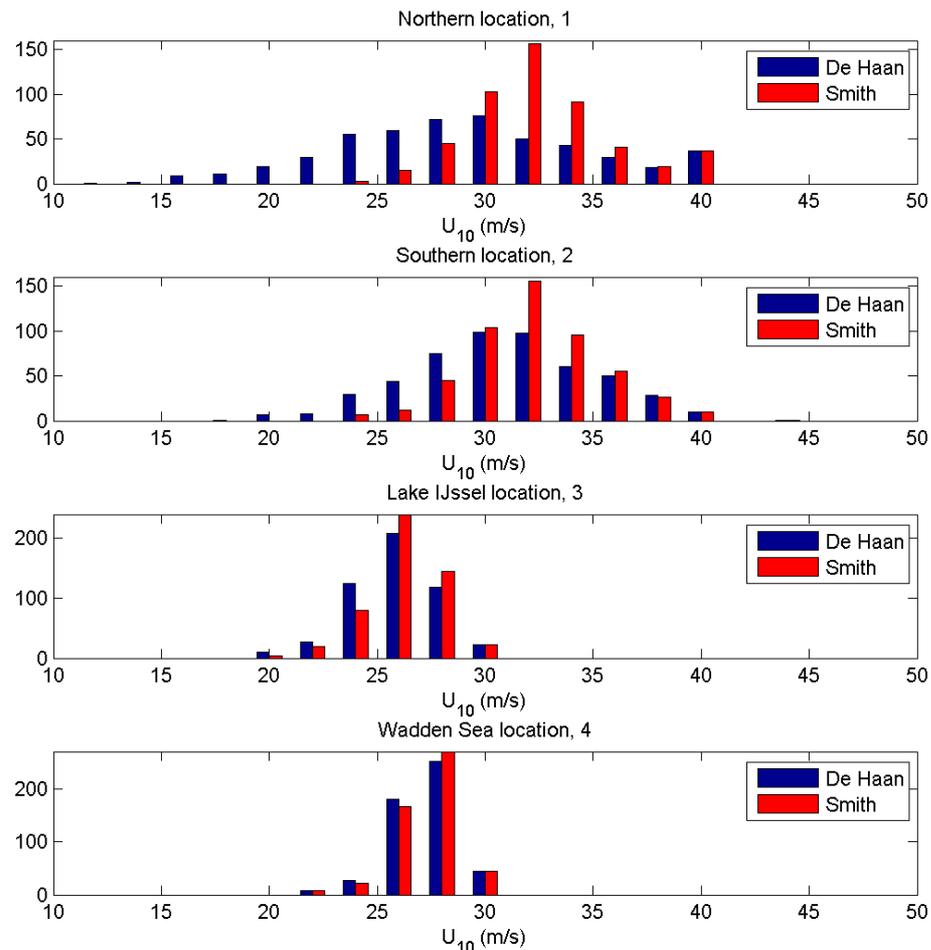


Figure 7. Histograms of the uplifted wind speeds at four locations. Optimal GPD estimates.

#### Choice of threshold

In this section the results of fitting an exponential distribution to the highest 100 wind speed POT values at each CFSR location and using these tail estimates in the methods are presented. These computations differ thus from those of the baseline results in the choice of threshold, since in this case no optimal thresholds have been used.

The uplifted wind speed fields show a spatial pattern closer to the one of the original CFSR fields (figures not shown), which suggests that they may be more realistic. Of course the opposite is true for the GPD estimates using 100 peaks at each location (not shown), since the large spatial variation of the shape parameter estimates leads to significant spatial variation in the uplifted fields (no spatially smooth propagating extreme wind fields).

Figure 8 presents the uplifted wind speed histograms for the Northern (reference), Southern, Lake IJssel and Wadden Sea locations, respectively. The histograms are similar to those of the baseline results, though the results at the Northern, Southern and Lake IJssel locations are higher and those at the Wadden Sea location are lower.

In conclusion it can be said that, relative to the baseline results, fitting an exponential distribution to the highest 100 peaks at each location leads to uplifted fields with spatial variations closer to those of the original fields.

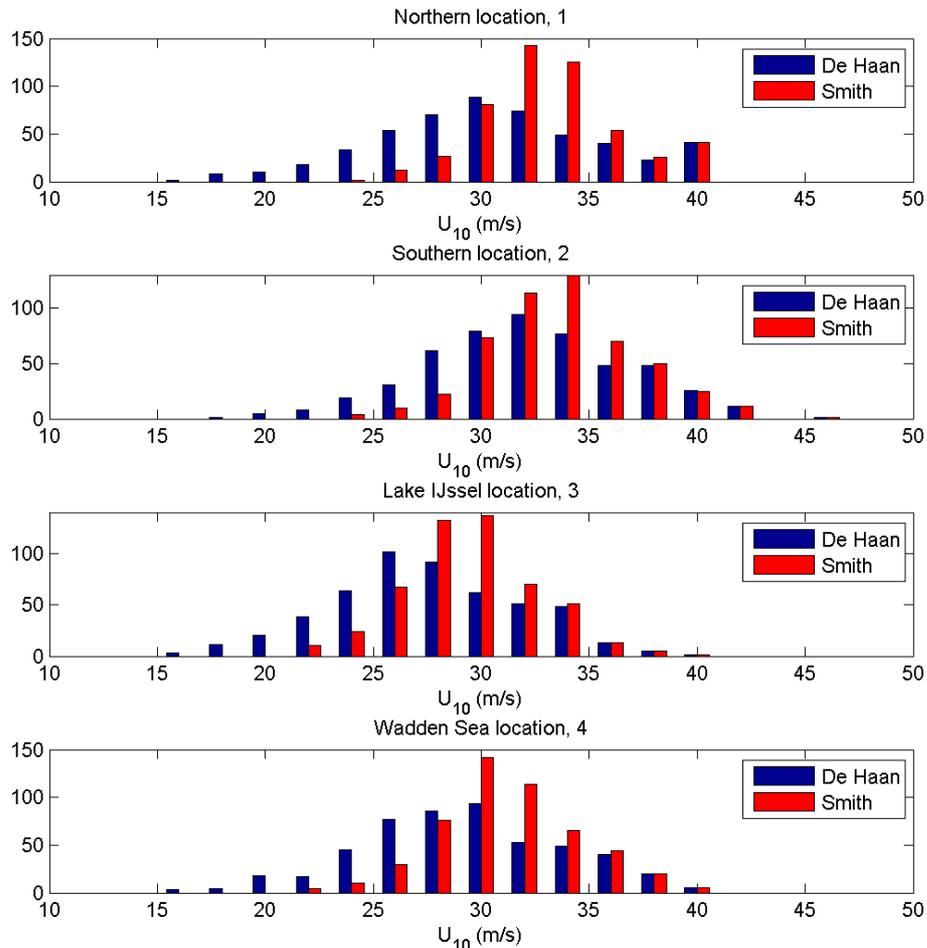


Figure 8. Histograms of the uplifted wind speeds at four locations. Exponential estimates based on local fits using 100 peaks.

#### Choice of reference location

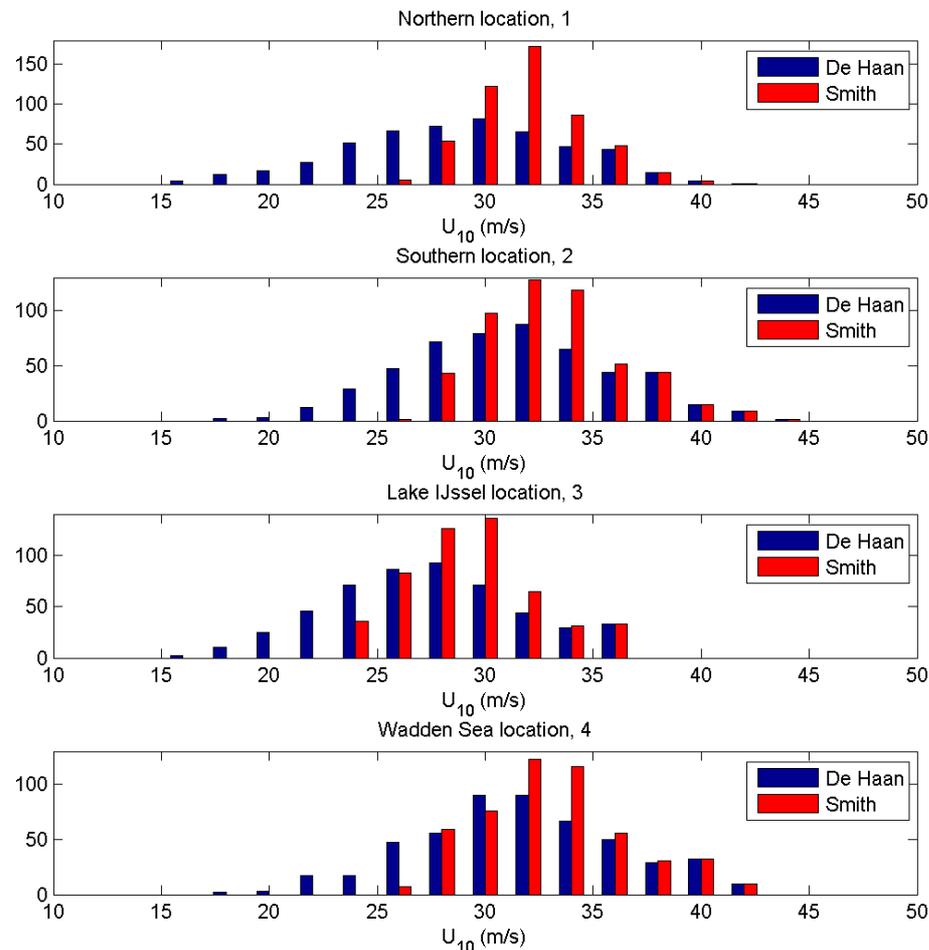
In this section, the effects of choosing a different reference location are analysed. The presented results were computed using the same fits and samples as those used in the computation of the baseline results, but using another reference location. The reference location used in these computations is the Lake IJssel location.

The dates of the identified highest four storm peaks (ranked from the highest to the lowest peak wind speeds before the lifting) are given in Table 1. The ranking of the storms differs thus from that of the baseline results, in which the Northern location is the reference location. The storms nr. 2 and nr. 3 in these results are respectively storms nr. 4 and nr. 1 in the baseline results. Comparing the uplifted fields of these storms (not shown) one can conclude that the spatial patterns are similar, but the magnitude of the uplifted fields is generally higher when the Lake IJssel location is used instead of the Northern location as reference.

Table 1. Ranking of the highest four storm peaks using North Sea and Lake IJssel locations as reference locations.		
	North Sea ref. loc. (Baseline)	Lake IJssel ref. loc.
Storm nr. 1	1-2-1983 12h	25-1-1990 18h
Storm nr. 2	19-12-1982 18h	14-1-1984 18h
Storm nr. 3	16-10-1987 12h	1-2-1983 12h
Storm nr. 4	14-1-1984 18h	18-1-2007 12h

Figure 9 shows the uplifted wind speed histograms for the Northern, Southern, Lake IJssel (reference) and Wadden Sea locations, respectively. The histograms are similar to those of the baseline results, but since the Northern location is in this case no longer the reference location, the histograms

contain values higher than the 10,000-yr return value at the Northern location and the histograms for the Lake IJssel no longer contain values above the its 10,000-yr return value.



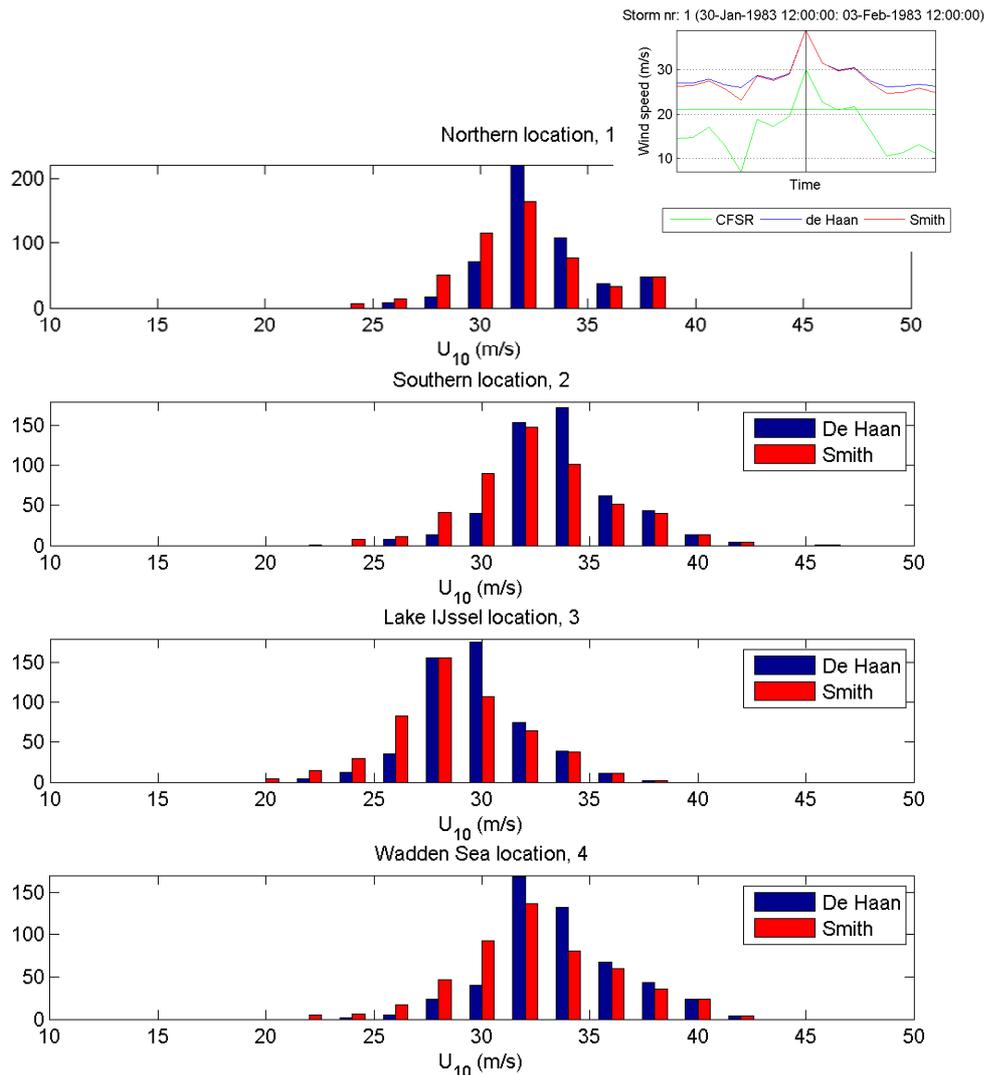
**Figure 9. Histograms of the uplifted wind speeds at four locations. Optimal exponential estimates using the Lake IJssel location as the reference location.**

The main conclusions to this analysis are that the choice of another reference location may influence the ranking of the selected storms, the severity of the uplifted fields and the distribution of the wave loads. Once restricting the number of considered storms to the 30 highest in terms of the POT values at the reference location, the population of storms obtained using difference reference locations may differ significantly.

#### *Augmentation of the extreme value distribution with the empirical distribution in Method de Haan*

The results of method de Haan presented so far have been obtained by applying the GPD distribution to observation both above and below the thresholds. Strictly speaking, the latter is invalid. Therefore the effects of augmenting the extreme value distribution with the empirical distribution of the observation below the threshold in the Method de Haan are analysed in this section.

Figure 10 presents the uplifted wind speed histograms for the Northern (reference), Southern, Lake IJssel and Wadden Sea locations, respectively. The wind speed time series at the reference location is plotted along the respective histogram. The figure shows that when also augmenting the exponential in the Method de Haan the results of the methods of Smith and de Haan are really close to each other. The observations below the local threshold are further lifted when augmenting the extreme value distribution in Method de Haan and the results of the two methods are close even away from the peak (cf. the wind speed time series in figures 4 and 10). We are not yet in a position to comment on which temporal evolution is more accurate and will, therefore, still consider both versions of the Method de Haan in further studies.



**Figure 10. Histograms of the uplifted wind speeds at four locations. The wind speed time series at the reference location is plotted along the respective histogram. Optimal exponential estimates augmenting the fitted location-specific exponential with the empirical distribution of the observations below the threshold.**

## CONCLUSIONS AND RECOMMENDATIONS

The main conclusions of this study are that:

- The uplifted wind fields show more spatial variation (or noise) than the original fields. This is probably due to the uncertainties in the extreme value analysis per location.
- The methods produce rather similar uplifted wind speed fields around the peak of the storm.
- When considering the uplifted wind speeds from 48 hours before until 48 hours after the peak of the storm, the Method Smith uplifted wind speeds have a higher mean and a more peaked histogram than those of the Method de Haan. More precisely, the uplifted wind speeds of the two methods coincide for values close to the peak values, further away from the peak (for lower original time series values) the results of the Method de Haan are lower than those of the Method Smith. In other words, the time evolution of the uplifted Method de Haan wind speeds show stronger increases and decays in time, whereas the Method Smith wind speeds remain at a higher level closer to the peak value. However, if in the Method de Haan the locally-fitted extreme value distributions are augmented below the threshold with the respective empirical distributions, then the time evolution of the uplifted wind speeds of both methods (including uplifted wind speeds originally below the threshold) are similar.

- The higher uncertainties associated with using the GPD distribution instead of the exponential distribution for the local fits lead to larger differences between the original and uplifted fields of both methods in terms of spatial smoothness. Furthermore, the lighter (the less extreme) the GPD tails at a certain location the closer the uplifted local wind speeds of the two methods.
- For both methods, when assuming a type I tail, different choices of the threshold may lead to uplifted wind speed fields with a spatial pattern smoother and closer to the one of the original fields.
- For both methods, different reference locations may lead to a different ranking of (and population of) chosen storms, uplifted fields with higher magnitudes and a different distribution of the associated wave loads.

To motivate our recommendations, we would like to point out that in this first application of the methods only omni-directional statistics have been determined. Because only omni-directional statistics have been determined, the magnitude of the lifting is fully determined by the omni-directional statistics. In other words, the magnitude of the lifting is the same for easterly wind directions as for westerly wind directions, which is not in accordance with reality since easterly conditions are much milder than westerly conditions in the Netherlands. In order to obtain more realistic estimates of extreme wind fields, directional estimates of the local fits need to be obtained.

Furthermore, in order to determine spatially and directionally smoother extreme statistics, more attention should be given to the spatial continuity of the estimates of the extreme value distribution parameters per location and variable. In principle the extremes of the wind and associated hydrodynamic variables are not homogeneous in direction (or time, or season) and the extreme value distribution estimations must be direction and probably also seasonal dependent. Given the number of analyses that need to be carried out, automatic techniques and methods that lead to smooth variations in space, direction and time of the estimate should therefore be investigated.

Another limiting aspect of this study is that the results of the methods of Smith and de Haan have only been assessed in terms of comparisons of the uplifted time series. A more thorough assessment of the results of these methods should involve using the uplifted fields to force the wave and hydrodynamic models used in the determination of the hydraulic boundary conditions and assessment of the model results.

Last but not least, an assessment of the results of these statistical methods in terms of the underlying physics (including their time evolution) should also to be carried out.

#### ACKNOWLEDGMENTS

This work is part of the Strength and Loads on Water Defences project commissioned by Rijkswaterstaat Centre for Water Management in the Netherlands. We are thankful to NOAA's National Centers for Environmental Prediction (NCEP) for the CFSR data. We are indebted to Prof. de Haan and Prof. Smith for providing us with the methods and to Prof. Anderson for his suggestions on the application of the methods.

#### REFERENCES

- Buishand, T.A., L. de Haan and C. Zhou. 2008. On spatial extremes with application to a rainfall problem. *The Annals of Applied Statistics*, 2(2), 624–642, DOI: 10.1214/08-AOAS159.
- Caires, S., 2009: *Extreme wind statistics for the Hydraulic Boundary Conditions for the Dutch primary water defences. SBW-Belastingen: Phase 2 of subproject 'Wind modelling'*. Deltares Report 1200264-005.
- Caires, S., L. de Haan and R. L. Smith. 2011. *On the determination of the temporal and spatial evolution of extreme events*. Deltares report 1202120-001-HYE-004 (for Rijkswaterstaat, Centre for Water Management).
- Saha, Suranjana, and Coauthors. 2010. The NCEP Climate Forecast System Reanalysis. *Bull. Amer. Meteor. Soc.*, 91, 1015.1057. doi: 10.1175/2010BAMS3001.1.