

EXTREME VALUE ANALYSIS: WAVE DATA

by

Sofia Caires

2011

JCOMM Technical Report No. 57

EXTREME VALUE ANALYSIS: WAVE DATA

by

Sofia Caires

2011

JCOMM Technical Report No. 57

NOTE

The designations employed and the presentation of material in this publication do not imply the expression of any opinion whatsoever on the part of the Secretariats of the Intergovernmental Oceanographic Commission (of UNESCO), and the World Meteorological Organization concerning the legal status of any country, territory, city or area, or of its authorities, or concerning the delimitation of its frontiers or boundaries.

CONTENTS

Introduction	1
Methodologies	2
Introduction	2
Extreme value theory	2
Block maxima	2
Peaks Over Threshold	4
r-largest	6
Other approaches	6
Choice of distribution	6
Initial-distribution approach	7
Estimation and diagnosis	8
Climate change and variability	9
Worked examples	10
Introduction	10
NDBC deep water data	10
Data availability	11
POT/GPD analysis	14
POT/Exponential analysis	15
POT/Weibull analysis	17
AM/GEV analysis	18
AM/Gumbel analysis	19
Estimation method	19
Stationarity	20
North Sea shallow water data	21
POT/GPD analysis	22
POT/Exponential analysis	24
POT/Weibull analysis	24
AM/GEV analysis	25
AM/Gumbel analysis	26
Estimation method	27
Stationarity	27
Software packages	29
Recommendations	30
Extreme value analysis	30
Further aspects	30
References	32

Introduction

Estimates of the m -year return value of significant wave height—the value which is exceeded on average once every m years—are needed for the safety control and design of ship, offshore, and coastal structures, and for the mapping of flood risk areas. The WMO Guide to Wave Analysis and Forecasting aims at providing guidance on how to obtain those estimates. In the design of ships and offshore platforms 1/20-yr to 1/100-yr return values are often used. In the control of the safety of the Netherlands sea defenses return values of up to 1/10,000-yr are used. In the mapping of flood risk area in the United Kingdom 1/1,000-yr return values are used. The longer time series of significant wave height available come from hindcasts and usually cover no more than 50 years, meaning that one generally needs to extrapolate well beyond the range of the available data and thus resort to extreme value analysis to obtain the required return value estimates.

In this report we begin by describing and discussing approaches that can be used to estimate such return values in Chapter 2. Approaches based on extreme value theory as well as *ad hoc* methods are considered. We then present in Chapter 3 some worked examples using two time series of significant wave height measurements, one in deep and the other in shallow waters. In Chapter 4 we provide an inventory of software packages available to carry out extreme value analyses. We finish in Chapter 5 with some guidelines / recommendations.

Methodologies

Introduction

This chapter briefly introduces the principles of extreme value theory and describes the methods used in extreme value analysis (EVA). Other approaches to estimate return values will also be considered. For further background information on extreme value theory and analyses we recommend the book of Stuart Coles (Coles, 2001), which is comprehensive, easy to read and presents many applications to environmental data.

Extreme value theory

Extreme value theory provides analogues of the central limit theorem for the extreme values in a sample. According to the central limit theorem, the mean of a large number of random variables, irrespective of the distribution of each variable, is distributed approximately according to a Gaussian distribution. For example, the sea surface elevation is often modelled as a sum of several individual random waves and accordingly its distribution is often assumed to be Gaussian. According to extreme value theory, the extreme values in a large sample have an approximate distribution that is independent of the distribution of each variable.

Block maxima

In order to explain the basic ideas, let us define $M_n = \max\{X_1, \dots, X_n\}$, where X_1, X_2, \dots is a sequence of independent random variables having a common distribution function F . In its simplest form, the *extremal types theorem* states the following: If there exist sequences of constants $\{\sigma_n > 0\}$ and $\{\mu_n\}$ such that $P\{\sigma_n M_n + \mu_n \leq z\} \rightarrow G(z)$ as $n \rightarrow \infty$, where G is a *non-degenerate cumulative distribution function*¹, then G must be a generalized extreme value (GEV) distribution, which is given by

$$G(z) = \begin{cases} \exp\left\{-\left[1 + \xi\left(\frac{z - \mu}{\sigma}\right)\right]^{-1/\xi}\right\}, & \text{for } \xi \neq 0 \\ \exp\left\{-\exp\left[-\left(\frac{z - \mu}{\sigma}\right)\right]\right\}, & \text{for } \xi = 0, \end{cases} \quad (2.1)$$

where z take values in three different sets according to the sign of ξ : $z > \mu - \sigma/\xi$ if $\xi > 0$ (the domain of z has a lower bound), $z < \mu - \sigma/\xi$ if $\xi < 0$ (the domain of z has an upper bound), and $-\infty < z < \infty$ if $\xi = 0$.

1. A distribution function is said to be degenerate if it allocates probability 1 to a single point.

In other words, if the sequence of distribution functions of (normalizations of) the maximum value in a random sample of size n converges to a (single) distribution function as n tends to infinity, then that distribution function must be a GEV distribution. Moreover, this and other results of extreme value theory *hold true even under general dependence conditions* (Coles, 2001).

In Eq. (2.1), the parameters μ , σ and ξ are called the location, scale, and shape parameters and satisfy $-\infty < \mu < \infty$, $\sigma > 0$ and $-\infty < \xi < \infty$. For $\xi = 0$ the GEV is the Gumbel distribution, for $\xi > 0$ it is the Fréchet distribution, and for $\xi < 0$ it is the Weibull distribution (of maxima). For $\xi > 0$ the tail of the GEV is “heavier” (i.e., decreases more slowly) than the tail of the Gumbel distribution, and for $\xi < 0$ it is “lighter” (decreases more quickly and actually reaches 0) than that of the Gumbel distribution. The GEV is said to have a type II tail for $\xi > 0$ and a type III tail for $\xi < 0$ ². The tail of the Gumbel distribution is called a type I tail.

The extremal types theorem gives rise to the *annual maxima* (AM) method of modelling extremes, in which the GEV distribution is fitted to a sample of block maxima (e.g. to annual maxima, though biannual, seasonal, monthly or even daily maxima can of course be used as well).

One of the main applications of extreme value theory is the estimation of the once per m year (1/ m -yr) return value, the value which is exceeded on average once every m years. The 1/ m -yr return value based on the AM method/GEV distribution, z_m , is given by

$$z_m = \begin{cases} \mu - \frac{\sigma}{\xi} \left(1 - \left\{ -\log \left(1 - \frac{1}{m} \right) \right\}^{-\xi} \right), & \text{for } \xi \neq 0 \\ \mu - \sigma \ln \left\{ -\log \left(1 - \frac{1}{m} \right) \right\}, & \text{for } \xi = 0. \end{cases} \quad (2.2)$$

The sample sizes of annual maxima data are usually small, so that model estimates, especially return values, have large uncertainties. This has motivated the development of more sophisticated methods that enable the modelling of more data than just block maxima. These methods are based on two well-known characterizations of extreme value distributions: one based on exceedances of a threshold, and the other based on the behaviour of the r largest, for small values of r , observations within a block. These are described in the following two sub-sections.

2. Please note that some articles (e.g. Hosking and Wallis, 1987) use another convention for the signal of the shape parameter: a negative shape parameter in those references corresponds to a type II distribution.

Peaks Over Threshold

The approach based on the exceedances of a high threshold, hereafter referred to as the POT (Peaks Over Threshold) method, consists of fitting the generalized Pareto distribution (GPD) to the peaks of clustered excesses over a threshold, the excesses being the observations in a cluster minus the threshold, and calculating return values by taking into account the rate of occurrence of clusters (see Pickands, 1971 and 1975, and Davidson and Smith, 1990). Under very general conditions this procedure ensures that the data can have only three possible, albeit asymptotic, distributions (the three forms of the GPD given below) and, moreover, that observations belonging to different peak clusters are (approximately) independent. In the POT method, the peak excesses over a high threshold u of a time series are assumed to occur in time according to a Poisson process with rate λ_u and to be independently distributed as a GPD, whose distribution function is given by

$$F_u(y) = \begin{cases} 1 - \left(1 + \xi \frac{y}{\sigma_u}\right)^{-1/\xi}, & \text{for } \xi \neq 0 \\ 1 - \exp\left(-\frac{y}{\sigma_u}\right), & \text{for } \xi = 0, \end{cases} \quad (2.3)$$

where $0 < y < \infty$, $\sigma_u > 0$ and $-\infty < \xi < \infty$. The two parameters of the GPD are called the scale (σ_u) and shape (ξ) parameters. When $\xi = 0$ the GPD is said to have a type I tail and amounts to the exponential distribution with mean σ_u ; when $\xi > 0$ it has a type II tail and it is the Pareto distribution; and when $\xi < 0$ it has a type III tail and it is a special case of the beta distribution. If $\xi < 0$, just as with the GEV, the support of the GPD was an upper-bound, $-\sigma_u/\xi$, which is called the *upper end-point* of the GPD. The significance of this upper end-point is that (because when $\xi < 0$ one must have $x < -\sigma_u/\xi$ in Eq. (2.3)) the excesses over u modelled by the GPD cannot take values greater than $-\sigma_u/\xi$, which in turn means that the exceedances of the variable of interest cannot exceed the value

$$x^* = u - \sigma_u/\xi. \quad (2.4)$$

This parameter x^* is to be thought of as the *upper-limit of the variable of interest* (e.g. of H_s).

The 1/m-yr return value based on a POT/GPD analysis, z_m , is given by

$$z_m = \begin{cases} u + \frac{\sigma_u}{\xi} \{ (\lambda_u m)^\xi - 1 \}, & \text{for } \xi \neq 0 \\ u + \sigma_u \log(\lambda_u m), & \text{for } \xi = 0. \end{cases} \quad (2.5)$$

Note that this expression is obtained from Eq. (2.3) by solving $(1 - F_u(y)) = \frac{1}{\lambda_u m}$ for y and then adding the threshold u to the result.

Just as block maxima have the GEV as their approximate distribution, the threshold excesses have a corresponding approximate distribution within the GPD family. Moreover, the parameters of the GPD of threshold excesses are uniquely determined by those of the associated GEV distribution of block maxima. In particular, the shape parameter is the same, and the scale parameters of the two distributions are related by

$$\sigma_u = \sigma + \xi(u - \mu). \quad (2.6)$$




Figure 2.1 Example of the POT sample selection. The green stars indicate the observations exceeding the threshold and the red stars the selected POT points.

The sample to be used in the POT method has to be 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 declustering 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 (so that they belong to more or less ‘independent storms’) are considered as belonging to the collection of POT points. Specifically, in the present applications we have treated cluster maxima at a distance of less than 48 h apart as belonging to the same cluster (storm). Figure 2.1 shows an example of the POT sample selection. The green stars indicate the observations exceeding the threshold and the red stars the selected POT points. Note that only the cluster maxima are retained and that the peak of the first and third clusters were in this case not considered because they are less than 48h apart from the peaks of the second and fourth clusters, respectively.

The choice of threshold (analogous to the choice of block size in the block maxima approach) represents a trade off between bias and variance: too low a threshold is likely to violate the asymptotic basis of the model, leading to bias; too high a threshold will generate fewer excesses with which to estimate the model, leading to high variance. An important property of the POT/GPD approach is the threshold stability property: if a GPD is a reasonable model for excesses of a threshold u_0 , then for a higher threshold u a GPD should also apply; the two GPD's have identical shape parameter and their scale parameters are related by $\sigma_u = \sigma_{u_0} + \xi(u - u_0)$, which can be reparameterized as

$$\sigma^* = \sigma_u - \xi u. \quad (2.7)$$

Consequently, if μ_0 is a valid threshold for excesses to follow the GPD then estimates of both σ^* and ξ should remain nearly constant above μ_0 . This property of the GPD can be used to find the minimum threshold at which a GPD model applies to the data.

r-largest

The characterizations of extreme value distributions based on the behaviour of the r largest observations within a block is known as the r -largest approach. This approach is not often used in practice and is therefore not described in detail nor applied in this study.

Briefly, it consists of collecting the r -largest values per year (instead of merely the annual maxima) and fitting the r -largest distribution (see, for instance, Coles, 2001, pp. 68) to the data. An example of the application of this method to estimate return values of significant wave height is given by Guedes Soares and Scotto (2004).

The choice of r is analogous to the choice of threshold in the POT method and the choice of block size in the block maxima approach.

Other approaches

Choice of distribution

The choice of distribution to fit AM and POT data is limited by the extreme value theory as mentioned above. However, mostly for historical reasons, in many studies the Weibull distribution of minima instead of the GPD is fitted to POT data.

This Weibull distribution is not the Weibull distribution of *maxima* referred to in Eq. (2.1) for $\xi < 0$, but the Weibull distribution of minima (a form of the GEV distribution for minima). In particular, while the latter has a Type III upper tail, the former has a Type I (exponential) upper tail. It was introduced by Weibull in connection with failure data because it is the approximate distribution of the *minimum* of many variables, which could be seen as the

weakest link among many links that can be broken in a structure. Popularized by reliability engineers, its use has spread to other areas, in particular to ocean engineering. It is given by

$$F_u(y) = 1 - \exp\left(-\left(\frac{y}{a}\right)^c\right). \quad (2.8)$$

Although not generally advisable, there are certain situations in which it makes sense to fit a Weibull distribution of minima to the peak excesses in place of the GPD model. Indeed, suppose that the data really follow a Type I tail, or at least that this has been convincingly demonstrated on the basis of some statistical analyses (which is often the case in deep waters; see Caires and Sterl, 2005). Then the asymptotic distribution of the excesses is exponential. Since the exponential is a special case of the Weibull distribution of minima, one might think that there would be no harm in fitting a Weibull rather than an exponential to the data. Now, if the data are truly exponential, this would actually entail more uncertainty in parameter estimates, which would be undesirable (intuitively, to know that the data are exactly exponential amounts to more information than knowing that they are Weibull). However, it may happen that, because the exponential is only valid asymptotically, the Weibull distribution will provide a better approximation to the data (since it has one more parameter and hence more flexibility), and in that case fitting the latter would provide better results than fitting the former. In any case, if one is to step outside the GPD domain one should do so on the basis of some justification.

Still in this connection, we should add that, because wave data most often exhibits a type I tail or a slightly lighter type III tail, studies comparing estimates from the GPD with estimates from the Weibull are often inconclusive as to what is the most appropriate tail/distribution (e.g. Van Vledder et al., 1993) since there are no statistically significant differences between the two models. However, a good reason for always considering the GPD for POT data and the GEV for AM data is that they have a substantial and solid theoretical basis stemming from asymptotic considerations.

Initial-distribution approach

Before the popularization of extreme value theory methodology for the extreme value analysis of wave data, a method that was often used to obtain return values estimates was the so-called Initial-distribution approach (Lapatoukhin et al., 2000 and Holthuijsen, 2007). It consists of fitting a distribution to the whole data (not only the extremes) and estimating the high percentiles of such a distribution. Two widely used distributions are the log-normal distribution and the Weibull distribution (see e.g. Guedes Soares and Henriques, 1996).

The Initial-distribution approach is not advisable. The arguments against, which are well exposed by Ferreira and Guedes Soares (1998, 2000) and Anderson et al. (2001), are:

- Due to dependence and nonstationarity, metocean time series violate the assumptions of independence and identity in distribution, which invalidates the application of the common statistical methods used (confidence intervals and tests) as well as the definition of return value.
- There is no scientific justification for using one particular distribution to fit the data being considered (e.g.: significant wave height data), and the usual goodness-of-fit diagnostics are not able (on the basis of realistic sample sizes and given the length of the required “prediction horizon”) to distinguish data with type I (exponential) tail, say,

from data with type II (heavier than exponential) tail. In contrast, if for example one concentrates on averages, maximum values, or excesses over a high threshold of very general variables, then statistical theory provides a scientific basis for the use of, respectively, the normal, GEV and generalized Pareto distributions.

Estimation and diagnosis

There are several numerical methods available for the estimation of the parameters of extreme value distributions. Most of them, for instance the methods of moments and of probability weighted moments (PWM), give explicit expressions for the parameter estimates. The maximum likelihood (ML) method tends to be the preferred estimation method since it is quite general and more flexible than other methods, especially when the number of parameters is increased as for instance when extending the extreme value approach to account for non-stationarity. However, in ordinary extreme value analyses like the ones we are concerned with in this report the flexibility provided by the ML method is not necessary, and for the range of tails typically found with wave data (not too heavy-tailed distributions) and for small to moderate sample sizes the method of PWM performs better than the ML method in the estimation of the GPD and GEV parameters (for details, see Hosking and Wallis, 1987, and Hosking et al., 1985).

Historically, before computers were so widely used, methods based on *probability paper* and a visual or least-square linear fits were used to estimate the parameters of the distributions. However, such estimating techniques have their shortcomings, are no longer needed, and are not advisable as pointed out in WMO (1998, p. 107).

When estimating it is not only important to obtain the (point) estimates, but also the uncertainty in the estimation. There are several methods for computing the confidence intervals (uncertainty) of the estimates (see e.g. Caires, 2007):

- The standard approach is based on the fact that the estimates are typically asymptotically normally distributed with the parameter values as their mean and a certain covariance matrix (the covariance matrices corresponding to the ML and PWM estimates are given by Hosking and Wallis, 1987) and on the *delta method* (Ferguson, 1996, p. 45)³. The (symmetric) confidence intervals obtained in this way will be called *asymptotic intervals*.
- The asymptotic intervals are often unsatisfactory because the actual distribution of the ML estimates, even for moderate sample sizes, can be quite skewed. Thus, for the computation of confidence intervals based on maximum likelihood estimates it has been found (e.g. Coles, 2001) that the *profile likelihood method* is usually preferable, as it accounts better for the skewness of the distribution of the ML estimates. This method is based on the likelihood ratio and is valid under certain regularity conditions (see Coles, 2001, and references therein for more details). The generally asymmetric confidence intervals obtained in such way are known as *profile likelihood intervals*.

3. *The delta method allows the determination of the asymptotic distribution of estimators that are functions of other estimators whose distribution is known by means of a Taylor expansion.*

- In some cases the delta method cannot be used to find explicit expressions for the variances of the estimators. In such cases, resampling methods like the bootstrap offer a simple and reliable alternative for estimating standard errors of estimators. Furthermore, the bootstrap method also allows one to compute *percentile confidence intervals* (Efron and Tibshirani, 1993) which also work asymptotically and can be asymmetric. However, Tajvidi (2003) investigated the performance of several bootstrap methods for constructing confidence intervals for the parameters and quantiles of the GPD and concluded that none of the bootstrap methods gives satisfactory intervals for small sample sizes. In addition, Coles and Simiu (2003) state that “it is well known that bootstrap procedures are not consistent for extreme value problems—there is a tendency for the bootstrap sample to generate shorter tails than the true sample distribution”. Coles and Simiu (2003) propose an ad-hoc method to correct/adjust the bootstrap estimates which consists of applying a bias correction to the bootstrap parameter estimates assuring that the bootstrap sample mean coincide with the parameter estimates. We shall refer to such confidence intervals as *adjusted bootstrap*.

Caires (2007) studied the coverage rate of confidence intervals of extreme value estimates based on the methods above and concluded the adjusted percentile bootstrap method generally produce the best confidence intervals from the point of view of coverage rates. Furthermore, that the quality of the coverage rate does not depend much on the bootstrap sample size; a bootstrap sample size of 1000 seems to be quite adequate for most practical purposes. Adjusted percentile bootstrap confidence intervals are to be given with the estimates presented in this report.

The parameter, λ_u , the yearly cluster rate, needed for the estimation of the return values, can be estimated by the average number of clusters/peak excesses per year. However, for yearly series with different numbers of observations (gaps) the estimation of λ_u should account for the gaps in the data (see Ferreira and Guedes Soares, 1998). λ_u should then be estimated by

$$\lambda_u = k^{-1} \sum_{i=1}^k N_i / p_i, \quad (2.9)$$

where k is the number of years considered, $p_i = n_i/n$, n_i is the number of observations available in the i th year, N_i is the corresponding number of peak excesses, and n is the maximum number of observations in a yearly series.

Model checking in extreme value analyses is usually done by means of probability plots, quantile plots and return level plots (see Coles, 2001). In this report we have chosen to illustrate the fits with return level plots.

Climate change and variability

In the methods described so far the extreme wave climate is assumed to be stationary. However, it is believed today that climate is not stationary, as the detection of both decadal variability and long term time trends in different climate variables, reported by several authors, indicates. Both the AM/GEV and POT/GPD approach can be extended to the non-stationary situation by making the parameters of the distributions functions of time (see Coles, 2001). An example of application to wave data of the non-stationary analogue of the AM/GEV approach is given in Wang and Swail (2006) and of the non-stationary analogue of the POT/GPD approach in Caires et al. (2006).

Worked examples

Introduction

To illustrate the EVA methodologies described can be applied, this chapter provides two examples of EVA of significant wave height observations. The first example deals with deep water measurements and the second with shallow water data.

NDBC deep water data

Measurements by buoy 46005 of the NOAA database (National Data Buoy Center, <http://seaboard.ndbc.noaa.gov/>) have been considered, see Figure 3.1. The buoy is located at 46.050° N 131.020° W at the Pacific Ocean offshore the US Washington state, at a water depth of 2780 m.




Figure 3.1 Left: Google Earth aerial view of the location of NDBC buoy 46005. Right: NDBC buoy 46005.

This dataset was selected because its quality is quite high, it is freely available on the internet and the measurements of this particular buoy have been used in other studies (e.g. Anderson et al., 2001 and Holthuisen, 2007). The buoy significant wave height measurements are available hourly from 20-minutes long records. These are assumed to describe sea states with durations of about three hours. The measurements have gone through some quality control by NOAA. It is however always recommended to perform some basic quality checks.

Data availability

All available buoy 46005 measurements were downloaded from the NDBC site⁴ and analysed. Data is available from 1976 until 2008. However, there are several gaps in the data. Gaps in buoy measurements are common, and they can be due to the buoy not being operational for a long period, to data processing problems, to the loss of a single observation because of communication problems, etc. For instance, this particular buoy went adrift on 18/12/2008 and recovered as of 13/3/2009. Table 3.1 presents the percentage of available hourly H_s measurements per month until 2008. Note that not only there are missing measurements as there are also repeated measurements (e.g. in July 1996 the percentage of available hourly data is 112%).

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
1976	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	20.56	32.93	33.19	33.20
1977	0.13	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	21.24	31.94	29.44
1978	30.78	31.10	32.53	32.50	32.26	29.86	32.39	32.26	32.64	33.06	33.33	33.20
1979	32.93	32.74	33.20	33.33	33.20	33.19	33.06	34.41	33.06	33.06	33.33	33.33
1980	32.66	33.05	33.20	33.06	32.26	29.31	82.80	99.73	98.89	100.00	100.00	99.60
1981	97.58	98.21	98.79	99.72	99.46	99.86	99.73	100.00	98.19	96.77	99.86	98.52
1982	97.98	99.85	98.12	99.44	99.46	100.00	99.73	99.06	99.31	96.24	86.25	48.52
1983	96.91	99.40	99.87	98.19	99.73	98.75	100.00	99.87	99.72	99.87	99.72	84.95
1984	97.85	98.99	98.79	100.00	99.33	99.44	99.19	99.06	99.44	99.33	98.89	98.92
1985	57.39	0.00	0.00	0.00	26.34	33.19	95.43	95.03	40.56	46.10	0.00	0.00
1986	0.00	70.98	99.06	99.72	99.06	99.58	99.73	99.19	90.83	0.00	0.00	69.09
1987	97.58	99.85	99.60	99.17	82.93	0.00	34.68	99.19	94.58	99.73	99.72	99.73
1988	98.66	99.43	99.73	99.58	99.73	99.86	99.87	99.06	35.00	33.06	5.42	52.69
1989	98.52	99.85	99.33	99.31	98.52	97.78	99.19	98.25	99.31	99.60	99.44	98.52
1990	98.66	99.85	99.46	99.17	56.18	33.06	33.06	16.53	0.00	0.00	0.00	0.00
1991	75.54	99.40	99.46	99.72	99.19	99.58	99.60	98.39	99.17	99.06	99.31	99.19
1992	98.12	98.56	97.85	99.58	99.19	99.58	98.92	98.39	98.06	97.85	96.94	98.25
1993	96.64	89.58	45.03	0.00	0.00	0.00	9.14	70.56	63.61	99.60	98.75	99.73
1994	98.25	95.39	98.39	99.72	99.06	99.44	93.95	98.12	97.78	98.66	99.17	99.60
1995	98.12	99.11	99.73	98.89	99.73	99.31	99.06	99.60	96.25	98.92	98.19	98.52
1996	86.56	93.82	83.47	101.67	97.72	98.33	111.96	6.45	0.00	0.00	0.00	0.00
1997	59.68	40.03	0.13	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
1998	95.16	98.66	96.24	95.97	94.76	95.28	97.45	104.84	96.67	99.46	98.89	98.92
1999	98.39	98.21	96.37	99.86	99.60	100.00	98.79	97.18	42.50	57.12	89.17	42.88
2000	51.21	63.79	95.70	96.11	97.85	98.75	97.98	99.60	98.89	98.52	98.61	99.46
2001	97.98	98.66	99.33	99.17	99.19	99.17	97.04	98.92	99.72	97.72	99.72	99.73
2002	97.45	97.62	97.58	97.78	98.25	98.89	98.92	94.76	96.53	99.06	99.86	99.19
2003	98.39	100.00	99.33	99.44	99.46	100.00	99.73	99.46	98.75	99.60	99.31	99.87
2004	98.79	100.00	99.60	99.44	99.73	99.58	99.73	99.19	97.92	99.87	99.86	81.32
2006	0.00	0.00	0.00	81.25	99.87	99.86	99.87	99.46	96.67	99.19	99.72	99.60
2007	98.66	98.96	99.46	98.61	99.19	99.44	99.06	99.73	99.58	99.46	96.25	100.00
2008	98.79	99.86	94.89	14.58	3.36	16.81	99.87	99.87	99.44	99.19	99.03	51.61

Table 3.1 Monthly percentage of available hourly significant wave height measurements of buoy 46005 (all data).

Figure 3.2 shows the time series of the available H_s and mean wave period (T_m) measurements.

4. http://www.ndbc.noaa.gov/station_history.php?station=46005

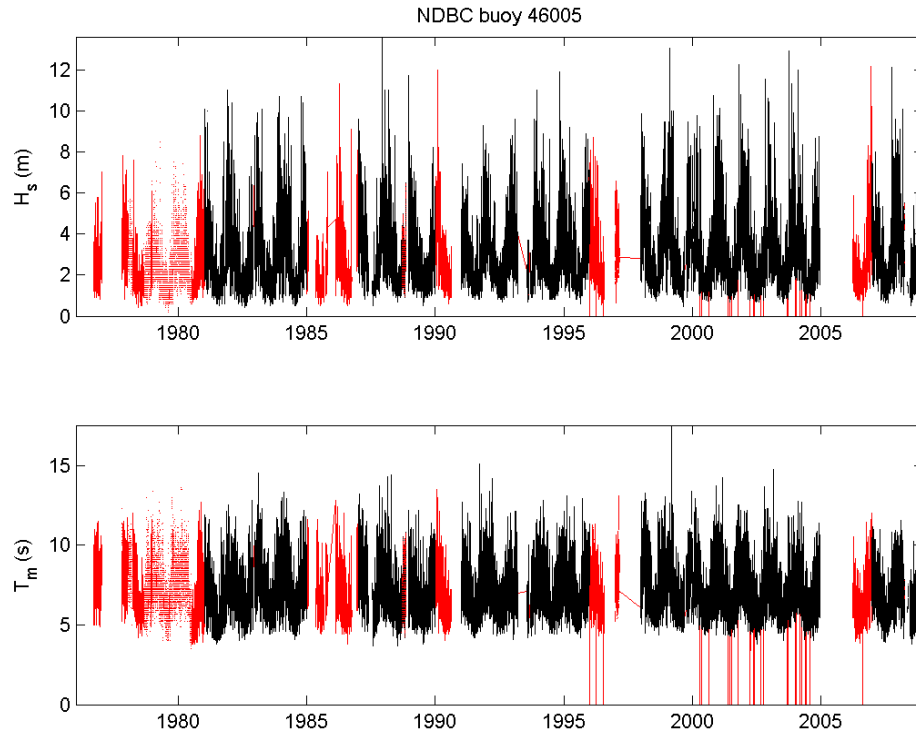


Figure 3.2 Time series of the available H_s and T_m measurements of buoy 46005.

In order to ensure that only good quality data is considered in the EVA, the following quality checks were carried out:

- Because our aim is to carry out an EVA of the data considering both POT and AM samples, we have discarded the years on which there are many missing measurements (only years for which at least half of the winter and autumn measurements were available were kept). Like this we try to avoid that the AM sample is contaminated by missing storms and the estimate of the rate of storms per year, λ_w , in the POT method is also more forthright.
- In order to have the same sampling variability in the considered data, years on which mostly only 3-hourly observations are available instead of hourly were discarded (1978-1980).
- Repeated measurements (i.e., having the same date) and measurements for which $H_s < 0.15$ were discarded.
- The time series was controlled for outliers: observations that deviate more than 7 times the standard deviation of the monthly data from its mean, or more than 3 times the standard deviation of the monthly data from the previous observation. There were no outliers identified.

Table 3.2 presents the percentage of selected hourly H_s measurements per month resulting from the above described quality checks. The dismissed measurements are plotted in red in Figure 3.2. Following the quality controls the original dataset covering 33 years has been reduced to a consistent dataset of 21 years. This dataset provides a reliable basis for the EVA. Figure 3.3 shows the density scatter of the selected H_s and T_m measurements. Note that, due to the swell, T_m extreme values do not always correspond to H_s extreme values. In

fact, on 20/3/1999 there is a T_m storm peak of 17.08 seconds while the H_s peak is 6.54 m (about half of the maximum H_s storm peak).

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
1981	97.58	98.21	98.79	99.72	99.46	99.86	99.73	100.00	98.19	96.77	99.86	98.52
1982	97.98	99.85	98.12	99.44	99.46	100.00	99.73	99.06	99.31	96.24	86.25	48.52
1983	96.91	99.40	99.87	98.19	99.73	98.75	100.00	99.87	99.72	99.87	99.72	84.95
1984	97.85	98.99	98.79	100.00	99.33	99.44	99.19	99.06	99.44	99.33	98.89	98.92
1987	97.58	99.85	99.60	99.17	82.93	0.00	34.68	99.19	94.58	99.73	99.72	99.73
1988	98.66	99.43	99.73	99.58	99.73	99.86	99.87	99.06	35.00	33.06	5.42	52.69
1989	98.52	99.85	99.33	99.31	98.52	97.78	99.19	98.25	99.31	99.60	99.44	98.52
1991	75.54	99.40	99.46	99.72	99.19	99.58	99.60	98.39	99.17	99.06	99.31	99.19
1992	98.12	98.56	97.85	99.58	99.19	99.58	98.92	98.39	98.06	97.85	96.94	98.25
1993	96.64	89.58	45.03	0.00	0.00	0.00	9.14	70.56	63.61	99.60	98.75	99.73
1994	98.25	95.39	98.39	99.72	99.06	99.44	93.95	98.12	97.78	98.66	99.17	99.60
1995	98.12	99.11	99.73	98.89	99.73	99.31	99.06	99.60	96.25	98.92	98.19	98.52
1998	95.16	98.66	96.24	95.97	94.76	95.28	97.45	95.16	96.67	99.46	98.89	98.92
1999	98.39	98.21	96.37	99.86	99.60	100.00	98.79	97.18	42.50	57.12	89.17	42.88
2000	51.21	63.79	95.70	96.11	97.85	98.75	97.98	99.60	98.89	98.52	98.61	99.46
2001	97.98	98.66	99.33	99.17	99.19	99.17	97.04	98.92	99.72	97.72	99.72	99.73
2002	97.45	97.62	97.58	97.78	98.25	98.89	98.92	94.76	96.53	99.06	99.86	99.19
2003	98.39	100.00	99.33	99.44	99.46	100.00	99.73	99.46	98.75	99.60	99.31	99.87
2004	98.79	100.00	99.60	99.44	99.73	99.58	99.73	99.19	97.92	99.87	99.86	81.32
2007	98.66	98.96	99.46	98.61	99.19	99.44	99.06	99.73	99.58	99.46	96.25	100.00
2008	98.79	99.86	94.89	14.58	3.36	16.81	99.87	99.87	99.44	99.19	99.03	51.61

Table 3.2 Monthly percentage of selected hourly significant wave height measurements of buoy 46005 (after quality checks).




Figure 3.3 Density scatter of the selected H_s and T_m measurements of buoy 46005. Note that both wind-sea and swell waves can be identified.

POT/GPD analysis

We start by analysing the data using the POT/GPD approach. We have used the threshold stability property mentioned in Section 2.2.2 to choose the most appropriate threshold for selecting a sample of peak excesses and fitting the GPD to it. More precisely, we have looked for threshold values around which the estimate of the shape parameter and σ^* seem to be stable before becoming rather variable due to reduction of the sample size. Figure 3.4 shows the threshold plot, i.e. the estimates of the shape parameter, of σ^* and of the 1/100-yr return value as functions of the threshold, obtained with the H_s data. The threshold that we have chosen is marked by a vertical line. The return value plot of the corresponding GPD fit is shown in Figure 3.5 and the model parameter estimates are presented in Table 3.3 (2nd column). The estimation was done using the PWM method.

The return value plot suggests that the GPD model is appropriate for the data. As can be seen in Table 3.3 (2nd column), the estimate of the shape parameter is negative but close to zero, which suggests that the data have a type I tail. In fact the confidence interval of the shape parameter estimates does not exclude any type of tail.




Figure 3.4 Variation of the estimates of ξ , σ^* and 1/100-yr H_s return values of the GPD model using the PWM method with the threshold used to collect the H_s POT sample.




Figure 3.5 Return value plot of the GPD model fitted to the H_s data obtained with the PWM method (solid black line) and associated adjusted bootstrap 95% confidence intervals (dashed black lines). The POT data are represented by the asterisks.

	POT/GPD	POT/Exponential	POT/Weibull	AM/GEV	AM/Gumbel
Sample size	119	119	119	21	21
u or $\hat{\mu}$ (m)	8.27	8.27	8.27	10.52 (9.85, 11.24)	10.29 (9.72, 10.93)
$\hat{\xi}$	-0.08 (-0.30, 0.16)	—	—	-0.35 (-0.71, -0.03)	—
$\hat{\sigma}$ or \hat{a} (m)	1.40 (1.04, 1.82)	1.29 (1.09, 1.52)	1.32 (1.11, 1.55)	—	—
$\hat{\sigma}$ (m)	—	—	—	1.53 (1.06, 1.95)	1.36 (1.00, 1.68)
\hat{c}	—	—	1.05 (0.93, 1.22)	—	—
\hat{H}_{s100} (m)	15.23 (12.98, 18.66)	16.47 (15.18, 17.91)	15.90 (14.38, 17.77)	14.04 (12.78, 15.67)	16.55 (14.98, 18.19)

Table 3.3 Parameter estimates and associated 95% confidence intervals of the models fitted to the H_s data.

POT/Exponential analysis

Given that the data may have a type I tail we have also fitted an exponential distribution to the POT data. In fact, the Anderson–Darling test (see, e.g., Stephens 1974) and the Gomes and van Montfort (1998) test do not reject the exponential distribution as model for the data.

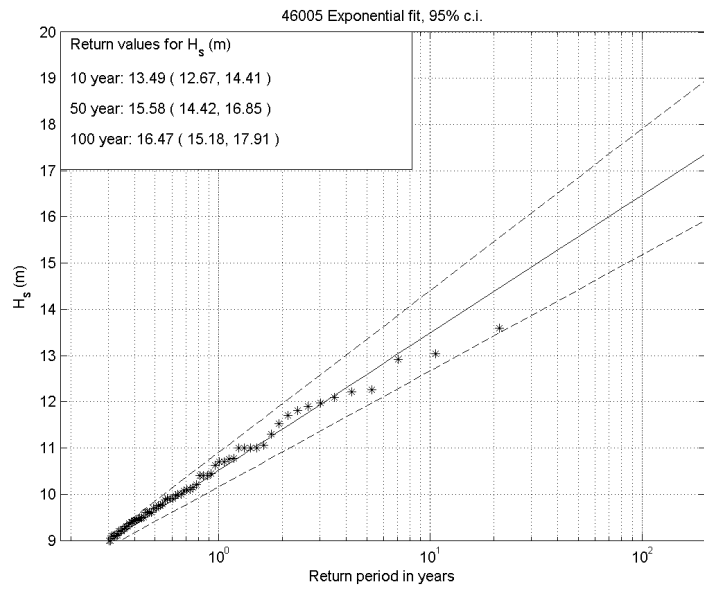


Figure 3.6 Return value plot of the Exponential fit (solid black line) and associated adjusted bootstrap 95% confidence intervals (dashed black lines). The POT data are represented by the asterisks.




Figure 3.7 Return value plot of the GPD (black) and Exponential (red) fits and associated adjusted bootstrap 95% confidence intervals (dashed lines). The POT data are represented by the asterisks.

The return value plot of the corresponding exponential fit is shown in Figure 3.6 and the model's parameter estimates are presented in Table 3.3 (3rd column). The two fits are presented in Figure 3.7. The return value estimates are higher than the estimates from the GPD, which was to be expected given that the estimated shape parameter of the GPD is less than 0, and the fit looks slightly worse (cf. Figure 3.7), but the estimates of the two models do not differ significantly (e.g. the confidence intervals of the GPD estimates

contain the point estimates of the exponential model). Note the smaller amplitude of the confidence intervals of the exponential estimates. This is not surprising, as the confidence intervals of the exponential model do not involve any uncertainty associated with the type of tail the data may have, since we have fixed the shape parameter to zero.

POT/Weibull analysis

We have also fitted the Weibull distribution of minima, Eq. (2.8), to the POT data. The return value plot of the corresponding Weibull fit is shown in Figure 3.8 and the model's parameter estimates are presented in Table 3.3 (4th column). Since the exponential distribution is a special case of the Weibull distribution of minima, and since the exponential model provides a reasonable fit to the data, it is not surprising that the parameter estimates of the two models are so similar (e.g. the shape parameter estimate of the Weibull is 1.05, while a shape parameter of 1 yields the exponential distribution). However, the return values computed from the Weibull are a bit smaller than those of the exponential, and the fit of the former distribution looks somewhat better in the return value plot, which is understandable in view of the greater flexibility (i.e. greater 'number of degrees of freedom', namely two, relative to the exponential distribution which has only one parameter) of the Weibull distribution.




Figure 3.8 Return value plot of the Weibull fit (solid black line) and associated adjusted bootstrap 95% confidence intervals (dashed black lines). The POT data are represented by the asterisks.

AM/GEV analysis

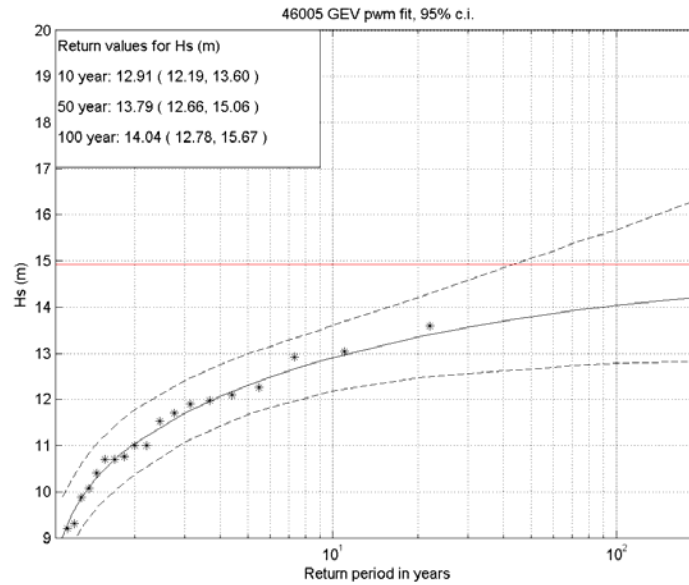


Figure 3.9 Return value plot of the GEV model fit to the AM H_s data obtained with the PWM method (solid black line) and associated adjusted bootstrap 95% confidence intervals (dashed black lines). The AM data are represented by the asterisks.

Figure 3.9 shows the return value plot of the GEV fit to the annual maxima of the H_s data. Table 3.3 (5th column) gives the corresponding parameter estimates. The estimation was done using the PWM method. Comparing these estimates with those obtained with the POT/GPD approach, one can conclude that the estimate of the GEV shape parameter is much smaller than that of the GPD, which may indicate that the GEV sample is not large enough to provide reliable estimates. In fact, looking at Figure 3.9 one can see that although the GEV fit to the AM data is rather good, the upper bound estimate of the GEV (red line in Figure 3.9) is even lower than the 1/100-yr return value estimates of the GPD. In view of this observation, it would be advisable in this case to use only the results of the GPD model.

AM/Gumbel analysis

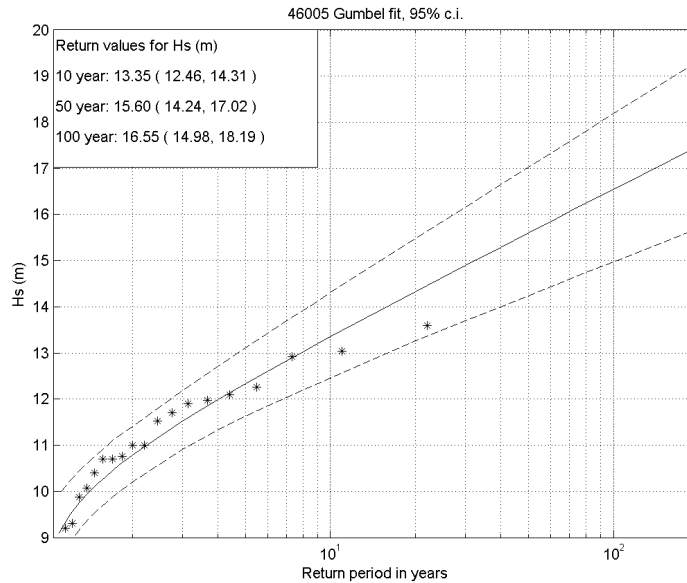


Figure 3.10 Return value plot of the Gumbel fit (solid black line) and associated adjusted bootstrap 95% confidence intervals (dashed black lines). The AM data are represented by the asterisks.

Given that the data may have a type I tail, even though this is not supported by the confidence intervals of GEV shape parameter estimates, we have also fitted a Gumbel distribution to the AM data. The return value plot of the corresponding Gumbel fit is shown in Figure 3.10 and the model's parameter estimates are presented in Table 3.3 (last column). The Gumbel return value estimates are rather close to those of the exponential (cf. Figure 3.6 and Figure 3.10), which illustrates the compatibility of the two approaches provided the estimation of the tail is correct.

Estimation method

The POT and GPD estimates provided above were obtained using the PWM method as advised by Hosking and Wallis (1987) and Hosking et al. (1985). In order to compare these estimates with those of the ML method, we have also computed the ML estimates. Figure 3.11 shows the return value plots of the GPD ML fit to the POT data and the GEV ML fit to the AM data. Table 3.4 compares the estimates of both methods. Comparing the estimates and the fits (cf. Figure 3.5 vs. Figure 3.11A and Figure 3.10 vs. Figure 3.11B), one can conclude that the ML fits seem less adequate and that the shape parameter (and consequently return value) estimates are lower than those of the PWM fits. These results support the recommendations of Hosking et al. to always use the PWM method for GPD or GEV estimation from relative short sets of data with not too heavy-tailed distributions.




Figure 3.11 Return value plots obtained using the ML method (solid black line) and associated adjusted bootstrap 95% confidence intervals (dashed black lines). A: GPD fit to the POT data. B: GEV model fit to the AM data.

	POT/GPD PWM	POT/GPD ML	AM/GEV PWM	AM/GEV ML
Sample size	119	119	21	21
u or $\hat{\mu}$ (m)	8.27	8.27	10.52 (9.85, 11.24)	10.57 (8.52, 11.50)
$\hat{\xi}$	-0.08 (-0.30, 0.16)	-0.15 (-0.33, 0.03)	-0.35 (-0.71, -0.03)	-0.38 (-1.15, -0.09)
$\hat{\sigma}$ or $\hat{\sigma}$ (m)	1.40 (1.04, 1.82)	1.49 (1.11, 1.93)	1.53 (1.06, 1.95)	1.46 (0.59, 2.46)
\hat{H}_{s100} (m)	15.23 (12.98, 18.66)	14.36 (12.81, 16.57)	14.04 (12.78, 15.67)	13.72 (10.57, 16.76)

Table 3.4 Parameter estimates using the PWM and ML methods and associated adjusted bootstrap 95% confidence intervals.

Stationarity

Given that the analyses reported above were carried out assuming stationarity of the H_s extremes, we have tried to test the validity of that assumption. Figure 3.12 shows linear fits to the H_s annual means, POT and AM data. Although with some scatter, in all cases the data shows a positive linear trend of 0.72, 0.57 and 4.2 cm per year, respectively.

We have therefore checked whether it would be more appropriate to use the non-stationary analogues of the POT/GPD and AM/GEV to analyse the data (cf. Caires et al., 2006). Models with linear trends in the location parameters were considered, and in both cases the results of the likelihood ratio test was that the trends were not significant and therefore that stationary models are good models for the data.




Figure 3.12 Annual trends in the H_s measurements of buoy 46005

North Sea shallow water data

The Dutch Ministry of Transport, Public Works and Water Management (VenW) maintains a network of nine wave monitoring stations offshore the Netherlands. The buoy locations are given in Figure 3.13. Time series of hourly significant wave height data are available from 1979 to 2002 from <http://www.golfklimaat.nl/>. Each hourly measurement is computed from a 20-minutes long record of the sea surface elevation, and is assumed to be representative of a 3-hour long sea state. When compiling the measured sequences missing information is filled-in using estimates made using a control and validation programme. We performed further quality checks on the data and the conclusion was that the quality of the data is good. There were no outliers identified and the data coverage is close to 100% for the full period 1970-2002.

The measurements by the Schiermonnikoog noord (SON) buoy are analysed here. The buoy is located at 53°35'44" N and 06°10'00" E in the North Sea, at a water depth of 19 m. Therefore, contrary to the 45006 buoy measurements, these are in shallow waters.



Figure 3.13 Measuring stations along the Dutch coast.

POT/GPD analysis

Again, we start by analysing the data using the POT/GPD approach. Figure 3.14 shows the threshold plot obtained with the SON H_s data. The threshold that we have chosen is marked by a vertical line. The return value plot of the corresponding GPD fit is shown in Figure 3.15 and the model parameter estimates are presented in Table 3.5 (2nd column). The estimation was done using the PWM method.

The return value plot suggests that the GPD model is appropriate for the data. The estimate of the shape parameter is negative, which suggests that the data have a type III tail. However, the confidence interval of the shape parameter estimates does not exclude a type I tail. The shape parameter estimate for this data (-0.13) is much smaller than that for the deep water data (-0.07), which is not surprising since waves in shallow waters are depth limited. The upper-limit estimate (c.f. Eq. (2.4)) for this data is 11.4 m, whereas for the deep water data was 22 m.




Figure 3.14 Variation of the estimates of ξ , σ^* and 1/100-yr H_s return values of the GPD model using the PWM method with the threshold used to collect the H_s POT sample.




Figure 3.15 Return value plot of the GPD model fitted to the H_s data obtained with the PWM method (solid black line) and associated adjusted bootstrap 95% confidence intervals (dashed black lines). The POT data are represented by the asterisks.

	POT/GPD	POT/Exponential	POT/Weibull	AM/GEV	AM/Gumbel
Sample size	108	108	108	24	24
u or $\hat{\mu}$ (m)	4.50	4.50	4.50	5.82 (5.42, 6.27)	5.75 (5.40, 6.13)
$\hat{\xi}$	-0.13 (-0.37, 0.09)	—	—	-0.17 (-0.50, 0.13)	—
$\hat{\sigma}$ or \hat{a} (m)	0.96 (0.70, 1.27)	0.86 (0.71, 1.01)	0.88 (0.73, 1.05)	—	—
$\hat{\sigma}$ (m)	—	—	—	0.92 (0.61, 1.21)	0.84 (0.62, 1.04)
\hat{c}	—	—	1.08 (0.93, 1.27)	—	—
\hat{H}_{s100} (m)	8.63 (7.24, 10.79)	9.74 (8.87, 10.71)	9.21 (8.23, 10.46)	8.74 (7.83, 10.02)	9.61 (8.52, 10.65)

Table 3.5 Parameter estimates and associated 95% confidence intervals.

POT/Exponential analysis

The Anderson–Darling test (see, e.g., Stephens 1974) and the Gomes and van Montfort (1998) test do not reject the exponential distribution as model for the data. We have therefore also fitted an exponential distribution to the POT data. The return value plot of the corresponding exponential fit is shown in Figure 3.16 and the model's parameter estimates are presented in Table 3.5 (3rd column). Again, the return value estimates are higher than the estimates from the GPD and the fit looks reasonable.

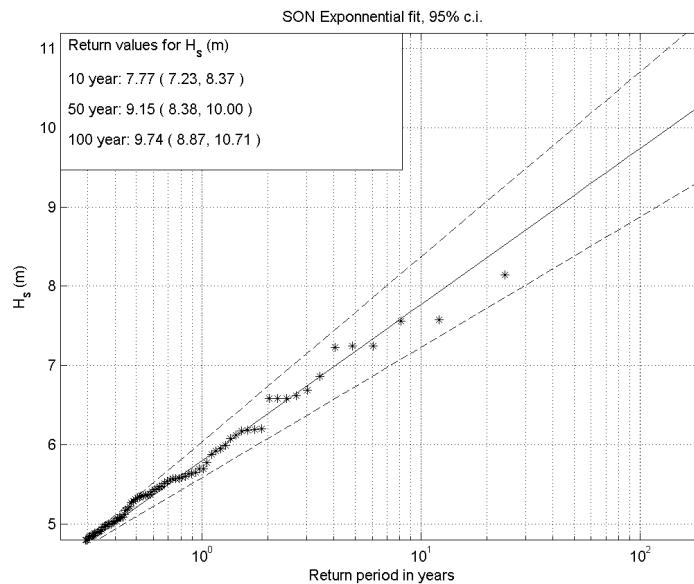


Figure 3.16 Return value plot of the Exponential fit (solid black line) and associated adjusted bootstrap 95% confidence intervals (dashed black lines). The POT data are represented by the asterisks.

POT/Weibull analysis

We have also fitted the Weibull distribution of minima, Eq. (2.8), to the POT data. The return value plot of the corresponding Weibull fit is shown in Figure 3.17 and the model's parameter estimates

are presented in Table 3.5 (4th column). Again the Weibull fit looks better than the exponential fit. The shape parameter estimate of the Weibull is 1.08, which is again very close to the exponential shape parameter yielding the exponential distribution.




Figure 3.17 Return value plot of the Weibull fit (solid black line) and associated adjusted bootstrap 95% confidence intervals (dashed black lines). The POT data are represented by the asterisks.

AM/GEV analysis

Figure 3.18 shows the return value plot of the GEV fit to the annual maxima of the H_s data. Table 3.5 (5th column) gives the corresponding parameter estimates. The estimation was done using the PWM method. Comparing these estimates with those obtained with the POT/GPD approach, one can conclude that the estimate of the GEV model are close to those of the GPD, which gives confidence in the estimates of both models.




Figure 3.18 Return value plot of the GEV model fit to the AM H_s data obtained with the PWM method (solid black line) and associated adjusted bootstrap 95% confidence intervals (dashed black lines). The AM data are represented by the asterisks.

AM/Gumbel analysis

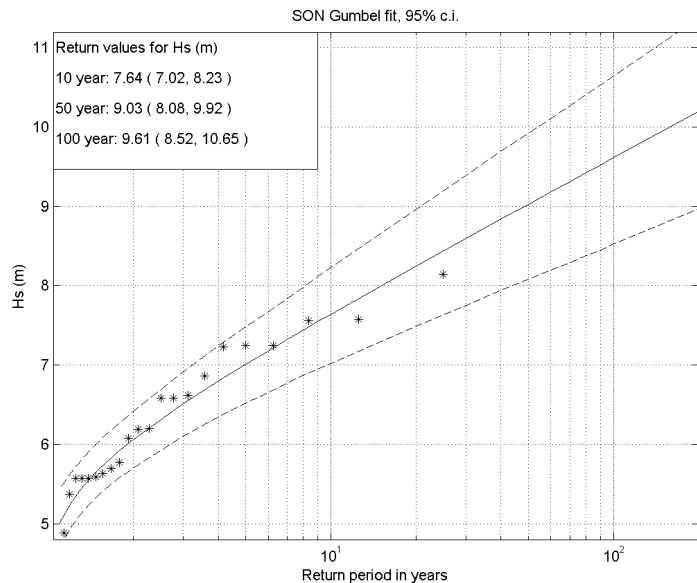


Figure 3.19 Return value plot of the Gumbel fit (solid black line) and associated adjusted bootstrap 95% confidence intervals (dashed black lines). The AM data are represented by the asterisks.

Figure 3.19 shows the return value plot of the Gumbel fit to the annual maxima of the H_s data. Table 3.5 (last column) gives the corresponding parameter estimates. Again, the Gumbel return value estimates are rather close to those of the exponential.

Estimation method

Figure 3.20 shows the return value plots of the GPD ML fit to the POT data and the GEV ML fit to the AM data. Table 3.4 compares the estimates of both methods. Comparing the estimates and the fits, one can conclude that the ML fits are worst and the shape parameter (and consequently the return value) estimates are lower than those of the PWM fits. The differences in the results of the two methods are, however, smaller than those computed for the deep water data, that is probably because the sample sizes of the shallow water data are larger.




Figure 3.20 Return value plots obtained using the ML method (solid black line) and associated adjusted bootstrap 95% confidence intervals (dashed black lines). Left: GPD fit to the POT data. Right: GEV model fit to the AM data.

	POT/GPD PWM	POT/GPD ML	AM/GEV PWM	AM/GEV ML
Sample size	108	108	24	24
u or $\hat{\mu}$ (m)	4.50	4.50	5.82 (5.42, 6.27)	5.86 (5.30, 6.42)
$\hat{\xi}$	-0.12 (-0.37, 0.09)	-0.14 (-0.30, 0.03)	-0.17 (-0.50, 0.13)	-0.24 (-1.20, 0.20)
$\hat{\sigma}$ or $\hat{\sigma}$ (m)	0.96 (0.70, 1.27)	0.98 (0.74, 1.25)	0.92 (0.61, 1.21)	0.90 (0.57, 1.44)
\hat{H}_{s100} (m)	8.63 (7.24, 10.79)	8.51 (7.59, 9.74)	8.74 (7.83, 10.02)	8.35 (6.62, 10.35)

Table 3.6 Parameter estimates using the PWM and ML methods and associated adjusted bootstrap 95% confidence intervals.

Stationarity

Figure 3.21 shows linear fits to the H_s annual means, POT and AM data. The estimated linear trends are of -0.09, 0.81 and 3.2 cm per year, respectively.

As done for the deep water data, we have checked whether it would be more appropriate to use the non-stationary analogues of the POT/GPD and AM/GEV to analyse the SON data. Models with linear trends in the location parameters were considered, and again in both cases the results of the likelihood ratio test was that the trends were not significant. The stationary model is therefore a good model for the data.

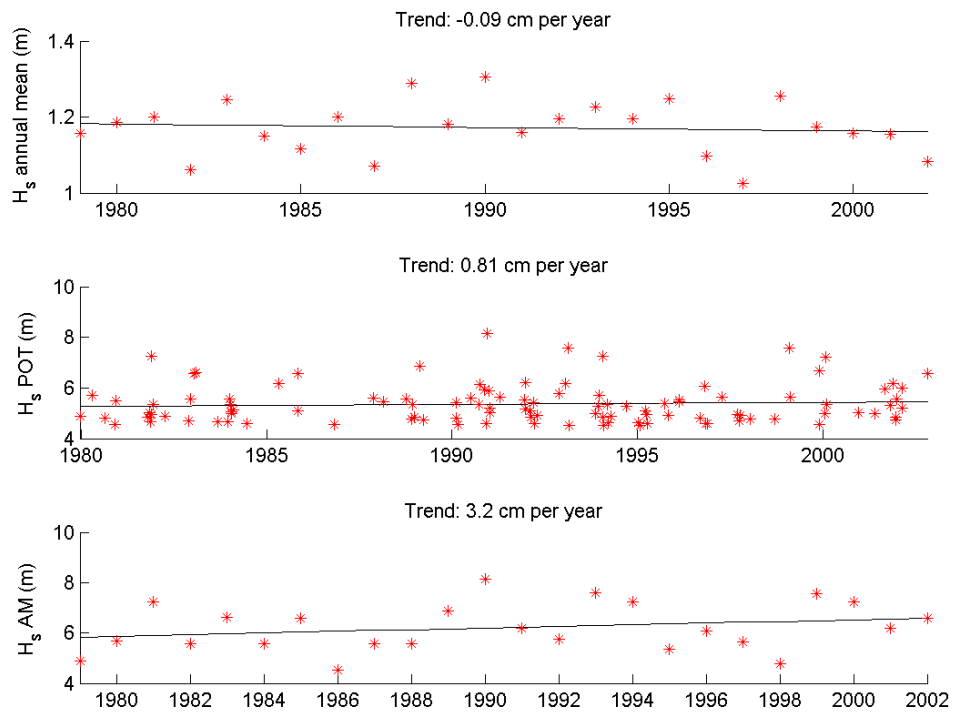


Figure 3.21 Annual trends in the H_s measurements of the SON buoy.

Software packages

In the last years, mostly the last decade, there has been an increased use of the extreme value methodology and an associated increased offer of software packages to perform extreme value analysis (see e.g. Stephenson and Gilleland, 2006). A lot of the available packages target financial instead of environmental data. Without the pretence of covering all that is available, Table 4.1 gives an overview of available extreme value analysis software packages directly suitable for the analysis of wave data. Some have command line (routine) and/or Graphical User Interfaces (GUI). GUIs make packages very easy to use, but when not accompanied by command line / routine options provide little flexibility for more experienced users.

Package name and reference	Programming language	Avail.	Interface
EVIM http://www.bilkent.edu.tr/%7Efaruk/evim.htm	MATLAB	free software	command line
EXTREMES http://mistis.inrialpes.fr/software/EXTREMES/accueil.html	C++ with a MATLAB GUI	free software	GUI
extRemes http://www.isse.ucar.edu/extremevalues/evtk.html Uses ismev	R	free software	command line and GUI
ismev http://cran.r-project.org/web/packages/ismev/index.html (The R routines are based on the S-Plus (http://www.insightful.com/products/splus/) routines of Coles, 2001)	S-Plus and R	free software	command line
ORCA http://www.widelft.nl/soft/chess/orca/index.html	MATLAB	commercial	command line and GUI
Statistics of Extremes http://lstat.kuleuven.be/Wiley/index.html	S-Plus and FORTRAN	free software	command line
WAFO http://www.maths.lth.se/matstat/wafo/about.html	MATLAB	free software	command line
Xtremes http://www.xtremes.de/xtremes/	Pascal	commercial	GUI

Table 4.1 Software packages available for EVA. Please note that web addresses may be liable to change.

Although a lot of researchers, consultants and engineers mostly do their analysis in FORTRAN, most available packages of EVA routines are in MATLAB (<http://www.mathworks.com/products/matlab/>) or R (<http://www.r-project.org/>). FORTRAN seem to prefer to develop their own routines instead of using packages. For instance, the FORTRAN routines used for GPD and GEV fits in the popular climate explorer site (<http://climexp.knmi.nl>) are freely available (<http://climexp.knmi.nl/source/fortran.zip>), but the available routine were develop solely with the goal of own use and not for comprehensive usage by others.

The analysis presented here were carried out using the MATLAB tool **ORCA** ("metOcean data tRansformation, Classification and Analysis tool").

Recommendations

Extreme value analysis

When carrying out an extreme value analysis of significant wave height data we suggest that the following steps are followed:

1. The selection of the samples of significant wave height POT and/or AM, paying attention to declustering, so that the data can be considered independent.
2. The selection of an extreme value distribution as a statistical model for the data produced in the first step. According to theory the generalized Pareto distribution (GPD) should be selected when the sample is collected by means of the POT method, while in the case of the AM method a generalized extreme value (GEV) distribution should be chosen.
3. The selection of a method to estimate the unknown parameters of the distribution chosen in Step 2. When possible, the PWM method is preferable.
4. When doing a POT/GPD analysis always produce threshold plots to guide the threshold choice.
5. On the basis of the estimated parameters (and thus the estimated distribution) the extreme values corresponding to one or more prescribed return period(s) can be estimated.
6. The selection of a method to quantify the uncertainty in the estimates of the distribution's parameters and in the associated return values. The uncertainty in the parameters and return values is usually quantified by a confidence interval of some significance level (here 95%). Adjusted bootstrap confidence intervals are considered optimal.

Furthermore, we have the following recommendations:

- Before carrying out an extreme value analysis, a rigorous data quality analysis should always be carried out.
- When the availability of data is limited give preference to a POT/GPD analysis instead of AM/GEV analysis. Doing both would help in diagnosis.
- If one is to choose a distribution other than GPD or GEV one should do so on the basis of some justification.
- The initial-distribution approach is not recommended.
- Always check that whether the climate can be assumed stationary.

Further aspects

In these analyses only omni-directional significant wave height data were considered. However, the wave period and direction also plays an important role in wave loads. Caires and van Gent (2008) have studied the incorporation of wave period data in the extreme value analysis of wave data. A important aspect when analysing the mean wave period data is to establish whether the overall (wind sea and swell) are of interest or only those associated with significant wave height extremes (wind sea). Figure 3.3 shows the type of relations between the significant wave height and the mean wave period in a given region. When doing EVA per directional sector it is important to properly divide the data into sectors. A sector should not consider only a proportion of a given population, since that will only increase the uncertainty. Furthermore, two apart and equally important, in terms of extremes, populations should not be mixed within a sector.

Furthermore, often individual parameters within a sea state are also of interest, such as the maximal individual wave height. In deep waters, these can be computed using the Rayleigh distribution and in shallow water using the Battjes and Groenendijk distribution (see Battjes and Groenendijk, 2000).

The data considered here were regularly sampled. That is not the case when considering satellite data, which makes the extreme value analyses of the data not straightforward. Anderson et al. (2001) provide suggestions for how to approach the problem of irregularly sampled data.

References

- Anderson, C. W., D. J. T. Carter, and P. D. Cotton, 2001: *Wave climate variability and impact on offshore design extremes*. Shell International and the Organization of Oil and Gas Producers Rep., 99 p.
- Battjes, J.A. and H.W. Groenendijk, 2000: Wave height distributions on shallow foreshores, *Coastal Engineering*, **40**, 161-182.
- Caires, S., 2007: *Extreme wave statistics. Confidence Intervals*. WL | Delft Hydraulics Report H4803.30.
- Caires, S. and A. Sterl, 2005: 100-year return value estimates for ocean wind speed and significant wave height from the ERA-40 data. *J. Climate*, 18(7), 1032-1048.
- Caires, S., V. Swail and X. L. Wang, 2006: Projection and Analysis of Extreme Wave Climate. *J. Climate*, 19(21), 5581-5605.
- Caires, S. and M.R.A. van Gent. Extreme wave loads. *Proc. 27th Int. Conf. on Offshore Mechanics and Arctic Eng. (OMAE2008)*, Lisbon, 15-20 June 2008.
- Coles, S., 2001: *An Introduction to the Statistical Modelling of Extreme Values*. Springer Texts in Statistics, Springer-Verlag: London.
- Coles, S., and E. Simiu, 2003: Estimating uncertainty in the extreme value analysis of data generated by a hurricane simulation model. *J. Engrg. Mech.*, 129 (11), 1288-1294.
- Davison, A. C., and R. L. Smith, 1990: Models for exceedances over high thresholds (with discussion). *J. Roy. Stat. Soc.*, 52B, 393-442.
- Efron, B. and Tibshirani, R.J., 1993: *An Introduction to the Bootstrap*. Monographs on Statistics & Applied Probability 57, Chapman and Hall/CRC, 436p.
- Ferguson, T.S., 1996: *A Course in Large Sample Theory*, Chapman and Hall, New York.
- Ferreira, J.A., and C. Guedes Soares, 1998: An application of the peaks over threshold method to predict extremes of significant wave height. *J. Offshore Mech. Arct. Eng.*, 120, 165-176.
- Ferreira, J. A., and C. Guedes Soares, 2000: Modelling distributions of significant wave height. *Coast. Eng.*, 40, 361-374.
- Gomes, M. Y., and M. A. J. van Montfort, 1986: Exponentiality versus generalized Pareto, quick tests. *Third Int. Conf. on Statistical Climatology*, Vienna, Austria.
- Guedes Soares, C. and A.C. Henriques, 1996: Statistical uncertainty in long-term predictions of significant wave height. *J. Offshore Mech. Arct. Eng.*, 11, 284-291.
- Guedes Soares, C. and M.G. Scotto, 2004: Application of the r largest-order statistics for long-term predictions of significant wave height, *Coastal Eng.*, 21 (5-6), pp. 387-394.
- Holthuijsen, L.H., 2007: *Waves in Oceanic and Coastal Waters*. Cambridge University Press, ISBN 0521860288, 387 p.
- Hosking, J.R.M. and J.R. Wallis, 1987: Parameter and quantile estimation for the Generalized Pareto Distribution. *Technometrics*, 29, 339-349.
- Hosking, J.R.M., J.R. Wallis, and E.F. Wood, 1985: Estimation of the generalized extreme-value distribution by the method of probability-weighted moments. *Technometrics*, 27, 251-261.
- Lopatoukhin, L.J., V.A. Rozhkov, V.E. Ryabinin, V.R. Swail, A.V. Boukhanovsky and A. B. Degtyarev, 2000: *Estimation of extreme wind wave heights*. WMO/TD-No. 1041, JCOMM Technical Report No. 9.
- Pickands, J., 1971: The two-dimensional Poisson process and extremal processes. *Journal of Applied Probability*, 8, 745-756.
- Pickands, J., 1975: Statistical inference using extreme order statistics. *Annals of Statistics*, 3, 119-131.
- Stephens, M. A., 1974: EDF statistics for goodness of fit and some comparisons. *J. Amer. Stat. Assoc.*, **69**, 730-737.
- Stephenson, A. and E. Gilleland, 2006: Software for the analysis of extreme events: The current state and future directions. *Extremes*, **8**, 87-109.
- Tajvidi, N., 2003: Confidence Intervals and Accuracy Estimation for heavy-tailed Generalized Pareto Distributions. *Extremes*, **6**, 111-123.
- Van Vledder, G., Y. Goda, P. Hawkes, E. Mansard, M.J. Martin, M. Mathiesen, E. Peltier, and E. Thompson, 1993: Case studies of extreme wave analysis: a comparative analysis. *Proc. WAVES'93 Conf.*, New Orleans, USA, 978-992.

- Wang, X.L. and V. Swail, 2006: Climate change signal and uncertainty in projections of ocean wave height, *Climate Dyn.*, 26, 109-126.
- WMO, 1998: *Guide to Wave Analysis and Forecasting*. WMO-No. 702, 2nd edition, Geneva, Switzerland, 159 p., <http://www.wmo.int/pages/prog/amp/mmop/documents/WMO%20No%20702/WMO702.pdf>.