NORGES TEKNISK-NATURVITENSKAPELIGE
UNIVERSITET

Combination of The Peaks-Over-Threshold and Bootstrapping Methods for Extreme Value Prediction

by

A. Naess and P.H. Clausen

PREPRINT
STATISTICS NO. 14/2001
ISSN 0804-9173

NORWEGIAN UNIVERSITY OF SCIENCE AND
TECHNOLOGY
TRONDHEIM, NORWAY

This report has URL
http://www.math.ntnu.no/preprint/statistics/2001/S14-2001.ps
Arvid Naess has homepage: http://www.math.ntnu.no/∼arvidn
E-mail: arvidn@math.ntnu.no
Address: Department of Mathematical Sciences, Norwegian University of Science and Technology,
N-7491 Trondheim, Norway.

Combination of The Peaks-Over-Threshold and Bootstrapping Methods for Extreme Value Prediction

A. Naess
Department of Mathematical Sciences
P. H. Clausen
Department of Structural Engineering
Norwegian University of Science and Technology
NO-7491 Trondheim, Norway

Abstract

The paper investigates the application of the peaks-over-threshold (POT) method in combination with the bootstrapping method for estimating extreme values, in particular, long return period characteristic values of environmental loads, e.g. wind loads, on the basis of observed data. Attention is focused on the effect different statistical estimators of the parameters inherent in the POT method have on the predicted characteristic values. The accuracy of the predicted long return period characteristic values provided by the different methods will be evaluated by application to synthetic data where the relevant characteristic values can be calculated. The bootstrapping method is advocated as a practical tool for estimating confidence intervals on the provided point estimates of the long return period values.

Keywords: Peaks-Over-Threshold method, Extreme value estimation, de Haan estimators, moment estimators, Bootstrap confidence intervals, 50-year wind speed

Introduction

The focus of this paper is the practical implementation of the peaks-over-threshold (POT) method (Smith 1989; Davison and Smith 1990) for the estimation of extreme values characterized by return periods that are much longer than the periods of data acquisition. Typically, the goal is to estimate extremes of an environmental process like for instance wind speed with a return period of say 50 or 100 years on the basis of 25 years of recorded data.

The traditional method adopted for estimating 50 or 100 year return period values is usually based on the assumption that the distribution of yearly extreme values can be approximated by a Gumbel or Type I extreme value distribution (Gumbel 1958; Castillo 1988). The general procedure is then to plot the data on Gumbel probability paper and carry out some form of linear regression on the data to fit a straight line. By restricting the database to comprise only yearly extreme values, a great amount of potentially useful information is discarded. This procedure can be improved as discussed by Naess (1998a). In recent years the POT method has also received considerable attention.

The POT method is based on utilizing all peak events of a given time series exceeding a specified threshold. The word 'event' is used deliberately, because in many practical applications, like wind speed time series, one would consider peak values with a time separation less than a certain value to belong to the same peak event. For example, the passage of a low pressure system with a duration of a few days is seen as an event in this context. By considering peak events instead of yearly maxima, the number of data points for statistical processing may be increased considerably. However, in this connection it is appropriate to recall the basic assumption of the POT method that the threshold is high relative to typical values. Excessive lowering of the threshold to obtain more data may lead to substantial bias. Gross et al. (1994) have performed numerical simulations that indicate that in samples taken from extreme value populations for wind speed, optimal results seem to be obtained by a choice of threshold corresponding to a number of exceedances of the order of ten per year on the average.

Naess (1998b) used a POT analysis based on estimators proposed by de Haan (1994) to analyse wind speed data recorded at 44 US weather stations. It was shown that the predicted 50 year values of wind speed obtained by the POT method were in general significantly below the corresponding predictions obtained by using the traditional Gumbel method. Naess and Clausen (1999) investigated the performance of the POT method using several different estimators. The data used for the analyses were synthetic data sampled from a stationary Gaussian process. For such processes the 50 year values can be accurately given. It was established that the POT method in combination with the chosen estimators significantly underestimated the correct values.

In the present paper two new estimators for use with the POT method will be introduced. It will be demonstrated for an extensive set of synthetic data that one of these estimators provides fairly good agreement with the target values, and has the desirable property, at least from a practical point of view, of being, on the average, conservative.

The problem of estimating confidence intervals for the obtained long return period values is also discussed, and it is shown that the bootstrapping method (Efron and Tibshirani 1993; Davison and Hinkley 1997) offers a practical solution.

The Peaks Over Threshold Method

The Generalized Pareto Distribution

The POT method is based on what is called the generalized Pareto (GP) distribution (defined below) in the following manner: It has been shown (Pickands 1975) that asymptotically, the excess values above a high level will follow a GP distribution if and only if the parent distribution belongs to the domain of attraction of one of the extreme value distributions. The assumption of a Poisson process model for the exceedance times combined with GP distributed excesses can be shown to lead to the generalized extreme value (GEV) distribution for corresponding extremes, see below. The expression for the GP distribution is

$$G(y) = G(y; a, c) = Prob[Y \leq y] = 1 - \left(1 + c\,\frac{y}{a}\right)_{+}^{-1/c} \tag{1}$$

Here $a > 0$ is a scale parameter and c $(-\infty < c < \infty)$ determines the shape of the distribution. $(z)_{+} = \max(0, z)$.

The asymptotic result referred to above implies that equation (1) can be used to represent the conditional cumulative distribution function of the excess $Y = X - u$ of the observed variate X over the threshold u, given that $X > u$ for u sufficiently large (Pickands 1975). The cases $c > 0$, $c = 0$ and $c < 0$ correspond to Fréchet (Type II), Gumbel (Type I), and reverse Weibull (Type III) domains of attraction, respectively, cf. section below. For $c = 0$, the expression between the parentheses in equation (1) is understood in a limiting sense as the exponential $\exp(-y/a)$.

The Generalized Extreme Value Distribution

When the analysis is based on the maximum observation during specified periods of time, like one year, the usual procedure is to assume that the resulting set of data are independent and identically distributed (iid) and follow a generalized extreme value (GEV) distribution for maxima with cumulative distribution function (CDF)

$$GEV(x; a, b, c) = \exp\left\{-\left(1 + c\,\frac{x-b}{a}\right)_{+}^{-1/c}\right\} \tag{2}$$

Here $a > 0$ and b are scale and location parameters, respectively. c is a parameter determining the shape of the distribution. The cases $c > 0$, $c = 0$ and $c < 0$ correspond to Fréchet (Type II), Gumbel (Type I), and reverse Weibull (Type III) extreme value distributions, respectively. The special case $c = 0$ is again interpreted as a limiting case, that is

$$GEV(x; a, b, 0) = \exp\left\{-e^{-(x-b)/a}\right\} \tag{3}$$

which is the standard expression for the Gumbel distribution.

Similarly as for the distribution of exceedances for high levels discussed in the previous section, the assumption that yearly maxima follow a GEV distribution is also based on a general asymptotic result, see e.g. Leadbetter et al. (1983). Of particular interest here is the observation that for a particular process, the c-parameter in the GP distribution representing the asymptotic distribution of exceedances, is the same as the c-parameter in the corresponding GEV distribution for the same process, explaining the use of the same symbol in the two cases. It is important to note, however, that this statement is valid only asymptotically. The practical significance of this, is that, in general, the c-parameter estimated on the basis of a finite set of data rarely equals that of the correct asymptotic extreme value distribution. Also relevant in this context is the observation made by Fisher and Tippet (1928) that the best fit to a finite set of extreme value data from a stationary Gaussian process is provided by a Type III distribution and not by the asymptotically correct Type I distribution.

Return Periods

The return period R of a given wind speed, in years, is defined as the inverse of the probability that the specified wind speed will be exceeded in any one year. If λ denotes the mean crossing rate of the threshold u per year (i.e., the average number of data points above the threshold u per year), then the return period R of the value of X corresponding to the level $x_R = u + y$ is given by the relation

$$R = \frac{1}{\lambda\, Prob[X > x_R]} = \frac{1}{\lambda\, Prob[Y > y]} \tag{4}$$

Hence, it follows that

$$Prob[Y \leq y] = 1 - 1/(\lambda R) \tag{5}$$

Invoking equation (1) for $c \neq 0$ leads to the result

$$x_R = u - a[1 - (\lambda R)^c]/c \tag{6}$$

Similarly, for $c = 0$, it is found that

$$x_R = u + a \ln(\lambda R) \tag{7}$$

where u is the threshold used in the estimation of c and a.

Statistical Estimators

The de Haan Estimators

Let n denote the total number of data points, while the number of observations above the threshold value u is denoted by k. The threshold u then represents the $(k + 1)$-th highest data point(s). An estimate for λ is $\hat{\lambda} = k/n_{yrs}$, where n_{yrs} denotes the length of the record in years. The highest, second highest, ... , k-th highest, $(k + 1)$-th highest variates are denoted by X^*_n, X^*_{n-1}, ...,$X^*_{n-(k+1)}$, $X^*_{n-k} = u$, respectively.

The parameter estimators proposed by de Haan (1994) are based on the following two quantities

$$H_{k,n} = \frac{1}{k} \sum_{i=0}^{k-1} \{\ln(X^*_{n-i}) - \ln(X^*_{n-k})\} \tag{8}$$

and

$$H^{(2)}_{k,n} = \frac{1}{k} \sum_{i=0}^{k-1} \{\ln(X^*_{n-i}) - \ln(X^*_{n-k})\}^2 \tag{9}$$

Estimators for a and c are then given by the relations

$$\hat{a} = \rho\, X^*_{n-k}\, H_{k,n} = \rho\, u\, H_{k,n} \tag{10}$$

and

$$\hat{c} = H_{k,n} + 1 - \frac{1}{2} \left\{ 1 - \frac{(H_{k,n})^2}{H^{(2)}_{k,n}} \right\}^{-1} \tag{11}$$

where $\rho = 1$ if $\hat{c} \geq 0$, while $\rho = 1 - \hat{c}$ if $\hat{c} < 0$.

Subject to general conditions on the underlying probability law, de Haan (1994) showed that $\hat{a} \to a$ and $\hat{c} \to c$ as $n \to \infty$ (in probability).

Closely related to the de Haan estimators are the Hill estimators. Their application to problems similar to the ones discussed in this paper were investigated by Naess and Clausen (1999). Their conclusion was that the Hill estimators lead to results that are quite similar to those provided by the de Haan estimators. Since the Hill estimators require considerably higher numerical efforts than the de Haan estimators, we have therefore chosen to exclude the Hill estimators from the present study.

The Moment Estimators

In terms of the the mean value $E(Y)$ and the standard deviation $s(Y)$ of the exceedance variate Y, it can be shown that (Hosking and Wallis 1987)

$$a = \frac{1}{2}E(Y)\{1 + [E(Y)/s(Y)]^2\} \tag{12}$$

and

$$c = \frac{1}{2}\{1 - [E(Y)/s(Y)]^2\} \tag{13}$$

Hence, empirical estimates of the first two moments of Y provide estimates of a and c. We shall refer to the resulting estimators as the moment estimators for a and c.

The c0 Estimators

Naess (1998a) showed that if the underlying long term probability density of the stochastic process $X(t)$ giving rise to the observations can be assumed of the form

$$f_X(x) = A\, \exp\{-\alpha(x)\} \tag{14}$$

where A is a positive constant, and $\alpha(x)$ denotes a well-behaved function that is strictly increasing for increasing large x, then, under reasonable assumptions, the data obtained by transforming the initial data using the function $\alpha(x)$ are Gumbel distributed with good approximation. That is, if $x_1, \ldots, x_n$ denotes the original set of extreme value or peak event data, then the transformed data $z_1, \ldots, z_n$, where $z_i = \alpha(x_i)$, will have $c \approx 0$. The estimators applied to the transformed data are denoted by c0 estimators to emphasize that for the transformed data only the a-parameter will be estimated. In the context of this paper, this leads to two $c0$ estimators. One is based on the de Haan estimator for a. This estimator will be referred to as the $c0$ de Haan estimator. The other is based on the moment estimator for a and is referred to as the $c0$ moment estimator.

Estimating Confidence Intervals by Bootstrapping

The bootstrapping method is a statistical technique of fairly recent origin that can be used for estimating e.g. confidence intervals on statistical estimates on the basis of a limited sample (Efron and Tibshirani 1993; Davison and Hinkley 1997). It is based on resampling from a distribution determined by the available sample of data. Despite the fact that the name of the method alludes to lifting oneself up by the bootstraps (Baron von Munchausen), the method appears to be quite effective for the specific purpose of e.g. estimating confidence bands. For convenience, a brief discussion of some basic features of the bootstrapping method will be provided here.

Assume that $\boldsymbol{u} = (u_1, u_2, \ldots, u_n)$ is a sample or vector consisting of n independent observations of a random variable U. There are in principle two available options for obtaining bootstrap estimates. One is the nonparametric approach, where a purely empirical distribution function (EDF) is established for U on the basis of the observed data by allocating a probability of $1/n$ to each of the observed data points. The other is the parametric bootstrap, which is obtained by assuming that U has a specified cumulative distribution function (CDF) $F_U(u;\theta) = Prob(U \leq u)$, where θ denotes a vector of unknown parameters, which determine the CDF. These parameters are then estimated from the observed data $\boldsymbol{u}$, giving $\hat{\theta}$, and $F_U(u;\hat{\theta})$ is adopted as the CDF of U.

In this paper only the nonparametric bootstrap will be used. The goal is to estimate some statistical quantity R given by the unknown CDF on the basis of the observed sample $\boldsymbol{u}$. Let $\hat{R}$ denote the estimate of R based on the given sample. The nonparametric bootstrapping technique for estimating confidence intervals on the point estimator $\hat{R}$ is based on resampling (with replacement) from the EDF provided by the observed sample $\boldsymbol{u}$.

This is done as follows: The EDF gives rise to an empirical random variable U^*. l bootstrap samples $\boldsymbol{u}_j^*$, $j = 1, \ldots, l$, with n independent observations of U^* in each sample are now generated. Each sample $\boldsymbol{u}_j^*$ gives rise to an estimate R_j^* of R.

A simple estimator for confidence intervals on $\hat{R}$ is derived by calculating the sample standard deviation s_R^*:

$$s_R^* = \sqrt{\frac{1}{l-1}\sum_{j=1}^{l}(R_j^* - \bar{R}^*)^2} \tag{15}$$

where $\bar{R}^* = (1/l)\sum_{j=1}^{l} R_j^*$. An approximate confidence interval at level $1-q$ is then obtained as

$$(\,\hat{R} - w_{q/2}\, s_R^*\,,\; \hat{R} + w_{q/2}\, s_R^*\,) \tag{16}$$

where $w_{q/2}$ denotes the $100(1-q/2)\%$ standard normal quantile. To get stable results for this estimator, usually 20-30 bootstrap samples are sufficient. To avoid making the assumption that the bootstrap estimates are normally distributed, the true distribution may be approximated by generating a large number of bootstrap samples, usually several thousand are needed, especially for small values of q. If l samples have been generated, the R_j^* are rearranged in increasing order. A $100(1-q)\%$ confidence interval is then

$$(\,R_L^*\,,\; R_M^*\,) \tag{17}$$

where $L = [ql/2]$ and $M = [(1-q/2)l]$ ($[a]$ means the integer part of a). Such estimates may be further improved as described by Davison and Hinkley (1997). However, we shall not enter such details here.

Numerical Results

Synthetic Data

The accuracy and behaviour of the various estimation procedures discussed above will be studied by application to synthetic data, for which the values to be estimated can be calculated accurately. In particular, it is assumed that the underlying stochastic process $X(t)$ is stationary and Gaussian with mean value zero and standard deviation equal to 1.0. It is also assumed that the mean zero upcrossing rate $\nu_X^+(0)$ is such that the product $\nu_X^+(0)\,T = 10^3$, where T= 1 year, which would be typical for the wind speed process. The distribution of the yearly extreme value of $X(t)$ is then calculated by the formula

$$F^{1yr}(x) = \exp\left\{-10^3\,\exp\left(-\frac{x^2}{2}\right)\right\} \tag{18}$$

The 50-year return period value x^{50yr} is then calculated from the relation $F^{1yr}(x^{50yr}) = 1 - 1/50$, which gives $x^{50yr} = 4.65$.

The corresponding data to be used for the POT analyses are generated by observing that the peak events extracted from measurements of the wind speed process are usually peak events that are separated by three to four days. This is done to obtain approximately independent data, cf. Naess (1998b). In accordance with this, the peak event data are generated from the extreme value distribution

$$F^{3d}(x) = \exp\left\{-10\,\exp\left(-\frac{x^2}{2}\right)\right\} \tag{19}$$

where $\nu_X^+(0)\,T = 10$, which implies that T corresponds to approximately 3.65 days.

For the POT analyses, 5000 time series of 2000 peak events in each time series were generated from $F^{3d}(x)$. Each time series then corresponds to 20 years of data.

The POT analyses were carried out as follows: First, a threshold level was chosen so that about 40 peak events exceeded this threshold. The subsequent threshold levels were chosen so that approximately 10 new peak events were added for each new threshold. Figure 1 displays the results for x^{50yr} from the POT analyses using the four different estimators: de Haan, moment, c0 de Haan and c0 moment. The results are based on one particular time series and the figure shows x^{50yr} as a function of the number k of peak events above the threshold. To reduce the variability and facilitate the practical choice of a suitable threshold to be used for the prediction of x^{50yr}, the results given in Figure 1 are smoothed by a running filter. Specifically, in Figure 2 are plotted the results for $\tilde{x}_j^{50yr}$, where $\tilde{x}_j^{50yr}$ = $(1/5)\sum_{i=-2}^{2} x_{j+i}^{50yr}$. For comparison, the result obtained by the Gumbel method for yearly extremes combined with transformation of data was 4.66, cf. Naess (1998a). The target value is 4.65.

Figure 3 shows the mean values $\bar{x}^{50yr}$ of the calculated estimates from the 5000 independent samples as a function of the number k above the threshold. Judging from the mean values, Figure 3 indicates that the POT analysis in combination with the de Haan and the $c0$ de Haan estimators lead to varying degrees of significant underestimation of the target value. In other words, these two estimation procedures seem to have a significant negative bias. The best performance is clearly offered by the moment and the $c0$ moment estimators.

On the average they are also more stable as a function of threshold level expressed through the parameter k. While the moment estimator leads on the average to a slight underestimation, the $c0$ moment estimator leads to a slight overestimation of the target value, which may be preferable from a practical point of view.

In Figure 4 are presented the mean values $\bar{x}^{50yr}$ of the calculated estimates from 5000 bootstrapped samples based on the one sample used to produce the results shown in Figure 1. This figure clearly demonstrates what must be expected, that bootstrapping cannot rectify inherent bias in the EDF. It appears that for estimating x^{50yr}, the running mean of Figure 2 is perhaps a preferable alternative to bootstrapping.

As mentioned in the Introduction, another topic worth discussing is the problem of estimating confidence intervals associated with the predicted long return period values. In particular, the performance of the bootstrapping method for this purpose shall be investigated. To elucidate on this, for the specific choice of $k = 145$, two empirical PDFs have been calculated for each one of the estimators for x^{50yr} discussed above. One is obtained from the 5000 independent samples, providing an accurate approximation to the target PDF. The other is obtained on the basis of one arbitrary sample for which 5000 nonparametric bootstrap samples have been generated. The calculated PDFs using each of the four estimators are shown in Figures 5 - 12.

These figures indicate that the bootstrap method will provide reasonably accurate estimates of the confidence intervals associated with the various estimators. In particular, there is excellent agreement between the target and bootstrap PDFs for the two $c0$ estimators. A feature of significant practical interest is the observation that the $c0$ estimators have much higher efficiency than the other two estimators. This is clearly connected to the fact that effectively only one estimated parameter enters the $c0$ estimator, viz. a, while estimates of both a and c enter the other two estimators. To summarize this discussion, Table 1 lists the 2.5% and 97.5% quantile values of the distributions given in Figures 5 - 12 together with the length of the corresponding 95% confidence intervals. This table, in fact, also points to another interesting feature. In the sense of agreement with the 5000-sample values, it can be observed that bootstrapping seems to work much better for the POT method in combination with the two c0 estimators than it does for the POT method with the other two estimators.

Table 1: 2.5% and 97.5% quantile values and 95% confidence interval lengths of distributions in Figures 5 - 12

Method	2.5% value	97.5% value	95% conf. int. length
de Haan	4.10	5.08	0.98
de Haan bootstr.	4.19	4.99	0.80
Moment	4.22	5.10	0.88
Moment bootstr.	4.33	4.92	0.59
c0 de Haan	4.32	4.76	0.44
c0 de Haan bootstr.	4.33	4.77	0.44
c0 moment	4.42	5.01	0.59
c0 moment bootstr.	4.43	5.03	0.60

Wind Speed Data

Wind speed data measured at two locations in Norway have been analysed in the same way as the synthetic data. Specifically, the time series of the wind speed observed at the lighthouses at Ferder and Utsira have been subjected to POT analyses. The results for the predicted values of x^{50yr} are presented in Figures 13 and 14. Also shown in these figures are the corresponding results obtained by the classical Gumbel method, but for transformed data, cf. Naess (1998a)

In both cases the $c0$ moment estimator leads to significantly higher estimates than the de Haan estimator, in agreement with what was observed for the synthetic data. However, it is now seen that both the $c0$ estimators give significantly higher estimates than the other two methods, in contrast to the average behaviour for the synthetic data. For Utsira, the two $c0$ estimators leads to fairly good agreement with the Gumbel result for transformed data. This is not the case for Ferder. However, closer scrutiny of the underlying data reveals that the transformation used, that is $\alpha(x) = x^2$, is not the appropriate one for the Ferder Lighthouse data. For this location, it is, in fact, found that $\alpha(x) \approx x^3$ fits the tail behaviour of the underlying PDF better. Figure 15 shows the results for Ferder with this last transformation. Now the $c0$ moment estimator provides fairly good agreement with the Gumbel result for transformed data, which is 36.3 m/s, while the $c0$ de Haan estimator tend to underestimate slightly relative to this result.

Conclusions

The main focus of this paper has been the performance of the POT method in combination with various estimators. The goal has been the prediction of characteristic values associated with long return periods, typically 50 or 100 years, on the basis of recorded data for shorter periods of time, e.g. 20 years. In particular, the attention has been directed toward the de Haan and moment estimators applied to the data as recorded, that is, to untransformed data. In addition, two new estimators have been discussed, which are related to transformed data. The transformation is derived from the underlying distribution of the instantaneous value of the stochastic process under study.

An extensive set of synthetic data have been used to test the performance of the various estimators. The overall conclusion is that the de Haan estimator for untransformed data and the $c0$ de Haan estimator for transformed data on the average lead to significant underestimation of the target values. On the other hand, the moment estimator for untransformed data and the proposed $c0$ moment estimator for transformed data, provide on the average much better agreement with the target values. It has also been pointed out that the $c0$ estimators are superior in terms of efficiency.

To deal with the problem of estimating confidence intervals on the provided point estimates of the characteristic extreme values, bootstrapping has been proposed. It has been demonstrated that this technique seems to be an attractive practical tool for this purpose.

Acknowledgements

The authors want to express their gratitude to the Norwegian Meteorological Institute for making available the wind speed data for Ferder and Utsira.

References

Castillo, E. (1988). *Extreme Value Theory in Engineering*. San Diego: Academic Press.

Davison, A. C. and D. V. Hinkley (1997). *Bootstrap Methods and their Applications*. London: Cambridge University Press.

Davison, A. C. and R. L. Smith (1990). Models of Exceedances Over High Thresholds. *Journal of the Royal Statistical Society 52*(Ser. B), 393–442.

de Haan, L. (1994). Extreme Value Statistics. In J. Galambos, J. A. Lechner, and E. Simiu (eds.), *Extreme Value Theory and Applications*. Dordrecht: Kluwer Academic Publishers.

Efron, B. and R. J. Tibshirani (1993). *An Introduction to the Bootstrap*. New York: Chapman and Hall.

Fisher, R. A. and L. H. C. Tippet (1928). Limiting forms of the frequency distribution of the largest or smallest member of a sample. *Proc. Camb. Phil. Soc. 24*, 180–190.

Gross, J. L. et al. (1994). Novel Extreme Value Procedures: Application to Extreme Wind Data. In J. Galambos, J. A. Lechner, and E. Simiu (eds.), *Extreme Value Theory and Applications*. Dordrecht: Kluwer Academic Publishers.

Gumbel, E. (1958). *Statistics of Extremes*. New York: Columbia University Press.

Hosking, J. R. M. and J. R. Wallis (1987). Parameter and Quantile Estimation for the Generalized Pareto Distribution. *Technometrics 29*, 339–349.

Naess, A. (1998a). Estimation of Long Return Period Design Values for Wind Speeds. *Journal of Engineering Mechanics, ASCE 124*(3), 252–259.

Naess, A. (1998b). Statistical Extrapolation of Extreme Value Data Based on the Peaks Over Threshold Method. *Journal of Offshore Mechanics and Arctic Engineering, ASME 120*, 91–96.

Naess, A. and P. H. Clausen (1999). Statistical Extrapolation and the Peaks Over Threshold Method. OMAE–99–6422. In *Proceedings 18th International Conference on Offshore Mechanics and Arctic Engineering*. New York: ASME.

Pickands, J. (1975). Statistical Interference Using Order Statistics. *Annals of Statistics 3*, 119–131.

Smith, R. L. (1989). Extreme Value Theory. In W. Ledermann, E. Lloyd, S. Vajda, and C. Alexander (eds.), *Handbook of Applicable Mathematics*, pp. 437–472. New York: John Wiley and Sons.

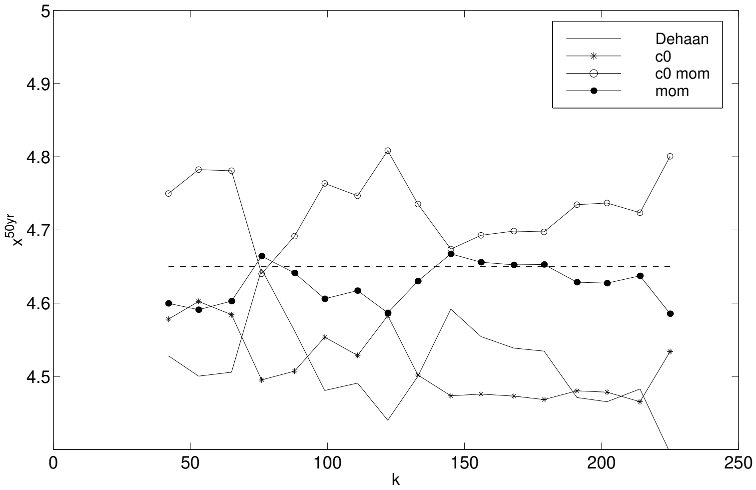

Figure 1:
The point estimate x^{50yr} of the 50-year return period value based on one sample as function of the number of peak events above threshold.

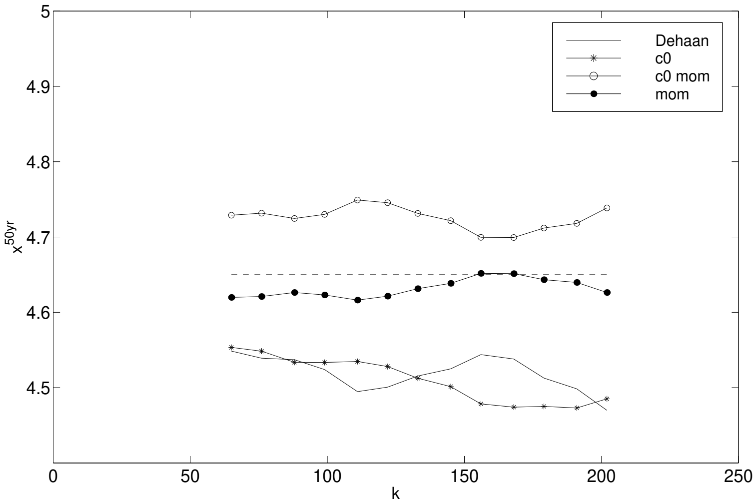

Figure 2:
The running mean value $\tilde{x}^{50yr}$ of the curves in Figure 1.

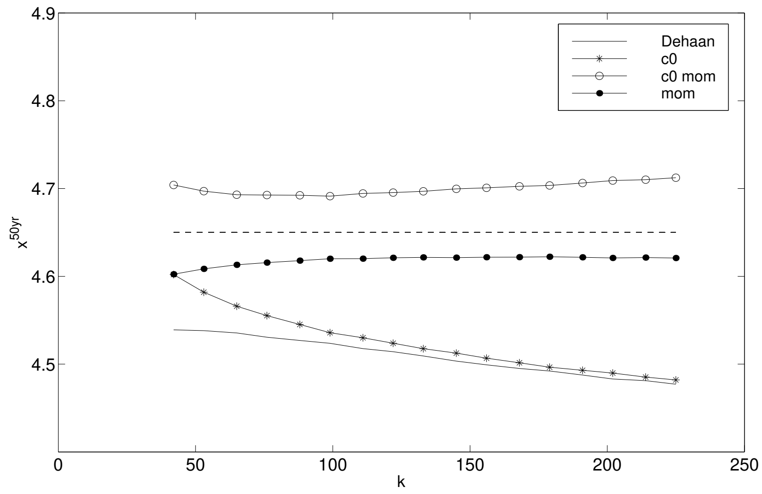

Figure 3:
The mean value $\bar{x}^{50yr}$ of the calculated estimates from 5000 independent samples.

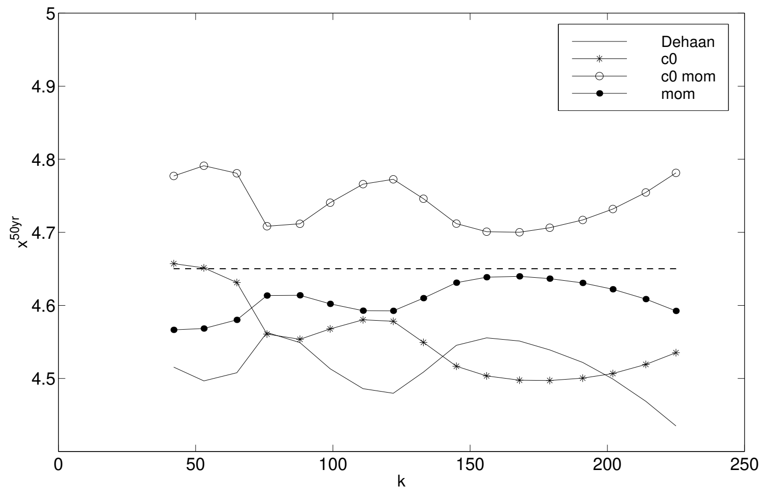

Figure 4:
The mean value $\bar{x}^{50yr}$ of the calculated estimates from 5000 bootstrap samples.

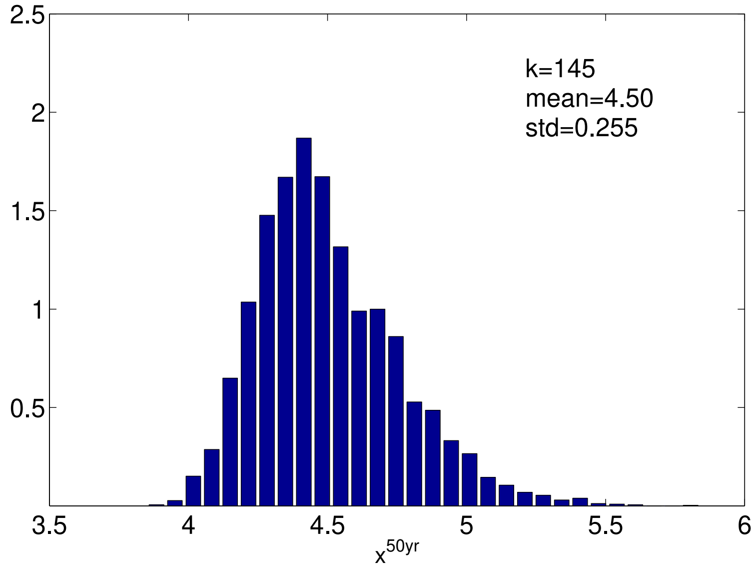

Figure 5:
Empirical PDF of x^{50yr} by de Haan's method from 5000 independent samples.

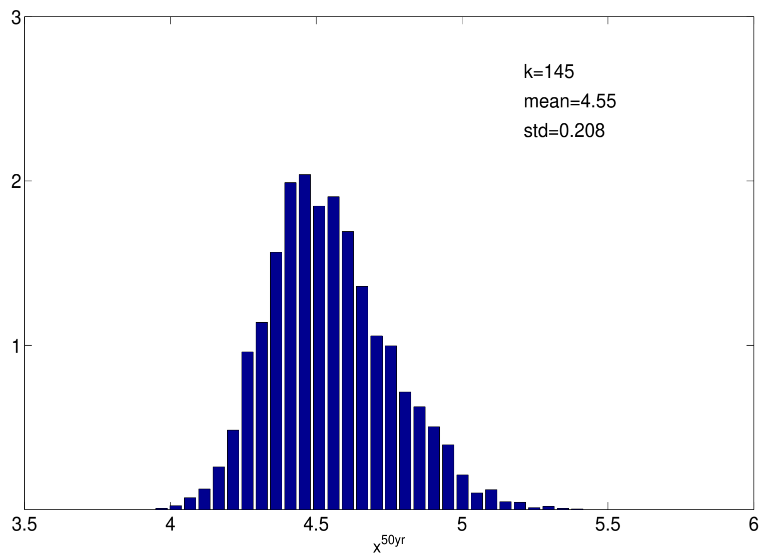

Figure 6:
Empirical PDF of x^{50yr} by de Haan's method from 5000 bootstrap samples.

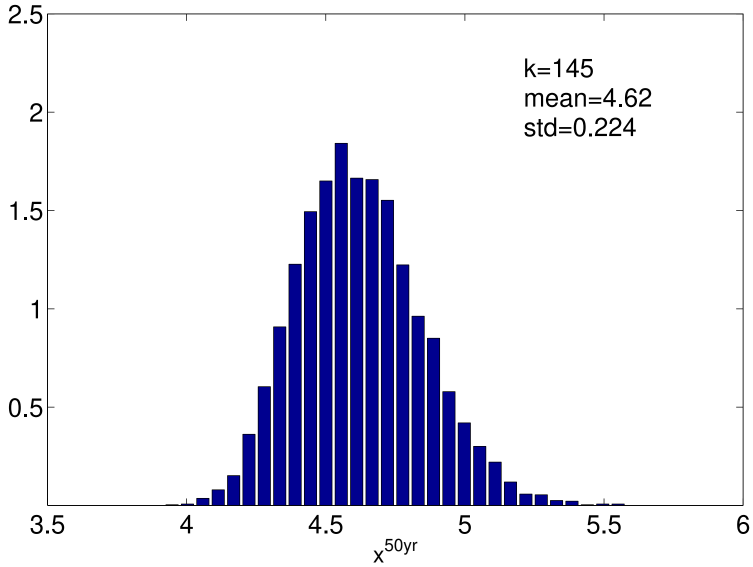

Figure 7:
Empirical PDF of x^{50yr} by the moment method from 5000 independent samples.

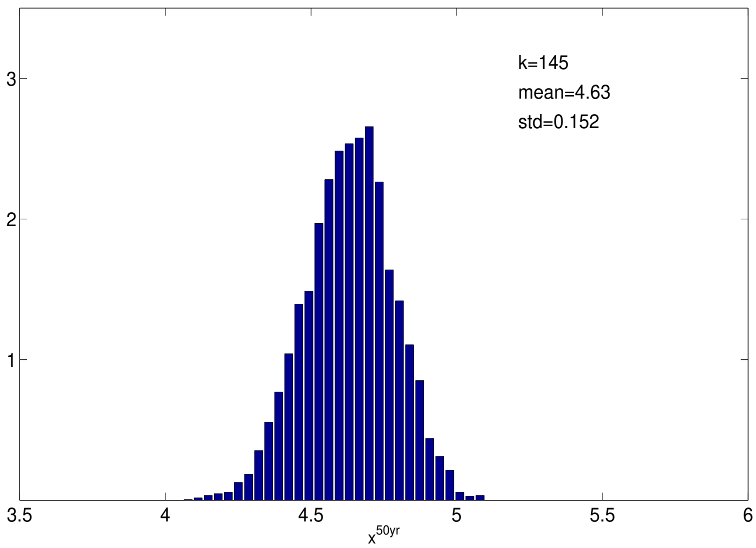

Figure 8:
Empirical PDF of x^{50yr} by the moment method from 5000 bootstrap samples.

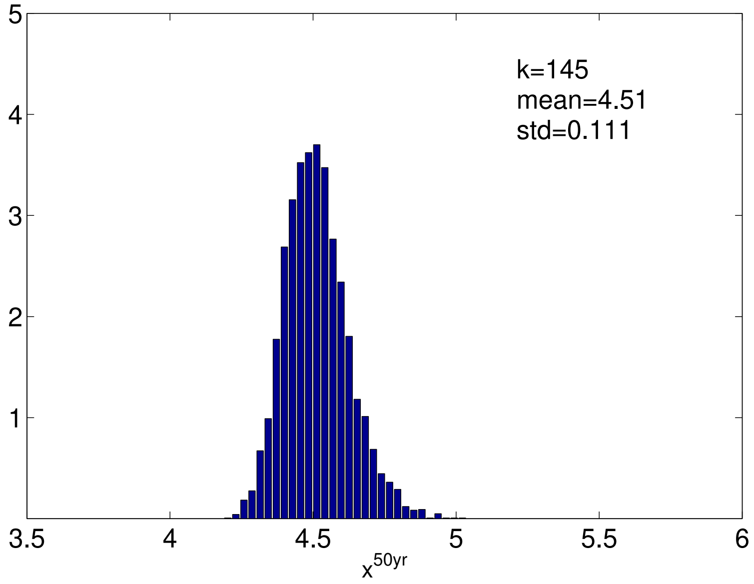

Figure 9:
Empirical PDF of x^{50yr} by the $c0$ de Haan method from 5000 independent samples.

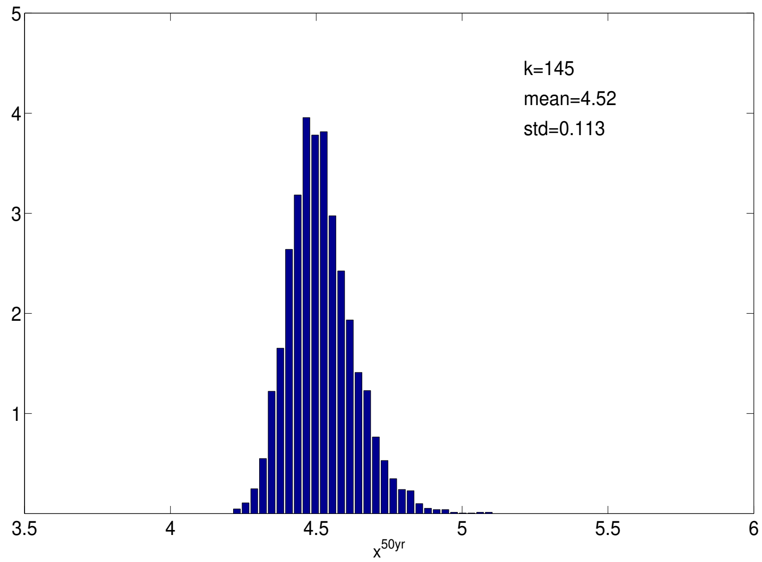

Figure 10:
Empirical PDF of x^{50yr} by the $c0$ de Haan method from 5000 bootstrap samples.

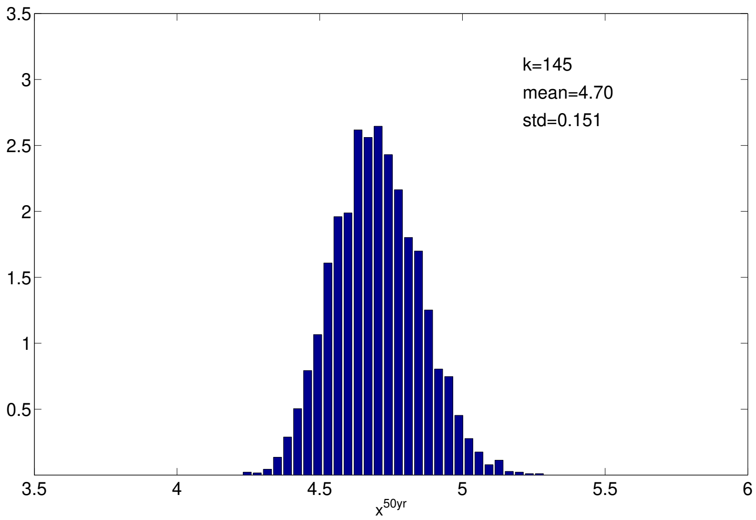

Figure 11:
Empirical PDF of x^{50yr} by the $c0$ moment method from 5000 independent samples.

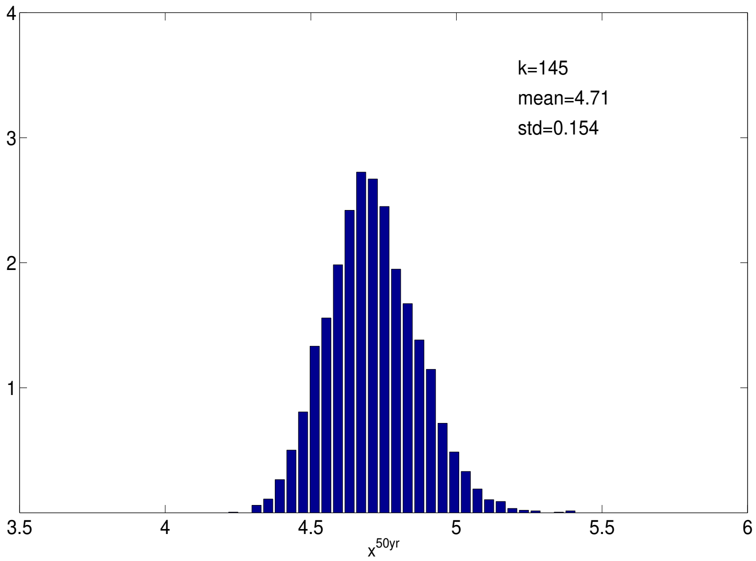

Figure 12:
Empirical PDF of x^{50yr} by the $c0$ moment method from 5000 bootstrap samples.

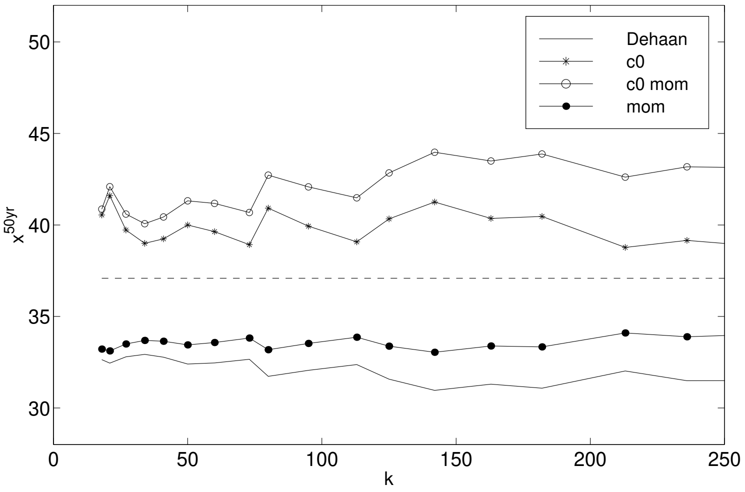

Figure 13:
Smoothed estimates of x^{50yr} for Ferder lighthouse. The broken line is the Gumbel estimate for transformed data.

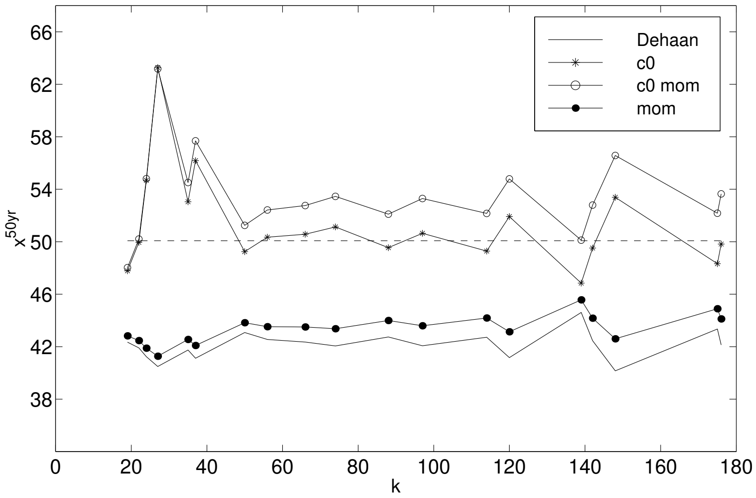

Figure 14:
Smoothed estimates of x^{50yr} for Utsira lighthouse. The broken line is the Gumbel estimate for transformed data.

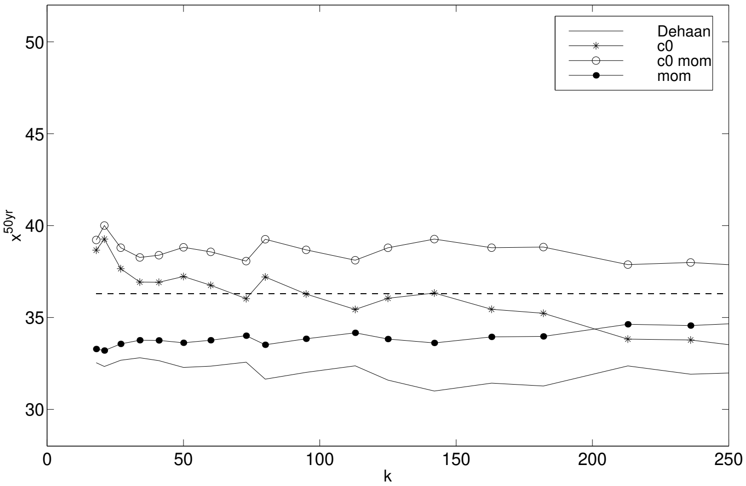

Figure 15:
Estimates of x^{50yr} for Ferder lighthouse with the x^3 transformation.