

GPD scale parameter

Tom Howard

August 31, 2021

0.1 Statistical Modelling of Extreme Values

To identify, for example, the 1000-year return level based solely on tide-gauge observations, some philosophy for making out-of-sample estimates is required. The usual approach is to exploit the most extreme observations, and theories concerning their behaviour, under some restrictive assumptions.

Annual Maxima

One popular and simple approach is fitting a Generalised Extreme Value (GEV) distribution to the annual maxima. The GEV distribution (GEVD) arises as the limiting case for block maxima as the block size tends to infinity. In the case of annual maxima, “block” means one year. The GEVD is characterised by three parameters. For readers unfamiliar with the GEVD, it may be helpful to picture the effect of these parameters in terms of a return-level curve, such as the ones shown in Fig. ???. The location parameter, μ , is comparable to an intercept. An increase in μ slides the whole curve up the Y-axis. μ is the Y-value (return level) evaluated at the one-year return period:

$$\mu = y \Big|_{L=0}$$

where $L = \log(\text{return period})$ and y is the return level. Notice that, though not particularly useful, this could be written

$$\mu = y \Big|_{y=\mu}$$

The GEV scale parameter, σ , is the gradient of the curve, evaluated at the one-year return period. This could either be written as

$$\sigma = \frac{dy}{dL} \Big|_{L=0} \tag{1}$$

or, for comparison with equation 6,

$$\sigma = \frac{dy}{dL} \Big|_{y=\mu}$$

Since y is a monotonic function of L , and in view of the first (unnumbered) equation, this is an alternative way to unambiguously define the point on the RL curve at which to evaluate the gradient. It's just specified in a non-standard way: in terms of the ordinate (y) instead of the usual specification in terms of the abscissa (L).

The shape parameter, ξ , determines the curvature. Negative ξ corresponds to a curve which flattens out at high return periods, approaching an upper bound as the return period tends to infinity. With positive ξ the curve has no upper bound, but has a lower bound as the return level decreases. When $\xi = 0$ the curve is a straight line and has neither lower nor upper bound. This follows the convention of [?] for the shape parameter. However, not all sources follow this convention. In CFB2018, "shape parameter" refers to the negative of our ξ . In the wider literature the "shape parameter" may refer to the negative or the reciprocal of our ξ . To make our shape parameter notation unambiguous: if Y is a random variable with GEV distribution, our shape parameter ξ is defined such that the distribution of Y is given by

$$P(Y < y) = \exp \left\{ - \left[1 + \xi \left(\frac{y - \mu}{\sigma} \right) \right]^{-1/\xi} \right\} \quad (2)$$

This can be more simply expressed as the corresponding return level curve, which is

$$\frac{y - \mu}{\sigma} = \frac{R^\xi - 1}{\xi} \quad (3)$$

where the average recurrence interval (or "return period") is R and the corresponding return level is y . The connection between equations 2 and 3 is seen by regarding exceedances of the R -year return level y as Poisson-distributed random occurrences, occurring at an average rate

$$\lambda = 1/R \quad (4)$$

The probability of no such occurrences in a given year is then given by standard Poisson statistics:

$$P(\text{no occurrences}) = P(Y < y) = \exp(-\lambda) \quad (5)$$

Combining 3, 4 and 5 gives equation 2. The particular case $\xi = 0$ is obtained by taking the limit as $\xi \rightarrow 0$.

Peaks over Threshold

The most extreme storm surges in the UK are caused by the storminess of the winter atmosphere, so the annual maximum event is always expected to occur in winter. Thus, an advantage of the annual-maxima approach described above is that the annual maxima are typically very well separated from each other and thus can be considered independent, particularly if the nominal year change is taken to be in the summer. A disadvantage of the approach is that it uses

only the annual maxima. On the other hand, the peaks-over-threshold (POT) approach uses all of the data exceeding a chosen threshold. This formed part of the approach taken by CFB2018. An advantage of this approach is that, if a low-enough threshold is used, it has the potential to exploit more of the available data (i.e. an average of more than one extreme event per year), whilst including only extreme events. Such exploitation of more data usually reduces the uncertainties in inferred statistics (e.g. the out-of-sample estimates). This is particularly desirable when short observational records limit the available extremes. However, if the threshold is too low, some of the data included can no longer be considered “extreme” and may bias the result. This is the well-recognised bias-variance trade-off. Another disadvantage is that including more than one event from a winter may compromise the independence of the events. (Skew surge can be evaluated for every high tide, and a weather system can generate a substantial skew surge on successive high tides.) Dependence is accommodated by CFB2018 using an extremal index... For a detailed comparison of the annual-maxima and POT approaches see...

The usual POT approach is to fit a Generalised Pareto Distribution (GPD) to the peaks. The GPD has two parameters. The shape parameter ξ is shared with the GEVD. The GPD scale parameter, $\tilde{\sigma}$, is the gradient of the plot of return level against log of return period at the return period of the chosen threshold, u ,

$$\tilde{\sigma} = \left. \frac{dy}{dL} \right|_{y=u} \quad (6)$$

As in the unnumbered equation following equation 1, the point on the RL curve at which to evaluate the gradient is specified in a non-standard way: in terms of the ordinate (y) instead of the usual specification in terms of the abscissa (L).

$\tilde{\sigma}$ is a property of both the extreme value distribution and the chosen threshold. The GEV scale parameter, σ , on the other hand, is a property of the extreme value distribution only and is thus a more fundamental parameter for making comparisons: it can be used in a like-for-like comparison of the results of different thresholds, or for comparison of GEV and GPD results. The two different scale parameters are related by $\sigma = \tilde{\sigma}\lambda_u^\xi$, where λ_u is the expected number of exceedances of u per year.

Though not formally a parameter of the GPD, a threshold must be chosen. CFB2018 tested 14 different thresholds and, finding no clear support for dismissal of any, elected to evaluate statistics based on each threshold and identify the median as the best estimate.