

Modelling non-stationary extremes of storm severity: comparing parametric and semi-parametric inference

Konzen, E., Neves, C. and Jonathan, P.

October 26, 2020

1 Introduction and Motivation

Statistical inference for extreme values has developed rapidly in recent years, offering considerable scope for practical application in science and engineering. Two seemingly divergent camps of statistical thought have emerged, proposing different approaches to extreme value modelling, yielding inferences which are not always obviously in agreement. The work presented in this paper seeks to identify potential points of contact between these so-called parametric and semi-parametric frameworks for extreme value inference, to encourage better common understanding and convergence of at least some practices, in particular for tackling non-stationary extremes.

Non-stationarity is commonplace in environmental extremes; physical processes generate extreme values which typically vary systematically with covariates, including space and time. For the peaks-over-threshold (POT) method, where only data exceeding a threshold are used for analysis [Balkema and de Haan, 1974, Pickands III, 1975], various models have been proposed to capture non-stationarity, including those of Davison and Smith [1990] and Leadbetter [1991]. In the parametric framework, non-stationarity can be incorporated within the appropriate distribution function for threshold exceedances by allowing the distribution's (shape and scale) parameters to vary with covariates [see e.g. Coles, 2001, Chavez-Demoulin and Davison, 2005, and references therein]. Important assumptions underpinning this approach are that the data generating process is locally stationary and that observations from the data generating process can be considered approximately independent given covariates.

Applications of non-stationary extreme value analysis are more numerous using parametric inference. For example, in an ocean engineering context, Forristall [2004] performs extreme value analysis of significant wave height for directional octants. This approach (i) accommodates directional non-stationarity (by conducting independent inferences per directional octant) and (ii) allows extreme quantiles for specific directional sectors to be estimated. Choice of number and widths of directional sectors is an open problem (see e.g. Ross et al. 2018, Folgueras et al. 2019). Moreover, non-stationary inference provides a number of challenges in practical assessment and interpretation of uncertainty and risk (e.g. Jonathan and Ewans 2007, Mackay et al. 2010, Feld et al. 2019).

The need to incorporate a non-stationary threshold within extreme value inference is well-recognised for environmental applications. Robinson and Tawn [1997] base their approach on a non-stationary threshold to characterise the evolution of extreme sea currents. Northrop and Jonathan [2011] propose a covariate-dependent threshold estimated using quantile regression, and Northrop et al. [2016] propose a cross-validation procedure for threshold selection. We introduce a method for selection of a non-stationary threshold amenable to both parametric and semi-parametric approaches to inference. The basic assumption is that on a narrow region of the directional domain there is a unique physical regime which dictates the tail heaviness of the underlying distribution although extreme events can become more (or less) frequent. In this work, we propose a procedure for non-stationary threshold selection, which incorporates a smooth weight function to describe the varying density of extreme observations with directional covariate on the circular domain. The smooth weight function avoids inconsistencies stemming from the sectoring approach identified by Forristall

[2004]. The threshold selection procedure is common to both parametric and semi-parametric approaches used for extreme value estimation.

Estimating the effects of periodic or circular covariates including direction and season in extreme value analysis is challenging in general. Coles and Walshaw [1994] consider extreme value analysis of wind speed on a circular domain, and Mendez et al. [2008] a seasonal model for significant wave height. Since extreme value parameters might be expected to vary smoothly with covariate, various different smooth basis representations (e.g. splines) for parameters of the conditional distribution of threshold exceedances, as a function of covariate as useful (see e.g. Ugarte et al. 2012, Oumow et al. 2012, Jones et al. 2016, Zanini et al. 2020). Covariate basis representations require careful choice of order of complexity to avoid over-parametrisation. Some representations are attractive since basis functions have compact support on the covariate domain, facilitating computationally efficient and stable inference (Randell et al. 2016)). In this work we consider a parametric estimation scheme using a penalised periodic B-spline covariate representation.

An alternative approach to non-stationary extreme value analysis is to estimate parametric and semi-parametric models which are local on the covariate domain. Local models avoid the complexity of global covariate representations, particularly problematic for covariate intervals where extreme events are rare. In this work we consider semi-parametric Moment and local parametric maximum likelihood models.

Objectives and layout

The goal in this paper is to explore the synergies between different parametric and semi-parametric approaches to inference on extreme value data non-stationary with respect to periodic covariate. The default parametric approach employed is maximum likelihood estimation for peaks over threshold assumed to follow a generalised Pareto distribution, with spline descriptions of model parameter variation with covariate on some covariate domain. For comparison, two other covariate-local stationary models are considered: the first uses local parametric maximum likelihood estimation, and the second local semi-parametric Moment estimation; to achieve the latter we extend the semi-parametric methodology so that inference for quantities of interest in ocean engineering is possible and meaningful. We develop (a) a new common local heuristic procedure for extreme value threshold selection for both local parametric and semi-parametric frameworks, and (b) local inference for generalised Pareto shape parameter, extreme value index and related quantities.

We then present a comparative study showing how the spline parametric, local parametric and semi-parametric approaches provide complementary inference for directional extremes in terms of (a) extreme value threshold selection, (b) extreme value model estimation and (c) extreme quantile estimation including the T -year value and right endpoint of the support of the underlying distribution when is finite. We illustrate the approaches in application to directional extreme value analysis of hindcast storm peak significant wave height recorded at a northern North Sea location offshore Norway. We conduct hypothesis tests to characterise variation in the max-domain of attraction of the distribution of storm peak significant wave height as a function of covariate. We hope to demonstrate that spline parametric, local parametric and semi-parametric approaches provide complementary non-stationary extreme value inference, and that local parametric and semi-parametric approaches may be of practical benefit to practitioners in coastal and ocean engineering and environmental sciences.

The remainder of the paper is organised as follows. The application to directional modelling of H_S^{SP} is introduced in Section 2. Section 3 provides theoretical motivation for non-stationarity analysis. Section 4 discusses non-stationary threshold estimation and extreme value model parameter estimation. Section 5 and Section 6 discuss estimation of extreme value parameters and extreme quantiles. Results of the North Sea application are presented in Section 7. Finally, Section 8 provides discussion and conclusions. Appendix A provides theoretical motivation for non-stationary inference on peaks over threshold, and Appendix B provides algorithms for the estimation of extreme value threshold, spline generalised Pareto inference and for the overall analysis.

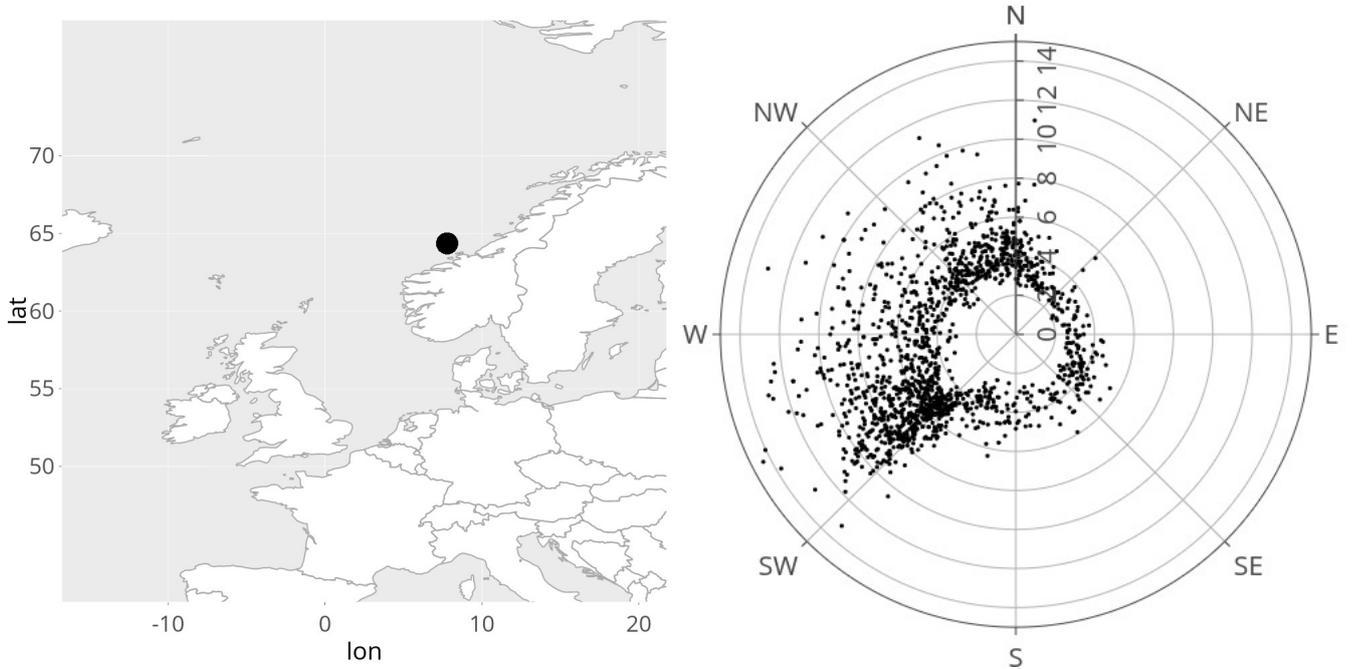


Figure 1: Left: map showing the North Sea location. Right: polar scatter-plot of H_S^{sp} on the direction θ from which waves propagate; radial scale of H_S^{sp} is in metres.

2 Motivating application

The data for the motivating application is described by Randell et al. [2016], consisting of observations of storm severity, in terms of storm peak significant wave height henceforth referred to as H_S^{sp} and storm direction θ at a location in the northern North Sea. Significant wave height H_S measures the roughness of the ocean surface, and can be defined as four times the standard deviation of the ocean surface elevation at a spatial location for a specified period of observation. H_S^{sp} is then the maximum value of H_S observed per time period corresponding to a storm event. Values of H_S^{sp} can be regarded as approximately independent given storm peak characteristics.

The application sample is taken from the WAM hindcast of Reistad et al. [2011], which provides time-series of significant wave height, (dominant) wave direction and season (defined as day of the year, for a standardised year consisting of 360 days) for three hour sea-states for the period September 1957 to December 2012 at a northern North Sea location in the vicinity of the black disk in the upper panel of Figure 1. A hindcast is a physical model of the ocean environment, incorporating pressure field, wind field and wind-wave generation models in particular; the hindcast model is calibrated to observations of the environment from instrumented offshore facilities, moored buoys and satellite altimeters in the neighbourhood of the location for a period of time, typically decades. Extreme seas in the North Sea are dominated by winter storms originating in the Atlantic Ocean and propagating eastwards across the northern part of the North Sea. Due to their proximity to the storms, sea states at northern North Sea locations are usually more intense than in the southern North Sea. Occasionally, the storms travel south-eastward and intrude into the southern North Sea producing large sea states. Directions of propagation of extreme seas vary considerably with location, depending on land shadows of the British Isles, Scandinavia, and the coast of mainland Europe, and fetches associated with the Atlantic Ocean, Norwegian Sea, and the North Sea itself. In the northern North Sea the main fetches are the Norwegian Sea to the North, the Atlantic Ocean to the west, and the North Sea to the south. Extreme sea states from the directions of Scandinavia to the east and the British Isles to the south-west are unlikely. The shielding by these land masses is more effective for southern North Sea locations, resulting in a similar directional distribution but reduced wave heights by comparison with northern North Sea locations.

At the location of interest, observations of storm peak significant wave height H_S^{sp} are isolated from the hindcast time-series using the procedure described in Ewans and Jonathan [2008] as follows. Briefly, contiguous intervals of significant wave height above a low peak-picking threshold are identified, each interval assumed to correspond to a storm event. The peak-picking threshold corresponds to a directional-seasonal quantile of H_S with specified non-exceedance probability, estimated using quantile regression. The maximum of significant wave height during the storm interval is taken as the storm peak significant wave height H_S^{sp} . The values of directional and seasonal covariates at the time of storm peak significant wave height are referred to as storm peak values of those variables. The resulting storm peak sample consists of 2941 values. With direction *from which* a storm travels expressed in degrees clockwise with respect to north, the lower panel of Figure 1 shows a polar plot of observations of H_S^{sp} (in metres) versus direction θ .

The land shadow of Norway (approximately the directional interval $(45^\circ, 210^\circ)$) has a considerable effect on the rate and size of occurrences with direction. In particular, there is a dramatic increase in both rate and size of occurrences with increasing direction at around 210° , corresponding to Atlantic storm events from the south-west able to pass the Norwegian headland. We therefore expect considerable directional variability in extreme value parameter estimates for the sample. In contrast, the magnitude of the rate of change of both rate and size of occurrences with respect to season (not shown, but see Randell et al. 2016) is lower. Winter storms (approximately from October to March) are more intense and frequent. Only winter storms with storm peak events occurring in October to March including have been considered further in this work, corresponding to a sample of 1521 storm peak values.

3 Theoretical motivation for the non-stationary case

In Section 3.1 we first summarise theoretical results and assumptions underpinning extreme value modelling for the stationary (or omni-directional, or covariate-free) case. In Section 3.2 we then present a framework for local modelling of non-stationary extremes with covariate over a circular directional domain. These form the basis of inference schemes presented in Section 4.

3.1 Stationary case

Suppose the available sample consists of realisations of n independent (or weakly dependent) and identically distributed random variables X_1, X_2, \dots, X_n . Since the random variables follow the same (unknown) distribution function, for brevity use the symbol X to refer to any of the random variables when there is no need to be specific. The common distribution function is then $F(x) = P(X \leq x)$, for every $x \in \mathbb{R}$. We denote by x^F the right endpoint of the support of F , namely the ultimate value which bounds all possible observations from above, $x^F = \sup\{x : F(x) < 1\}$. We note that x^F may be less than or equal to infinity.

The extreme value (or extreme types) theorem (Fisher and Tippett 1928, Gnedenko 1943, de Haan 1970) establishes that the limit distribution of linearly normalised partial maxima $M_n = \max(X_1, \dots, X_n)$ (informally “block maxima”), with normalising constants $a_n > 0$ and $b_n \in \mathbb{R}$, must be one of only possible three extreme value distributions: Fréchet, Gumbel or Weibull. These three types can be nested in the one-parameter Generalised Extreme Value (GEV) distribution. Specifically, if a_n and b_n exist such that,

$$\lim_{n \rightarrow \infty} P\{a_n^{-1}(M_n - b_n) \leq x\} = \lim_{n \rightarrow \infty} F^n(a_n x + b_n) = G(x), \quad (3.1)$$

for every continuity point of (non-degenerate) G , then G must be a GEV distribution with distribution function given by:

$$G_\xi(x) = \exp(-(1 + \xi x)^{-1/\xi}), \quad (3.2)$$

for all x such that $1 + \xi x > 0$. We then say that F belongs to the max-domain of attraction of the GEV, for some $\xi \in \mathbb{R}$, and write $F \in \mathcal{D}(G_\xi)$. Parameter ξ is conventionally referred to as the shape parameter in parametric literature, and the extreme value index (EVI) in semi-parametric literature of extremes. A max-domain of attraction is thought of as a class enclosing all distribution functions which, in the limit, are

associated with a particular extreme value index ξ . For example, both Cauchy and Student's t -distribution with one degree of freedom are found in the same max-domain of attraction $\mathcal{D}(G_\xi)$, with $\xi = 1$. The Fréchet max-domain of attraction, corresponding to $\xi > 0$, contains all distributions exhibiting polynomial decay, such as the Pareto, again Cauchy and Student's t and the Fréchet itself. These distributions have infinite right endpoint x^F . All distribution functions belonging to $\mathcal{D}(G_\xi)$ with $\xi < 0$, referred to as the Weibull max-domain of attraction, are short-tailed with finite x^F ; examples include the uniform and beta distribution. For $\xi = 0$, a continuity argument gives $G_{\xi=0}(x) = \exp(-e^{-x})$, $x \in \mathbb{R}$; the corresponding Gumbel max-domain of attraction ($\mathcal{D}(G_\xi)$ with $\xi = 0$, or $\mathcal{D}(G_0)$), of particular interest in many applied fields due to both simplicity of inference and the great variety of distributions possessing an exponential tail (with either $x^F < \infty$ or $x^F = \infty$). The normal, gamma and log-normal distributions are members of $\mathcal{D}(G_0)$.

Extreme value inference using block maxima has been practised for many years, from the time of Gumbel [1958] in application to hydrology. However, in this paper, inference is based on analysis of threshold exceedances or peaks over threshold (POT), motivated by characterisation of the max-domains of attraction above in terms of exceedances of high threshold. Balkema and de Haan [1974] and Pickands III [1975] established that the max-domain of attraction condition (3.1) is equivalent to the assertion that the (conditional) distribution of X given $X > u$, with u near the right endpoint x^F , converges to the generalised Pareto distribution (GPD) with distribution function $1 + \log G_\xi$. Analysis of threshold exceedances is potentially statistically more efficient than that of block maxima: in the former, all large values of threshold exceedance in the sample are admitted, including multiple occurrences of large values belonging to same block, which might be excluded in a block maximum analysis.

The parametric and semi-parametric approaches to extreme value analysis of threshold exceedances considered in this paper have strong conceptual similarities, and both approaches involve estimation of three quantities corresponding to threshold, generalised Pareto shape and scale in parametric inference, and threshold, extreme value index (EVI) and associated scale in semi-parametric inference. There is a large literature on estimation of shape parameter in parametric inference, and the extreme value index in semi-parametric inference. In this work we use Maximum likelihood (ML) (Coles 2001, Smith 1987, Davison and Smith 1990) and Moment (M) estimators [Dekkers et al., 1989] as typical parametric and semi-parametric estimators, due to their popularity. The maximum likelihood estimator assumes that the limiting generalised Pareto distribution provides an accurate description of the largest values in the sample, as though the exceedances over a fixed threshold were drawn from an actual generalised Pareto distribution; this provides a highly flexible framework for inference. The Moment estimator builds on the concept of max-domain of attraction, and in this sense is model-free; it is more robust: there are seldom any assurances that exceedances of high threshold actually follow the generalised Pareto form. The Moment estimator can be viewed as intuitive, since its form involves summary statistics like averages. When goodness of fit can be assessed with reasonable accuracy, the maximum likelihood estimator can be shown to be efficient for the shape parameter. There is concern however that the maximum likelihood may be influenced by deviations of the true underlying distribution function from the assumed generalised Pareto, resulting in extrapolation bias. For this reason, it is sometimes preferable to consider a robust estimator building on moments to lessen the importance of such deviations, at the expense of increased variance.

Viewed parametrically or semi-parametrically, choosing an appropriate extreme value threshold is fundamental to inference for peaks over threshold, and involves a typical bias-variance trade-off. If the threshold is set too high, the set of threshold exceedances is small resulting in increased variance of parameter estimates. If the threshold is set too low, the variance of the parameter estimate is reduced at the cost of increased bias. Sensitivity to the extreme value threshold choice is a common critical feature of parametric and semi-parametric approaches. Confirming the relative stability of estimated shape parameter (or extreme value index) and high quantiles (such as the T -year level) with near-zero exceedance probability with respect to threshold is a key diagnostic test for peaks over threshold analysis, addressed further in Section 6.

Given threshold, inference for peaks over threshold in the parametric setting involves estimation of generalised Pareto shape and scale parameters, the latter being threshold-dependent; for estimation of T -year levels, we also need to estimate the rate of threshold exceedance. In the semi-parametric setting, we

focus on estimation of extreme value index; scale and location normalising functions (akin to $a > 0$ and $b \in \mathbb{R}$ in (A.5)) required for extreme quantile calculations, are estimated separately.

3.2 Non-stationary case

In the non-stationary setting, we assume that the covariate domain \mathcal{S} is a circle, partitioned into m intervals, with centroids θ_j , and write Θ for the set $\{\theta_j\}_{j=1}^m$ of centroids; for directional analysis considered here, we set $\mathcal{S} = [0^\circ, 360^\circ)$, and $\theta_j = j - 1$ for $j = 1, \dots, m = 360$. For each θ_j , we suppose there are n independent copies $\{X_i(\theta_j)\}_{i=1}^n$ of random variable $X(\theta_j)$ with distribution function $F_{\theta_j}(x) = P\{X(\theta_j) \leq x\}$, $x \in \mathbb{R}$. We also assume that $X(\theta_j)$ and $X(\theta_{j'})$ are independently distributed when $j \neq j'$. We then assume the extreme value theorem holds on covariate neighbourhood $\mathcal{N}(\theta_j, h)$ of half-width $h > 0$ centred at θ_j (see Section 4.1). This means that condition (3.1) holds with respect to F_{θ_j} and the shape parameter (or extreme value index) $\xi(\theta_j)$ governs the tail behaviour of the underlying F_{θ_j} . We also assume throughout that changes in the tail dynamics of the data generating process are smooth, and that $\xi(\theta)$ is a smooth (unknown) function of $\theta \in \mathcal{S}$. These assumptions justify the non-stationary parametric and semi-parametric inference procedures used in this work.

The theoretical motivation for the non-stationary case is given in Appendix A, and summarised here. Suppose that for each $\theta \in \mathcal{S}$, neighbourhood $\mathcal{N}(\theta, h)$ contains $N(\theta)$ random occurrences of X . We start by selecting $k(\theta)$ threshold exceedances on that neighbourhood, assuming a limiting generalised Pareto distribution holds for sufficiently large $N(\theta)$. Given the initial assumption of (weak) dependence on \mathcal{S} , tail observations from F_θ can be viewed as peaks over threshold at each θ as follows, for any $\theta \in \mathcal{N}(\theta, h) \subset \mathcal{S}$. With the notation $N = n \times m$, let $X_{N-k(\theta):N}$ denote the $(N - k(\theta))^{\text{th}}$ neighbourhood-specific threshold having selected fixed number $k(\theta)$ for that neighbourhood. To ensure integrity of statistical analysis on \mathcal{S} , the number $k(\theta)$ is defined using a weighting function $\omega(\theta_j) > 0$ so that $k(\theta) := [k \times \omega(\theta)]$, with $[\cdot]$ indicating ‘‘integer part of’’, and k is the total number of largest observations retained for inference across the whole of \mathcal{S} . For a finite number m of directional sectors θ_j , the smooth function $\omega(\theta)$ satisfies $\sum_{j=1}^m \omega(\theta_j)/m = 1$.

Given threshold $u(\theta) = X_{N-k(\theta):N}$, we perform non-stationary extreme value inference for peaks over threshold. In the parametric setting, inference involves estimation of generalised Pareto shape and scale parameters which vary on the covariate domain; for estimation of T -year levels, we also need to estimate the rate of threshold exceedance which also varies with covariate. The semi-parametric approach is concerned with extrapolation beyond the sample through estimation of extreme quantiles with very small tail probabilities; semi-parametric estimators are typically smooth functionals (averages, quantiles, etc.) of sample upper order statistics, estimating extreme value index and scale function, in terms of the extreme sample fraction $k(\theta)/N$.

4 Non-stationary threshold selection

Here we describe an approach to non-stationary threshold selection useful for all of parametric and semi-parametric settings. The approach exploits a heuristic criterion to select a number $k = k(\theta)$ of upper order statistics, and hence a random threshold on the covariate domain. Algorithm 1 of Appendix B provides a simple implementation of the threshold selection procedure. In Section 5 we then discuss how the estimated threshold is used in extreme value estimation using each of the parametric, local parametric and semi-parametric approaches. Note that although threshold estimation is made here using either the semi-parametric Moment estimator (Section 4.2) or maximum likelihood estimation (Section 4.3), any other consistent estimator would be equally applicable.

4.1 Threshold selection

Let $\mathcal{N}(\theta_j, h)$ be the covariate neighbourhood with centre at $\theta_j \in \Theta$ and fixed radius $h > 0$ defined by

$$\mathcal{N}(\theta_j, h) = \{\theta^* \in \mathcal{S} : 0 \leq d(\theta^*, \theta_j) \leq h\}, \quad (4.1)$$

for every θ_j for an appropriate choice of distance d : in the directional application, we adopt wrapped-Euclidean distance $d(\theta^*, \theta_j) := \min\{|\theta^* - \theta_j|, 360 - |\theta^* - \theta_j|\}$ for $\theta^* \in \mathcal{S} = [0, 360)$. For each θ_j , the number $N(\theta_j)$ of realisations of $X_i(\theta_{j'})$, $i = 1, \dots, n$, $\theta_{j'} \in \Theta$, in the neighbourhood within lag h of centroid θ_j is

$$N(\theta_j) = |\mathcal{C}(\theta_j, h)|$$

and $\mathcal{C}(\theta_j, h) = \{X_i(\theta_{j'}) \mid \theta_{j'} \in \mathcal{N}(\theta_j, h), i = 1, 2, \dots, n\}$. A judicious choice of h ensures a large enough number $N(\theta_j)$ is present, so that the extreme value theorem holds on $\mathcal{N}(\theta_j, h)$, for every θ_j .

We now consider a heuristic threshold selection procedure exploiting the relative density of extreme values in $\mathcal{N}(\theta_j, h)$. We seek the optimal number k_j^* of the k largest observations in $\mathcal{N}(\theta_j, h)$ as threshold exceedances. Then k_j^* is the solution of $S_\phi(k_j^*) = \min_k S_\phi(k)$ with

$$S_\phi(k) = \frac{1}{k} \sum_{i \leq k} i^\phi |\hat{\xi}_i(\theta_j) - \text{median}(\hat{\xi}_1(\theta_j), \hat{\xi}_2(\theta_j), \dots, \hat{\xi}_k(\theta_j))| \quad (4.2)$$

where $0 \leq \phi < 0.5$, and $\hat{\xi}_k(\theta_j)$ is a locally consistent estimator of $\xi(\theta_j)$ using the k largest observations in $\mathcal{N}(\theta_j, h)$. The heuristic procedure (4.2) was introduced in the first edition of Reiss and Thomas (the most recent edition being Reiss and Thomas 2007) and then studied in detail in Neves and Fraga Alves [2004]. It facilitates automatic threshold choice, understood intuitively as follows. For small k , the weighted deviations in (4.2) tend to be large due to the inherently large variance of $\hat{\xi}_k(\theta_j)$. As k increases, the summands in (4.2) are expected to decrease until bias sets in and overrides the variance from which point $S_\phi(k)$ is expected to increase again. Minimising the weighted empirical distance (4.2) is equivalent to optimising the bias-variance trade-off by exploiting the settled behaviour of estimates $\{\hat{\xi}_k(\theta_j) : k < N(\theta_j)\}$ for appropriate k .

We specify a covariate weight $\omega(\theta_{j'}; \theta_j)$ of observations for every $\theta_{j'} \in \mathcal{N}(\theta_j, h) \cap \Theta$ relative to centroid $\theta_j \in \Theta$ using a von Mises kernel [cf. Pewsey et al., 2014], such that

$$\omega(\theta_{j'}; \theta_j) := \frac{\mathcal{K}_\eta(\theta_{j'}; \theta_j)}{\sum_{\theta_{j'} \in \mathcal{N}(\theta_j, h)} \mathcal{K}_\eta(\theta_{j'}; \theta_j)}, \quad (4.3)$$

with von Mises kernel $\mathcal{K}_\eta(\theta; \theta_j) := \exp\{\eta \cos(\theta - \theta_j)\} / (2\pi B_0(\eta))$ for $\theta \in \mathcal{S}$, where B_0 is the modified Bessel function of the first kind of order zero $B_0(\eta) = \pi^{-1} \int_0^\pi \exp[\eta \cos(\theta)] d\theta$. Concentration parameter $\eta > 0$ controls kernel spread: as η increases, the local density becomes more concentrated about θ_j ; thus η plays a role similar to bandwidth h in (4.1), regulating degree of directional smoothness. The values of η and h need to be specified carefully during applications to reflect the characteristics of the data generating process, and the resulting choice of k_j^* reflects the extent of stationarity on $\mathcal{N}(\theta_j, h)$. The corresponding optimal estimate for non-stationary threshold is denoted $u(\theta_j)$ for each $\theta_j \in \Theta$.

In this section we indicate the dependence of $\omega(\theta_{j'}; \theta_j)$ on centroid θ_j explicitly for clarity. In the discussion in Appendix A, we write $\omega(\theta_{j'})$ for notational convenience, with the dependence on θ_j understood. We do not claim that this heuristic approach to extreme value threshold specification is the most effective (e.g. with respect to every estimator one might devise for extreme value index), but we have found it useful in the application to estimation of ξ in Section 7.

4.2 Semi-parametric estimation

Now we outline semi-parametric estimation of $\xi_k(\theta_j)$ required for threshold selection above. For each $\theta_j \in \Theta$, we propose an adaptation of the extended version of the Moment estimator (Dekkers et al. 1989) for $\xi(\theta_j)$, which incorporates directional weights $\omega(\theta_{j'}; \theta_j)$ defined in (4.3). Let $h > 0$ be a bandwidth such that $h = h(N(\theta_j)) \rightarrow 0$ and $k(\theta)h \rightarrow \infty$, as $N(\theta_j) \rightarrow \infty$. The Moment estimator for the extreme value index is defined as

$$\hat{\xi}_k^M(\theta_j) := M^{(1)}(\theta_j) + 1 - \frac{1}{2} \left(1 - \frac{(M^{(1)}(\theta_j))^2}{M^{(2)}(\theta_j)} \right)^{-1}, \quad (4.4)$$

with

$$M^{(l)}(\theta_j) := \sum_{\mathcal{C}(\theta_j, h)} \omega(\theta_{j'}; \theta_j) (\log X_i(\theta_{j'}) - \log X_{N(\theta_j) - k, N(\theta_j)})^l \mathbb{1}_{\{X_i(\theta_{j'}) > X_{N(\theta_j) - k, N(\theta_j)}\}}, \quad l = 1, 2, \quad (4.5)$$

where $X_{N(\theta_j) - k, N(\theta_j)}$ denotes the $(k + 1)^{\text{th}}$ largest value in the observed sample in $\mathcal{N}(\theta_j, h)$. Note only the largest k observations in $\mathcal{N}(\theta_j, h)$ contribute in (4.5). Further discussion of the Moment estimator is given in Section 5.1.

4.3 Maximum likelihood estimation

The local maximum likelihood estimator $\hat{\xi}_k(\theta_j)$ for $\theta_j \in \Theta$ for threshold selection above is obtained by maximising the directionally-weighted log-likelihood

$$\mathcal{L}(\xi(\theta_j), \sigma_u(\theta_j)) := \sum_{\mathcal{C}(\theta_j, h)} \omega(\theta_{j'}; \theta_j) \ell(\xi(\theta_j), \sigma_u(\theta_j) | X_i(\theta_{j'}) - X_{N(\theta_j) - k, N(\theta_j)}) \mathbb{1}_{\{X_i(\theta_{j'}) - X_{N(\theta_j) - k, N(\theta_j)} > 0\}}, \quad (4.6)$$

with respect to the parameter pair $(\xi(\theta_j), \sigma_u(\theta_j)) \in (-1, \infty) \times \mathbb{R}^+$, with weights $\omega(\theta_{j'}; \theta_j)$ as in Section 4.2, and

$$\ell(\xi(\theta_j), \sigma_u(\theta_j) | y) = -\log \sigma_u(\theta_j) - (1 + 1/\xi(\theta_j)) \log(1 + \xi(\theta_j) y / \sigma_u(\theta_j)) \quad (4.7)$$

for $y = X_i(\theta_{j'}) - X_{N(\theta_j) - k, N(\theta_j)}$ when $\xi(\theta_j) \neq 0$ and $\ell(\sigma_u(\theta_j) | y) = -\log \sigma_u(\theta_j) - y/\sigma(\theta_j)$ otherwise.

5 Non-stationary extreme value estimation

This section considers estimation of extreme value parameters and extreme quantiles. We first present adapted versions of widely-used semi-parametric estimators for extreme value index and associated scale function based on a known moving threshold $u(\theta), \theta \in \mathcal{S}$. Then we discuss two classes of parametric maximum likelihood estimators for generalised Pareto shape and scale. The first is a minor modification of the local maximum likelihood estimator introduced in Section 4.3. The second is motivated by a spline representation of the variation of generalised Pareto shape and scale with respect to covariate. Section 6 then addresses estimation of extreme quantiles using the estimated parameters.

In this and subsequent sections, we consider the non-stationary extreme value threshold $u(\theta)$ to be determined and fixed; our focus is therefore on characterising threshold exceedances. Regardless of the approach to extreme value estimation, the heuristic procedure of Section 4 is used to estimate $u(\theta)$.

Descriptions of the semi-parametric, local parametric and spline parametric estimation schemes are given in Sections 5.1-5.3 below.

5.1 Local semi-parametric moment estimation

Given threshold $u(\theta_j)$ for $\theta_j \in \Theta$, we adopt the Moment estimator for the extreme value index given by

$$\hat{\xi}_k^M(\theta_j) := M^{(1)}(\theta_j) + 1 - \frac{1}{2} \left(1 - \frac{(M^{(1)}(\theta_j))^2}{M^{(2)}(\theta_j)} \right)^{-1} \quad (5.1)$$

where moments $M^{(l)}(\theta_j), l = 1, 2$ are calculated using

$$M^{(l)}(\theta_j) = \frac{1}{k(\theta_j)} \sum_{\mathcal{C}(\theta_j, h)} (\log X_i(\theta_{j'}) - \log u(\theta_j))^l \mathbb{1}_{\{X_i(\theta_{j'}) > u(\theta_j)\}} \quad (5.2)$$

where $k(\theta_j)$ is the number of threshold exceedances in $\mathcal{N}(\theta_j, h)$. Equation 5.2 corresponds to an unweighted version of the Moment estimator used for threshold estimation in Equation 4.5.

This Moment estimator is motivated by the asymptotic behaviour of the $k(\theta_j)^{\text{th}}$ upper order statistics associated with the sample of independent random variables $X_i(\theta_{j'}) - X_{N(\theta_j) - k(\theta_j), N(\theta_j)}$ on $\mathcal{N}(\theta_j, h)$ as established in the theory of extremes for threshold excesses. Conditionally on $X_{N(\theta_j) - k(\theta_j), N(\theta_j)} = u(\theta_j)$, their underlying distribution function tends to $F^{[u(\theta_j)]}(t) = P(X(\theta_j) - u(\theta_j) > t | X(\theta_j) > u(\theta_j))$, for all $t > 0$ [cf. de Haan and Ferreira, 2006, page 90]. Smith (1987) defined the analogous maximum likelihood estimator in terms of excesses over a high threshold u . We refer the reader to Appendix A for a precise definition of $F^{[u]}$, and further details on how it approaches the GPD distribution function.

An associated scale estimator (e.g. de Haan and Ferreira 2006, Section 4.2) necessary for semi-parametric estimation of extreme quantiles in Section 6 is then given by

$$\hat{a}_{\theta_j}^M \left(\frac{N}{k(\theta_j)} \right) := u(\theta_j) M^{(1)}(\theta_j) \frac{1}{2} \left(1 - \frac{(M^{(1)}(\theta_j))^2}{M^{(2)}(\theta_j)} \right)^{-1}, \quad (5.3)$$

The Moment estimator for the scale function $a_{\theta_j}^*(N/k(\theta_j))$ is discussed in Appendix A(ii).

5.2 Local maximum likelihood estimation

Given threshold $u(\theta_j)$ for $\theta_j \in \Theta$, local maximum likelihood estimation described of generalised Pareto shape $\xi(\theta_j)$ and scale $\sigma_u(\theta_j)$ is performed by the procedure described in Section 4.3, with all directional weights ω set to unity. For completeness, the log-likelihood to be maximised is given by

$$\mathcal{L}(\xi(\theta_j), \sigma_u(\theta_j)) := \sum_{\mathcal{C}(\theta_j, h)} \ell(\xi(\theta_j), \sigma_u(\theta_j) | X_i(\theta_{j'}) - u(\theta_j)) \mathbf{1}_{\{X_i(\theta_{j'}) - u(\theta_j) > 0\}}. \quad (5.4)$$

This formulation incorporates a semi-parametric approach to threshold selection within a local parametric approach to parameter estimation: conditioned on the random number $K = k(\theta_j)$ of exceedances, random exceedances of threshold $u(\theta_j)$ can again be viewed as independent identically distributed random variables with distribution function converging asymptotically to the generalised Pareto (Appendix A).

5.3 Spline-based maximum likelihood estimation

Given non-stationary threshold $u(\theta_j)$ for $\theta_j \in \Theta$, we can also estimate a non-stationary generalised Pareto model using the full sample of threshold exceedances, by allowing generalised Pareto shape ξ and scale σ_u to be functions of θ on \mathcal{S} . This is the type of analysis typically undertaken in environmental and ocean engineering settings (see e.g. Northrop et al. 2016). Specifically, we assume a periodic cubic B-spline representation for the variation of generalised Pareto shape and scale parameters with covariate (see e.g. Chapter 5 of Wood 2017, Zanini et al. 2020). We then estimate spline coefficients using maximum penalised likelihood estimation, regulating the roughness of shape and scale with covariate to optimise predictive performance assessed using cross-validation.

On the index set Θ of covariate values θ_j , $j = 1, 2, \dots, m$, we relate the values $\xi(\theta_j), \sigma_u(\theta_j)$ of generalised Pareto shape and scale to the periodic B-spline basis via basis matrix \mathbf{B} with elements B_{jb} such that

$$\xi(\theta_j) = \sum_{b=1}^{n_b} B_{jb} \beta_b^{(1)} \quad \text{and} \quad \log \sigma_u(\theta_j) = \sum_{b=1}^{n_b} B_{jb} \beta_b^{(2)}, \quad \theta_j \in \Theta, \quad (5.5)$$

where n_b is the number of basis functions, and the β s are basis coefficients. The sample log-likelihood \mathcal{L} is

$$\mathcal{L}(\boldsymbol{\beta}) = \sum_{j=1}^m \sum_{i=1}^{n(\theta_j)} \ell(\xi(\theta_j), \sigma_u(\theta_j) | X_i(\theta_j) - u(\theta_j)) \mathbf{1}_{\{X_i(\theta_j) > u(\theta_j)\}}, \quad (5.6)$$

with $\ell(\xi(\theta_j), \sigma_u(\theta_j) | X_i(\theta_j) - u(\theta_j))$ defined in (4.7). $n(\theta_j)$ denotes the number of observations in the sample at covariate θ_j , $\boldsymbol{\beta}^{(a)} = (\beta_1^{(a)}, \beta_2^{(a)}, \dots, \beta_{n_b}^{(a)})^\top$ for $a = 1, 2$, and $\boldsymbol{\beta} = (\boldsymbol{\beta}^{(1)\top}, \boldsymbol{\beta}^{(2)\top})^\top$. We set the number of spline

knots on \mathcal{S} to be more than sufficient to capture the anticipated parameter variability with covariate, and then penalise parameter roughness globally to obtain a model with good predictive performance. Penalisation is performed using first-order difference penalties for the coefficients in (5.5),

$$P^{(a)} = \sum_{b=1}^{n_b-1} \left(\beta_{b+1}^{(a)} - \beta_b^{(a)} \right)^2 + \left(\beta_1^{(a)} - \beta_n^{(a)} \right)^2 = \boldsymbol{\beta}^{(a)\top} \mathbf{D}^\top \mathbf{D} \boldsymbol{\beta}^{(a)}, \quad a = 1, 2,$$

with difference matrix \mathbf{D} given by:

$$\mathbf{D} = \begin{bmatrix} -1 & 1 & 0 & 0 & \cdot & \cdot \\ 0 & -1 & 1 & 0 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ 0 & 0 & \cdot & 0 & -1 & 1 \\ 1 & 0 & \cdot & 0 & 0 & -1 \end{bmatrix}.$$

The penalised log-likelihood is then:

$$\mathcal{L}_{pen}(\boldsymbol{\beta}) = \mathcal{L}(\boldsymbol{\beta}) - \lambda P^{(1)} - \kappa P^{(2)} \quad (5.7)$$

where λ and κ are smoothing parameters chosen to maximise cross-validated predictive likelihood. In the application illustrated in Section 7, cross-validation is applied as follows. Each iteration of the cross-validation consists in using a bootstrap resample (of the original sample of threshold exceedances) as training set, and observations omitted from the bootstrap resample as test set. Note that the test sets corresponding to different bootstrap resamples may therefore overlap. The training set is used to estimate the model parameters for a grid of combinations of λ and κ , and the test set to assess prediction performance using predictive log likelihood. The choices of λ and κ provide best predicted performance are adopted for subsequent inferences. The procedure is outlined in Algorithm 2 in Appendix B.

6 Estimation of non-stationary extreme quantile

In this section we use extreme value estimates from Section 5 to calculate extreme quantiles, including the right endpoint of the distribution of $X(\theta)$ when this is finite, in the semi-parametric setting (Section 6.1) and parametric setting (Section 6.2). The stepwise inference procedure is given in Algorithm 3 of Appendix B.

A common motivation applies for estimation of extreme quantiles in both parametric and semi-parametric settings. Suppose that F_θ is the actual distribution function of the data generating process at direction $\theta \in \mathcal{S}$ and Q_θ the corresponding quantile function; that is, $Q_\theta = F_\theta^{\leftarrow}$, with the left arrow indicating a generalised inverse. Conditions (A.2) and (A.3) of Appendix A ensure that an extreme quantile $F_\theta^{\leftarrow}(1-p)$, with $1-p > F_\theta(u(\theta))$ for some high threshold $u(\theta)$ and hence very small probability p , depends only on the tail of the distribution function F_θ . Consequently, we can estimate an extreme quantile using linear functional $\hat{b}_\theta + \hat{a}_\theta Q_\theta^H(1-p)$, with Q_θ^H relating to a generalised Pareto distribution function H_θ . This result applies for both parametric and semi-parametric approaches, with small differences. The most obvious difference is that in the parametric approach, normalising constants a_θ and b_θ are provided by the scale and location of the limiting generalised Pareto distribution, whereas in the semi-parametric approach these are estimated as functionals of the sample analogue of distribution function F_θ .

6.1 Semi-parametric-based estimation

6.1.1 Extreme quantile

In the semi-parametric setting, we assume that the number $k(\theta)$ of exceedances of threshold $u(\theta)$ is such that $k(\theta) \rightarrow \infty$ and $k(\theta)/N(\theta) \rightarrow 0$ as $N(\theta) \rightarrow \infty$. That is, the number of threshold exceedances $k(\theta)$

remains negligible compared to the total number $N(\theta)$ of observations in $\mathcal{N}(\theta, h)$. This condition can be written in terms of the covariate-specific sample size n , requiring that $n \rightarrow \infty$ implies $N(\theta) \rightarrow \infty$. The proposed Moment estimator for the extreme quantile with small probability $p \rightarrow 0$, $N(\theta)p/k(\theta) \rightarrow 0$ as $n \rightarrow \infty$, conditioned on threshold $u(\theta)$, is

$$\hat{x}_p^M(\theta) = u(\theta) + \hat{a}_\theta^M \left(\frac{N}{k(\theta)} \right) \frac{\left(\frac{k(\theta)}{N(\theta)p} \right)^{\hat{\xi}_k^M(\theta)} - 1}{\hat{\xi}_k^M(\theta)}, \quad (6.1)$$

with estimated extreme value index $\hat{\xi}_k^M(\theta)$ and associated scale parameter $\hat{a}_\theta^M \left(\frac{N}{k(\theta)} \right)$ given in Section 5.1. The associate scale in (6.1) is a function of the top sample fraction $k(\theta)/N(\theta)$.

6.1.2 Finite right endpoint

Semi-parametric estimators for the right endpoint have received considerable attention in the literature [e.g. Chapter 4 of de Haan and Ferreira, 2006, and connected references]. We propose a model-free, data-driven estimator of the covariate-dependent right endpoint $x^{F_\theta} < \infty$ for $\theta \in \mathcal{S}$, motivated by the general endpoint estimator of Fraga Alves and Neves [2014] with non-stationary threshold $u(\theta)$. We proceed by formulating an extreme value condition which partitions the Gumbel max-domain of attraction (with extreme value index $\xi = 0$) into two classes of distributions, one with finite right endpoint, and the other right unbounded. (An example of a distribution function in the former class is the Negative Fréchet, with distribution function $F_{\alpha, \beta}(x) = 1 - \exp\{-(\alpha - x)^{-\beta}\}$, $x \leq \alpha$, $\alpha \in \mathbb{R}$, $\beta > 0$; simple calculations show that $F_{\alpha, \beta}$ has $\xi = 0$ despite having finite right endpoint $x^F = \alpha < \infty$). We seek a right endpoint estimator for distributions in this Gumbel subclass.

With $\{Y_i\}_{i=1}^{N(\theta)}$ representing the random variables X_i in a particular neighbourhood $\mathcal{N}(\theta, h)$, defined in (4.1), we denote by $Y_{1, N(\theta)} \leq \dots \leq Y_{N(\theta)-k(\theta), N(\theta)} \leq \dots \leq Y_{N(\theta), N(\theta)}$ the corresponding ascending order statistics, and define the general endpoint estimator of x^{F_θ} , assumed finite, for every $\theta \in \mathcal{S}$,

$$\hat{x}_0^G(\theta) := Y_{N(\theta), N(\theta)} + \frac{1}{\log 2} \sum_{i=0}^{k(\theta)-1} \log \left(\frac{k(\theta) + i + 1}{k(\theta) + i} \right) (Y_{N(\theta)-k(\theta), N(\theta)} - Y_{N(\theta)-k(\theta)-i, N(\theta)}). \quad (6.2)$$

This estimator is valid for any $\xi(\theta) \leq 0$, and does not require prior estimation of extreme value index [cf. Fraga Alves et al., 2017], in contrast to competitor estimators for the right endpoint proposed in the literature.

6.2 Parametric-based estimation

6.2.1 Extreme quantile

In the parametric setting, level $x_p(\theta)$ with small exceedance probability p corresponds to the quantile of the distribution of threshold exceedances for direction θ . This suggests defining the $1/p$ extreme quantile as the value $x_p(\theta) = Q_\theta^H(1 - p)$, with normalising constants set to the scale and location parameters of a generalised Pareto function H_θ . Hence, for local parametric and spline parametric approaches, estimates of $x_p(\theta)$ of the form

$$\hat{x}_p^{\text{ML}}(\theta) = u(\theta) + \hat{\sigma}_u(\theta) \frac{\left(\frac{\varphi(u, \theta)}{p} \right)^{\hat{\xi}(\theta)} - 1}{\hat{\xi}(\theta)}, \quad (6.3)$$

are obtained, where $\varphi(u, \theta) = k(\theta)/N(\theta)$ is the sample fraction of exceedances of $u(\theta)$ within $\mathcal{N}(\theta, h)$. Estimation of $\hat{\xi}(\theta)$ and $\hat{\sigma}_u(\theta)$ (for $\theta \in \Theta$) using local parametric and spline parametric approaches is discussed in Sections 5.2 and 5.3. The extreme level estimate $\hat{x}_p^{\text{ML}}(\theta)$ can be interpreted as a T -year level used in ocean engineering and hydrology: that is, the value exceeded with probability $1/T$ per annum. Expressions

(6.3) and (6.1) show obvious similarities, and also distinctive traits of the semi-parametric Moment, local parametric and spline parametric approaches.

For local maximum likelihood estimation (Section 5.2), the values of $k(\theta_j)$ and $N(\theta_j)$ are estimated directly in Section 4. In the spline maximum likelihood approach (Section 5.3), a smoothed version of $\varphi(u, \theta_j)$ is estimated as the probability of threshold exceedance for $\theta_j \in \Theta$ using logistic regression, with log-likelihood

$$\mathcal{L}_{\text{Lgs}}(\boldsymbol{\beta}) = \sum_{j=1}^m \tau_j \log[\nu_j] + (1 - \tau_j) \log[1 - \nu_j], \quad (6.4)$$

where $\tau_j = k(\theta_j)/N(\theta_j)$ is the sample proportion of threshold exceedances of $u(\theta_j)$ at $\theta_j \in \Theta$. $\nu_j = (1 + \exp[-\eta_j])^{-1}$ is the probability of threshold exceedance at θ_j , with $\boldsymbol{\eta} = \{\eta_j\}_{j=1}^m = \mathbf{B}\boldsymbol{\beta}$ for B-spline basis matrix \mathbf{B} and parameter vector $\boldsymbol{\beta}$ to be estimated. Roughness penalisation of $\boldsymbol{\eta}$, with optimal roughness coefficient μ estimated by cross-validation, ensures good predictive performance. The penalised log-likelihood thus takes the form $\mathcal{L}_{\text{pen}}(\boldsymbol{\beta}) = \mathcal{L}(\boldsymbol{\beta}) - \mu\mathbf{P}$, with $\mathbf{P} = \boldsymbol{\beta}^\top \mathbf{D}^\top \mathbf{D}\boldsymbol{\beta}$ (c.f. Section 5.3).

6.2.2 Finite right endpoint

In case $\xi(\theta) < 0$, the limiting generalised Pareto distribution function has finite right endpoint. A consistent estimator for this right endpoint follows from (6.3) by setting $p = 0$

$$\hat{x}_0^{\text{ML}}(\theta) = u(\theta) - \frac{\hat{\sigma}_u(\theta)}{\hat{\xi}(\theta)}. \quad (6.5)$$

7 Application to storm peak significant wave height

The discussion above provides a suite of complementary approaches to non-stationary extreme value analysis incorporating elements of both parametric and semi-parametric inference, featuring a common non-stationary extreme value threshold capturing covariate dependence of threshold exceedances. In this section, we apply the methodology to estimation of T -year values from the sample of storm peak significant wave height H_S^{sp} on storm direction θ introduced in Section 2. The mechanics of inference for extreme quantiles, and right endpoint if appropriate, is given in Algorithm 3 of Appendix B. For brevity henceforth, we refer to semi-parametric, local parametric and spline parametric estimators as ‘‘M’’, ‘‘local ML’’ and ‘‘spline ML’’, with ‘‘M’’ indicating a Moment estimator, and ‘‘ML’’ referring to maximum likelihood.

Exploratory analysis of the sample, supported by previous analysis by Randell et al. [2016], suggests that the covariate domain can be partitioned into five directional sectors assumed approximately homogeneous in terms of the characteristics of H_S^{sp} . Referring to Figure 1, directional sectors corresponding to the following intervals of θ were identified. Sector 1 corresponds to $\theta \in (0^\circ, 50^\circ] \cup (320^\circ, 360^\circ]$, for storms propagating from the Norwegian Sea to the North; Sector 2 for $\theta \in (50^\circ, 140^\circ]$ corresponds to the ‘‘land shadow’’ of Norway, with fetch-limited storms propagating from the coast with a more northerly direction relative to the normal to the coast; Sector 3 is $\theta \in (140^\circ, 210^\circ]$, again for the Norwegian land-shadow, but with storms propagating from a more southerly direction; Sector 4 is $\theta \in (210^\circ, 270^\circ]$ corresponding to storms from the Atlantic potentially ‘‘funnelled’’ by the Norwegian coast; and Sector 5 with $\theta \in (270^\circ, 320^\circ]$, for more northerly Atlantic storms. Further information about the underlying physics is given in Section 2. The partitioned sample is summarised in Figure 2, using so-called ‘‘violin’’ plots which add kernel density estimates to a box-whisker representation. The long-tailed behaviour of storms from the Atlantic is clear in Sectors 4 and 5, compared to the fetch-limited characteristics in storms from Sectors 2 and 3. Although Sector 4 exhibits the largest values of threshold exceedances in Figure 2, there is evidence from the kernel density plots that Sector 5 has a relatively long tail. In this section we seek to quantify tail-heaviness by estimating $\xi(\theta)$ using both parametric and semi-parametric approaches, and hence estimate extreme quantiles.

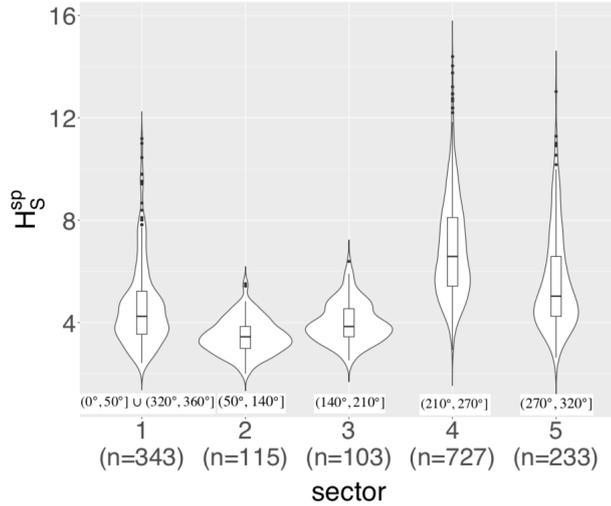


Figure 2: Violin plots of H_S^{sp} observations to the total of $N = 1521$ data points, with n here referring to the number of observations in each of the 5 sectors considered.

Using the approach in Section 4, non-stationary thresholds $u(\theta)$ were estimated using both the local ML and M estimator for lag $h = 60^\circ$, concentration $\eta = 10$ and parameter $\phi = 0.35$; a sensitivity study was conducted to ensure that threshold inference was not overly sensitive to these choices. Estimates are shown together in Figure 3. The general trends shown by the two estimates are in agreement across the covariate domain; for example, larger values of $u(\theta)$ are observed for storms from the Atlantic (SW to W), and smaller values for fetch-limited directions (E to SE). Subsequent parametric and semi-parametric inference for exceedance characteristics therefore has a relatively common starting point.

Estimates $\hat{\xi}(\theta_j)$ for $\theta_j \in \Theta$ from each of the M, local ML and spline ML approaches described in Section 5 are displayed in Figure 4, in terms of bootstrap means and 95% confidence intervals. Note that the M estimate of ξ was obtained using the corresponding M estimate of u ; analogously the local ML and spline ML estimates of ξ adopt the local ML estimate of u . There is reasonable overall agreement between the three estimates for ξ , and the estimates are also qualitatively plausible given other analyses of these data (Randell et al. 2016) and general physical considerations. Effects of land shadows (e.g. $\theta \in (80, 150)$) resulting in lower values for ξ are clear. The spline ML estimator is smoothest with respect to θ and its confidence band does not include zero for any direction. In contrast, $\xi(\theta) = 0$ is suggested by semi-parametric and local parametric estimators, most obviously for the M estimator, implying that the data generating distribution may lie in the Gumbel max-domain of attraction.

Identifying domains of attraction

Given the evidence in Figure 4, it is interesting to consider whether we can identify intervals with $\xi(\theta) < 0$, $\xi(\theta) = 0$ (and $\xi(\theta) > 0$ in principle) for $\theta \in \mathcal{S}$ (c.f. choice of domain of attraction, see Fraga Alves et al. 2017). We attempt this using two sets of hypothesis tests. In the first test, we adopt null hypothesis $F_\theta \in \mathcal{D}(G_0)$ versus the two-sided alternative $F \in \mathcal{D}(G_\xi)_{\xi \neq 0}$. Four different test statistics are considered, three of which

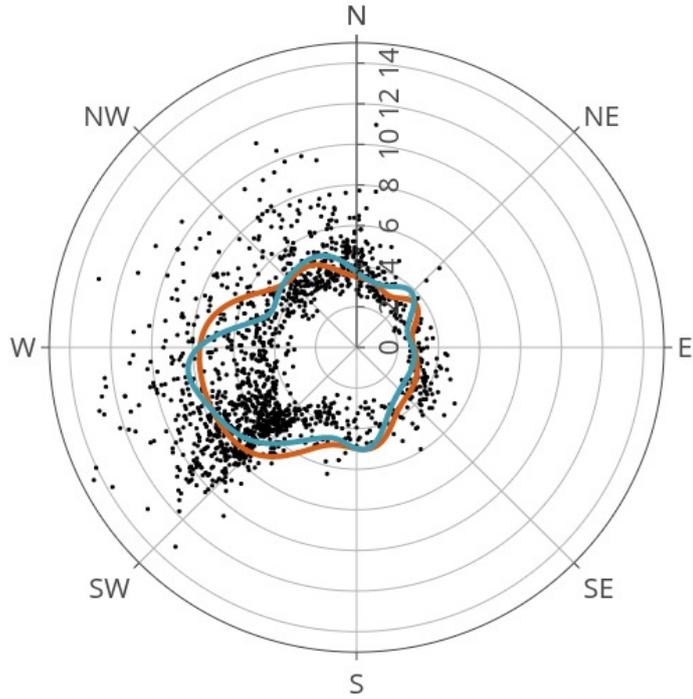


Figure 3: Polar plot of extreme value thresholds $u(\theta)$ from the M (orange) and ML (blue) estimators. Storm direction θ is the direction from which waves propagate; radial scale of H_S^{sp} is in metres.

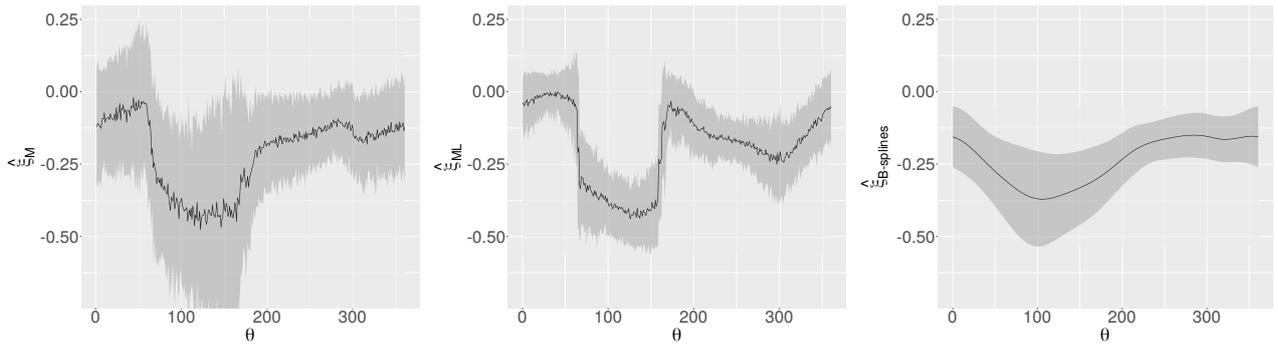


Figure 4: Bootstrap mean and corresponding 95% confidence intervals for $\hat{\xi}(\theta_j)$, $\theta_j \in \Theta$ based on: M (*left*), local ML (*middle*) and spline ML (*right*) estimators.

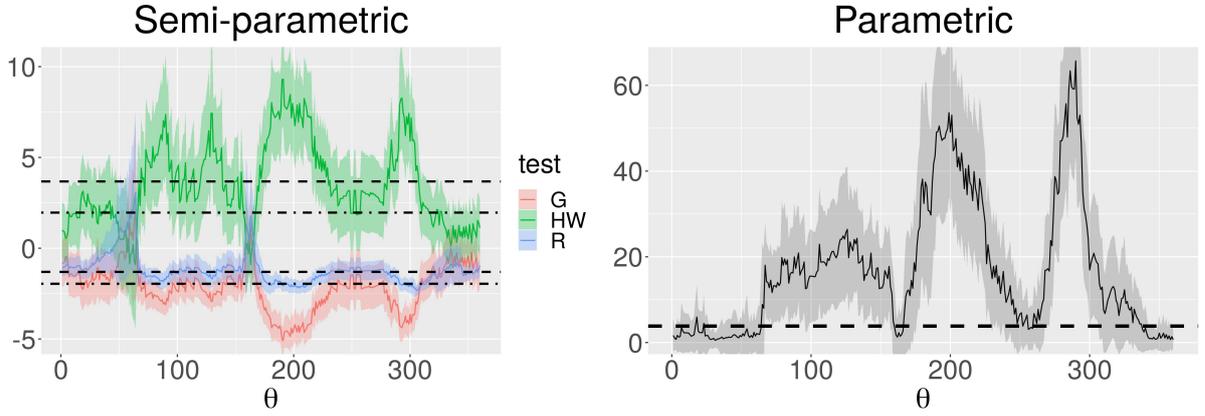


Figure 5: Bootstrap mean and corresponding 95% confidence intervals for the ratio (R , R), Greenwood-type (G , G) and Hasofer-Wang (HW , W) (left) and likelihood ratio (LR , D) (right) test statistics for $\xi = 0$. Horizontal lines in both panels indicate asymptotic critical values at significance level $\alpha = 5\%$. In the left panel: dashed lines for R test (-1.305 and 3.676); dot-dashed lines for G and W (± 1.96). In the right panel: dashed line at 3.84 .

are semi-parametric

$$\begin{aligned}
 R_{k,n}(\theta) &:= \frac{\max_{\mathcal{C}(\theta,h)} (X_i(\theta_j) - u(\theta_j))}{E^{(1)}(\theta)} - \log k(\theta) \\
 G_{k,n}(\theta) &:= \sqrt{\frac{k(\theta)}{4}} \left(\frac{E^{(2)}(\theta)}{(E^{(1)}(\theta))^2} - 2 \right) \\
 W_{k,n}(\theta) &:= \sqrt{\frac{k(\theta)}{4}} \left(\frac{1}{1 + [E^{(2)}(\theta)/(E^{(1)}(\theta))^2 - 2]} - 1 \right)
 \end{aligned}$$

based on moment spacings

$$E^{(l)}(\theta) = \frac{1}{k(\theta)} \sum_{\mathcal{C}(\theta,h)} (X_i(\theta_j) - u(\theta_j))^l \mathbf{1}_{\{X_i(\theta_j) > u(\theta_j)\}} \quad \text{for } l = 1, 2$$

in the notation of Section 5; these test statistics are discussed in Fraga Alves et al. [2017] and references therein. The ratio test (R) assesses the relative contribution of the maximum to the average of the excesses over the given threshold u . For the R test, the null hypothesis $F_\theta \in \mathcal{D}(G_0)$ is rejected at level $\alpha \in (0, 1)$ if $R_{k,n}(\theta) < g_{\alpha/2}$ or $R_{k,n}(\theta) > g_{1-\alpha/2}$, where g_ε is the ε -quantile of the Gumbel distribution, $g_\varepsilon = -\log(-\log \varepsilon)$. The Greenwood-type test (G) and the Hasofer and Wang test (HW) are moment-based, and associated critical regions follow from the central limit theorem. The critical region for the two-sided test of nominal size α for $T^* \in \{G, W\}$ is given by $|T_{k,n}^*(\theta)| > z_{1-\alpha/2}$, with z_ε denoting the ε -quantile of the standard normal distribution.

When $\xi(\theta) > -1/2$, a parametric test statistic based on the generalised Pareto distribution is available (c.f. Section 3.1 and Appendix A). Asymptotic normality of maximum likelihood estimators for peaks over threshold gives rise to a likelihood ratio (LR) test using deviance statistic

$$D(\theta) = 2 \left\{ \mathcal{L}(\hat{\xi}(\theta), \hat{\sigma}_u(\theta)) - \mathcal{L}(0, \hat{\sigma}_u(\theta)) \right\} \quad (7.1)$$

to test the hypothesis, where \mathcal{L} is the log-likelihood in (5.4). At level α , the null hypothesis is rejected when the observed value for D exceeds the $(1 - \alpha)$ -quantile of the χ_1^2 distribution [see e.g. Coles, 2001].

Results using the four test statistics are illustrated in Figure 5. The values for test statistics G and HW lie generally on or beyond their critical values as a function of θ , providing evidence to reject the null hypothesis that $\xi(\theta) = 0$ at the significance level $\alpha = 5\%$. We note that upward crossings of the indicated $\alpha = 5\%$ critical barriers in the HW test sample path evidence an underlying short tailed distribution belonging the Weibull max-domain, whereas for the G-test, the Weibull max-domain is reflected in the downward crossings of corresponding $\alpha = 5\%$ critical barriers. The value of the R test statistic lies generally on or near its lower critical value, again providing some support for the Weibull max-domain of attraction; we note that this test statistic is recognised as being somewhat more conservative than G and HW, particularly for near-rejection in favour of the Weibull domain.

The semi-parametric approach provides a further possibility (Neves and Pereira 2010, Fraga Alves and Neves 2014) for hypothesis testing, specifically to test for a Gumbel max-domain of attraction with infinite right endpoint as null hypothesis, versus an alternative hypothesis of any max-domain of attraction with $\xi(\theta) \leq 0$ and $x^{F_\theta} < \infty$. No similar parametric test exists, since in the parametric setting a generalised Pareto distribution with $\xi(\theta) = 0$ corresponds to the exponential distribution which is right unbounded. In this sense, the semi-parametric approach provides a better framework for tackling the statistical assessment of a finite upper bound. For every $\theta \in \mathcal{S}$, we test $H_0 : (F_\theta \in \mathcal{D}(G_0) \text{ and } x^{F_\theta} = \infty)$ versus $H_1 : (F_\theta \in \mathcal{D}(G_\xi)_{\xi \leq 0} \text{ and } x^{F_\theta} < \infty)$ using the test statistic T given by

$$T_{k,n}(\theta) := \frac{1}{k(\theta)} \sum_{\mathcal{C}(\theta,h)} \frac{X_i(\theta_j) - u(\theta_j) - \hat{a}_\theta(N/k(\theta))}{\max_{\mathcal{C}(\theta,h)}(X_i(\theta_j) - u(\theta_j))} \mathbb{1}_{\{X_i(\theta_j) > u(\theta_j)\}}, \quad (7.2)$$

with M estimator \hat{a} provided in (5.3). At the significance level $\alpha \in (0, 1)$, the null hypothesis is rejected whenever $\sqrt{k} \log k |T_{k,n}(\theta)| > z_{1-\alpha/2}$, where $z_\alpha = \Phi^\leftarrow(\alpha)$ denotes the α -quantile of the standard normal distribution function. Results using the test statistic T are shown in Figure 6. The null hypothesis of an infinite upper bound is rejected at the $\alpha = 5\%$ level for the majority of the directional domain, except notably for storms emanating from approximately $(310^\circ, 360^\circ)$, around 50° and around 160° . From a physical perspective, these observations can be explained by occurrences of storms from the Norwegian Sea $((310^\circ, 360^\circ))$, funnelling from the north-east along the Norwegian coast (50°) , and rare occurrences of dominant wave directions from the Norwegian coast (160°) . We summarise the results of hypothesis tests concerning choice of max-domain of attraction in Figure 7, and note directional intervals of $\xi(\theta) = 0, \xi^{F_\theta} < \infty$ corresponding to storms from the Atlantic funnelled along the Norwegian coast from south-west to west, and along the Norwegian coast from north to north-east.

Estimating extreme quantiles and right endpoints

We use expressions (6.1) and (6.3) to estimate quantiles $x_p(\theta)$ using M, local ML and spline ML estimators. Assuming N_0 storm occurrences in observation period T_0 , the T -year directional level corresponds to the value $x_p(\theta)$ such that $P\{X(\theta) > x_p(\theta)\} = (T_0/N_0) \times T^{-1}$. Figure 8 gives a matrix rose-plots of estimated directional 100-year and 10,000-year levels with accompanying 95% bootstrap confidence bands. There is general qualitative agreement between the three estimates for 100-year level (top row) and 10,000-year level (bottom row), in terms of bootstrap mean. Uncertainties from the M estimation are somewhat larger, as might be expected recalling the evidence of ξ estimation in Figure 4. Not surprisingly, estimated extreme levels for directions with short fetches ($\theta \in [60, 140]^\circ$) are low, whereas those corresponding to long fetches from the Atlantic Ocean and Norwegian Sea ($\theta \in [310, 360] \cup [0, 60]^\circ$) are high. These findings are broadly in line with those from hypothesis testing in Figure 7. Estimated 100-year levels for H_S^{SP} are between 15m and 20m the most severe sectors, in terms of bootstrap mean and confidence bands. The same is true of bootstrap means at the 10,000-year level except for M-related estimates which exceed 20m. Randell et al. [2016] report 100-year values from a directional-seasonal parametric analysis of approximately 14m for the Atlantic sector.

Figure 7 indicates there is significant evidence in the H_S^{SP} sample for concluding that the underlying distribution is right bounded on much of the directional domain. This suggests that estimation of the right

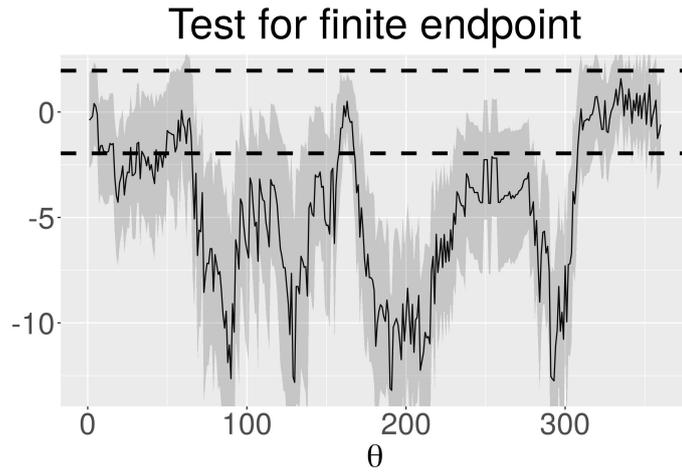


Figure 6: Bootstrap mean and corresponding 95% confidence intervals for the T-test statistic. The horizontal dashed lines correspond to the asymptotic critical values at the significance level $\alpha = 5\%$, i.e. the normal quantiles which are solutions of $|z_\varepsilon| = 1.96$.

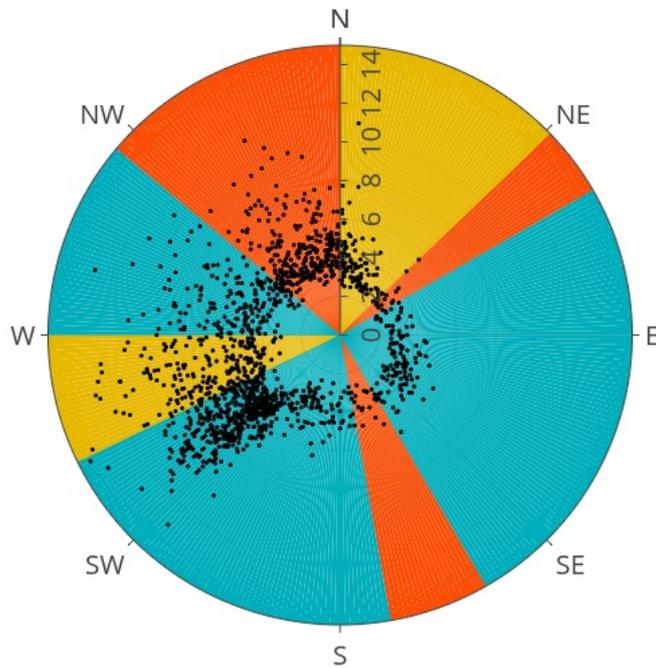


Figure 7: Polar plot showing combined deductions from tests of hypotheses for max-domain of attraction of H_S^{sp} at the $\alpha = 5\%$ significance level. Orange: $\xi(\theta) = 0, x^{F_\theta} = \infty$; yellow: $\xi(\theta) = 0, x^{F_\theta} < \infty$; blue: $\xi(\theta) < 0$. The original directional sample of H_S^{sp} on direction θ is also shown as black dots.

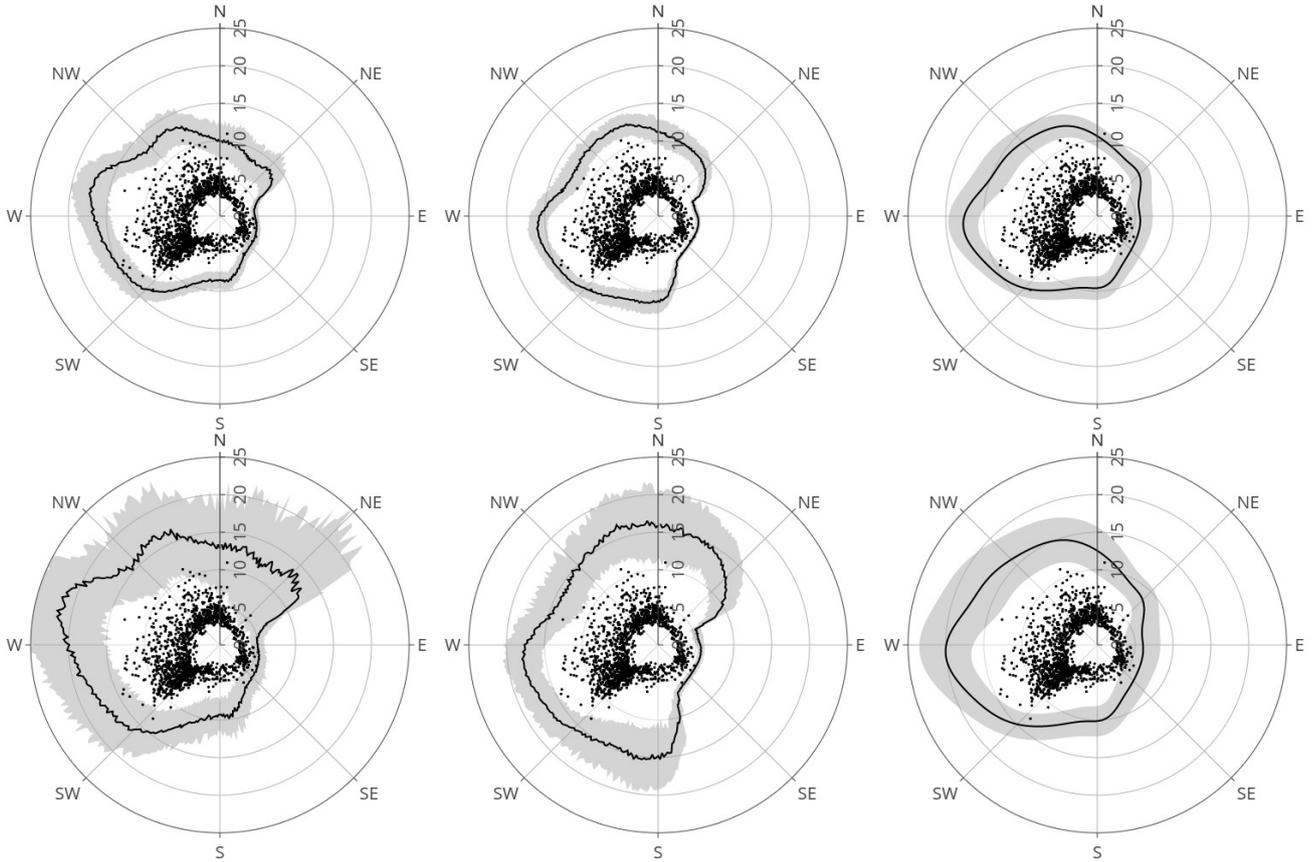


Figure 8: Rose diagrams for the 100-year (top) and 10,000-year (bottom) levels based on M (*left*), local ML (*middle*) and spline ML (*right*) estimators. Bootstrap means and corresponding 95% confidence bands are displayed as a function of direction measured clockwise from north.

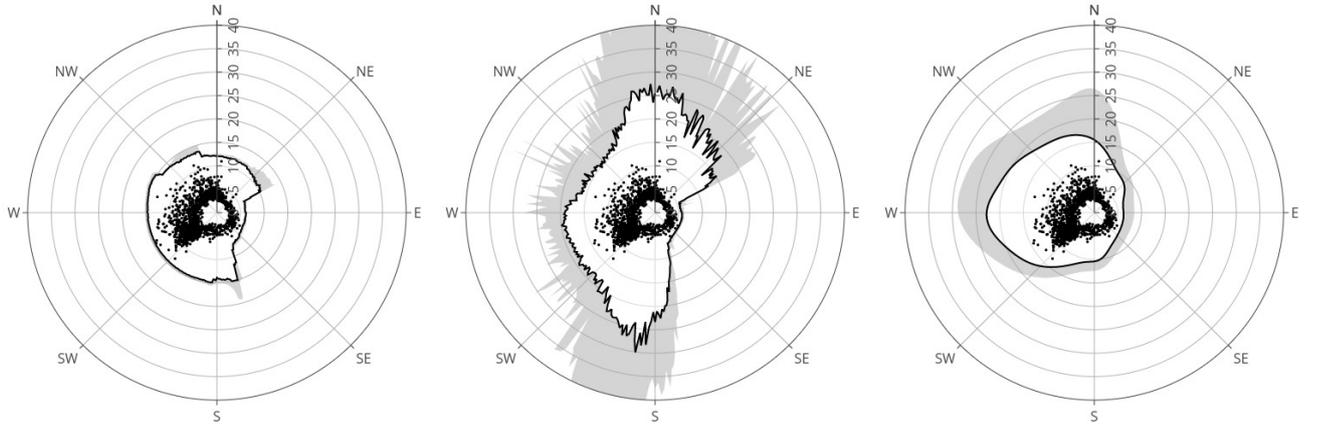


Figure 9: Rose diagrams for the finite right endpoint of H_S^{sp} . Three estimators are used: the general estimator (*left*), the local ML (*middle*) and spline ML (*right*) estimators. Bootstrap means and corresponding 95% confidence bands are displayed as a function of direction measured clockwise from north.

endpoint x^{F_θ} might be usefully attempted. Three estimators are considered, as described in Section 6. The local ML and spline ML estimators for x^{F_θ} are only appropriate when $\xi(\theta) < 0$, whereas the semi-parametric general estimator does not require estimation of $\xi(\theta)$ at all, but imposes the constraints $\xi(\theta) \leq 0$, $x^{F_\theta} < \infty$ by construction. To ensure fair comparison of the estimates for x^{F_θ} , we therefore repeated the local ML and spline ML generalised Pareto fits of Section 5, with the domain of ξ restricted to $\xi(\theta) < 0$. The left panel of Figure 9 shows estimates using general estimator (6.2). The centre and right panels show estimates from local ML and spline ML estimation, with $\xi(\theta)$ constrained to be negative everywhere. Clearly, maximum likelihood estimation becomes more challenging as the true shape parameter ξ approaches zero from below, with numerical optimisation routines more than often experiencing convergence issues [e.g. Gomes and Neves, 2008, for a comparison between M and ML within the univariate semi-parametric setting]. Estimates from the general estimator (left panel) are surprisingly low compared with those for local ML and spline ML, and with semi-parametric estimates of the 100- and 10,000-year values using the Moment estimator, shown in Figure 8. Agreement between the three estimators for the finite right endpoint cannot be expected to be nearly as good as that observed for T -year level, since the estimators are rather different in nature.

Local ML and spline ML estimates (Figure 9, centre and right panels) show some agreement particularly for westerly, and north to north-easterly directions, where Figure 7 suggests $\xi(\theta) = 0$ and $x^{F_\theta} < \infty$; the local ML estimator provides considerably more uncertain estimates, because of input estimates for $\xi(\theta)$ near zero. The general endpoint estimator exhibits small variance in comparison, a distinctive trait that might be attributed to the simplicity in form not requiring external estimation of ξ ; the small variance of this semi-parametric estimator has been proved both theoretically and via extensive simulation studies (Fraga Alves and Neves 2014). Increased uncertainty in the general endpoint estimator (Figure 9) corresponds to regions in Figure 7 for which $\xi(\theta) = 0$ and $x^{F_\theta} = \infty$. A study of the sensitivity of the general endpoint estimator to specification of h and $k(\theta)$ was undertaken, particularly due to the low estimates it yields. The study revealed that the general endpoint estimator is in fact remarkably stable with respect to both h and $k(\theta)$. We must of course keep in mind that all of the estimates in Figure 9 assume a finite right endpoint, whereas the evidence from Figure 7 is that this is not the case for all directions. Note further that the asymptotic distribution for the general endpoint estimator is non-normal, and therefore the estimated bootstrap 95% confidence bounds offer a limited view of its variance. A similar remark applies to the local-ML estimator for the right endpoint of the generalised Pareto distribution for the irregular case for $\xi \in (-1, -1/2]$.

In summary, application of the parametric and semi-parametric methodologies developed in Sections 4-6 above to the sample of directional storm peak significant wave height suggests that estimates for ξ , and 100-year and 10,000-year levels are in good qualitative agreement. Where differences occur, they can be explained in terms of specific modelling assumptions made, rather than in terms of fundamental differences

in the underlying approaches to extreme value analysis.

8 Discussion and conclusions

This paper presents a framework for inference on non-stationary peaks over threshold, reconciling approaches from semi-parametric and parametric extreme value analysis, in application to directional ocean storm severity. The key components of the framework are (a) estimation of non-stationary extreme value threshold, and (b) estimation of tail characteristics from threshold exceedances, including extreme quantiles and finite right endpoint when appropriate. Threshold estimation is performed using a non-stationary extension of a heuristic approach for semi-parametric Moment (M) and parametric maximum likelihood (ML) estimators. Tail characteristics and extreme quantiles are then estimated, based on semi-parametric M, local parametric ML and spline ML estimators. Given threshold, the two streams of extreme value inference mirror each other. In the spline ML approach, a penalised B-splines representation with compact support is used: each basis function is non-zero on a specific interval of the covariate domain. This feature plays a similar role to bandwidth of directional neighbourhoods in the semi-parametric M and parametric local ML approaches.

We also consider estimators for the non-stationary finite right endpoint, including extension of the semi-parametric general estimator of Fraga Alves and Neves [2014]. Hypothesis tests, used to characterise the max-domain of attraction of storm peak significant wave height as a function of direction, suggest that the distributions of storm peak significant wave height for storms from the Atlantic and Norwegian Sea are unbounded above.

Inferences regarding directional thresholds for storm peak significant wave height are in good agreement over the covariate domain. Estimates for 100 and 10,000-year levels are also in reasonable agreement. Estimates for the right endpoint are more different across approaches, and are influenced by the specifics of modelling assumptions made associated with the different estimation strategies. For the application considered, both parametric and semi-parametric inference provides similar characterisations of extreme non-stationary ocean environments. Indeed, we illustrate how ideas from the semi-parametric and parametric schools of thought can be used in tandem to exploit the desirable features of the approaches, whilst overcoming some obvious pitfalls. For example, threshold estimation (used for both semi-parametric and parametric analysis) is motivated by an inherently non-parametric heuristic in Section 4. Quantile regression is a popular approach to non-stationary extreme value threshold [e.g. Northrop and Jonathan, 2011]. Evaluation of threshold estimates from the current heuristic approach, in comparison of estimates from a quantile regression, will be an interesting future study.

Parametric approaches to non-stationary extremes are relatively well-studied due in part to the wide range of flexible covariate representations for generalised Pareto parameters for threshold exceedances, and associated methods for regression and assessment of model fit. In contrast, from a semi-parametric perspective, an exact generalised Pareto correspondence need not be assumed, offering greater flexibility. However, avoiding a particular model choice from the outset generally results in increased uncertainty of estimates for ξ and high quantiles. Nevertheless, we show in this work that semi-parametric and parametric approaches perform rather similarly when set up reasonably. We hope that the fusion of ideas from parametric and semi-parametric approaches to extreme value analysis outlined in this article provide a basis for increased understanding and quantification of extreme phenomena in environmental and engineering applications.

References

- A. A. Balkema and L. de Haan. Residual life time at great age. *Ann. Probab.*, 2(5):792–804, 1974.
- V. Chavez-Demoulin and A. C. Davison. Generalized additive modelling of sample extremes. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 54(1):207–222, 2005.
- S.G. Coles. *An Introduction to Statistical Modeling of Extreme Values*. Springer, 2001.

- S.G. Coles and D. Walshaw. Directional modelling of extreme wind speeds. *Applied Statistics*, 43:139–157, 1994.
- Anthony C Davison and Richard L Smith. Models for exceedances over high thresholds. *Journal of the Royal Statistical Society: Series B (Methodological)*, 52(3):393–425, 1990.
- A. L. M. Dekkers, J. H. J. Einmahl, and L. de Haan. A moment estimator for the index of an extreme-value distribution. *The Annals of Statistics*, 17(4):1833–1855, 1989.
- K. C. Ewans and P. Jonathan. The effect of directionality on northern North Sea extreme wave design criteria. *J. Offshore. Arct. Eng.*, 130:041604:1–041604:8, 2008.
- G. Feld, D. Randell, and P. Jonathan. On the estimation and application of directional design criteria. *Proc. 38th Int. Conf. on Ocean, Offshore & Arctic Engineering, Scotland (draft at www.lancs.ac.uk/~jonathan)*, 2019.
- R. A. Fisher and L. H. C. Tippett. Limiting forms of the frequency distribution of the largest and smallest member of a sample. *Cambridge Philosophical Society. Mathematical Proceedings*, 24:180–190, 1928.
- Pedro Folgueras, Sebastián Solari, and Miguel Ángel Losada. The selection of directional sectors for the analysis of extreme wind speed. *Natural Hazards and Earth System Sciences*, 19(1):221–236, 2019.
- George Z Forristall. On the use of directional wave criteria. *Journal of waterway, port, coastal, and ocean engineering*, 130(5):272–275, 2004.
- I. Fraga Alves, C. Neves, and P. Rosário. A general estimator for the right endpoint with an application to supercentenarian women’s records. *Extremes*, 20(1):199–237, 2017.
- Maria Isabel Fraga Alves and Cláudia Neves. Estimation of the finite right endpoint in the Gumbel domain. *Statistica Sinica*, 24(4):1811 – 1835, 2014.
- B. V. Gnedenko. Sur la distribution limite du terme maximum d’une série aléatoire. *Ann. Math.*, 44:423–453, 1943.
- M. Ivette Gomes and Cláudia Neves. Asymptotic comparison of the mixed moment and classical extreme value index estimators. *Statistics & Probability Letters*, 78(6):643 – 653, 2008. ISSN 0167-7152.
- E. J. Gumbel. *Statistics of Extremes*. Columbia University Press, New York, 1958.
- L. de Haan. *On regular variation and its application to the weak convergence of sample extremes*. Mathematisch Centrum Amsterdam, 1970.
- Laurens de Haan and Ana Ferreira. *Extreme Value Theory: An Introduction*. Springer, 2006.
- Laurens de Haan, Albert Klein Tank, and Cláudia Neves. On tail trend detection: modeling relative risk. *Extremes*, 18(2):141–178, 2015.
- P. Jonathan and K. C. Ewans. The effect of directionality on extreme wave design criteria. *Ocean Eng.*, 34: 1977–1994, 2007.
- Matthew Jones, David Randell, Kevin Ewans, and Philip Jonathan. Statistics of extreme ocean environments: Non-stationary inference for directionality and other covariate effects. *Ocean Engineering*, 119: 30–46, 2016.
- M. R. Leadbetter. On a basis for Peaks over Threshold modeling. *Statistics & Probability Letters*, 12(4): 357–362, 1991.

- E. B. L. Mackay, P. G. Challenor, and A. S. Bahaj. On the use of discrete seasonal and directional models for the estimation of extreme wave conditions. *Ocean Eng.*, 37:425–442, 2010.
- F J Mendez, M Menendez, A Luceno, R Medina, and N E Graham. Seasonality and duration in extreme value distributions of significant wave height. *Ocean Eng.*, 35:131–138, 2008.
- C. Neves and A. Pereira. Detecting finiteness in the right endpoint of light-tailed distributions. *Statistics and Probability Letters*, 80:437–444, 2010.
- Cláudia Neves and Maria Isabel Fraga Alves. Reiss and Thomas’ automatic selection of the number of extremes. *Computational Statistics & Data Analysis*, 47(4):689 – 704, 2004.
- P. J. Northrop and P. Jonathan. Threshold modelling of spatially dependent non-stationary extremes with application to hurricane-induced wave heights. *Environmetrics*, 22(7):799–809, 2011.
- P. J. Northrop, P. Jonathan, and D. Randell. Threshold modeling of nonstationary extremes. In Dipak K Dey and Jun Yan, editors, *Extreme Value Modeling and Risk Analysis: Methods and Applications*, chapter 5, pages 87–108. Chapman and Hall / CRC, 2016.
- B. Oumow, M. de Carvalho, and A. C. Davison. Bayesian P-spline mixture modeling of extreme forest temperatures. Retrieved from www.mat.puc.cl/~mdecarvalho, 2012.
- A. Pewsey, M. Neuhauser, and G.D. Ruxton. *Circular Statistics in R*. 2014.
- J. Pickands III. Statistical inference using extreme order statistics. *Ann. Statist.*, 3(1):119–131, 1975.
- D. Randell, K. Turnbull, K. Ewans, and P. Jonathan. Bayesian inference for non-stationary marginal extremes. *Environmetrics*, 27:439–450, 2016.
- Rolf-Dieter Reiss and Michael Thomas. *Statistical Analysis of Extreme Values*. Springer, 3rd edition, 2007.
- M. Reistad, O. Breivik, H. Haakenstad, O. J. Aarnes, B. R. Furevik, and J.-R. Bidlot. A high-resolution hindcast of wind and waves for the North sea, the Norwegian sea, and the Barents sea. *J. Geophys. Res.*, 116:1–18, 2011.
- M. E. Robinson and J. A. Tawn. Statistics for extreme sea currents. *Appl. Statist.*, 46:183–205, 1997.
- E Ross, S Sam, D Randell, G Feld, and P Jonathan. Estimating surge in extreme North Sea storms. *Ocean Engineering*, 154:430–444, 2018.
- Richard L. Smith. Estimating tails of probability distributions. *Ann. Statist.*, 15(3):1174–1207, 09 1987. doi: 10.1214/aos/1176350499. URL <https://doi.org/10.1214/aos/1176350499>.
- M D Ugarte, T Goicoa, J Etxeberria, and A F Militino. Projections of cancer mortality risks using spatio-temporal P-spline models. *Stat. Methods in Med. Res*, 21:545–560, 2012.
- Simon N Wood. *Generalized Additive Models: An Introduction With R*. CRC Press, 2017.
- E. Zanini, E. Eastoe, M. J. Jones, D. Randell, and P. Jonathan. Flexible covariate representations for extremes. *Environmetrics*, 31(5):e2624, 2020.

A . Basis for inference on directional extremes via the POT method

The contents of this appendix build on Appendix B of de Haan et al. [2015]. This paper offers added flexibility to the latter by allowing $\xi = \xi(s)$ to vary with $s \geq 0$, where s might represent direction θ , or time or some other covariate. As a consequence, the right endpoint does not need to be assumed constant in s . We will not delve into the theoretical details in terms of explicit smoothness and boundedness conditions needing to be in place particularly by assuming $h = h_n > 0$. These are clearly beyond the scope of this paper, but we envisage that the probabilistic underpinning to this work will stem from Chapter 9 of de Haan and Ferreira [2006].

For direction $s \geq 0$, let $(X_1(s), X_2(s), \dots, X_n(s))$ be a vector of i.i.d. random variables with common distribution function $F_s(x)$, for all $x \in \mathbb{R}$, absolutely continuous with right endpoint $x^{F_s} \leq \infty$. Assume $F_s \in \mathcal{D}(G_{\xi(s)})$, for some $\xi(s) \in \mathbb{R}$ and for every $s \geq 0$, i.e., that condition (3.1) holds locally for each s . In this setting, Theorem 1.1.6 of de Haan and Ferreira [2006] ascertains that it is possible to replace n with t running over the real line in such a way that (3.1) becomes equivalent to the following extreme value condition: there exists a positive function a_s^* such that

$$\lim_{t \uparrow x^{F_s}} \frac{1 - F_s(t + xa_s^*(t))}{1 - F_s(t)} = \frac{\log G_{\xi(s)}(x)}{\log G_{\xi(s)}(0)}, \quad (\text{A.1})$$

for all x with $1 + \xi(s)x > 0$. The limit in (A.1) is the tail distribution function (also known as survival function) of the generalised Pareto (GP) distribution with shape parameter $\xi(s) \in \mathbb{R}$, given by $(1 + \xi(s)x)^{-1/\xi(s)}$. The extreme value condition (A.1) is often key for describing rare events' behaviour in lieu of the dual max-domain of attraction characterisation $F \in \mathcal{D}(G_\xi)$.

(i) Parametric approach

Taking a parametric view, the peaks over threshold (POT)-domain of attraction condition (A.1) prescribes the GP distribution as the proper fit to the normalised exceedances given these are above a certain high threshold near the right endpoint x^{F_s} . With minimal notational changes around a fixed (deterministic) threshold u , it is straightforward to see that condition (A.1) implies

$$\lim_{u \uparrow x^{F_s}} \left| P\{X_1(s) \leq x + u \mid X_1(s) > u\} - H_{\xi(s), u, \sigma(u)}(x) \right| = 0, \quad (\text{A.2})$$

locally uniformly in $x > u$, for each $s \geq 0$, with $\sigma(u) > 0$ (we omit the subscript s on σ and μ below) for simplicity of notation) and $H_{\xi, \mu, \sigma}(x) := 1 - (1 + \xi(x - \mu)/\sigma)^{-1/\xi}$, for all x such that $0 < H_{\xi, \mu, \sigma}(x) < 1$, with location $\mu \in \mathbb{R}$ and scale $\sigma > 0$. Informally, $F_s^{[u]}(x) := P\{X_1(s) \leq x + u \mid X_1(s) > u\} \approx H_{\xi(s), u, \sigma(u)}(x)$, for all $x > 0$ and u near the right endpoint x^{F_s} , with the scale parameter implicitly defined in terms of s through the threshold $u(s)$.

For each s , the limiting relation (A.2) provides the probabilistic underpinning for fitting a GP distribution function to the unconditional tail distribution function $\overline{F}_s(x) := 1 - F_s(x)$ with x sufficiently large. This becomes more evident since

$$F_s(x) = (1 - F_s(u))F_s^{[u]}(x) + F_s(u),$$

from which

$$\overline{F}_s(x) = (1 - F_s^{[u]}(x))(1 - F_s(u)) \approx (1 - H_{\xi(s), u, \sigma(u)}(x))(1 - F_s(u)), \quad (\text{A.3})$$

for $x > u$, as $u \rightarrow x^{F_s}$. Finally, we note that

$$(1 - H_{\xi(s), u, \sigma(u)}(x))(1 - F_s(u)) = H_{\xi(s), \mu^*, \sigma^*(u)}(x), \quad (\text{A.4})$$

where $\mu^* - u = \sigma(u)U_H(1 - F_s(u))$, $\sigma^*(u) = \sigma(u)(1 - F_s(u))^{\xi(s)}$, and U_H standing for the tail quantile function pertaining to the standard GPD, that is

$$U_H(t) := \left(\frac{1}{1 - H_{\xi(s),0,1}} \right)^{\leftarrow} (t) = \frac{t^{\xi(s)} - 1}{\xi(s)},$$

for all $t \geq 1$ (the left arrow indicates the left-continuous inverse). The representation (A.4) facilitates the view that, in practice, changes in the threshold (e.g. through covariates) will be reflected in the scale parameter. In turn, the approach for inference is reflected in the way we go choose to go about the term $1 - F_s(u)$ for this to become statistically meaningful. In order to able to perform large sample inference drawing on the POT-GPD framework streamlined above, we now make the threshold dependent on the sample size n and $u = u(n)$ will naturally become larger as n goes to infinity. A parametric approach typically advocates for a large enough threshold to be fixed and inference to be conducted on the basis of the resulting POT framework, whereby the expected number of exceedances above the selected threshold is a random number K_s , say, satisfying $(n/K_s)(1 - F_s(u_n)) \rightarrow 1$ in probability, as $n \rightarrow \infty$. This suggests estimation of $(1 - F_s(u))$ via the analogous tail empirical distribution function (stepping up by $1/n$ at each observation) evaluated at u , adding up to $1 - K_s/n$ in the above, associated with the random number K_s of exceedances of u at direction (or location) s . Hence, for a given (fixed) k_s , the location and scale parameters in (A.4) become:

$$\begin{aligned} \mu^*(s) &= u + \sigma(u)U_H\left(1 - \frac{k_s}{n}\right), \\ \sigma^*(u) &= \sigma(u)\left(1 - \frac{k_s}{n}\right)^{\xi(s)}. \end{aligned}$$

Therefore, the crux of parametric inference for extremes values lays in the estimation of the shape and scale parameters, respectively, $\xi(s)$ and $\sigma_s = \sigma^*(u)$.

(ii) Semi-parametric approach

It will be notationally cleaner to express the argument (A.1) in terms of the pertaining tail quantile function $U_s := (1/(1 - F_s))^{\leftarrow}$. Note that $U_s(t)$ is non-decreasing and provides a straightforward link to an extreme quantile, with the right endpoint representing the ultimate quantile: $\lim_{t \rightarrow \infty} U_s(t) = U_s(\infty) = x^{F_s}$. To this effect, we make t in (A.1) depend on the (possibly unknown) sample size n at each location s through replacing it by $U_s(n/k_s)$, where k_s is an intermediate sequence of positive integers such that $k_s = k_s(n) \rightarrow \infty$ and $k_s/n \rightarrow 0$, as $n \rightarrow \infty$. This is possible because (A.1) holds uniformly in x . Hence, we have for the left hand-side of (A.1):

$$\frac{1 - F_s(U_s(n/k_s) + x a_s(U_s(n/k_s)))}{1 - F_s(U_s(n/k_s))} = \frac{n}{k_s} (1 - F_s(U_s(n/k_s) + x a_s^*(n/k_s))),$$

with $a_s^*(n/k_s) = a_s(U_s(n/k_s))$.

For simplicity, we consider regularly spaced independent vectors $(X_1(s), X_2(s), \dots, X_n(s))$, $s = 1, 2, \dots, m, \dots$ with i.i.d. components, with partial tally of $N = n \times m \in \mathbb{N}$ observations across the whole system, and where n is potentially unknown (without affecting inference on extremes), yet assumed large ($n \rightarrow \infty$). In this setting, the basic extreme value condition is (cf. (A.3)):

$$\lim_{n \rightarrow \infty} \frac{N}{k \omega(s)} \left\{ 1 - F_s \left(U_s \left(\frac{N}{\omega(s)k} \right) + x a_s^* \left(\frac{N}{\omega(s)k} \right) \right) \right\} = 1 - H_{\xi(s),0,1}(x), \quad (\text{A.5})$$

for all x with $1 + \xi(s)x > 0$, uniformly in $s = 1, 2, \dots$, with $\omega \geq 0$ a continuous and smooth function satisfying $(1/m) \sum_{s=1}^m \omega(s) \rightarrow 1$, as $m \rightarrow \infty$, i.e. the resulting sequence of weights $\{\omega(m)\}_{m \in \mathbb{N}}$ is Cesàro summable. The latter is to maintain integrity, also ensuring that the stationary case is well-defined. In particular,

the case of complete stationarity, corresponding to omni-directional data in the context of this paper, is recovered if $\omega(s)$ is identically one over the stipulated range for s . The interest lies in the estimation of the various extreme value indices $\xi(s)$, and the scale and location terms, respectively $a_s^*(N/k_s)$ and $U_s(N/k_s)$, now with $k_s := [\omega(s) \times k]$ and $[\cdot]$ standing for integer part. Since the N -th order statistic $X_{N-k_s:N}(s)$ is close to $U_s(N/k_s)$, if $k_s \rightarrow \infty$, $k_s/N \rightarrow 0$, as $N \rightarrow \infty$, we shall adopt it, as the usual estimator for the threshold $\widehat{U}_s(N/k_s) = X_{N-k_s:N}(s)$. The random adaptive threshold in this setting emulates the non-stationarity mirrored in the scale $\sigma_s > 0$ which features the parametric setting (i).

The role of the of the (deterministic) function ω can perhaps be better understood if the domain of attraction condition (A.5), characterising non-stationary POT, is formulated in terms of their respective inverse functions intervening in both hand-sides. To this end, we use $t \rightarrow \infty$ in place of $N/k \rightarrow \infty$, so that (A.5) translates into:

$$\lim_{t \rightarrow \infty} t \left\{ 1 - F_s \left(U_s \left(\frac{t}{\omega(s)} \right) + x a_s^* \left(\frac{t}{\omega(s)} \right) \right) \right\} = \omega(s) (1 - H_{\xi(s), 0, 1}(x)),$$

which, by inversion in terms of $t/y = 1 / \{ 1 - F_s(U_s(t/\omega(s)) + x a_s^*(t/\omega(s))) \}$, and similarly for $y/\omega(s) = (1 + \xi(s))^{-1/\xi(s)}$ on the right hand side, yields the desired condition:

$$\lim_{t \rightarrow \infty} \frac{U_s(t/y) - U_s(t/\omega(s))}{a_s^*(t/\omega(s))} = \frac{(\omega(s)/y)^{\xi(s)} - 1}{\xi(s)},$$

for all $y > 0$. In particular, putting $y = 1$, and by virtue of a_s^* being a regularly varying function at infinity with index $\xi(s)$, uniformly in $s \in \mathcal{S}$ [cf. de Haan and Ferreira, 2006], i.e. $\lim a_s^*(tx)/a_s^*(t) = x^{\xi(s)}$, $x > 0$, we obtain the following condition on the ‘‘moving’’ threshold:

$$\lim_{t \rightarrow \infty} \frac{U_s(t/\omega(s)) - U_s(t)}{a_s^*(t)} = \frac{(\omega(s))^{-\xi(s)} - 1}{\xi(s)}. \quad (\text{A.6})$$

In case $\xi(s) = 0$, for some s , then the limit in the above reads as $\log(\omega(s))$. Furthermore, if the right endpoint is finite, i.e. $x^{F_s} < \infty$, then $a^* > 0$ satisfies $a_s^*(t) \rightarrow 0$, as $t \rightarrow \infty$, and $F_s \in \mathcal{D}(G_{\xi(s)})$, with $\xi(s) \leq 0$, which entails that $U_s(t)$ is asymptotic to $U_s(t/\omega(s))$, for $\omega(s) \in (0, \infty)$. The smooth function $\omega(s)$ thus primarily reflects the amplitude in extremes rather than their magnitude or severity. Indeed, there is a clear connection between how extreme spacings amplify and the tail behaviour of the underlying distribution function.

B . Algorithms

This appendix provides algorithms for estimation of local non-stationary threshold and optimal spline representation. It also provides an algorithm for the complete extreme value analysis given non-stationary extreme value threshold.

B.1 Local non-stationary threshold selection

B.2 Penalised maximum likelihood B-spline estimation

B.3 Estimation of T -year level and finite right endpoint

The procedure for estimation of an extreme level and of the right endpoint, if appropriate, for any of the three methods is summarised in the following algorithm for clarity. The algorithm assumes that the non-stationary threshold $u(\theta)$, $\theta \in \mathcal{S}$ has been pre-determined through Algorithm 1 in connection with the adopted local estimators for ξ .

Algorithm 1 Non-stationary threshold estimation for local estimators

- 1: Specify lag h , concentration η and parameter ϕ ;
 - 2: **for** $\theta_j \in \Theta$ **do**
 - 3: Estimate weights $\omega(\theta_{j'})$ for $\theta_{j'} \in \Theta$;
 - 4: **for** $k = 1, \dots, N(\theta_j) - 1$ **do**
 - 5: Use the k largest values in $\mathcal{N}(\theta_j, h)$ to obtain $\hat{\xi}_k(\theta_j)$ using (4.6) or (4.4);
 - 6: Calculate $S_\phi(k)$;
 - 7: **end for**
 - 8: Set $k_j^* = \arg \min_k S_\phi(k)$;
 - 9: Identify threshold $u(\theta_j)$ with the $(k_j^* + 1)^{\text{th}}$ largest value with direction in $\mathcal{N}(\theta_j, h)$;
 - 10: **end for**
-

Algorithm 2 Maximum penalised likelihood for non-stationary periodic B-spline representation

- 1: Evaluate B-spline basis functions on index set Θ ;
 - 2: Specify sets of values of smoothing penalty coefficients λ and κ to consider;
 - 3: **for** each choice of λ **do**
 - 4: **for** each choice of κ **do**
 - 5: **for** each of a number of bootstrap resamples **do**
 - 6: Estimate model parameters using bootstrap resample;
 - 7: Use estimated model to predict test observations (not occurring in bootstrap resample);
 - 8: Calculate the predictive log-likelihood;
 - 9: **end for**
 - 10: Accumulate total predictive log-likelihood;
 - 11: **end for**
 - 12: **end for**
 - 13: Select pair of values of λ and κ with best predictive performance, and evaluate spline coefficients for these choices of roughness coefficients.
-

Algorithm 3 Extreme quantile (including T -year level) and finite right endpoint estimation using M, local ML and spline ML based estimators.

- 1: Specify data sample; non-stationary threshold $u(\theta)$; period of sample T_0 , period T ;
 - 2: Isolate set of directional threshold exceedances;
 - 3: **if** M **then**
 - 4: Specify window half-width h ;
 - 5: **else if** local ML **then**
 - 6: Specify window half-width h ;
 - 7: **else if** spline ML **then**
 - 8: Specify values of smoothing coefficients λ and κ to consider;
 - 9: Specify details for B-spline basis function construction;
 - 10: **end if**
 - 11: **for** each of a large number of bootstrap resamples **do**
 - 12: Generate bootstrap resample from sample of threshold exceedances;
 - 13: **if** M **then**
 - 14: Count number $k(\theta)$ of threshold exceedances in $\mathcal{N}(\theta, h)$;
 - 15: Count number $N(\theta)$ of observations in $\mathcal{N}(\theta, h)$;
 - 16: Estimate $\xi(\theta)$ on $\mathcal{N}(\theta, h)$ using (5.1);
 - 17: Estimate a_θ on $\mathcal{N}(\theta, h)$ using (5.3);
 - 18: Estimate high quantile with $p \ll 1/N(\theta)$ value using (6.1);
 - 19: Estimate finite right endpoint ($x^{F_\theta} < \infty$) using (6.2);
 - 20: **else if** local ML **then**
 - 21: Count number $k(\theta)$ of threshold exceedances in $\mathcal{N}(\theta, h)$;
 - 22: Count number $N(\theta)$ of observations in $\mathcal{N}(\theta, h)$;
 - 23: Estimate $\xi(\theta)$, $\sigma_u(\theta)$ on $\mathcal{N}(\theta, h)$ using (5.4);
 - 24: Estimate T -year value using (6.3);
 - 25: Estimate right endpoint (when $\xi(\theta) < 0$) using (6.5);
 - 26: **else if** spline ML **then**
 - 27: Estimate optimal smoothing parameters λ , κ and hence estimate $\xi(\theta)$, $\sigma_u(\theta)$ (Algorithm 2);
 - 28: Estimate optimal smoothing parameter μ and fraction $\tau(\theta)$ of threshold exceedances by logistic regression using (6.4);
 - 29: Estimate T -year value using (6.3), with $\hat{\tau}(\theta)$ in place of $k(\theta)/N(\theta)$;
 - 30: Estimate right endpoint (when $\xi(\theta) < 0$) using (6.5);
 - 31: **end if**
 - 32: Accumulate bootstrap estimates for parameters, extreme levels or quantiles and for the right endpoint;
 - 33: **end for**
 - 34: Calculate bootstrap means and confidence intervals for parameters, T -year values and endpoint.
-