Ocean Engineering 142 (2017) 315–328

Contents lists available at ScienceDirect

Ocean Engineering

journal homepage: www.elsevier.com/locate/oceaneng

Efficient estimation of return value distributions from non-stationary marginal extreme value models using Bayesian inference

Emma Ross[a], David Randell[b], Kevin Ewans[c], Graham Feld[d], Philip Jonathan[a,*]

[a] Shell Projects & Technology, London SE1 7NA, United Kingdom
[b] Shell Projects & Technology, 1031 HW Amsterdam, The Netherlands
[c] Metocean Research Limited, New Plymouth 4310, New Zealand
[d] Shell Projects & Technology, Aberdeen AB12 3FY, United Kingdom

ARTICLE INFO

Keywords:
Extreme
Significant wave height
Non-stationary
Splines
Bayesian inference
Return value
Numerical integration

ABSTRACT

Extreme values of an environmental response can be estimated by fitting the generalised Pareto distribution to a sample of exceedances of a high threshold. In oceanographic applications to responses such as ocean storm severity, threshold and model parameters are typically functions of physical covariates. A fundamental difficulty is selection or estimation of an appropriate threshold or interval of thresholds, of particular concern since inferences for return values vary with threshold choice. Historical studies suggest that evidence for threshold selection is weak in typical samples.

Hence, following Randell et al. (2016), a piecewise gamma-generalised Pareto model for a sample of storm peak significant wave height, non-stationary with respect to storm directional and seasonal covariates, is estimated here using Bayesian inference. Quantile regression (for a fixed quantile threshold probability) is used to partition the sample prior to independent gamma (body) and generalised Pareto (tail) estimation. An ensemble of independent models, each member of which corresponds to a choice of quantile probability from a wide interval of quantile threshold probabilities, is estimated. Diagnostic tools are then used to select an interval of quantile threshold probabilities corresponding to reasonable model performance, for subsequent inference of extreme quantiles incorporating threshold uncertainty.

The estimated posterior predictive return value distribution (for a long return period of the order of 10,000 years) is a particularly useful diagnostic tool for threshold selection, since this return value is a key deliverable in metocean design. Estimating the distribution using Monte Carlo simulation becomes computationally demanding as return period increases. We present an alternative numerical integration scheme, the computation time for which is effectively independent of return period, dramatically improving computational efficiency for longer return periods.

The methodology is illustrated in application to storm peak and sea state significant wave height at a South China Sea location, subject to monsoon conditions, showing directional and seasonal variability.

1. Introduction

1.1. Threshold selection

Threshold selection in practical application of extreme value analysis is almost always problematic. Even in the absence of covariate effects, it is rarely clear where the threshold should be set, or indeed if setting a single threshold is even desirable. A review of threshold selection for extreme value analysis is given by Scarrott and MacDonald (2012). The generalised Pareto (GP) model for peaks-over-threshold is motivated by asymptotic arguments: the threshold needs to be set high enough so that a generalised Pareto model fits threshold exceedances reasonably, to reduce bias. Yet the threshold should be set low enough that there are sufficient exceedances to estimate generalised Pareto model parameters, to reduce variance: a typical bias-variance trade-off. Graphical techniques such as the mean excess plot (Ghosh and Resnick, 2010) can be of some use in aiding a sensible threshold choice, as can inspection of the stability of generalised Pareto shape parameter estimate or other key inferences, such as estimates for return values or other structure variables as a function of threshold. Here, a structure variable is any variable defined in terms of one or more responses modelled by extreme value analysis. For example, the

* Corresponding author.
E-mail address: philip.jonathan@shell.com (P. Jonathan).

http://dx.doi.org/10.1016/j.oceaneng.2017.06.059

load on an offshore structure can be considered a structure variable, defined in terms of extreme value responses including significant wave height, current speed, wind speed etc. Graphical techniques for threshold selection are rarely conclusive however. Some authors, including Sanchez-Archilla et al. (2008), Thompson et al. (2009), Northrop and Coleman (2014) and Wadsworth (2016), have proposed procedures for estimating good thresholds, but these all contain subjective elements. In the presence of covariate effects, threshold selection is even more problematic. Typically, the threshold is set as a local (covariate-dependent) quantile of the response, and the problem of threshold selection transformed into one of specifying the appropriate threshold quantile level for the covariates used.

1.2. Threshold estimation

One approach to overcoming the need to specify an extreme value threshold ψ before extreme value inference is to make ψ a model parameter to be estimated. To achieve this, the extreme value model must be extended so that it describes part or all of the body of the sample, as well as extreme value threshold exceedances, such as in Tancredi et al. (2006), Wadsworth et al. (2010), MacDonald et al. (2011) and Randell et al. (2016). Yet in such a model, incorporating extreme value inference and estimation of ψ, it is not clear whether is desirable that estimation of ψ be influenced by model fit to threshold non-exceedances: adequate fit of the generalised Pareto tail should take priority, since tail estimation is our primary concern. From this perspective, models for which ψ is pre-specified with no regard to threshold non-exceedances would seem advantageous.

1.3. Incorporating threshold uncertainty

Values $q(\psi)$ for return values (and other structure variables of interest) can be found corresponding to any threshold ψ. A final "preferred" return value q_1 might then be selected corresponding a single "best" threshold $\hat{\psi}$, found either from inspection of diagnostics (or model estimation), such that

$$q_1 \triangleq q(\hat{\psi}).$$

Alternatively, we can provide final return values q_2 by integrating over suitable values of ψ specified by some density $f(\psi)$, itself inferred from inspection of diagnostics (or directly estimated). Then

$$q_2 \triangleq \int_\psi q(\psi) f(\psi) d\psi.$$

The advantage of using q_2 over q_1 is that uncertainty in ψ is propagated through to return values. When pre-specification or estimation of $\hat{\psi}$ is problematic, it seems reasonable to prefer q_2 over q_1. When pre-specification or estimation of $\hat{\psi}$ is straightforward, we expect q_2 and q_1 to be similar since $f(\psi)$ provides probability mass only around $\hat{\psi}$. Northrop et al. (2017) discuss cross-validatory threshold selection, including a method for incorporating threshold uncertainty.

In this work, we choose to employ estimates of the form q_2 for return values, but note that there is still considerable subjectivity in the choice of $f(\psi)$ to be used. We seek to inform this choice by consideration of various model diagnostics. We thereby hopefully introduce some rationality, but make no claim to have removed all subjectivity.

1.4. Non-stationary Bayesian extreme value modelling

In Randell et al. (2016), a model for the distribution of independent observations of peaks-over-threshold of a response such as storm peak significant wave height given multidimensional covariates is developed, incorporating the generalised Pareto distribution for exceedances of some estimated non-stationary threshold. A truncated Weibull distribution characterises values below the threshold. The model is used primarily to estimate distributions of return values corresponding to long return periods, for use in the design and reliability assessment of marine and coastal structures. The methodology is intended to be easy to use, and computationally efficient for full-scale oceanographic applications for sample sizes from 10^3 to 10^7 with at least a two-dimensional (e.g. directional-seasonal) covariate domain. Practical applications' experience using the model suggests that (a) there is often little or no evidence in the sample to inform threshold selection, that (b) the model form itself is perhaps too restrictive to facilitate easy threshold estimation using Bayesian inference, and that (c) the gamma distribution provides an equally flexible form for the distribution of non-exceedances, together with more stable inference when likelihood gradients are exploited (Section 3.5).

We have concluded, for typical applications to ocean storm severity, that there is little value in seeking to estimate extreme value threshold directly. Instead, making inferences using an ensemble of piecewise gamma-generalised Pareto models, each of which corresponds to a specific threshold choice, is less problematic inferentially and more useful in practice. Ensemble members (and hence thresholds) are selected to give reasonable performance, as judged by inspection of relevant model diagnostics. Specifically, ensemble threshold choices typically correspond to an interval of threshold probabilities for a non-stationary quantile regression threshold, as in Randell et al. (2015). The model is described in Section 3.

1.5. Return value estimation

Return value estimates are the key deliverable of metocean design. In the presence of covariates and multiple extreme value thresholds, closed form expressions for return values are not available. As a result, Monte Carlo simulation is typically used. This involves simulating under the estimated model, generating thousands of realisations of sets of extreme values corresponding to the return period of interest. Monte Carlo simulation also provides an intuitive framework for estimating return value distributions for "dissipated" or sea state significant wave height from an extreme value model for storm peak events. The computational burden increases approximately linearly with return period. Parallel computation is useful, but nevertheless return value estimation by Monte Carlo simulation remains computationally intensive, accounting for the vast majority of computing resource required for a typical study (Section 4).

We have developed numerical integration algorithms to replace all Monte Carlo simulations previously performed, resulting in a huge reduction in time required for analysis. The approach is described in Section 4.

1.6. Outline of article

The objective of this article is to improve the applicability of the methodology presented in Randell et al. (2016), in two major respects. Firstly, we simplify the extreme value model by performing prior non-stationary quantile regressions to estimate possible extreme value thresholds, and use each of these thresholds in turn to partition the sample into "body" and "tail". Independent gamma and generalised Pareto models are next estimated for body and tail respectively using Bayesian inference. Finally, an ensemble of models (for different plausible threshold choices) is adopted for return value inference. Secondly, we replace return value inference using computationally-intensive Monte Carlo simulation by a computationally-efficient numerical integration scheme, which amongst other things, improves the