

OMAE2009-79915

Published. J. Offshore Mechanics Arctic Engineering 2011 133 / 011601-1.

A SPATIO-DIRECTIONAL MODEL FOR EXTREME WAVES IN THE GULF OF MEXICO

Philip Jonathan

Shell Technology Centre Thornton
P.O. Box 1
Chester
United Kingdom.
philip.jonathan@shell.com

Kevin Ewans

Shell International Exploration and Production
P.O. Box 60
2280 AB Rijswijk
The Netherlands.
kevin.ewans@shell.com

ABSTRACT

The characteristics of extreme waves in hurricane dominated regions vary systematically with a number of covariates, including location and storm direction. Reliable estimation of design criteria requires incorporation of covariate effects within extreme value models. We present a spatio-directional model for extreme waves in the Gulf of Mexico, motivated by the non-homogeneous Poisson model for peaks over threshold. The model is applied to storm peak significant wave height H_S for arbitrary geographic areas from the proprietary GOMOS hindcast for the US region of the Gulf of Mexico for the period 1900-2005. At each location, directional variability is modelled using a non-parametric directional location and scale; data are standardised (or "whitened") with respect to local directional location and scale to remove directional effects. For a suitable choice of threshold, the rate of occurrence of threshold exceedences of whitened storm peak H_S with direction per location is modelled as a Poisson process. The size of threshold exceedences is modelled using a generalised Pareto form, the parameters of which vary smoothly in space, and are estimated within a roughness penalised likelihood framework using natural thin plate spline forms in two spatial dimensions. By re-parameterising the generalised Pareto model in terms of asymptotically independent parameters, an efficient back-fitting algorithm to estimate the natural thin plate spline model is achieved. The algorithm is motivated in an appendix. Design criteria, estimated by simulation, are illustrated for a typical neighbourhood of

17 × 17 grid locations. Applications to large areas consisting of more than 2500 grid locations are outlined.

1 INTRODUCTION

The availability of comprehensive metocean data allows the effect of the heterogeneity of extremes with respect to direction, season and location to be accommodated in estimation of design criteria.

Capturing covariate effects of extreme sea states is important when developing design criteria. Design criteria derived from a model that adequately incorporates covariate effects can be materially different from a model which ignores those effects (e.g. Jonathan et al. 2008). In previous work (e.g. Jonathan and Ewans 2007, Ewans and Jonathan 2008) it has been shown that omni-directional storm peak H_{S100} derived from a directional model can be heavier tailed than that derived from a direction-independent approach, indicating that large values of storm peak H_S are more likely than we might anticipate were we to base our beliefs on estimates which ignore directionality. Similar effects have been demonstrated for seasonal covariates (e.g. Anderson et al. 2001, Jonathan and Ewans 2008).

There is a large body of statistics literature regarding modelling of covariate effects in extreme value analysis; see, e.g., Davison and Smith [1990] or Robinson and Tawn [1997]. The case for adopting an extreme value model incorporating covariate effects is clear, unless it can be demonstrated statistically that a model ignoring covariate effects is no less appropriate. Chavez-Demoulin and Davison [2005] and Coles [2001] provide straight-forward descriptions of a non-homogeneous Poisson model in which occurrence rates and extremal properties are modelled as functions of covariates. Scotto and Guedes-Soares [2000] describe modelling using non-linear thresholds. A Bayesian approach is adopted Coles and Powell [1996] using data from multiple locations, and by Scotto and Guedes-Soares [2007]. Spatial models for extremes (Coles and Casson [1998], Casson and Coles [1999]) have also been used, as have models (Coles and Tawn [1996, 2005]) for estimation of predictive distributions, which incorporate uncertainties in model parameters. Ledford and Tawn [1997] and Heffernan and Tawn [2004] discuss the modelling of dependent joint extremes. Chavez-Demoulin and Davison [2005] also describe the application of a block bootstrap approach to estimate parameter uncertainty and the precision of extreme quantile estimates, applicable when dependent data from neighbouring locations are used. Guedes-Soares and Scotto [2001] discuss the estimation of quantile uncertainty. Eastoe [2007] and Eastoe and Tawn [2009] illustrate an approach to removing covariate effects from extremes data prior to model estimation.

One of the first examinations of the spatial characteristics of extreme wave heights in the Gulf of Mexico was reported by Haring and Heideman [1978]. They performed extremal analysis of the ODGP hurricane hindcast data base (Ward et al. 1978) at a number of continental shelf locations from Mexico to Florida, and concluded that there was not practical difference between the sites, but they did observe a gradual reduction in extreme wave heights with decreasing water depth. Chouinard et al. [1997] took the opportunity to re-examine the spatial behaviour of extremes in the Gulf of Mexico, when the GUMSHOE hindcast data base became available. They found strong support for the existence of so called hurricane alleys in which regions of more severe hurricanes coinciding with regions of elevated near-surface water temperatures and confirmed the need for a spatially-dependent hurricane severity probability density; they

proposed a function with a spatial scale of 150 km. As a consequence the API [2007] have released interim guidance on hurricane conditions in the Gulf of Mexico that contains criteria for four separate regions.

Here we introduce a novel spatio-directional model for extremes and apply it to data for neighbourhoods of the Gulf of Mexico. The model incorporates non-parametric standardisation of extremes to remove certain covariate effects, and a natural thin plate spline formulation to characterise the smooth variation of extreme value parameters in 2-D space within a penalised maximum likelihood framework. We believe the main contributions of the work to be the development of a straight-forward approach to extreme value analysis with multiple covariates, and provision of an intuitive alternative to site pooling.

The layout of the article is as follows. In section 2 we describe the present application and illustrate the data. In section 3 we outline the extreme value modelling procedure and describe results. In section 4 we discuss estimation of design criteria for a typical neighbourhood of locations. In Section 5, we compare modelling results with earlier work using a parametric directional extremes approach, and outline Gulf-wide analyses. Conclusions are drawn and recommendations made. An informal outline of the natural thin plate spline generalised Pareto model is given in Appendix A. Some supporting illustrations are relegated to Appendix B for clarity.

2 DATA

Data examined are significant wave height H_S values from the proprietary GOMOS Gulf of Mexico hindcast Study (Oceanweather 2005), for the period September 1900 to September 2005 inclusive, at 30-minute intervals. For a total of 4363 grid locations, data are available at a grid spacing of 0.125 degrees in both latitude and longitude. We chose to retain 2658 "non-boundary" locations defined as follows for analysis. At a non-boundary location, it is possible to place a square box of dimensions 11×0.125 degrees centred at the location, such that all locations within the box belong to the full hindcast. In this way, non-boundary locations do not include coastal US regions and locations near to Mexican water. A total of 315 storm events were isolated, common to all non-boundary grid points. For each storm period for each grid point, we isolated storm peak significant wave height for subsequent analysis. We also extract the corresponding vector mean direction of the sea state at the time of the peak significant wave height, henceforth referred to as the storm peak direction. This quantity is not necessarily aligned with the storm track direction or the local wind at the time of the storm peak significant wave height at the grid point concerned.

We motivate the model development using a typical square neighbourhood N of 17×17 grid locations (corresponding to 2 degrees in both longitude and latitude). For reasons of confidentiality, we withhold the true co-ordinates of neighbourhood N , and refer to locations within this neighbourhood in terms of longitude and latitude relative to the centre of the neighbourhood. We also re-scale the values of storm peak H_S (in illustrations only) by an arbitrary multiplicative factor, the value of which is also withheld. For clarity, we refer to re-scaled values as H_S^* .

The directional and seasonal dependence of extreme events in GOMOS has already been illustrated (e.g. Jonathan and Ewans 2007, Jonathan and Ewans 2008). The spatial variability of storm severity in the Gulf of Mexico is also widely reported (e.g. Chouinard 1992,




Figure 1. Contours of the maximum value of re-scaled storm peak H_S^* for a typical neighbourhood of locations, N. The background grey-scale is graduated from the lowest (darkest) to the highest (lightest) values per location.

Chouinard et al. 1997). For the current work, a contour map of the maximum value of re-scaled storm peak H_S^* per location is given in Figure 1.

3 EXTREME VALUE MODELLING

We have values for storm peak significant wave heights $\{X_{ij}\}_{i=1,j=1}^{n,p}$ for $n = 315$ storms at $p = 2658$ locations in the Gulf of Mexico (GoM), with corresponding storm peak directions $\{\theta_{ij}\}_{i=1,j=1}^{n,p}$ occurring in some period P_0 . We seek to develop a spatio-directional model which will account for both directional and spatial variation in extreme value characteristics. The modelling procedure is comprised of the following elements:

1. At each location j , we characterise the variation of $\{X_i\}_{i=1}^n$ with respect to direction using a directional standardisation procedure. The resulting "whitened" data $\{W_{ij}\}_{i=1,j=1}^{n,p}$ exhibit little directional variability in local "location" (e.g. the median value) and "band" (e.g. a chosen inter-quantile range). The directional standardisation procedure is described in section 3.1.
2. At each location j , we select an appropriate threshold u_j above which the values of whitened data $\{W_{ij}\}_{i=1}^n$ can be reasonably fit with a generalised Pareto model. In the work reported here, the threshold is set to a fixed quantile of the data per location, for all $\{u_j\}_{j=1}^p$, for simplicity.
3. At each location j , we use the whitened data $\{W_{ij}\}_{i=1}^n$ to estimate the rate of occurrence $\rho_j(\theta)$ of exceedences of u_j , as a function of storm peak direction θ , using a Poisson model.
4. For all whitened data at all locations, we fit a spatial generalised Pareto (GP) model to threshold exceedences. The spatial GP model is estimated using roughness-penalised maximum likelihood, with a natural thin plate spline form for model parameters in space. The spatial GP model formulation is described in section 3.2.
5. Finally, a Monte Carlo simulation based on the fitted model is performed to estimate extreme quantiles, such as omni-directional

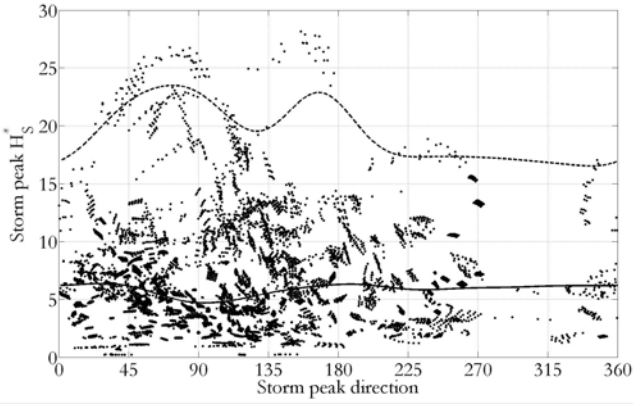


Figure 2. Re-scaled storm peak H_S^* data with storm peak direction for a typical 5×5 grid of neighbouring locations, with estimates of the local median (solid) and 0.99 quantile with respect to direction. There is considerable variation in the values of re-scaled storm peak H_S^* with direction. Storm peak direction gives the direction from which storms emanate, measured clockwise, with North at 0 degrees.

100-year return period events for the spatial neighbourhood of interest.

3.1 Directional standardisation or Whitening

At location j , the objective of directional standardisation is to transform the data $\{X_i\}_{i=1}^n$ so that they have an approximately constant location and scale with respect to direction, borrowing from the work of Tawn and colleagues (e.g. Dixon et al. 1998, Eastoe 2007). In this sense, the standardisation procedure removes directional "colour" from the data, and "whitens" it. The form of the transformation to be used is somewhat arbitrary. In the current work we have adopted the simple form:

$$W_{ij} = \frac{X_{ij} - \mu(\theta_{ij})}{\eta(\theta_{ij})}$$

where $\mu_j(\theta)$ and $\eta_j(\theta)$ are local estimates of data location and scale with respect to direction. For any direction θ_k , let I_{θ_k} be a narrow interval of directions centred at θ_k . Then we set $\mu_j(\theta_k)$ to the median value (corresponding to $q_L = 0.5$) of the set $X_{h_j s.t. \theta_{h_j} \in I_{\theta_k}}$, and we set $\eta_j(\theta_k)$ to be the difference between a high quantile q_U (e.g. the 0.99 quantile) and the median value of the same set. We refer to the interval of quantiles $[q_L, q_U]$ used to define η as the "whitening band" for convenience. In the current work, the same values of q_L and q_U are used for all locations for simplicity.

In practice, since the number of storm peaks per location occurring in a narrow interval of storm peak direction is small, it was necessary to pool data from a local 5×5 neighbourhood to obtain more precise estimates of $\mu(\theta)$ and $\eta(\theta)$. In addition, some smoothing

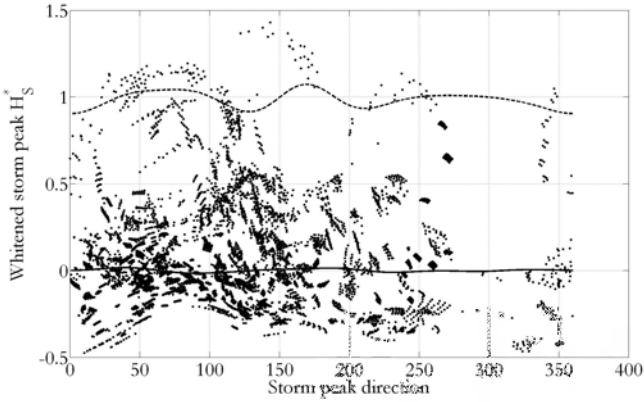


Figure 3. Whitened storm peak H_S data with storm peak direction for the 5×5 grid of neighbouring locations in Figure 2, with estimates of local median (solid) and 0.99 quantile with respect to direction. There is considerably less variation in the values of whitened storm peak H_S with direction, with respect to the two quantiles concerned.

of estimates was performed. Figures 2 and 3 illustrate data for a typical location, before and after directional standardisation. Before whitening there is clear variability in values of storm peak H_S with direction, which is reduced by directional standardisation.

The quality of directional standardisation can be assessed by statistical testing, e.g. conformity of whitened data to a random sample from a generalised Pareto distribution with constant parameters. A sensitivity study was also performed to explore the effect of selection of q_U and local directional smoothing of $\mu(\theta)$ and $\eta(\theta)$ on the characteristics of the whitened data. It was found that a large value $q_U > 0.9$ was necessary to obtain relatively stable results. However, it should be noted that further work is necessary to improve the reliability and stability of directional standardisation for routine application.

3.2 Generalised Pareto model

The whitened data $\{W_{ij}\}_{i=1, j=1}^{n, p}$ correspond to p samples of extreme values of size n . At location j , we model these data using a generalised Pareto form with distribution function:

$$F_{W_j|u_j}(x) = 1 - \left(1 + \frac{\gamma_j}{\sigma_j}(x - u_j)_+\right)^{-\frac{1}{\gamma_j}}$$

where γ_j is the shape parameter and σ is the scale. We estimate the sets of estimates $\{\gamma_j\}_{j=1}^p$ and $\{\sigma_j\}_{j=1}^p$ simultaneously for all locations using penalised maximum likelihood, such that the GP parameters are required to vary smoothly in space.

As has been noted by numerous authors, re-parameterising the GP model in terms of γ and ν , where $\nu = \sigma(1 + \gamma)$, yields a diagonal

Fisher information matrix, thereby yielding uncorrelated estimates (at least asymptotically) and simplifying the computation of estimates for the information matrix. One consequence of the (γ, ν) parameterisation is a simplification in the implementation of the natural thin plate spline (NTPS) algorithm for the spatial GP model. We therefore adopt the re-parameterised GP form for computations below.

At location j , the un-penalised negative log likelihood for observation x_{ij} takes the form:

$$\begin{aligned} l_{ij} &= \log \sigma_j + \left(\frac{1}{\gamma_j + 1} \right) \log \left(1 + \frac{\gamma_j}{\sigma_j} (x_{ij} - u_j) \right)_+ \\ &= \log \frac{1 + \gamma_j}{\nu_j} + \frac{1}{\gamma_j + 1} \log \left(1 + \frac{\gamma_j (1 + \gamma_j) (x_{ij} - u_j)}{\nu_j} \right)_+ \end{aligned}$$

The corresponding penalised negative log likelihood for all observations at all locations is:

$$l^* = \sum_{i=1}^n \sum_{j=1}^p l_{ij} + \frac{\lambda_\gamma}{2} R_\gamma + \frac{\lambda_\nu}{2} R_\nu$$

where R_γ and R_ν correspond to the spatial roughness of γ and ν respectively, which can be expressed as quadratic forms in the parameters of the NTPS expressions for γ and ν . This attractive form for parameter roughness makes maximum likelihood estimation possible using a so-called *back-fitting* algorithm. Moreover, since we have re-parameterised the problem in terms of asymptotically independent parameters, the estimation of γ can be made independently of ν , thereby greatly simplifying computation. Maximum likelihood (minimum log likelihood) is applied to find sets of values $\{\gamma_j\}_{j=1}^p$ and $\{\nu_j\}_{j=1}^p$ which are spatially smooth. The natural thin plate spline fitting procedure is given in outline in Appendix A.

The values of the roughness coefficients λ_γ and λ_ν dictate the smoothness of the solutions obtained for γ and ν . Various approaches are available to set appropriate values for these parameters, including cross-validation (see, e.g. Ewans and Jonathan [2008]). However, partly to the computational burden involved for larger neighbourhoods, an alternative approach was taken here. For a typical region of the GoM, we constructed a spatial GP model (M_0 , say) which appeared to give a realistic characterisation of the data for that region. We then refit various spatial GP models to different realisations from spatial model M_0 , and explored which intervals of values for λ_γ and λ_ν give acceptable fits. These intervals were taken as preferred starting values for the two roughness coefficients. We then explored the effect of varying λ_γ and λ_ν for the large spatial domains (including all 2658 locations), to ensure that estimates were not overly sensitive to choice of roughness coefficients. The spatial estimates for γ and σ for neighbourhood N (with whitening band $[q_L = 0.5, q_U = 0.99]$ incorporating moderate directional smoothing, and extreme value threshold set at the 0.75-quantile per location) is shown in Figures 4




Figure 4. Estimated value of generalised Pareto shape parameter using spatial model, based on whitened data with whitening band $[q_L=0.5, q_U=0.99]$, incorporating moderate directional smoothing, and extreme value threshold set at the 0.75 quantile per location.




Figure 5. Estimated value of generalised Pareto scale parameter using spatial model, based on whitened data with whitening band $[q_L=0.5, q_U=0.99]$, incorporating moderate directional smoothing, and extreme value threshold set at the 0.75 quantile per location.

and 5. These estimates are of no direct physical relevance, since they correspond to whitened data which must be directionally "coloured" prior to interpretation.

Setting the threshold u_j per location j is an important and usually challenging precursor to any extreme value modelling. In the current work, threshold selection was made by inspecting numerous diagnostic plots of the behaviour of generalised Pareto parameter estimates as a function of threshold, independently for different locations. Figure 6 shows the variation of shape estimate with empirical non-exceedence probability for 6 locations selected from different regions of the GoM, and is typical of the results found. Threshold $u_j = 0$ at all locations corresponds to admitting half the sample at each location for modelling, since we have whitened using the median, and a non-exceedence probability of 0.5. Setting the threshold per location based on a non-exceedence probability of 0.25 would appear plausible, but clearly choice of threshold is problematic.

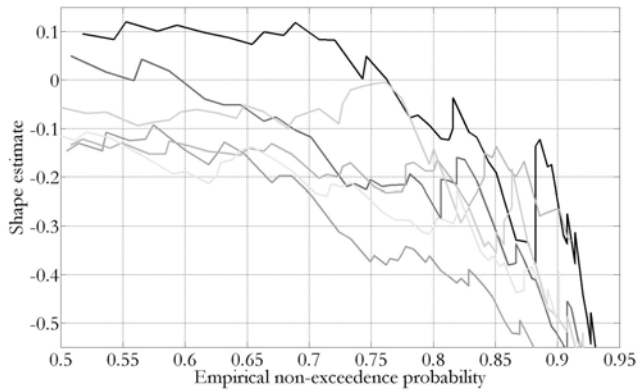


Figure 6. Variation of generalised Pareto shape parameter with threshold for 6 locations selected from different regions of the GoM. Selection of a suitable threshold for extreme value analysis is problematic.

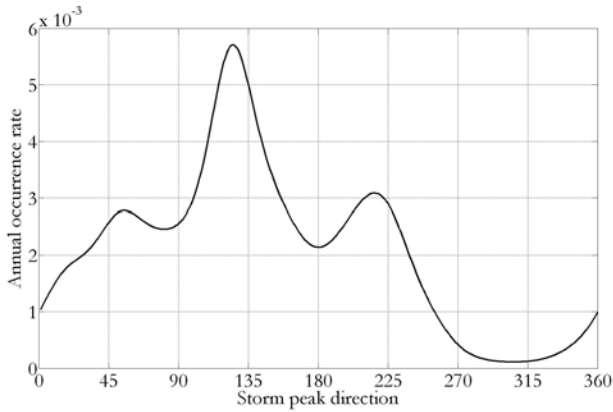


Figure 7. Annual occurrence rate for a typical location estimated using the Poisson model.

3.3 Poisson model

Per location j , the annual rate of exceedence $\rho_j(\theta)$ of u_j as a function of storm peak direction for whitened data is estimated as a roughness-penalised Poisson model as described in Jonathan and Ewans [2008]. A typical estimate for the annual rate of occurrence (per degree storm peak direction) is shown in Fig 7. The figure suggests that the occurrence rate of storms with directions in $[270,360)$ (i.e. emerging from the North West) is relatively small, as would be expected.

4 DESIGN CRITERIA

Monte Carlo simulations as outlined in Section 3 yield estimates for design conditions. The simulation proceeds as follows for the default spatio-directional model. First, 1000 realisations of 100 years of whitened storm peak events are simulated from the NTPS model, using the Poisson occurrence rate $\rho_j(\theta)$ per location j to estimate the number of occurrences and the distribution of those




Figure 8. Contours of median re-scaled H_{S100}^* , estimated using the NTPS model on whitened data. A [0.5,0.99] quantile whitening band was used with moderate directional smoothing, with a 0.75 quantile threshold for model estimation.

occurrences with respect to storm peak direction per location. The whitened data per location are then coloured using the corresponding directional standardisation parameters $\mu_j(\theta)$ and $\eta_j(\theta)$ for that location. Empirical estimates for the distribution of the 100-year storm peak maximum event are then accumulated. We illustrate the results of these simulation studies using contour plots of the median re-scaled storm peak H_{S100} for neighbourhood N, generated using MATLAB contouring software. The contour map for the spatio-directional model using a [0.5,0.99] quantile whitening band with moderate directional smoothing, with a 0.75 quantile threshold for model estimation using NTPS is given in Figure 8. We see that the gross features of Figure 1 are reproduced. For example, the maximum at around relative location [-0.4,0.5] extending North-East, and an area of low values in the bottom left quadrant are clear in both figures.

We compare estimates for this design condition obtained using this spatio-directional model with alternative approaches, namely:

1. Using directional standardisation with independent GP fits per location, based on local 5×5 pools of neighbouring grid points. This comparison allows us to assess the relative merits of independent GP estimation and the NTPS spatial model.
2. Using original (un-whitened) storm peak H_S data in the NTPS spatial model, allowing assessment of the effect of directional standardisation on design criteria.
3. Using original (un-whitened) storm peak H_S and independent GP fits per location (based on 5×5 pools), which might be thought of as similar to a default current engineering approach.

When the directional standardisation step is omitted, estimates for the median re-scaled storm peak H_{S100} can be obtained directly from theory in close form. Comparison of close-for estimates with simulation-based equivalents allows us to validate the simulation procedure. Note that using the original (un-whitened) data is equivalent to ignoring directional variability of extremes, precisely because the directional standardisation step has been omitted.

For the recommended spatio-directional model, we also provide estimates for the 0.25 and 0.75 quantiles of the distribution of




Figure 9. Contours of the 1st quartile (0.25 quantile) of re-scaled H_{S100}^* , estimated using the NTPS model on whitened data. A [0.5,0.99] quantile whitening band was used with moderate directional smoothing, with a 0.75 quantile threshold for model estimation.




Figure 10. Contours of the 3rd quartile (0.75 quantile) of re-scaled H_{S100}^* , estimated using the NTPS model on whitened data. A [0.5,0.99] quantile whitening band was used with moderate directional smoothing, with a 0.75 quantile threshold for model estimation.

median re-scaled storm peak H_{S100} . These are shown in Figures 9 and 10.

There is insufficient data per location to estimate the GP model reliably using data solely from that location. One effect of pooling data from neighbouring locations is to smooth the resulting estimates of design conditions; the spatial extent of the smoothing depends on the extent of the pooling employed. Figure 11 (corresponding to alternative 1 above using a 5×5 pool) is very similar in terms of gross features to Figure 8, suggesting that there is consistency between the spatial spline and independent generalised Pareto estimates (based on pooling); this was indeed found to be the case in general, provided that values for the roughness coefficients of the NTPS correspond approximately to the extent of pooling employed. Yet, as will be discussed later in Section 5, we believe the adoption of the spatial model is more intuitive than pooling.

For design conditions based on whitened data, the effect of varying the upper limit q_U of the whitening band over a relatively narrow

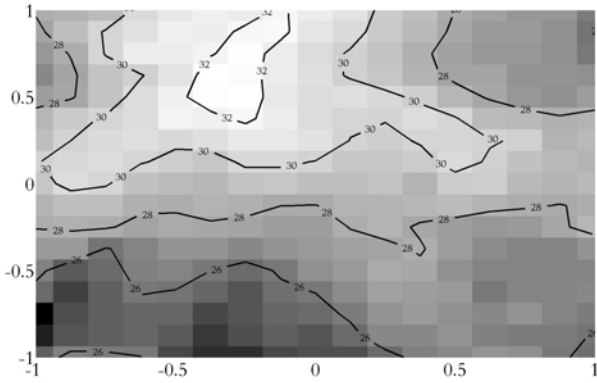


Figure 11. Contours of median re-scaled H_{S100}^* , estimated using independent GP fits per location on whitened data. A [0.5,0.99] quantile whitening band was used with moderate directional smoothing, with a 0.75 quantile threshold for model estimation.

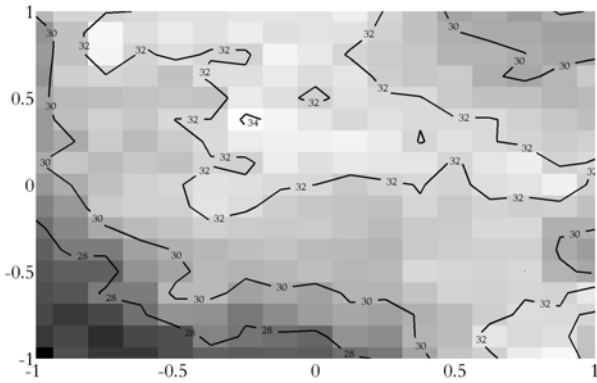


Figure 12. Contours of median re-scaled H_{S100}^* , estimated using the NTPS model on original (un-whitened) data. A [0.5,0.99] quantile whitening band was used with moderate directional smoothing, with a 0.75 quantile threshold for model estimation.

range produces stable results. For example, estimates in Figure 8 (for $q_U = 0.99$) are relatively consistent with those in Figure 16 (in Appendix B, for $q_U = 0.9$). Furthermore, the extent of directional smoothing during whitening does not unduly affect estimates for design conditions. This can be seen, for example, from comparison of Figure 8 (for moderate smoothing) with Figure 17 (in Appendix B, for heavier directional smoothing). For analysis based on whitened and original (un-whitened) data, the effect of extreme value threshold (expressed as a quantile per location) is relatively small within a relatively narrow range. For example, estimates using 0.75 and 0.90 quantiles (not shown) were found to be consistent.

There are differences between estimates for design conditions based on whitened (e.g. Figure 8 or 11) and un-whitened data (e.g. Figure 12 or 13). These differences appear consistently regardless of parameter variation for whitening, directional smoothing and extreme value threshold over reasonable ranges. This suggests systematic differences due to directional effects, and will be discussed in




Figure 13. Contours of median re-scaled H_{S100}^* , estimated using independent GP fits per location on original (un-whitened) data. A 0.75 quantile threshold for model estimation per location.

the next section.

5 DISCUSSION AND CONCLUSIONS

In the current implementation of the spatio-directional extremes model, the main sources of uncertainty in estimates of return value appear to be the specification of *hyper*-parameters, particularly the upper quantile q_U of the whitening band, and the extreme value threshold per location. Threshold selection presents a challenge in any extreme value analysis, especially one involving multiple threshold selections. We have adopted a constant quantile level to specify threshold per location. Our sensitivity studies indicate that estimates of return values are relatively stable with respect to small changes in q_U and threshold. Nevertheless, reliable spatio-directional modelling incorporating simultaneous fitting of generalised Pareto distributions over large numbers of locations demands systematic and realistic specification of hyper-parameters. Perhaps the most hopeful approach would be a Bayesian averaging with respect to plausible intervals of values for both q_U and extreme value threshold.

A major advantage of the current spatio-directional model is that data pooling is eliminated. Instead, neighbouring locations are constrained to have more similar values of extreme value parameters than locations which are far apart. To illustrate, first consider pooling data over a 5×5 grid of locations in order to achieve a sample of extremes large enough for extreme value analysis. We are implicitly assuming that the marginal extreme value distributions for the 25 locations pooled are identical, which is possibly not the case. Furthermore, if we set the extreme value threshold too high, then it is likely that the sample will be dominated by multiple (dependent) occurrences from a small number of large storms, biasing the extreme value estimation. Moreover, where we next to use the same pooling approach at an adjacent grid location (in either longitude or latitude), then 20 of the original locations would be included in the sample for extreme value analysis. Yet we might estimate a different value for the extreme value parameters. In the spatio-directional model, the situation is considerably more intuitive. The model assumes that extreme value shape and scale vary slowly with location.

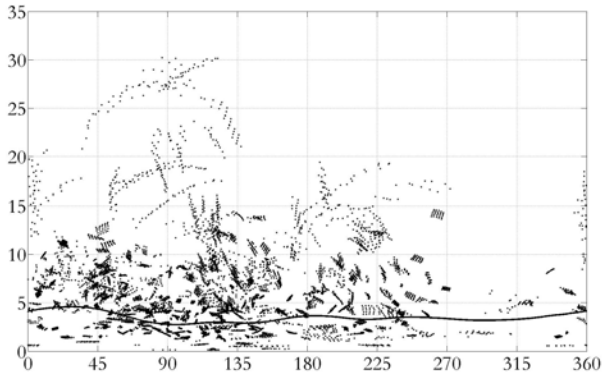


Figure 14. Re-scaled storm peak H_{S100}^* data with storm peak direction, for location L with relative longitude = 0.5 and latitude -0.75 within neighbourhood N. Also shown is a 0.25 quantile directional threshold used to fit a Fourier directional GP model.

Each occurrence of a threshold exceedence (regardless of location) is used exactly once for modelling. Marginally for each location, independent observations storm peaks are used for extreme value modelling, yet the extreme value shape and scale parameter estimates for neighbouring locations are constrained to be similar by their roughness-penalised spatial spline form.

The NTPS generalised Pareto model currently estimates spatial forms for both extreme value shape and scale parameters. The variability of parameters in space is controlled by the sizes of the roughness coefficients for shape and scale, set by inspection or a procedure similar to that suggested in Section 3. In careful application, it is essential to set these as rigorously as possible, possibly by independently varying spatial roughness coefficients for extreme value shape and scale over wide intervals using a cross-validation approach. Alternatively we might examine a set of nested models of increasing complexity, using hypothesis testing to justify the adoption of the most appropriate level of model complexity.

For the current neighbourhood N, the most noticeable difference between estimates for the median H_{S100} occurs at relative location $[0.5, -0.75]$, at which models based on both NTPS and independent pooled analysis of original (un-whitened) data (e.g. Figure 12 and Figure 13), ignoring direction effects, yield higher estimates than those incorporating directional whitening to accommodate directional variability of extremal behaviour (e.g. Figure 8 and Figure 11). To explore this difference further, a Fourier directional model was estimated (using the method described in Ewans and Jonathan [2008]) at this location. Using the 0.25 quantile of the data sample (comprised of a 5×5 pool of neighbouring grid location, see Figure 14) to set the extreme value threshold, the functional forms for extreme value shape and scale parameters shown in Figure 15 were obtained. Subsequent Monte Carlo simulation gave estimates for median re-scaled H_{S100}^* of 28.5, which compares favourably with the spatio-directional estimate at the same location. However, the forms of shape and scale parameter estimates with direction are rather sensitive to choice of extreme value threshold; for larger thresholds, estimates for median re-scaled H_{S100}^* closer to that obtained using a model based on the original (un-whitened) data, ignoring directionality, are obtained.




Figure 15. Estimates for generalised Pareto shape (grey) and scale (black) at location L .

Extensions to finite window natural thin plate splines (e.g. Green and Silverman 1994) would yield spatial solutions for extreme values parameters which are more variable near boundaries. This might be desirable in principle, but does not appear to cause difficulty in the present work. In particular, if there were to be concern about unrealistic smoothness near boundaries of neighbourhoods of the sizes discussed here, then the spatio-directional model could be applied to a wider spatial domain, retaining estimates for extreme value shape and scale corresponding to the original central region of the extended neighbourhood. For applications to bounded regions (e.g. the whole Gulf) we might also expect the variability of extreme value parameter near boundaries to be smaller. Indeed, the model introduced here has also been applied to Gulf-wide estimation for all non-boundary locations (see Section 2). We have estimated the median value of H_{S100} per location for 2658 grid locations in non-coastal regions, obtaining good agreement with location storm peak maxima for the period of the GOMOS hindcast. Illustrations are withheld for reasons of confidentiality. In current work we are considering modelling larger spatial domains in more detail.

Note that the form of the likelihood used here (e.g. l_{ij} and l^* in Section 3) ignores dependence between storm peak events. Per location, this is acceptable. Spatially, however, likelihoods of this form are only appropriate to estimate extremal behaviour marginally per location. Specifically, design criteria from the current analysis would be valid per location only. To incorporate the spatial dependence of events would require more sophisticated forms of joint likelihood (e.g. Davison and Gholamrezaee 2009), or the adoption of an approach similar to that of Heffernan and Tawn [2004] for joint modelling. With these techniques, valid estimates for design criteria for spatial neighbourhoods could be obtained. We are currently examining these approaches.

ACKNOWLEDGMENT

The authors thank Vince Cardone for his insight and guidance, particularly concerning the GOMOS hindcast. We acknowledge useful discussions with Idris Eckley, Jonathan Tawn and Michael Vogel, and thank Joost de Haan for his assistance with data handling. The authors further acknowledge the support of Shell International Exploration and Production and Shell Research Ltd.

REFERENCES

- C.W. Anderson, D.J.T. Carter, and P.D. Cotton. *Wave Climate Variability and Impact on Offshore Design Extremes*. Report commissioned from the University of Sheffield and Satellite Observing Systems for Shell International, 2001.
- API. *Interim guidance on hurricane conditions for the Gulf of Mexico*. API Bulletin, 2007.
- E. Casson and S. G. Coles. Spatial regression models for extremes. *Extremes*, 1:449–468, 1999.
- V. Chavez-Demoulin and A.C. Davison. Generalized additive modelling of sample extremes. *J. Roy. Statist. Soc. Series C: Applied Statistics*, 54:207, 2005.
- L. E. Chouinard. A statistical method for regional design wave heights in the Gulf of Mexico. *Offshore Technology Conference, OTC 6832*, 1992.
- L. E. Chouinard, C. Liu, and C. K. Cooper. Model for severity of hurricanes in Gulf of Mexico. *J. Wtrwy., Harb., and Coast. Engrg.*, May/June:120–129, 1997.
- S. Coles. *An introduction to statistical modelling of extreme values*. Springer, 2001.
- S. G. Coles and E. Casson. Extreme value modelling of hurricane wind speeds. *Structural Safety*, 20:283–296, 1998.
- S. G. Coles and E. A. Powell. Bayesian methods in extreme value modelling: a review and new developments. *International Statistics Review*, 64:119–136, 1996.
- S. G. Coles and J. A. Tawn. A Bayesian analysis of extreme rainfall data. *Applied Statistics*, 45:463–478, 1996.
- S. G. Coles and J. A. Tawn. Bayesian modelling of extreme sea surges on the uk east coast. *Phil. Trans. R. Soc. A*, 363:1387–1406, 2005.
- A. C. Davison. *Statistical models*. Cambridge University Press, 2003.
- A. C. Davison and M. M. Gholamrezaee. Geostatistics of extremes. Preprint. EPFL, `stat.epfl.ch`, 2009.
- A.C. Davison and R. L. Smith. Models for exceedances over high thresholds. *J. R. Statist. Soc. B*, 52:393, 1990.
- J. M. Dixon, J. A. Tawn, and J. M. Vassie. Spatial modelling of extreme sea-levels. *Environmetrics*, 9:283–301, 1998.
- E.F. Eastoe. *Statistical models for dependent and non-stationary extreme events*. Ph.D. Thesis, University of Lancaster, U.K., 2007.
- E.F. Eastoe and J.A. Tawn. Modelling non-stationary extremes with application to surface level ozone. *Appl. Statist.*, 58:22–45, 2009.
- K. C. Ewans and P. Jonathan. The effect of directionality on Northern North Sea extreme wave design criteria. *J. Offshore Mechanics Arctic Engineering*, 130:10, 2008.
- P. J. Green and B.W. Silverman. *Nonparametric regression and generalised linear models: A roughness penalty approach*. Chapman and Hall, London, UK, 1994.
- C. Guedes-Soares and M. Scotto. Modelling uncertainty in long-term predictions of significant wave height. *Ocean Engineering*, 28:329, 2001.
- R. R. Haring and J. C. Heideman. Gulf of Mexico rare wave return periods. *Offshore Technology Conference, OTC3229*, 1978.

- T. Hastie, R. Tibshirani, and J. Friedman. *The elements of statistical learning. Data mining, inference and prediction*. Springer, 2001.
- J. E. Heffernan and J. A. Tawn. A conditional approach for multivariate extreme values. *J. R. Statist. Soc. B*, 66:497, 2004.
- P. Jonathan and K. C. Ewans. The effect of directionality on extreme wave design criteria. *Ocean Engineering*, 34:1977–1994, 2007.
- P. Jonathan and K. C. Ewans. On modelling seasonality of extreme waves. In *Proc. 27th International Conf. on Offshore Mechanics and Arctic Engineering, 4-8 June, Estoril, Portugal, 2008*.
- P. Jonathan, K. C. Ewans, and G. Z. Forristall. Statistical estimation of extreme ocean environments: The requirement for modelling directionality and other covariate effects. *Ocean Engineering*, 35:1211–1225, 2008.
- A. W. Ledford and J. A. Tawn. Modelling dependence within joint tail regions. *J. R. Statist. Soc. B*, 59:475–499, 1997.
- Oceanweather. *GOMOS Gulf of Mexico hindcast*. Oceanweather Inc., 2005.
- M. E. Robinson and J. A. Tawn. Statistics for extreme sea currents. *Applied Statistics*, 46:183–205, 1997.
- M.G. Scotto and C. Guedes-Soares. Modelling the long-term time series of significant wave height with non-linear threshold models. *Coastal Engineering*, 40:313, 2000.
- M.G. Scotto and C. Guedes-Soares. Bayesian inference for long-term prediction of significant wave height. *Coastal Engineering*, 54:393, 2007.
- E. G. Ward, L. E. Borgman, and V. J. Cardone. Statistics of hurricane waves in the Gulf of Mexico. *Offshore Technology Conference*, 1978.

A APPENDIX

This section provides an overview of the natural thin plate spline (NTPS) model. The interested reader is referred to the books of Green and Silverman [1994], Hastie et al. [2001] and Davison [2003] for an introduction to roughness-penalised likelihood models using splines.

A.1 A natural cubic spline generalised Pareto model in one dimension

A natural cubic spline on an interval consists of a sequence of cubic polynomial pieces joined together to form a continuous function, whose first and second derivatives are also continuous on the whole interval. Moreover, the second and third derivatives are zero at the ends of the interval. For *distinct* locations $\{r_i\}_{i=1}^n$ on a straight line, one way of defining a natural cubic spline $f(r)$ is:

$$f(r) = a_1 + a_2 r + \sum_{i=1}^n \delta_i (r - r_i)^3 \quad (1)$$

subject to the constraints $\sum_{i=1}^n \delta_i = \sum_{i=1}^n \delta_i r_i = 0$.

For a sample of data $\{x_i\}_{i=1}^n$ measured at $\{r_i\}_{i=1}^n$, we estimate this model by minimising the roughness-penalised negative log likelihood l^* :

$$\begin{aligned} l^* &= \sum_{i=1}^n l_i^*(\lambda_\gamma, \lambda_\nu) \\ &= \sum_{i=1}^n l_i(r_i) + \frac{\lambda_\gamma}{2} \int \gamma'^2(r) dr + \frac{\lambda_\nu}{2} \int \nu'^2(r) dr \end{aligned}$$

where $l_i(r_i)$ is the generalised Pareto likelihood (see Section 3) and $\{\gamma_i\}_{i=1}^n = \underline{\gamma}$ and $\{\nu_i\}_{i=1}^n = \underline{\nu}$ are the n values of spline coefficients to be estimated for each of γ and ν .

Since we adopt natural cubic splines forms for γ and ν , we can write $\int \gamma'^2(r) dr = \underline{\gamma}' \underline{K} \underline{\gamma}$ and $\int \nu'^2(r) dr = \underline{\nu}' \underline{K} \underline{\nu}$ where \underline{K} is a symmetric matrix whose elements are fixed and easily computed. These quadratic forms for roughness suggest the approach to solution. To minimise l^* , we take derivatives with respect to the elements of $\underline{\gamma}$ and $\underline{\nu}$, yielding the set of score equations:

$$\begin{aligned} \frac{\partial l}{\partial \gamma_i} - \lambda_\gamma \underline{K} \underline{\gamma} &= 0 \\ \frac{\partial l}{\partial \nu_i} - \lambda_\nu \underline{K} \underline{\nu} &= 0 \end{aligned}$$

which can be solved using a procedure based on Taylor expansion, similar to Newton-Raphson, known in the statistics literature as back-fitting. The complexity of the solution scheme is greatly reduced by the adoption of the (γ, ν) parameterisation of the generalised Pareto model, for which $E(\frac{\partial^2 l}{\partial \gamma_i \partial \nu_j}) = 0 \quad \forall i, j$, decoupling the system into separate schemes for $\underline{\gamma}$ and $\underline{\nu}$. When there are multiple events at one or more locations, this scheme is easily modified by inclusion of an incidence matrix.

A.2 A natural thin plate spline generalised Pareto model in two dimensions

A natural thin plate spline in two dimensions is a function $f(\underline{r})$ of $\underline{r} = (r_{(1)}, r_{(2)}) \in \mathbb{R}^2$:

$$f(\underline{r}) = a_0 + a_1 r_{(1)} + a_2 r_{(2)} + \sum_{i=1}^n \delta_i \zeta(\|\underline{r} - \underline{r}_i\|)$$

for *distinct* locations $\{r_i\}_{i=1}^n$ subject to the constraints $\sum_{i=1}^n \delta_i = 0$ and $\sum_{i=1}^n \delta_i r_i = 0$, where the function ζ takes the form $\zeta(z) = \frac{1}{16\pi} z^2 \log_e(z^2)$.

The NTPS is parameterised in terms of $n + 3$ parameters $\{a_j\}_{j=0}^2 = \underline{a}$ and $\{\delta_i\}_{i=1}^n = \underline{d}$. The similarity in form of the NTPS in 2-D and the natural cubic spline in 1-D is clear. If we define the roughness of the NTPS as:

$$R(f) = \int \int \left(\frac{\partial^2 f}{\partial r_{(1)}^2} + \frac{\partial^2 f}{\partial r_{(1)} \partial r_{(2)}} + \frac{\partial^2 f}{\partial r_{(2)}^2} \right) dr_{(1)} dr_{(2)}$$

then the roughness $R(f)$ also takes a simple quadratic form in the parameters $\underline{\delta}$, $R(f) = \underline{\delta}' \underline{E} \underline{\delta}$, where $E_{ik} = \zeta(|r_i - r_k|)$, $\forall i, k = 1, 2, \dots, n$. Analogously to the natural cubic spline generalised Pareto model in 1-D, the NTPS 2-D form is solved by minimising the roughness-penalised likelihood l^* :

$$l^* = \sum_{i=1}^n l_i + \frac{\lambda_\gamma}{2} R_\gamma + \frac{\lambda_\nu}{2} R_\nu$$

to obtain estimates for $\{a_{\gamma j}\}_{j=0}^2 = \underline{a}_\gamma$, $\{\delta_{\gamma i}\}_{i=1}^n = \underline{d}_\gamma$, $\{a_{\nu j}\}_{j=0}^2 = \underline{a}_\nu$ and $\{\delta_{\nu i}\}_{i=1}^n = \underline{d}_\nu$.

We solve the penalised likelihood equation by minimising l^* with respect to the four sets of parameters, following a back-fitting procedure similar to that used for the natural cubic spline model. Once more, the (γ, ν) parameterisation decouples the score equations to simplify the computational scheme, and since there are typically multiple events per location, we introduce an incidence matrix. Fitting the NTPS model in Section 3 requires specification of the extreme value threshold across all locations, itself a spatial variable. In the current work, extreme threshold has been estimated independently per location, but this could clearly be improved if necessary. The NTPS weighting function ζ is also scale-dependent; note that $\zeta(kz) = k^2(\zeta(z) + \zeta(k))$. The scaling of spatial variables therefore must also be taken into consideration, alongside choice of roughness coefficients.

B APPENDIX

This appendix illustrates a sensitivity analysis for estimates of median re-scaled storm peak $H_{\zeta 100}^*$ with respect to varying parameters for whitening and NTPS model estimation, referred to in Section 4.




Figure 16. Contours of median re-scaled H_{S100}^* , estimated using the NTPS model on whitened data. A [0.5,0.9] whitening band was used with moderate directional smoothing, with a 0.75 quantile threshold for model estimation. There is reasonable correspondence with Figure 8




Figure 17. Contours of median re-scaled H_{S100}^* , estimated using the NTPS model on whitened data. A [0.5,0.99] whitening band was used with heavy directional smoothing, with a 0.75 quantile threshold for model estimation. There is reasonable correspondence with Figure 8




Figure 18. Contours of median re-scaled H_{S100}^* , estimated using the NTPS model on whitened data. A [0.5,0.99] whitening band was used with considerable directional smoothing, with a 0.90 quantile threshold for model estimation. There is reasonable correspondence with Figure 8