

Return level estimation from non-stationary spatial data exhibiting multidimensional covariate effects

Philip Jonathan, David Randell, Yanyun Wu

Shell Research Ltd., Manchester, M22 0RR, UK.

Kevin Ewans

Sarawak Shell Bhd., 50450 Kuala Lumpur, Malaysia.

Abstract

Careful modelling of non-stationarity is critical to reliable specification of marine and coastal design criteria. We present a spline based methodology to incorporate spatial, directional, temporal and other covariate effects in extreme value models for environmental variables such as storm severity. For storm peak significant wave height events, the approach uses quantile regression to estimate a suitable extremal threshold, a Poisson process model for the rate of occurrence of threshold exceedances, and a generalised Pareto model for size of threshold exceedances. Multidimensional covariate effects are incorporated at each stage using penalised (tensor products of) B-splines to give smooth model parameter variation as a function of multiple covariates. Optimal smoothing penalties are selected using cross-validation, and model uncertainty is quantified using a bootstrap re-sampling procedure. The method is applied to estimate return values for large spatial neighbourhoods of locations, incorporating spatial and directional effects. Extensions to joint modelling of multivariate extremes, incorporating extremal spatial dependence (using max-stable processes) or more general extremal dependence (using the conditional extremes approach) are outlined.

Keywords: offshore design; extreme; return value; covariate; B-spline; storm severity;

1. Introduction

Availability of comprehensive met-ocean data allows the effect of heterogeneity (or non-stationarity) of extremes with respect to direction, season and location to be accommodated in estimation of design criteria. Jonathan and Ewans (2013) review statistical modelling of extremes for marine design.

Capturing covariate effects in extreme sea states is important when developing design criteria. In previous work (e.g Jonathan and Ewans (2007a), Ewans and Jonathan (2008)) it has been shown that omnidirectional design criteria derived from a model that adequately incorporates directional covariate effects can be materially different from a model which ignores those effects (e.g. Jonathan et al. 2008). Directional return

*philip.jonathan@shell.com

values derived from a directional model can be heavier tailed than that derived from a direction-independent approach, indicating that large values of extreme events 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 et al. 2008).

There is a large body of statistical literature regarding modelling of covariate effects in extreme value analysis; for example, 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. Jonathan and Ewans (2007b) use block bootstrapping to evaluate uncertainties associated with extremes in storm peak significant wave heights in the Gulf of Mexico. 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 a sample of extremes prior to model estimation.

One of the first examinations of spatial characteristics of extreme wave heights was reported by Haring and Heideman (1978) for the Gulf of Mexico. 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 no 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. Jonathan and Ewans (2011b) used thin-plate splines to model the spatial characteristics of events in the Gulf of Mexico. Extending the thin-plate spline formalism to include other (possibly periodic) covariates is difficult; instead, the sample is typically pre-processed to remove the influence of all covariates other than the (2-D) spatial, prior to model estimation using thin-plate splines. Models estimated in this way suffer from the fact that interactions between the various modelling steps (and the parameters estimated therein) cannot be easily quantified.

Characterising the joint structure of extremes for different environmental variables is also important for improved understanding of those environments. Yet many applications of multivariate extreme value analysis adopt models that assume a particular form of extremal dependence between variables without justification, or restrict attention to regions in which all variables are extreme. The conditional extremes model of Heffernan and Tawn (2004) provides one approach avoiding these particular restrictions. Extremal dependence characteristics of environmental variables also typically vary with covariates. Reliable descriptions of extreme environments should also therefore characterise any non-stationarity. Jonathan et al. (2013) extends the conditional extremes model of Heffernan and Tawn to include covariate effects, using Fourier representations of model parameters for single periodic covariates.

The last decade has seen the emergence of useable statistical models for spatial extremes based on max-stable processes, at least in academia. The application of max-stable processes is complicated due to unavailability of the full multivariate density function. Padoan et al. (2010) develops inferentially practical, likelihood-based methods for fitting max-stable processes derived from a composite likelihood approach.

The procedure is sufficiently reliable and versatile to permit the simultaneous modelling of marginal and dependence parameters in the spatial context at a moderate computational cost. Davison and Gholamrezaee (2012) describes an approach to flexible modelling for maxima observed at sites in a spatial domain, based on fitting of max-stable processes derived from underlying Gaussian random process models. Generalised extreme value (GEV) margins as assumed throughout the spatial domain, and models incorporate standard geo-statistical correlation functions. Estimation and fitting are performed through composite likelihood inference applied to observations from pairs of sites. Davison et al. (2012) also provides a good introduction and review. Erhardt and Smith (2011) uses approximate Bayesian computation to circumvent the need for a joint likelihood function by instead relying on simulations from the (unavailable) likelihood avoiding the need to construct composite likelihoods at higher computational cost.

In this work, we apply a marginal model for spatio-directional extremes to a sample of data for storm severity on the north west continental shelf of Western Australia. The model (developed in Section 2) adopts a penalised B-spline formulation to characterise smooth variation of extreme value parameters spatially and directionally. The North West Shelf application is then presented in Section 3. In Section 4, we discuss model extension to incorporate appropriate spatial extremal dependence, and also outline a non-stationary extension of the conditional extremes model of Heffernan and Tawn (2004).

2. Model

The objective is to estimate design criteria for individual locations within a spatial neighbourhood, accounting for spatial and storm directional variability of extremal characteristics.

Model components

Following the work of Jonathan and Ewans (2008) and Jonathan and Ewans (2011b), summarised in Jonathan and Ewans (2013), we model storm peak significant wave height H_S^{sp} , namely the largest value of significant wave height H_S observed at each location during the period of a storm event. We assume that each of \dot{n}_S independent storm peak events is observed at all of n_L locations within the neighbourhood under consideration. We therefore start with a total of $\dot{n} = \dot{n}_S \times n_L$ observations of H_S^{sp} . We refer to these as the sample $\{\dot{z}_i\}_{i=1}^{\dot{n}}$ of \dot{n} storm peak significant wave heights observed at locations $\{\dot{x}_i, \dot{y}_i\}_{i=1}^{\dot{n}}$ with dominant wave directions $\{\dot{\theta}_i\}_{i=1}^{\dot{n}}$ (corresponding to the time of occurrence of H_S^{sp} , henceforth “storm directions”). We then proceed using the peaks over threshold approach as follows.

Extreme value threshold: We first estimate a threshold function ϕ above which observations \dot{z} are assumed to be extreme. The threshold varies smoothly as a function of covariates ($\phi \triangleq \phi(\theta, x, y)$) and is estimated using quantile regression. As the result of thresholding, the same number of observations of smaller storm peak events is eliminated from all locations by construction. The number of storms peak events remaining per location reduces from \dot{n}_S to n_S , so that the total number of observations (of threshold exceedances) for extreme value modelling is $n = n_S \times n_L$. We refer to these as the set of n threshold exceedances $\{z_i\}_{i=1}^n$ observed at locations $\{x_i, y_i\}_{i=1}^n$ with storm peak directions $\{\theta_i\}_{i=1}^n$.

Rate of threshold exceedance: We next estimate the rate of occurrence ρ of threshold exceedance using a Poisson process model with Poisson rate $\rho(\triangleq \rho(\theta, x, y))$.

Size of threshold exceedance: We estimate the size of occurrence of threshold exceedance using a generalised Pareto (henceforth GP) model. The GP shape and scale parameters ξ and σ are also assumed to vary smoothly as functions of covariates.

This approach to extreme value modelling follows that of Chavez-Demoulin and Davison (2005) and is equivalent to direct estimation of a non-homogeneous Poisson point process model (e.g., Dixon et al. 1998,

Jonathan and Ewans (2013)).

Parameter estimation

Extreme value threshold: For quantile regression, we seek a smooth function ϕ of covariates corresponding to non-exceedance probability τ of H_S^{sp} given any combination of θ, x, y . We choose to estimate ϕ by minimising the quantile regression lack of fit criterion

$$\ell_\phi = \left\{ \tau \sum_{i, r_i \geq 0}^n |r_i| + (1 - \tau) \sum_{i, r_i < 0}^n |r_i| \right\}$$

for residuals $r_i = z_i - \phi(\theta_i, x_i, y_i; \tau)$. We regulate the smoothness of the quantile function by penalising lack of fit for parameter roughness R_ϕ (with respect to all covariates), by minimising the revised penalised criterion

$$\ell_\phi^* = \ell_\phi + \lambda_\phi R_\phi$$

where the value of roughness coefficient λ_ϕ is selected using cross-validation to provide good predictive performance.

Rate of threshold exceedance: For Poisson modelling, we use penalised likelihood estimation. The rate ρ of threshold exceedance is estimated by minimising the roughness-penalised (negative log) likelihood

$$\ell_\rho^* = \ell_\rho + \lambda_\rho R_\rho$$

where R_ρ is parameter roughness with respect to all covariates, λ_ρ is again evaluated using cross-validation. The Poisson (negative log) likelihood is given by

$$\ell_\rho = - \sum_{i=1}^n \log \rho(\theta_i, x_i, y_i) + \int \rho(\theta, x, y) d\theta dx dy$$

which is approximated by

$$\hat{\ell}_\rho = - \sum_{j=1}^m c_j \log \rho(j\Delta) + \Delta \sum_{j=1}^m \rho(j\Delta)$$

where $\{c_j\}_{j=1}^m$ are counts of numbers of threshold exceedances per degree longitude, latitude and storm direction, per annum on an index set of m ($\gg 1$) bins on a regular lattice partitioning the covariate domain into intervals of constant volume Δ .

Size of threshold exceedance: The generalised Pareto model of size of threshold exceedance is estimated in a similar manner by minimising the roughness penalised (negative log) GP likelihood

$$\ell_{\xi, \sigma}^* = \ell_{\xi, \sigma} + \lambda_\xi R_\xi + \lambda_\sigma R_\sigma$$

where R_ξ and R_σ are parameter roughnesses for GP shape and scale with respect to all covariates, and roughness coefficients λ_ξ and λ_σ are evaluated using cross-validation. The GP (negative log) likelihood is given by

$$\ell_{\xi, \sigma} = \sum_{i=1}^n \log \sigma_i + \left(\frac{1}{\xi_i} + 1 \right) \log \left(1 + \frac{\xi_i}{\sigma_i} (z_i - \phi_i) \right)$$

where $\phi_i = \phi(\theta_i, x_i, y_i)$, $\xi_i = \xi(\theta_i, x_i, y_i)$ and $\sigma_i = \sigma(\theta_i, x_i, y_i)$, and a similar expression is used when $\xi_i = 0$ (see Jonathan and Ewans 2013).

Penalised likelihood optimisation for both Poisson and generalised Pareto inference is performed using the so-called back-fitting algorithm exploiting problem-specific structure, discussed in outline in the Appendix.

Return values

Simulation under the fitted model is used to estimate directional design values per location, optionally incorporating the effects of storm directional dissipation. In outline, we simulate a large number R (≈ 1000) of realisations, each realisation corresponding to a simulated set of storms for return period T of interest, for all locations in the spatial domain. In each realisation, we simply store the maximum value observed for directional sectors of interest per location. Over all R realisations, we accumulate R estimates of directional and omnidirectional extremes per location, with which we can estimate distributional statistics of the return value, including the median return value, the central 95% uncertainty band, or the full cumulative distribution function of return value. In the current work, we report omnidirectional design values, and design values for directional octants of equal size, centred on cardinal and inter-cardinal storm directions (see Section 3), estimated from simulation. The simulation procedure is given in Algorithm 1 below.

```

input:  $\phi$ ,  $\rho$ ,  $\xi$  and  $\sigma$  defined on index set of covariate combinations;  $\rho$  in units of number of
occurrences per location per 1-degree storm direction per annum; return period of interest,
 $T$ ; number of realisations,  $R$ , each corresponding to  $T$  years of events
output: per location: median and 95% uncertainty for return value omnidirectionally and for 8
directional sectors of equal size centred on cardinal and semi-cardinal directions
foreach fixed location  $(x, y)$  in turn
    estimate the expected total (omnidirectional) number  $S(x, y)$  of threshold exceedances in return
period (using  $S(x, y) = T \int \rho(\theta, x, y) d\theta$ );
    foreach realisation of  $T$  years in turn
        sample an actual number of threshold exceedances for return period (at random from a
Poisson distribution with mean  $S(x, y)$ );
        sample storm directions for threshold exceedances (from  $\rho(\theta, x, y)$ );
        sample sizes for threshold exceedances (from the GP distribution with parameters  $\xi(\theta, x, y)$ ,
 $\sigma(\theta, x, y)$  and  $\phi(\theta, x, y)$  for the previously-sampled values of storm directions  $\theta$ );
        foreach directional sector in turn
            accumulate maximum value observed, incorporating dissipation if required (see next
section on “Directional dissipation”);
        end
    end
    store return value statistics for the location;
end

```

Algorithm 1: Procedure for estimation of return values by simulation

Directional dissipation

In this work, for extreme value analysis, we characterize a storm in terms of its storm peak significant wave height, storm peak direction and location only. However, the full temporal evolution of a storm at a location is more fully described as a time-series of values of significant wave height and corresponding wave direction for consecutive sea states in the storm. In estimating a return value for arbitrary directional sector S , it is critical to account for the effects of all storms on S , including (1) storms whose storm peak

directions correspond to S , and (2) storms whose storm peak directions do not correspond to S , but whose storm evolution includes sea states with wave directions in S . To achieve this, we use the concept of storm dissipation, introduced by Jonathan and Ewans (2007a). The directional dissipation $\zeta(\theta)$ of a storm, for direction θ , is the value of significant wave height, expressed as a fraction of storm peak significant wave height, for that θ . For a given directional sector S , therefore, the directional dissipation $\zeta(S)$ of a storm is the largest impact of the storm in S , expressed as a fraction of the storm peak significant wave height. Directional dissipation captures the effects on a directional sector of storms whose peaks lie outside the sector. Incorporating directional dissipation increases return values in general, for all directional sectors other than the omnidirectional (which already includes the effects of all storms).

Directional dissipation is estimated directly from the directional evolution of significant wave height (or directional trajectory) per storm in the hindcast sample. During estimation of design values using simulation, each simulated storm peak significant wave height and storm direction is matched to a hindcast directional trajectory with similar storm peak characteristics. The simulated storm peak event is then assumed to exhibit the same directional dissipation as the matching hindcast storm event. The influence of the simulated storm event on arbitrary directional sectors is thereby included in the simulation.

Parameter smoothness

Physical considerations suggest we should consider parameters ϕ, ρ, ξ and σ to be smooth functions of covariates θ, x, y . For estimation, this can be achieved by expressing each parameter in terms of an appropriate basis for the domain D of covariates, where $D = D_\theta \times D_x \times D_y$. $D_\theta = [0, 360)$ is the (marginal) domain of storm peak directions, and D_x, D_y are the domains of x- and y-values (e.g. longitudes and latitudes) under consideration.

For each covariate (and marginal domain) in turn, we first calculate a B-splines basis matrix for an index set (of size $m \ll n$) of covariate values; potentially we could calculate the basis matrix for each of the n observations, but usually avoid this for computation efficiency. Specifically, for D_θ , we calculate basis matrix B_θ ($m_\theta \times p_\theta$) such that the value of any function at each of the m_θ points in the index set for storm direction can be expressed as $B_\theta \beta_\theta$ for some vector β_θ ($p_\theta \times 1$) of basis coefficients. Note that periodic marginal bases can be specified if appropriate (e.g. for D_θ).

Then we define a basis matrix for the three-dimensional spatio-directional domain D using Kronecker products of marginal basis matrices. Thus

$$B = B_\theta \otimes B_x \otimes B_y$$

provides a $(m \times p)$ basis matrix (where $m = m_\theta m_x m_y$, and $p = p_\theta p_x p_y$) for modelling each of ϕ, ρ, ξ and σ on the corresponding ‘‘spatio-directional’’ index set (of size m). Any of ϕ, ρ, ξ and σ (η , say, for brevity) can then be expressed in the form $\eta = B\beta$ for some $(p \times 1)$ vector β of basis coefficients. Model estimation therefore reduces to estimating appropriate sets of basis coefficients for each of ϕ, ρ, ξ and σ . The values of p_θ, p_x, p_y are functions of the number of spline knots for each marginal domain, and also depend on whether spline bases are specified to be periodic (e.g. D_θ) or not (e.g. D_x and D_y).

The roughness R of any function can be easily evaluated on the index set (at which $\eta = B\beta$). Following the approach of Eilers and Marx (e.g. Eilers and Marx 2010), writing the vector of differences of consecutive values of β as $\Delta\beta$, and vectors of second and higher order differences using $\Delta^k\beta = \Delta(\Delta^{k-1}\beta)$, $k > 1$, the roughness R of β is given by

$$R = \beta' P \beta$$

where $P = (\Delta^k)'(\Delta^k)$ for differences of order k . We use $k = 1$ throughout this work. With this choice of k ,

heavy roughness penalisation results in stationarity of parameters with respect to periodic and aperiodic covariates.

For readers less familiar with spline modelling, we recommend the excellent motivation and summary of Eilers and Marx (2010), which provides key references to the modelling literature, including Marx and Eilers (1998), Ruppert et al. (2003), Currie et al. (2006) and Eilers et al. (2006). Green and Silverman (1994) and Hastie et al. (2001) provide statistical motivations. de Boor (2001) provides technical background to splines.

There are many possible parameterisations of covariate effects potentially suitable for applications such as the current. For a single periodic covariate (e.g. direction), Fourier forms for model parameters might be considered. Polynomial or Legendre polynomial functions in longitude and latitude might be useful for spatially-varying parameters. However, we have adopted B-splines for two specific reasons. Firstly, B-splines (in marked contrast to Fourier) have local support facilitating efficient and stable estimation of spline coefficients. Secondly, spline models generally are easily and consistently extendible to describe multidimensional covariates using Kronecker products of one-dimensional spline bases. There are many competing spline models that may be considered (see, e.g., Lee and Durban 2011, Lee et al. 2013, Stephenson and Gilleland 2006 and VGAM software). A random field representation for covariate effects is a particularly attractive alternative to the spline formulation due to the algorithmic similarities of the two approaches.

Computational considerations

Quantile regression estimation is performed by direct minimisation of the criterion ℓ_ϕ^* from a good starting solution using a linear programming approach (e.g. Koenker 2005 and Bollaerts 2009). A starting solution is estimated by fitting a smoothing spline to estimates of the spatio-directional quantile with non-exceedance probability τ at each of the m covariate combinations on the index set. Poisson and generalised Pareto estimation is achieved using iterative back-fitting (e.g., Davison 2003). Good starting solutions are found to be essential for GP minimisation in particular. These are obtained by estimating local GP models at each of the m members of the index set (or combinations of neighbours thereof to increase sample size), then fitting smoothing spline models for each of GP shape ξ and scale σ . Using algorithms developed for generalised linear array models (Currie et al. 2006), direct computation of Kronecker products of the form $B_\theta \otimes B_x \otimes B_y$ (and some other computationally-demanding operations) can be avoided, providing large reductions in computer memory requirements and execution times. For larger problems, it is also computationally advantageous to adopt an appropriate model cost-complexity criterion (such as AIC) as an alternative to cross-validation, thereby avoiding the need for repeated model estimation. We refer to Chavez-Demoulin (1999) and Chavez-Demoulin and Davison (2005) who note the need for orthogonality of parameters when criteria such as AIC as used for multi-parameter models such as the generalised Pareto.

3. Application

Data and regional climatology

The application sample corresponds to storm peak significant wave height H_S^{sp} and dominant wave direction (at H_S^{sp}) for 6156 hindcast storm events at 1089 locations on a 33 x 33 regular grid over the north west continental shelf of Western Australia for the period 1970-2007. Note that storm peak events H_S^{sp} per location can be assumed to be independent in time, since they are isolated from time-series of temporally-dependent significant wave height H_S using a physically-motivated de-clustering procedure described by Jonathan and Ewans (2011a).

The regional climate is monsoonal, displaying two distinct seasons, “winter” from April to September and “summer” from October to March, with very rapid transitions between these, generally in April and September/October. The winter “dry” is the result of a steady air flow from the east (South East Trade Winds) originating from the Australian mainland, propagating over the Timor Sea. The summer “wet” is the result of the North West Monsoon, a steady, moist flow of air predominantly from the west to south west (and to a lesser extent from the north west). Tropical cyclones occur during summer months and are clearly the most important for extreme met-ocean criteria. Tropical cyclones originate from south of the equator in the eastern Indian Ocean and Timor and Arafura Seas. The most severe cyclones often occur in December and March-April, when sea surface temperatures are highest. In the region under consideration, most storms are tropical lows or developing storms, but can be very severe nevertheless, as exemplified by tropical cyclones Thelma (1998) and Neville (1992). Most storms passing through the region head in a westerly or south westerly direction before turning southwards.

[Figure 1 about here.]

[Figure 2 about here.]

The prevailing wave climate comprises contributions from Indian Ocean swell, winter easterly swell, westerly monsoonal swell, tropical cyclone swell, and locally generated wind-sea. Indian Ocean swell is a perennial feature typically, propagating from the south-west through north-west. The largest sea states are wind generated sea states associated with tropical cyclones. Figure 1 shows a scatter plot of H_S^{sp} against direction, pooled over locations; H_S^{sp} is non-stationary with respect to storm direction. There are a large number of small events from $260 - 310^\circ$, seen also in the rose plots of Figure 2; the rate of occurrence of large H_S^{sp} is also non-stationary with respect to storm direction. Figure 3 shows the spatial distribution of storm direction and storm peak significant wave height for the largest two storms in the hindcast. Both events are seen to be spatially localised, and to show clockwise rotation of wind field as expected. Figure 4 shows the directional evolution of significant wave height at the central location. At this location, the largest storm events emanate from the eastern and south-eastern octants, corresponding to cyclones. However, by far the largest number of events emanate from the western octant, associated with Indian Ocean Swell. Cyclone events (e.g. eastern octant) tend to dissipate over a wide angular range, whereas swell is directionally more restricted.

[Figure 3 about here.]

[Figure 4 about here.]

Spline parameterisation

We fit the spatio-directional spline model assuming an index set of 32 directional bins x 33 longitude bins x 33 latitude bins. Cubic splines are used for each dimension, the directional spline basis being the only periodic. Longitude and latitude domains are characterised by 15 evenly spaced knots, and the directional domain by 32 evenly spaced knots. Therefore, a total of $32 \times 18 \times 18$ spline parameters need to be estimated.

Roughness parameter estimation

Roughness parameter estimation is necessary for estimation of each of extreme value threshold, rate of occurrence of threshold exceedance, shape and scale parameters of GP-distributed sizes of threshold exceedance. In each case, the roughness parameter is estimated using 10-fold cross-validation over a domain

of plausible values of the roughness parameter, within which a clear minimum of the predicted lack of fit criterion (for the withheld data) is observed. Observations of storm peak events are temporally independent but dependent spatially. Here we are concerned only with marginal extreme value analysis per spatial location, and do not attempt to model extremal spatial dependence explicitly. As discussed in Chavez-Demoulin (1999) and Chavez-Demoulin and Davison (2005), the effect of using spatially dependent data to model marginal spatial extremes is small in terms of the estimated maximum likelihood parameters. However, parameter uncertainty is underestimated. As discussed in Wang and Wahba (1995) and Chavez-Demoulin and Davison (2005), block bootstrapping schemes are recommended to estimate model uncertainty in this situation. Sandwich estimators (see, e.g. Chandler and Bate (2007) and Northrop and Jonathan (2011)) are also available.

Extreme value threshold

We estimate threshold ϕ , using spline quantile regression, above which observations z of H_S are assumed to be extreme. A number of different quantile non-exceedance thresholds were examined; a 50% threshold was adopted. The optimal threshold, estimated using 10-fold block cross-validation, was found to be effectively constant with respect to location and direction. This is not surprising given the small sample size per location, and the strong spatial dependence across locations. A constant threshold of 3.4m was therefore adopted for subsequent analysis.

Rate of threshold exceedance

Next we estimate the rate of occurrence ρ of threshold exceedances using a Poisson process model with rate $\rho(\triangleq \rho(\theta, x, y))$. Figure 5 shows a spatio-directional plot of estimated ρ . The left hand plot shows the direction from which the estimate is largest per location. The 8 right-hand plots show estimates spatially for the 8 cardinal and inter-cardinal directions. Rate ρ is seen to be relatively constant spatially, but directionally ρ is higher for events from the west (250° - 290°), consistent with inferences from Figures 1, 2 and 4. That is, the rate of occurrence of storm peak events is dominated by Indian Ocean Swell. Note that the right-hand plots in Figures 5–7 show spatial variation for events from 8 specific semi-cardinal directions (as opposed to directional octants). Contrast this with the right-hand plots in Figures 10 and 11, and the plots in Figure 4, which present estimates pertaining to directional octants (as opposed to specific directions).

[Figure 5 about here.]

Size of threshold exceedance

Now we model the sizes of threshold exceedance using a GP model, again estimated using 10-fold block cross-validation. Figure 6 shows a spatio-directional plot for GP shape parameter, ξ , seen to be largest in general for events from the north - north east and south - south west. Figure 7 shows the corresponding plot for GP scale, σ , seen to be higher in general for events with directions from the east - south east. Because of the dependence between estimates of ξ and σ , care should be taken not to over-interpret these plots; inspection of estimates for return values (below) is preferred. Generalised Pareto maximum likelihood estimation with covariate-dependent parameters is computationally challenging. A back-fitting algorithm (see Appendix) is used to minimise the negative log-likelihood as efficiently as possible. A good starting solution also greatly aids efficient minimisation. The procedure used here to obtain a reasonable starting solution is outlined in Section 2 (Computational considerations).

[Figure 6 about here.]

[Figure 7 about here.]

Model validation

Model diagnostics are essential to demonstrate adequate model fit. Of primary concern is that the estimated H_S^{sp} extreme value model generates directional distributions of H_S^{sp} consistent with observed data at all spatial locations. To quantify this, we use the simulation procedure to generate 1000 realisations of storm peak events for all locations, each realisation for the same period as the original data. We then construct 95% uncertainty bands for cumulative distribution functions (cdfs) of H_S^{sp} , corresponding to different partitions of the spatio-directional covariate domain. We can then confirm that the actual sample lies within the 95% uncertainty band for each covariate partition. We illustrate this using two sets of partitions in Figures 8 and 9. Figure 8 compares omnidirectional cdfs for 9 locations on a 3×3 grid of locations across the domain, and indicates good agreement. Figure 9 compares directional octant return values aggregated over the 9 locations defined in Figure 8. Again, agreement is good.

[Figure 8 about here.]

[Figure 9 about here.]

Return value estimation

The characteristics of 100-year return value were estimated by simulating 1000 realisations of 100-years of storm events across the whole spatial neighbourhood (with or without directional dissipation effects). Figure 10 shows median 100-year storm peak return values H_{S100}^{sp} estimated by simulation under the fitted model, ignoring directional dissipation. Both omnidirectional return values and return values corresponding to directional sectors centred on the cardinal and inter-cardinal directions are given. Figure 11 shows analogous median 100-year return values H_{S100} incorporating directional dissipation (which plays no role in omnidirectional estimation). Generally, lower return values are seen in the south east corner nearest land. The highest storm peak return values at most locations are generally from the east. Return values are generally lower for storms from the south west. Omnidirectional return values are not influenced by the incorporation of directional dissipation. However, return values for most directional sectors are larger once directional dissipation is incorporated.

[Figure 10 about here.]

[Figure 11 about here.]

4. Discussion

In the paper, we introduce a marginal spatio-directional model for extreme storm peak significant wave height, applied to estimation of spatio-directional design values for a neighbourhood of locations off the North West Shelf of Australia. The model uses the peaks over threshold approach, incorporating estimation of an extreme value threshold and the rate and sizes of threshold exceedance. Model parameters are smooth spatio-directional functions. Block cross-validation is used to estimate appropriate parameter smoothness in each case; model cost-complexity measures (e.g. AIC) are also available, but must be used with care in the presence of dependence or non-orthogonal model parameterisations. Re-sampling techniques such as bootstrapping can be used to estimate the uncertainty of model parameters and estimates of return values and other structure variables. The model yields parameter estimates and design values consistent with physical intuition and previous estimates. The main advantage of the approach is that marginal spatial and directional variation of extreme value characteristics is incorporated in a rational, consistent, scalable

and computationally-efficient manner eliminating the need for ad-hoc procedures such as site pooling. In isolating storm peak events, we also estimate the directional dissipation (e.g. Jonathan and Ewans 2007a) of storms across locations. This allows us also to estimate design criteria for arbitrary directional sectors for a given location together with the omnidirectional estimate, in a consistent manner. For interested parties, further specifics of the modelling procedure are available from the authors on request.

Estimates of median 100-year return values were also estimated by pooling all locations, assuming spatial independence of observations and ignoring all directional and spatial effects. Figure 12 shows omnidirectional estimates, and estimates for storms from storm directions 90° - 135° , as a function of the non-exceedance probability of the constant extreme value threshold used. Omnidirectional estimates are low, but estimates for the directional sector 90° - 135° , from which most severe storms emanate, are broadly consistent with the more severe covariate combinations in the spatio-directional model.

[Figure 12 about here.]

Spline representations are also useful in non-stationary conditional extremes modelling based on the approach of Heffernan and Tawn (2004). Jonathan et al. (2014) introduces a general-purpose approach, common to all inference steps in conditional extremes inference. Non-crossing quantile regression estimates appropriate non-stationary marginal quantiles simultaneously (for a range of non-exceedance probabilities) as functions of covariate; these are necessary as thresholds for extreme value modelling, and for standardisation of marginal distributions prior to application of the conditional extremes model. Marginal extreme value and conditional extremes modelling is performed within a roughness-penalised likelihood framework, with cross-validation to estimate suitable model parameter roughness. A bootstrap re-sampling procedure, encompassing all inferences, quantifies uncertainties in, and dependence structure of, parameter estimates and estimates of conditional extremes of one variate given large values of another. The approach is validated using simulations from known joint distributions, the extremal dependence structures of which change with covariate. The approach is illustrated in application to joint modelling of storm peak significant wave height and associated storm peak period for extra-tropical storms at northern North Sea and South Atlantic Ocean locations, with storm direction as covariate.

The marginal model presented here can also be extended to incorporate multivariate spatial dependence, using composite likelihood and censored likelihood methods, so that joint characteristics of extremes of storm peak significant wave height across multiple locations can also be estimated and studied. At present, however, spatial extremes methods suffer from a number of restrictions which may cause biased inferences. Firstly, models are typically developed using block maxima rather than threshold exceedances, since spatial extremes theory is motivated by consideration of component-wise maxima, with full likelihoods replaced with approximations constructed as weighted sums of pairwise likelihoods. Censored composite likelihood approaches permit more efficient analysis of threshold exceedances (e.g. Huser and Davison 2014). Secondly, spatial extremes models only admit certain types of extremal dependence structure (e.g. Ledford and Tawn (1997)), namely perfect independence or asymptotic dependence; they do not admit asymptotic independence in particular. An extension to incorporate asymptotic independence has recently been proposed by Wadsworth and Tawn (2012). The generalised Pareto process is another promising emerging description for spatial extremes (Ferreira and de Haan 2014). In general, the extremal dependence structure of spatial extremes will itself be non-stationary; estimation of covariate effects in the dependence model will therefore also be necessary, using the spline approach outlined above. For example, the dependence structure of the simplest max-stable process for two spatial locations, known as the Smith process (Smith 1990), is parameterised in terms of a bivariate Gaussian covariance matrix (e.g. Jonathan and Ewans 2013), therefore requiring the estimation of three parameters (two variances and one covariance), all of which in principle may vary spatio-directionally. Other more realistic max-stable processes typically have larger numbers of dependence parameters (e.g., Davison et al. 2012). Spatial extremes methods are

nevertheless potentially of great value in met-ocean design, since they provide a framework within which extremal behaviour of complete ocean basins can be modelled, incorporating appropriate marginal and dependence structure and avoiding the need for site pooling in particular. It might be possible that in future, only one extreme modelling task would be necessary per hindcast. That model, for the whole ocean basin, could then be interrogated routinely to estimate design values for one location, or joint design values for an arbitrary number of arbitrary locations. The ocean engineer would then no longer in principle need to perform further site-specific extreme value analysis.

5. Acknowledgement

The authors gratefully acknowledge discussions with Graham Feld and Laks Raghupathi from Shell, contributions from Emma Ross and Elena Zanini of Lancaster University, UK, and thoughtful comments from two reviewers.

Appendix: Back-fitting for Poisson and generalised Pareto inference

A back-fitting algorithm (see, e.g., Davison (2003)) is used for optimisation of penalised (negative log) likelihoods for both Poisson and generalised Pareto model inference. Here, we outline the algorithm used for the generalised Pareto case, since it is somewhat more involved. Poisson inference proceeds along similar lines.

On the index set of m combinations of values of spatio-directional covariates, we express the generalised Pareto negative log sample likelihood ℓ in terms of asymptotically independent parameter pairs $\{\xi_i, \nu_i\}_{i=1}^m$ (where $\nu_i = \sigma_i(1 + \xi_i)$, $i = 1, 2, \dots, m$). B-spline parameterisations are $\xi = B\beta_\xi$, $\nu = B\beta_\nu$, for the $m \times 1$ vectors ξ and ν , where B is an appropriate (spatio-directional) B-spline basis matrix, and β_ξ and β_ν are $p \times 1$ parameter vectors to be estimated. The penalised (negative log) sample likelihood ℓ^* is then

$$\ell^* = \ell + \lambda_{\beta_\xi} \beta_\xi^t P \beta_\xi + \lambda_{\beta_\nu} \beta_\nu^t P \beta_\nu$$

where P is a symmetric (first-order) difference penalty matrix on the B-spline coefficients, and λ_{β_ξ} and λ_{β_ν} are roughness penalty coefficients estimated by cross-validation.

Back-fitting proceeds by applying a Newton-Raphson scheme to the partial derivatives of the penalised likelihood with respect to each of the $2p$ parameters $\{\beta_{\xi j}, \beta_{\nu j}\}_{j=1}^p$, known as penalised likelihood scores. For each element $\beta_{\xi j}$, $j = 1, 2, \dots, p$ of the generalised Pareto shape parameter vector β_ξ , with out loss of generality, a Taylor expansion of penalised likelihood score $u_{\beta_{\xi j}}^*$ about point $(\beta_\xi^o, \beta_\nu^o)$ yields

$$u_{\beta_{\xi j}}^* \approx u_{\beta_{\xi j}}^* \Big|_o + \sum_{l=1}^p \frac{\partial u_{\beta_{\xi j}}^*}{\partial \beta_{\xi l}} \Big|_o (\beta_{\xi l} - \beta_{\xi l}^o) + \sum_{l=1}^p \frac{\partial u_{\beta_{\xi j}}^*}{\partial \beta_{\nu l}} \Big|_o (\beta_{\nu l} - \beta_{\nu l}^o) \text{ for } j = 1, 2, \dots, p$$

where $|_o$ indicates quantities evaluated at $(\beta_\xi^o, \beta_\nu^o)$. Since the generalised Pareto parameters are asymptotically independent, we ignore the third term on the right hand side as explained below, greatly simplifying computation. If we set the left-hand-side to zero for each j (corresponding to minimum negative log likelihood), we can use the right-hand-side to iterate from starting solution $(\beta_\xi^o, \beta_\nu^o)$ to the minimum. This is the Newton-Raphson algorithm.

Penalised likelihood scores are written in terms of the penalised B-spline representation

$$\begin{aligned}
u_{\beta_{\xi j}}^* &= \frac{\partial \ell^*}{\partial \beta_{\xi j}} \\
&= \sum_{k=1}^m \frac{\partial \ell}{\partial \xi_k} \frac{\partial \xi_k}{\partial \beta_{\xi j}} + \sum_{k=1}^m \frac{\partial \ell}{\partial \nu_k} \frac{\partial \nu_k}{\partial \beta_{\xi j}} + \lambda_{\beta_{\xi j}} \{P\beta_{\xi j}\}_j \\
&= \{B'u_{\xi} + \lambda_{\beta_{\xi}} P\beta_{\xi}\}_j, \quad j = 1, 2, \dots, p
\end{aligned}$$

where u_{ξ} is a $m \times 1$ vector with elements $\frac{\partial \ell}{\partial \xi_k}$, $k = 1, 2, \dots, m$. The second term on the right-hand-side of the second line is identically zero, since $\frac{\partial \nu_k}{\partial \beta_{\xi j}} = 0$ always. In matrix notation

$$u_{\beta_{\xi}}^* = B'u_{\xi} + \lambda_{\beta_{\xi}} P\beta_{\xi}$$

An analogous expression is easily derived for scores $u_{\beta_{\nu}}^*$. To use the Newton-Raphson algorithm, we further need expressions for the (elements of the) partial derivatives of $u_{\beta_{\xi}}^*$ and $u_{\beta_{\nu}}^*$ with respect to (the elements of) β_{ξ} and β_{ν} . For example

$$\begin{aligned}
\frac{\partial u_{\beta_{\xi j}}^*}{\partial \beta_{\xi l}} &= \frac{\partial^2 \ell^*}{\partial \beta_{\xi j} \partial \beta_{\xi l}} \\
&= \frac{\partial}{\partial \beta_{\xi l}} (u_{\beta_{\xi j}}^*) \\
&= \frac{\partial}{\partial \beta_{\xi l}} (\{B'u_{\xi} + \lambda_{\beta_{\xi}} P\beta_{\xi}\}_j) \\
&= \{B'H_{\xi} B + \lambda_{\beta_{\xi}} P\}_{jl} \\
&\approx -\{B'\bar{H}_{\xi} B - \lambda_{\beta_{\xi}} P\}_{jl} \text{ for } j, l = 1, 2, \dots, p
\end{aligned}$$

where $\bar{H}_{\xi jl} = -E \frac{\partial^2 \ell^*}{\partial \xi_j \partial \xi_l}$, $j, l = 1, 2, \dots, p$ is diagonal. A similar expression is easily derived for $\frac{\partial u_{\beta_{\nu j}}^*}{\partial \beta_{\nu l}}$. Importantly, $\frac{\partial u_{\beta_{\xi j}}^*}{\partial \beta_{\nu l}} \approx \frac{\partial u_{\beta_{\nu j}}^*}{\partial \beta_{\xi l}} \approx 0$ always since $E \frac{\partial^2 \ell^*}{\partial \xi_j \partial \nu_l} = 0$ from asymptotic independence of the generalised Pareto parameters.

Substituting these expressions into the Newton-Raphson step above for $u_{\beta_{\xi}}^*$ we have, in matrix notation

$$0 = (B'u_{\xi}^o + \lambda_{\beta_{\xi}} P\beta_{\xi}^o) - (B'\bar{H}_{\xi}^o B - \lambda_{\beta_{\xi}} P)(\beta_{\xi} - \beta_{\xi}^o)$$

where the $\lambda_{\beta_{\xi}} P\beta_{\xi}^o$ terms cancel on the right-hand-side, from which

$$(B'\bar{H}_{\xi}^o B - \lambda_{\beta_{\xi}} P)\beta_{\xi} = B'u_{\xi}^o + \lambda_{\beta_{\xi}} P\beta_{\xi}^o + (B'\bar{H}_{\xi}^o B - \lambda_{\beta_{\xi}} P)\beta_{\xi}^o$$

where u_{ξ}^o and \bar{H}_{ξ}^o are u_{ξ} and \bar{H}_{ξ} evaluated at β_{ξ}^o . Thus the vector Newton-Raphson step for β_{ξ} becomes

$$\beta_{\xi} = (B'\bar{H}_{\xi}^o B - \lambda_{\beta_{\xi}} P)^{-1} (B'u_{\xi}^o + B'\bar{H}_{\xi}^o B\beta_{\xi}^o)$$

The analogous expression for β_{ν} is

$$\beta_{\nu} = (B'\bar{H}_{\nu}^o B - \lambda_{\beta_{\nu}} P)^{-1} (B'u_{\nu}^o + B'\bar{H}_{\nu}^o B\beta_{\nu}^o)$$

where u_{ν}^o and \bar{H}_{ν}^o , u_{ν} and \bar{H}_{ν} evaluated at β_{ν}^o . Alternating fits using the last two equations allow iteration from starting solution $(\beta_{\xi}^o, \beta_{\nu}^o)$ (typically corresponding to parameter stationarity) to the optimum

maximum penalised likelihood solution. The elements of u_ξ^o , u_ν^o , \bar{H}_ξ^0 and \bar{H}_ν^o change at each iteration.

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.
- Kaatje Bollaerts. *Statistical models in epidemiology and quantitative microbial risk assessment applied to salmonella in pork*. PhD thesis, Hasselt University, Belgium, 2009.
- E. Casson and S. G. Coles. Spatial regression models for extremes. *Extremes*, 1:449–468, 1999.
- R. E. Chandler and S. Bate. Inference for clustered data using the independence loglikelihood. *Biometrika*, 94:167–183, 2007.
- V. Chavez-Demoulin. *Two problems in environmental statistics: capture-recapture analysis and smooth extremal models*. PhD Thesis. Department of Mathematics, Swiss Federal Institute of Technology, Lausanne, 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, 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.
- I. D. Currie, M. Durban, and P. H. C. Eilers. Generalized linear array models with applications to multidimensional smoothing. *J. Roy. Statist. Soc. B*, 68:259–280, 2006.
- A. C. Davison. *Statistical models*. Cambridge University Press, 2003.
- A. C. Davison and M. M. Gholamrezaee. Geostatistics of extremes. *Proc. R. Soc. A*, 468:581–608, 2012.
- A. C. Davison, S. A. Padoan, and M. Ribatet. Statistical modelling of spatial extremes. *Statistical Science*, 27:161–186, 2012.
- A.C. Davison and R. L. Smith. Models for exceedances over high thresholds. *J. R. Statist. Soc. B*, 52:393, 1990.
- C. de Boor. *A practical guide to splines*. Springer-Verlag, New York (Applied Mathematical Sciences Revised Edition), 2001.
- 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, UK, 2007.
- E.F. Eastoe and J.A. Tawn. Modelling non-stationary extremes with application to surface level ozone. *Appl. Statist.*, 58:22–45, 2009.
- P H C Eilers and B D Marx. Splines, knots and penalties. *Wiley Interscience Reviews: Computational Statistics*, 2:637–653, 2010.
- P H C Eilers, I D Currie, and M Durban. Fast and compact smoothing on multi-dimensional grids. *Computational statistics and data analysis*, 50:61–76, 2006.
- R. J. Erhardt and R. L. Smith. Approximate Bayesian computing for spatial extremes. *Available from <http://arxiv.org/pdf/1109.4166.pdf>*, 2011.
- 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.
- A. Ferreira and L. de Haan. The generalised Pareto process. (*Accepted for Bernoulli Journal, preprint at www.bernoulli-society.org*), 2014.
- 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 Eng.*, 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*, 66:497–546, 2004.
- R. Huser and A. C. Davison. Space-time modelling of extreme events. *J. Roy. Statist. Soc. B*, 76:439–461, 2014.
- P. Jonathan and K. C. Ewans. The effect of directionality on extreme wave design criteria. *Ocean Eng.*, 34:1977–1994, 2007a.
- P. Jonathan and K. C. Ewans. Uncertainties in extreme wave height estimates for hurricane dominated regions. *J. Offshore Mechanics Arctic Engineering*, 129:300–305, 2007b.
- 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 and K. C. Ewans. Modelling the seasonality of extreme waves in the Gulf of Mexico. *ASME J. Offshore Mech. Arct. Eng.*, 133:021104, 2011a.
- P. Jonathan and K. C. Ewans. A spatiodectional model for extreme waves in the Gulf of Mexico. *ASME J. Offshore Mech. Arct. Eng.*, 133:011601, 2011b.
- P. Jonathan and K. C. Ewans. Statistical modelling of extreme ocean environments with implications for marine design : a review. *Ocean Engineering*, 62:91–109, 2013.

- 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 Eng.*, 35:1211–1225, 2008.
- P. Jonathan, K. C. Ewans, and D. Randell. Joint modelling of environmental parameters for extreme sea states incorporating covariate effects. *Coastal Engineering*, 79:22–31, 2013.
- P. Jonathan, K. C. Ewans, and D. Randell. Non-stationary conditional extremes of northern North Sea storm characteristics. *Environmetrics*, 25:172–188, 2014.
- R. Koenker. *Quantile regression*. Cambridge University Press, New York., 2005.
- A. W. Ledford and J. A. Tawn. Modelling dependence within joint tail regions. *J. R. Statist. Soc. B*, 59: 475–499, 1997.
- D.-J. Lee and M Durban. P-spline ANOVA type interaction models for spatio temporal smoothing. *Statistical Modelling*, 11:49–69, 2011.
- D.-J. Lee, M Durban, and P Eilers. Efficient two-dimensional smoothing with P-spline ANOVA mixed models and nested basis. *Computational Statistics and Data Analysis*, 61:22–37, 2013.
- B. D. Marx and P. H. C. Eilers. Direct generalised additive modelling with penalised likelihood. *Computational statistics and data analysis*, 28:193–209, 1998.
- P. Northrop and P. Jonathan. Threshold modelling of spatially-dependent non-stationary extremes with application to hurricane-induced wave heights. *Environmetrics*, 22:799–809, 2011.
- S. A. Padoan, M. Ribatet, and S. A. Sisson. Likelihood-based inference for max-stable processes. *J. Am. Statist. Soc.*, 105:263–277, 2010.
- M. E. Robinson and J. A. Tawn. Statistics for extreme sea currents. *Applied Statistics*, 46:183–205, 1997.
- D. Ruppert, M. P. Wand, and R. J. Carroll. *Semiparametric regression*. Cambridge University Press, 2003.
- M.G. Scotto and C. Guedes-Soares. Modelling the long-term time series of significant wave height with non-linear threshold models. *Coastal Eng.*, 40:313, 2000.
- M.G. Scotto and C. Guedes-Soares. Bayesian inference for long-term prediction of significant wave height. *Coastal Eng.*, 54:393, 2007.
- R. L. Smith. Max-stable processes and spatial extremes. *Unpublished article, available electronically from www.stat.unc.edu/postscript/rs/spatex.pdf*, 1990.
- A Stephenson and E Gilleland. Software for the analysis of extreme events: The current state and future directions. *Extremes*, 8:87–109, 2006.
- VGAM. Vector generalized linear and additive models. URL www.stat.auckland.ac.nz/~yee.
- J.L. Wadsworth and J.A. Tawn. Dependence modelling for spatial extremes. *Biometrika*, 99:253–272, 2012.
- Y Wang and G Wahba. Bootstrap confidence intervals for smoothing splines and their comparison to bayesian confidence intervals. *J. Statist. Comput. Simul.*, 51:263–279, 1995.
- E. G. Ward, L. E. Borgman, and V. J. Cardone. Statistics of hurricane waves in the Gulf of Mexico. *Offshore Technology Conference*, 1978.

List of Figures

1	H_S^{sp} against direction for all 1089 locations. Clear variability with direction, largest H_S^{sp} from $90^\circ - 130^\circ$. Large number of small events from $260^\circ - 310^\circ$	20
2	Rose histogram plots of H_S^{sp} . Size and colour of bin shows proportion of interval of H_S^{sp} values. Left-hand plot shows pooled rose for all 1089 locations. Right-hand plots show direction histograms for individual locations in the (from left to right, top to bottom) NW, N, NE, W, Central, E, SW, S and SE locations respectively. All right-hand plots (except “central”) correspond to locations at the boundary of the spatial domain in the given direction; the “central” plot corresponds to the centre of the spatial domain. In left hand plot, large number of smaller storms from west, smaller number of larger storms from east north-east.	21
3	Spatial distributions of storm direction (arrow direction) and relative size of storm peak significant wave height (arrow length) for the largest two cyclone events.	22
4	Directional evolution of significant wave height (metres) for all cyclone events at the (single) central location, emanating from 8 directional octants centred (from left to right, top to bottom) on NW, N, and NE; W and E; SW, S and SE respectively.	23
5	Spatio-directional plot for Poisson rate of threshold exceedance ρ . The 8 right-hand plots show ρ spatially for 8 directions (from left to right, top to bottom: storms from NW, N, and NE; W and E; SW, S and SE respectively). The left-hand plot shows the direction from which ρ is largest for each location. Units of ρ are number of threshold exceedances per degree longitude, latitude and direction, per annum.	24
6	Spatio-directional plot for estimates of GP shape, ξ . The 8 right-hand plots show ξ spatially for 8 directions (from left to right, top to bottom: storms from NW, N, and NE; W and E; SW, S and SE respectively). The left-hand plot shows the direction from which ξ is largest, for each location.	25
7	Spatio-directional plot for estimates of GP scale, σ . The 8 right-hand plots show σ spatially for 8 directions (from left to right, top to bottom: storms from NW, N, and NE; W and E; SW, S and SE respectively). The left-hand plot shows the direction from which σ is largest, for each location. Colour scale in metres.	26
8	Validation of spatio-directional model for H_S^{sp} . Comparison of cumulative distribution functions (cdfs) for omnidirectional return values of the original sample (black) with those from 1000 realisations under the model corresponding to the same time period as the original sample, shown in terms of the median (solid grey) and 95% uncertainty band (dashed grey). Panels correspond to a 3×3 grid of locations in the spatial domain, the relative longitudes and latitudes of which are given in the panel titles.	27
9	Validation of spatio-directional model for H_S^{sp} . Comparison of cumulative distribution functions (cdfs) for directional octant return values of the original sample (black) with those from 1000 realisations under the model corresponding to the same time period as the original sample, shown in terms of the median (solid grey) and 95% uncertainty band (dashed grey) aggregated over the 9 locations used in Figure 8. Panels correspond to directional octants centred (from left to right, top to bottom) on NW, N, and NE; W and E; SW, S and SE respectively.	28

10	100-year return values H_{S100}^{sp} without directional dissipation, using simulation. The left-hand plot shows median omnidirectional return values for each of the 1089 locations. The 8 right-hand plots show median directional sector H_{S100}^{sp} per location for 8 directional octants centred (from left to right, top to bottom) on storms from NW, N, and NE; W and E; SW, S and SE respectively. Colour scale in metres.	29
11	100-year return values H_{S100} with directional dissipation, using simulation. The left-hand plot shows median omnidirectional return values for each of the 1089 locations. The 8 right-hand plots show median directional sector H_{S100} per location for 8 directional octants centred (from left to right, top to bottom) on storms from NW, N, and NE; W and E; SW, S and SE respectively. Colour scale in metres.	30
12	Median omnidirectional 100-year return values H_{S100}^{sp} (solid black) from pooling of all locations assuming spatial independence ignoring all covariate effects, as a function of non-exceedance probability of constant extreme value threshold, with 95% uncertainty band (dashed black). Corresponding estimates for storms from directional sector 90° - 135° only are shown in grey.	31

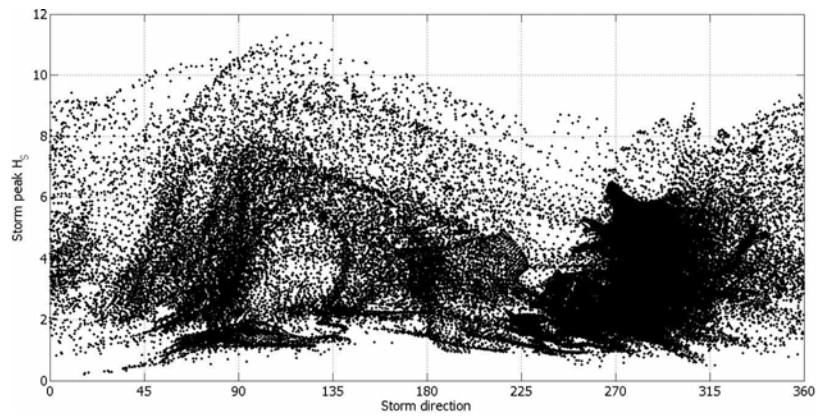


Figure 1: H_S^{sp} against direction for all 1089 locations. Clear variability with direction, largest H_S^{sp} from $90^\circ - 130^\circ$. Large number of small events from $260^\circ - 310^\circ$.

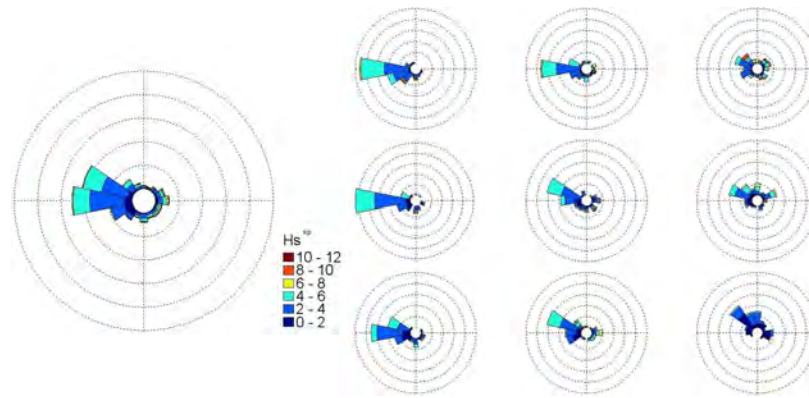


Figure 2: Rose histogram plots of H_S^{sp} . Size and colour of bin shows proportion of interval of H_S^{sp} values. Left-hand plot shows pooled rose for all 1089 locations. Right-hand plots show direction histograms for individual locations in the (from left to right, top to bottom) NW, N, NE, W, Central, E, SW, S and SE locations respectively. All right-hand plots (except “central”) correspond to locations at the boundary of the spatial domain in the given direction; the “central” plot corresponds to the centre of the spatial domain. In left hand plot, large number of smaller storms from west, smaller number of larger storms from east north-east.

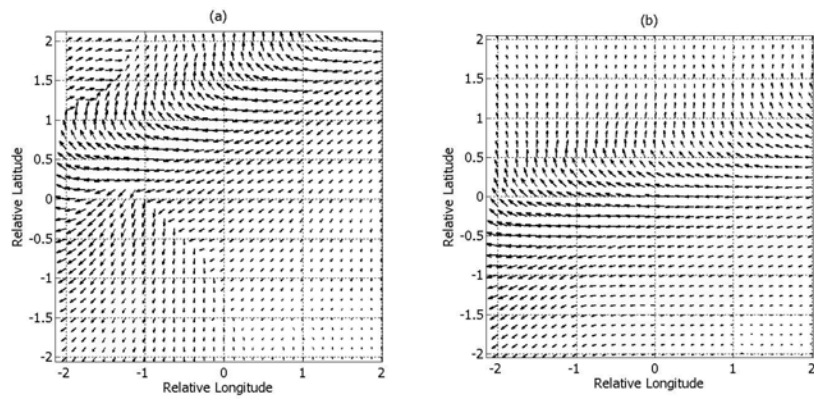


Figure 3: Spatial distributions of storm direction (arrow direction) and relative size of storm peak significant wave height (arrow length) for the largest two cyclone events.

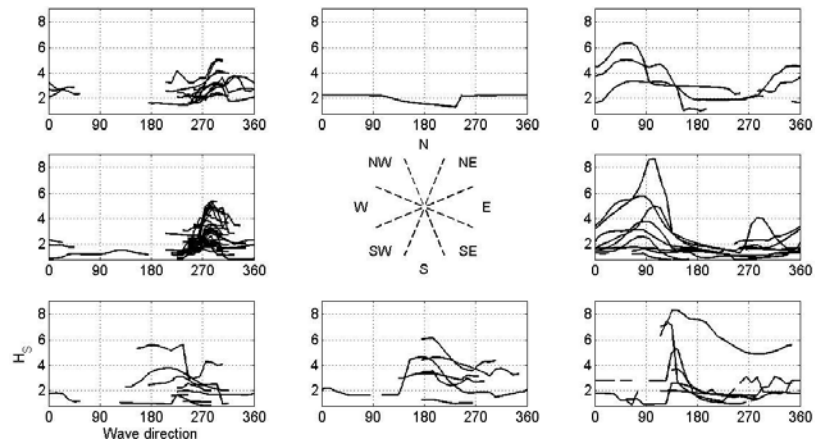


Figure 4: Directional evolution of significant wave height (metres) for all cyclone events at the (single) central location, emanating from 8 directional octants centred (from left to right, top to bottom) on NW, N, and NE; W and E; SW, S and SE respectively.

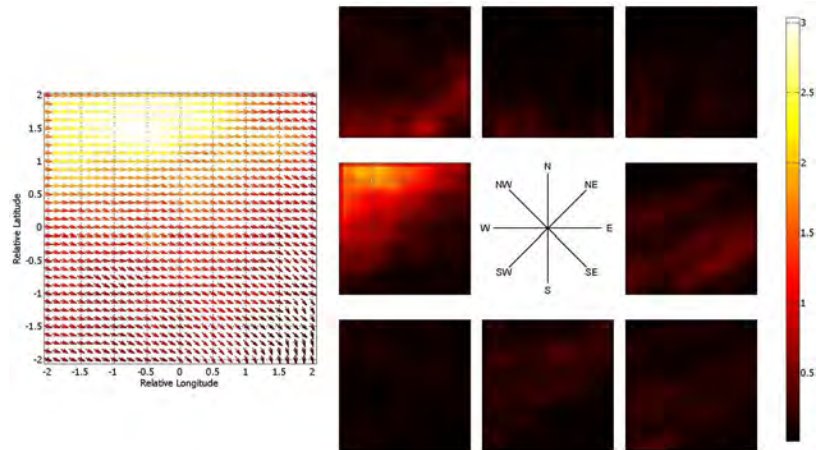


Figure 5: Spatio-directional plot for Poisson rate of threshold exceedance ρ . The 8 right-hand plots show ρ spatially for 8 directions (from left to right, top to bottom: storms from NW, N, and NE; W and E; SW, S and SE respectively). The left-hand plot shows the direction from which ρ is largest for each location. Units of ρ are number of threshold exceedances per degree longitude, latitude and direction, per annum.

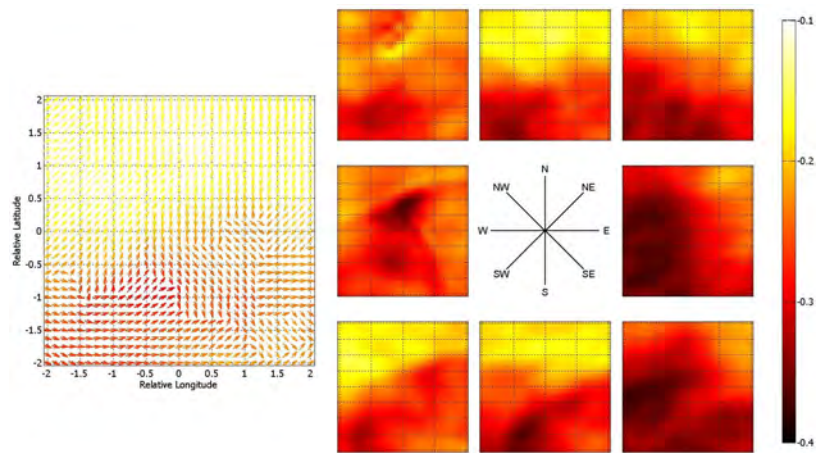


Figure 6: Spatio-directional plot for estimates of GP shape, ξ . The 8 right-hand plots show ξ spatially for 8 directions (from left to right, top to bottom: storms from NW, N, and NE; W and E; SW, S and SE respectively). The left-hand plot shows the direction from which ξ is largest, for each location.

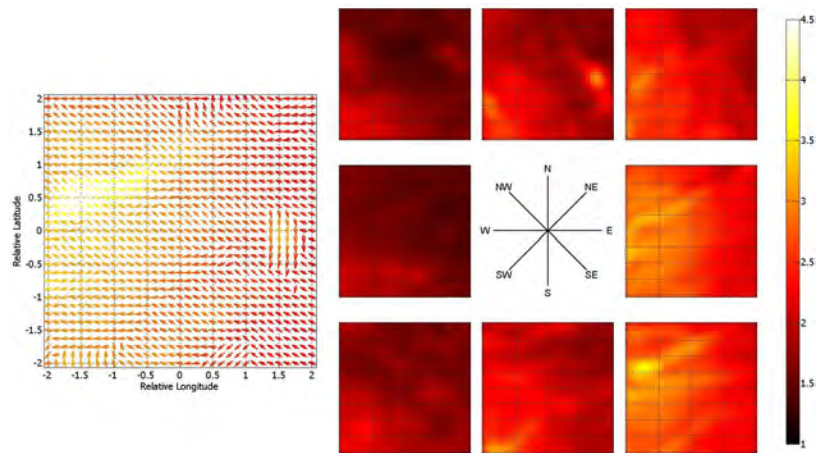


Figure 7: Spatio-directional plot for estimates of GP scale, σ . The 8 right-hand plots show σ spatially for 8 directions (from left to right, top to bottom: storms from NW, N, and NE; W and E; SW, S and SE respectively). The left-hand plot shows the direction from which σ is largest, for each location. Colour scale in metres.

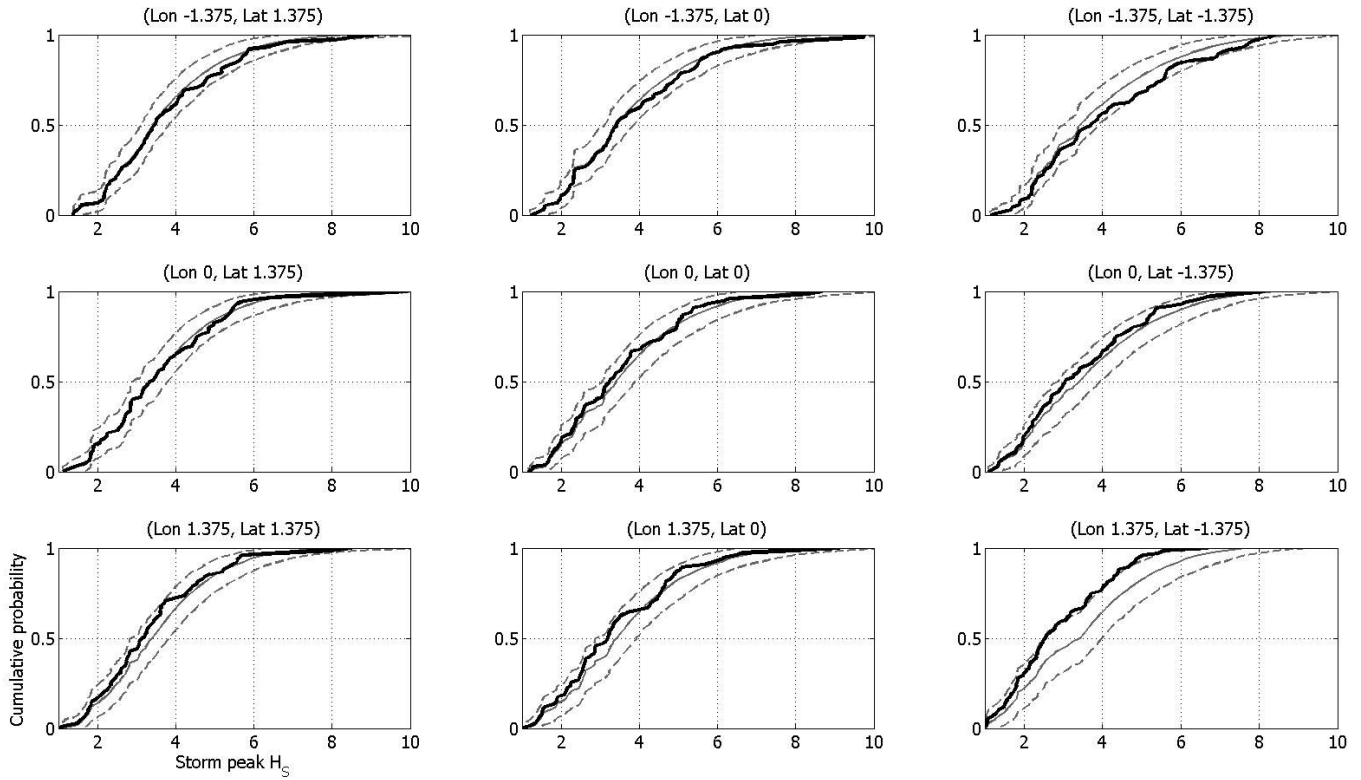


Figure 8: Validation of spatio-directional model for H_S^{sp} . Comparison of cumulative distribution functions (cdfs) for omnidirectional return values of the original sample (black) with those from 1000 realisations under the model corresponding to the same time period as the original sample, shown in terms of the median (solid grey) and 95% uncertainty band (dashed grey). Panels correspond to a 3×3 grid of locations in the spatial domain, the relative longitudes and latitudes of which are given in the panel titles.

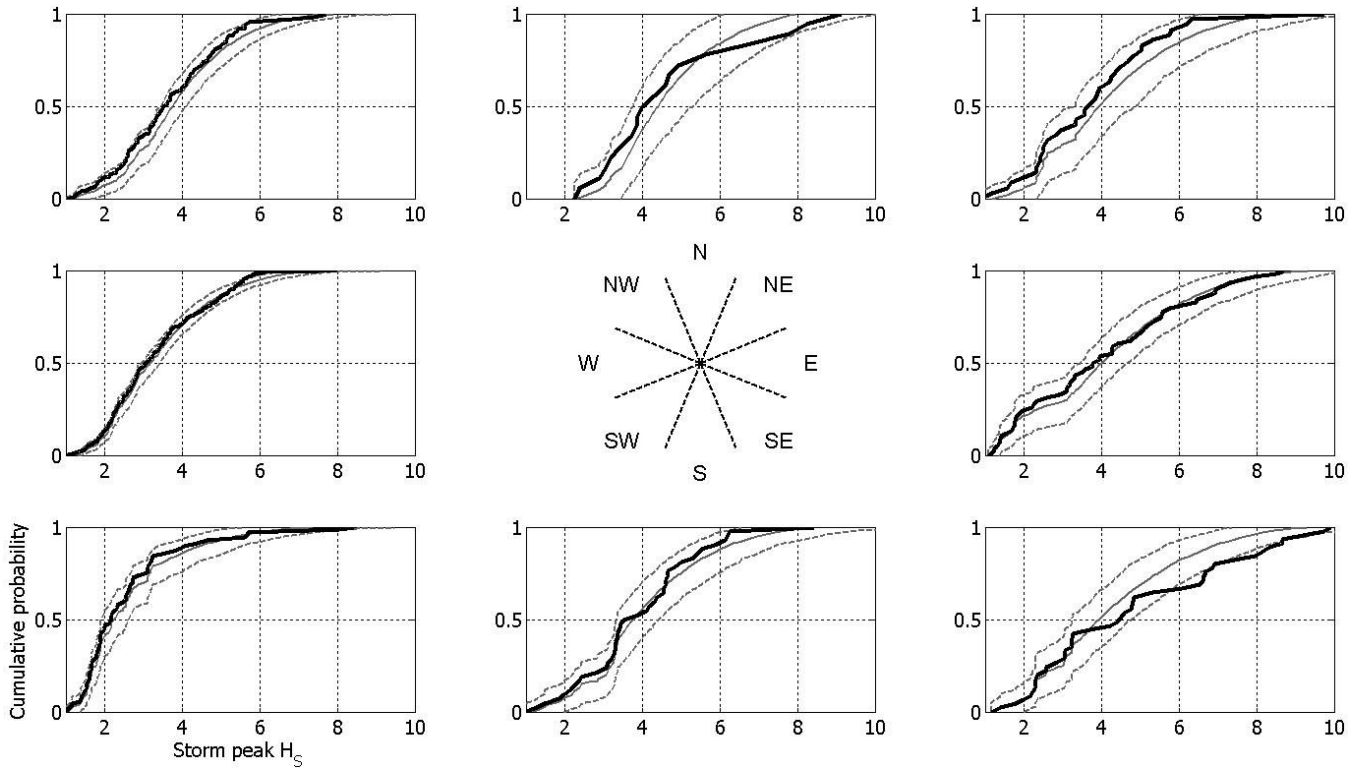


Figure 9: Validation of spatio-directional model for H_S^{sp} . Comparison of cumulative distribution functions (cdfs) for directional octant return values of the original sample (black) with those from 1000 realisations under the model corresponding to the same time period as the original sample, shown in terms of the median (solid grey) and 95% uncertainty band (dashed grey) aggregated over the 9 locations used in Figure 8. Panels correspond to directional octants centred (from left to right, top to bottom) on NW, N, and NE; W and E; SW, S and SE respectively.




Figure 10: 100-year return values H_{S100}^{sp} without directional dissipation, using simulation. The left-hand plot shows median omnidirectional return values for each of the 1089 locations. The 8 right-hand plots show median directional sector H_{S100}^{sp} per location for 8 directional octants centred (from left to right, top to bottom) on storms from NW, N, and NE; W and E; SW, S and SE respectively. Colour scale in metres.

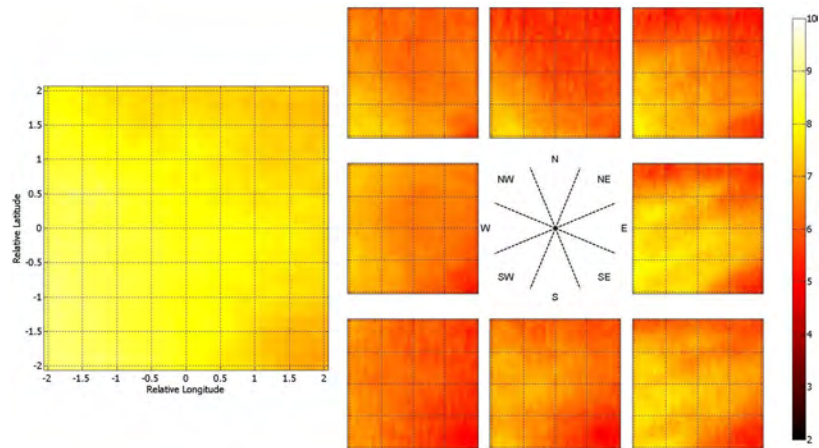


Figure 11: 100-year return values H_{S100} with directional dissipation, using simulation. The left-hand plot shows median omnidirectional return values for each of the 1089 locations. The 8 right-hand plots show median directional sector H_{S100} per location for 8 directional octants centred (from left to right, top to bottom) on storms from NW, N, and NE; W and E; SW, S and SE respectively. Colour scale in metres.




Figure 12: Median omnidirectional 100-year return values H_{S100}^{sp} (solid black) from pooling of all locations assuming spatial independence ignoring all covariate effects, as a function of non-exceedance probability of constant extreme value threshold, with 95% uncertainty band (dashed black). Corresponding estimates for storms from directional sector 90° - 135° only are shown in grey.