Covariate effects in oceanographic applications of marginal, conditional and spatial extremes

Slides at: *www.lancs.ac.uk/∼jonathan*

Lancaster STORi extremes workshop

Acknowledgement

- Shell stats : Philip Jonathan, **Mirrelijn van Nee**, Laks Raghupathi, David Randell, **Emma Ross**, Yanyun Wu
- Shell metocean : Kevin Ewans, Graham Feld
- STOR-i: **Monika Kereszturi**, *Kathryn Turnbull*, Elena Zanini
- ...

Motivation

- Rational and consistent design and assessment of marine structures
  - Reduce bias and uncertainty in estimation of structural integrity
  - Quantify uncertainty as well as possible
- Non-stationary marginal, conditional and spatial extremes
  - Multiple locations, multiple variables, time-series
  - Multidimensional covariates
- Improved understanding and communication of risk
  - Incorporation within established engineering design practices
  - Knock-on effects of improved inference
- Other current applications in Shell
  - Earthquake hazards
  - Corrosion and fouling

Motivation

- Environmental extremes vary smoothly with multidimensional covariates
  - Model parameters are functions of covariates
- Uncertainty quantification for whole inference
  - Data acquisition (simulator or measurement)
  - Data pre-processing (storm peak identification)
  - Extreme value threshold
  - Model form (marginal measurement scale effect, spatial extremal dependence)
- Statistical and computational efficiency
  - Slick algorithms
  - Parallel computation
  - Bayesian inference

Motivation: storm model

$H_S \approx 4\times$ standard deviation of ocean surface time-series at specific location corresponding to a specified period (typically three hours)

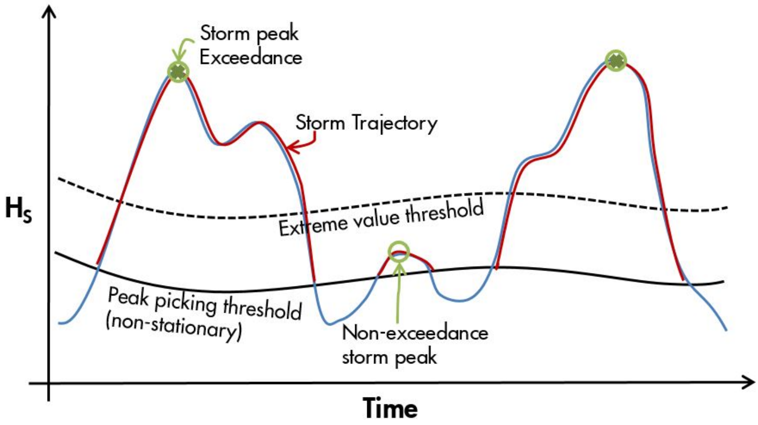

Outline

Covariate effects in:

- Marginal models
  - Simple introductory example (directional model)
  - Storm peak H_S with 2D, 3D and 4D covariates
- Conditional extremes models
  - Associated values of other wave field parameters given extreme stork peak H_S
- Spatial extremes models
  - Directional dependence in max-stable process parameters for storm peak H_S

North Sea example used as “connecting theme”; other examples to embellish

Outline: North Sea application

Storm peak H_S from gridded NEXTRA *winter* storm hindcast for North Sea locations; directional variability in storm severity; "strips" of locations with different orientations; central location for directional model

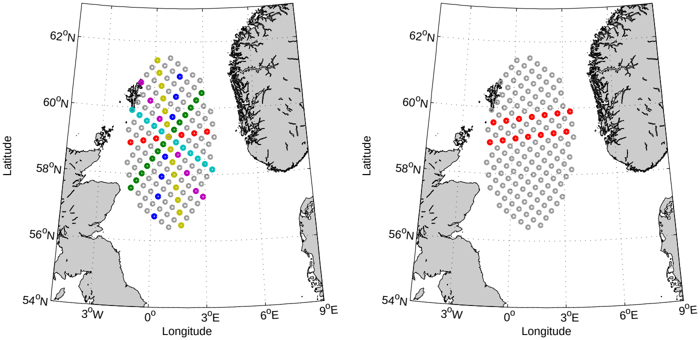

Marginal: simple gamma-GP model

- Sample of peaks over threshold *y*, with covariates θ
  - θ is 1D in motivating example : directional
  - θ is *nD* later : e.g. $4D$ spatio-directional-seasonal
- Below threshold ψ
  - *y* follows truncated gamma with shape α, scale $1/\beta$
  - Hessian for gamma better behaved than Weibull
- Above ψ
  - *y* follows generalised Pareto with shape ξ, scale σ
- ξ, σ, α, β, ψ all functions of θ
- ψ for pre-specified threshold probability τ
  - Generalise later to estimation of τ

- Frigessi et al. [2002], Behrens et al. [2004], MacDonald et al. [2011]
- Randell et al. [2016]

Marginal: simple gamma-GP model

- Density is $f(y|\xi, \sigma, \alpha, \beta, \psi, \tau)$

$$= \begin{cases} \tau \times f_{TG}(y|\alpha, \beta, \psi) & \text{for } y \leq \psi \\ (1-\tau) \times f_{GP}(y|\xi, \sigma, \psi) & \text{for } y > \psi \end{cases}$$

- Likelihood is $\mathcal{L}(\xi, \sigma, \alpha, \beta, \psi, \tau | \{y_i\}_{i=1}^n)$

$$\begin{aligned} = & \prod_{i:y_i \leq \psi} f_{TG}(y_i|\alpha, \beta, \psi) \prod_{i:y_i > \psi} f_{GP}(y_i|\xi, \sigma, \psi) \\ \times & \quad \tau^{n_B}(1-\tau)^{(1-n_B)} \text{ where } n_B = \sum_{i:y_i \leq \psi} 1. \end{aligned}$$

Estimate all parameters as functions of θ

Marginal: count rate c

- Whole-sample rate of occurrence ρ modelled as Poisson process given counts c of numbers of occurrences per covariate bin

- Chavez-Demoulin and Davison [2005]

Marginal: P-splines

- Physical considerations suggest $\alpha, \beta, \rho, \xi, \sigma, \psi$ and τ vary smoothly with covariates θ
- Values of $\eta \in \{\alpha, \beta, \rho, \xi, \sigma, \psi, \tau\}$ on some index set of covariates take the form $\boldsymbol{\eta} = \boldsymbol{B}\boldsymbol{\beta}_\eta$
  - For *nD* covariates, $\boldsymbol{B}$ takes the form of tensor product $\boldsymbol{B}_{\theta_n} \otimes ... \otimes \boldsymbol{B}_{\theta_\kappa} \otimes ... \otimes \boldsymbol{B}_{\theta_2} \otimes \boldsymbol{B}_{\theta_1}$
- Spline roughness with respect to each covariate dimension κ given by quadratic form $\lambda_{\eta\kappa}\beta'_{\eta\kappa}\boldsymbol{P}_{\eta\kappa}\boldsymbol{\beta}_{\eta\kappa}$
- $\boldsymbol{P}_{\eta\kappa}$ is a function of stochastic roughness penalties $\delta_{\eta\kappa}$
- Brezger and Lang [2006]

Marginal: priors and conditional structure

Priors

$$\begin{aligned} \text{density of } \boldsymbol{\beta}_{\eta\kappa} &\propto \exp\left(-\frac{1}{2}\lambda_{\eta\kappa}\boldsymbol{\beta}'_{\eta\kappa}\boldsymbol{P}_{\eta\kappa}\boldsymbol{\beta}_{\eta\kappa}\right) \\ \lambda_{\eta\kappa} &\sim \text{gamma} \\ (\text{ and } \tau &\sim \text{beta, when } \tau \text{ estimated }) \end{aligned}$$

Conditional structure

$$\begin{aligned} f(\tau|\boldsymbol{y}, \Omega \setminus \tau) &\propto f(\boldsymbol{y}|\tau, \Omega \setminus \tau) \times f(\tau) \\ f(\boldsymbol{\beta}_\eta|\boldsymbol{y}, \Omega \setminus \boldsymbol{\beta}_\eta) &\propto f(\boldsymbol{y}|\boldsymbol{\beta}_\eta, \Omega \setminus \boldsymbol{\beta}_\eta) \times f(\boldsymbol{\beta}_\eta|\boldsymbol{\delta}_\eta, \boldsymbol{\lambda}_\eta) \\ f(\boldsymbol{\lambda}_\eta|\boldsymbol{y}, \Omega \setminus \boldsymbol{\lambda}_\eta) &\propto f(\boldsymbol{\beta}_\eta|\boldsymbol{\delta}_\eta, \boldsymbol{\lambda}_\eta) \times f(\boldsymbol{\lambda}_\eta) \end{aligned}$$

$$\Omega = \{\alpha, \beta, \rho, \xi, \sigma, \psi, \tau\}$$

Marginal: inference

- Elements of $\boldsymbol{\beta}_\eta$ highly interdependent, correlated proposals essential for good mixing
- "Stochastic analogues" of IRLS and back-fitting algorithms for maximum likelihood optimisation used previously
- Estimation of different penalty coefficients for each covariate dimension

- Gibbs sampling when full conditionals available
- Otherwise Metropolis-Hastings (MH) within Gibbs, using suitable proposal mechanisms
  - mMALA where possible

- Roberts and Stramer [2002], Girolami and Calderhead [2011], Xifara et al. [2014]

Marginal: posterior parameter

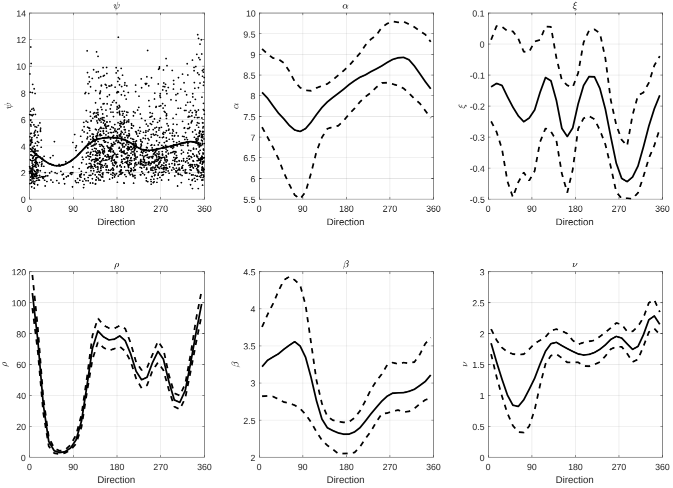

Marginal: posterior roughness penalty

Different scales so must be careful : rate is roughest, GP shape is smoothest

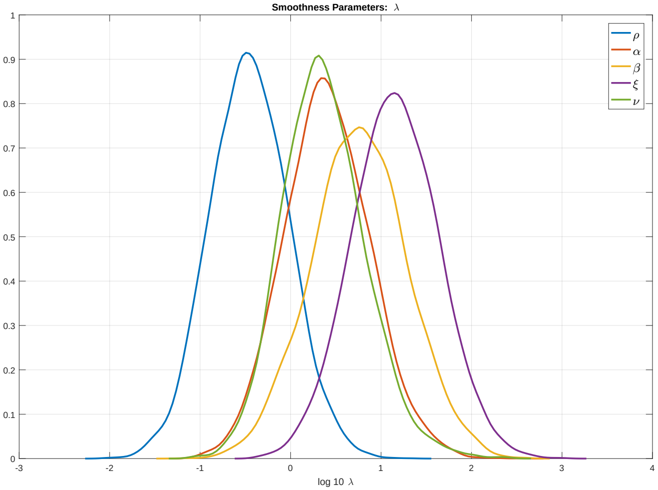

Marginal: validation

Compare sample with simulated values on partitioned covariate domain

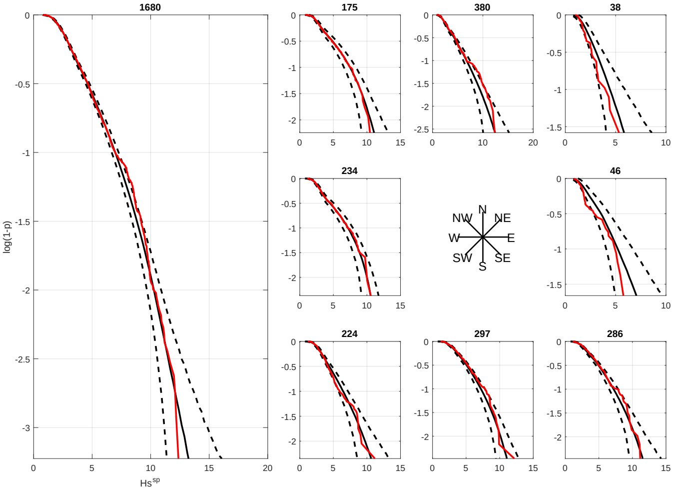

Marginal: extension to 2D

Directional-seasonal model for location in northern North Sea; τ estimated; land-shadow effect of Norway obvious; Randell et al. [2016]

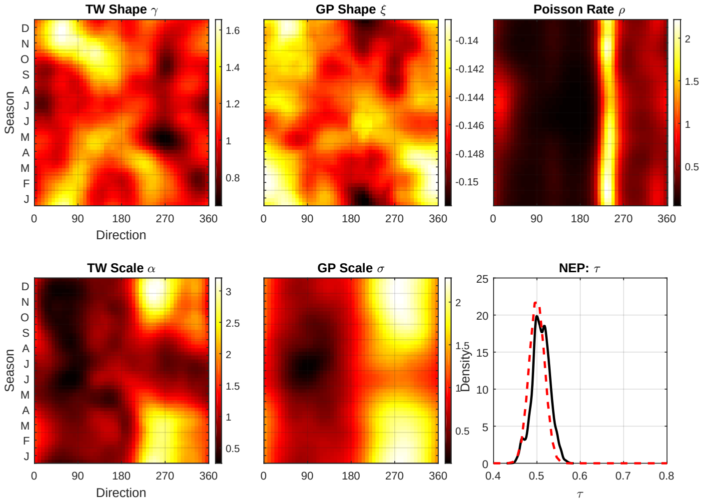

Marginal: extension to 2D

Summary statistics for return value distributions

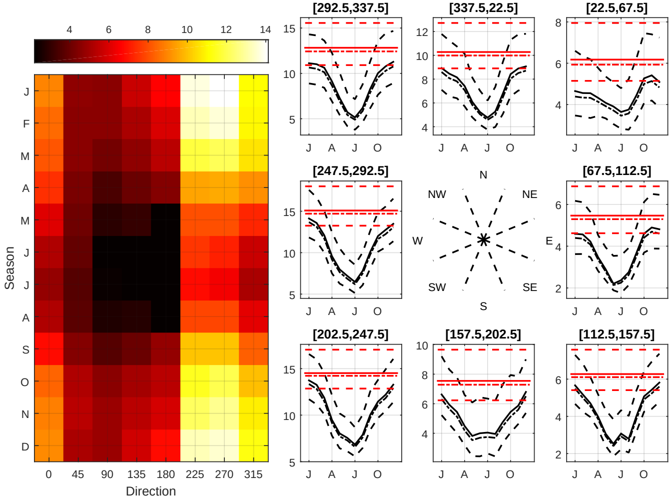

Marginal: extension to 4D

Spatio-directional-seasonal model for location in South China Sea; ML/CV/BS estimation; bootstrap median estimate after integration over season; clear spatial and directional effects; Raghupathi et al. [2016]

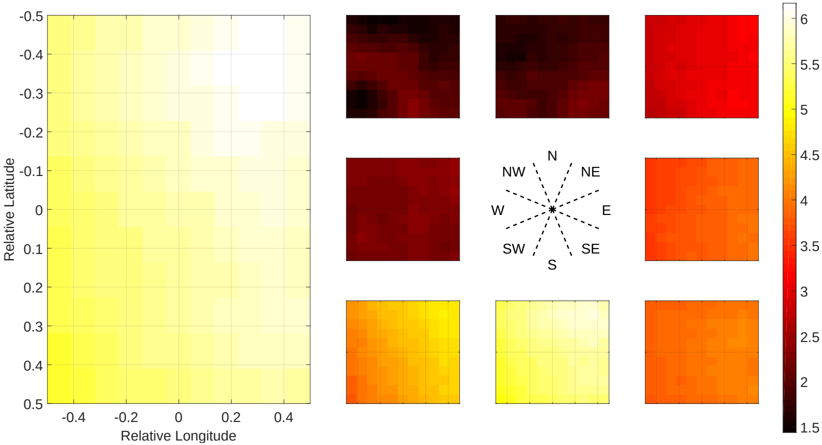

Marginal: extension to 4D

Bootstrap median estimate after integration over direction; clear spatial and seasonal effects

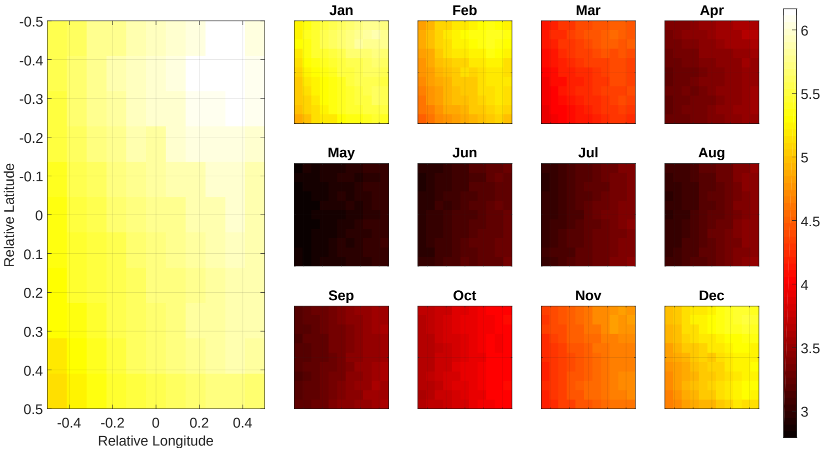

Conditional: summary

- Heffernan and Tawn [2004] and derivatives
- Evidence for covariate effects in conditional extremes of sea-state and storm peak variables
  - Marginal non-stationary extreme value model
  - Marginal transformation to standard scale removing marginal covariate dependence
  - Conditional dependence structure showing covariate effects
- Examples
  - Wave peak period | Significant wave height
  - Ocean current at one depth | Current at another depth
  - Significant wave height | Wind speed
  - Weather-vaning

Conditional: $T_P|H_S$ example

On **Gumbel** scale, extend with covariates θ

$$(Y_2|Y_1 = y, \theta) = \alpha_\theta y + y^{\beta_\theta}(\mu_\theta + \sigma_\theta Z) \text{ for } y > \psi_\theta(\tau)$$

where

- $\psi_\theta(\tau)$ is a high non-stationary quantile of Y_1 on Gumbel scale, for non-exceedance probability τ, above which the model fits well
- $\alpha_\theta \in [0, 1]$, $\beta_\theta \in (-\infty, 1]$, $\sigma_\theta \in [0, \infty)$
- Z is a random variable with **unknown** distribution G, assumed Normal for estimation

Application

- Estimate spectral peak wave period T_P for storm sea states with extreme severity (energy) H_S
- In T_P, H_S case, $\psi = \theta_j = \theta_k$
- Jonathan et al. [2014]

Conditional: $T_P|H_S$ example

ML/CV/BS inference; uncertainty bands capture uncertainty from marginal and dependence estimation; in conditional model, only α shows directional effect; reduction in conditional return value

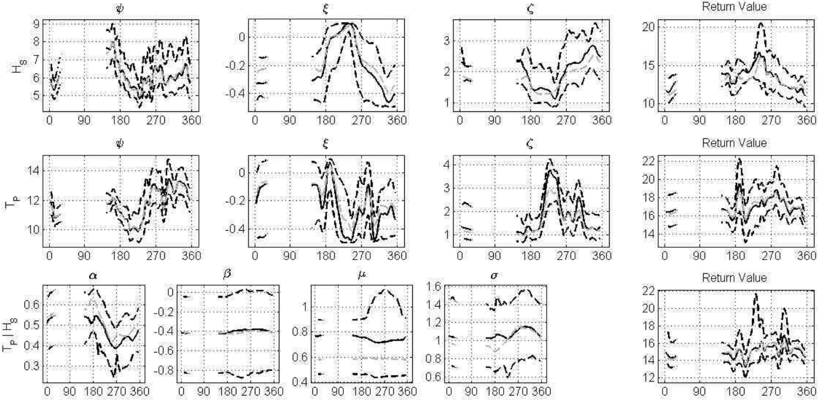

Spatial: outline

Why do spatial extremes?

- Improved inference at one location using data from spatial neighbourhood
- Insurance risk of damage to multiple structures from single "event"

Evidence for covariate effects in spatial extremes of storm peak significant wave height

- Neighbourhood of spatial locations
- Storm peak events corresponding to storm events observed at all spatial locations
- Marginal transformation per location to standard scale removing marginal covariate dependence
- Extremal spatial dependence structure showing anisotropy and location effect

Spatial: North Sea application

Storm peak H_S from gridded NEXTRA winter storm hindcast for North Sea locations; directional variability in storm severity; "strips" of locations with different orientations on bicycle wheel; multiple strips with same orientation

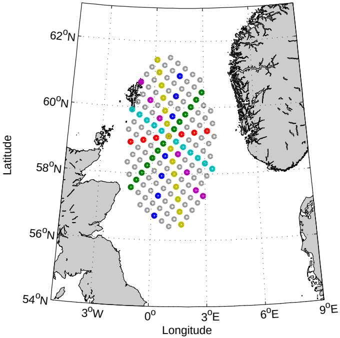

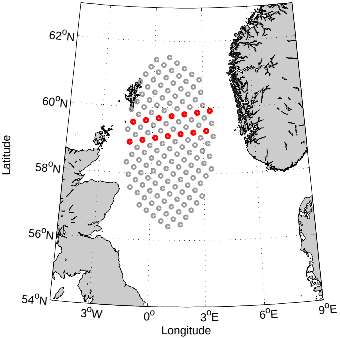

Spatial: storm on physical and Frechet scales

Storm peak H_S on physical and Frechet scales; marginal effects important

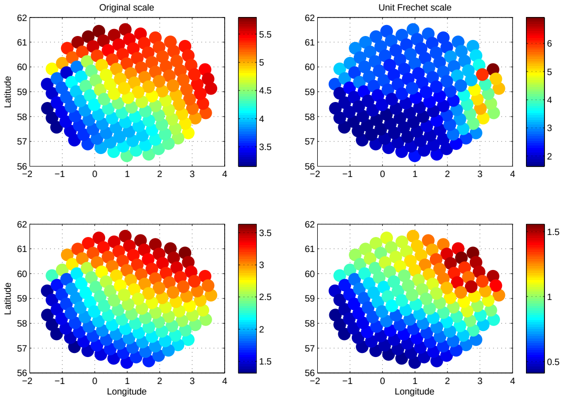

Spatial: Diagnosing dependence - simulated samples

Estimates for (a) η and (b) $\chi(x)$ against estimates for Spearman's ρ for sample size $n = 10^6$, from the Smith (magenta), Schlather (red), Brown-Resnick (blue), extremal-t (green) and Gaussian (black) processes, and the inverted logistic distribution (cyan). Estimation methods use model (19) for η with $q = 0.99$, and the empirical estimate for $\chi(x)$ with $x = 100$. Solid lines are median estimates from 1000 sample replications, dashed lines give 2.5% and 97.5% quantiles.

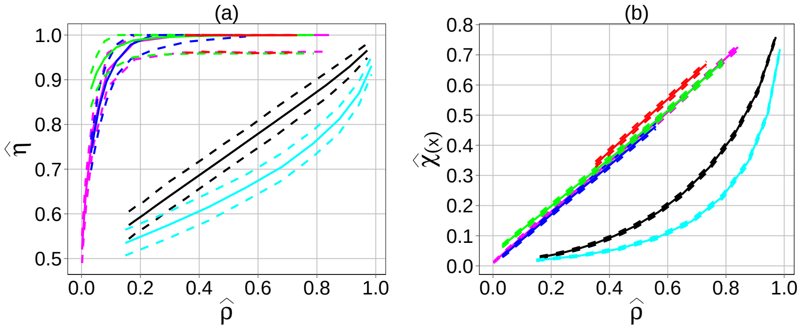

- $\eta = 1$ AD, $\chi = 0$ AI
- AD: Smith (magenta), Schlather (red), Brown-Resnick (blue)
- AI: extremal-t (green), Gaussian (black)
- Kereszturi et al. [2016]

Spatial: Diagnosing dependence - all sea states

Estimates of η with (a) $q = 0.90$ and (c) $q = 0.99$, and $\chi(x)$ with (b) $x = 10$ and (d) $x = 100$, plotted against Spearman's ρ for sea-state H_S sample of size $n = 58585$. Coloured points identify estimates from corresponding strip. Lines identify estimates using simulated samples of same size from Smith (black) and Gaussian (red) processes, and from the inverted logistic distribution (green);Kereszturi et al. [2016]

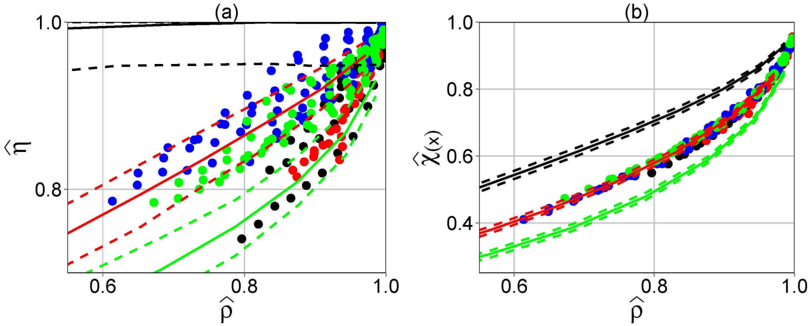

- η for $q = 0.9$, $\chi(x)$ for $x = 10$; $n = 58585$ individual sea states
- AD: Smith (black)
- AI: Gaussian (red), inverted logistic (green)

Spatial: Diagnosing dependence - all sea states

Estimates of η with (a) $q = 0.90$ and (c) $q = 0.99$, and $\chi(x)$ with (b) $x = 10$ and (d) $x = 100$, plotted against Spearman's ρ for sea-state H_S sample of size $n = 58585$. Coloured points identify estimates from corresponding strip. Lines identify estimates using simulated samples of same size from Smith (black) and Gaussian (red) processes, and from the inverted logistic distribution (green);Kereszturi et al. [2016]

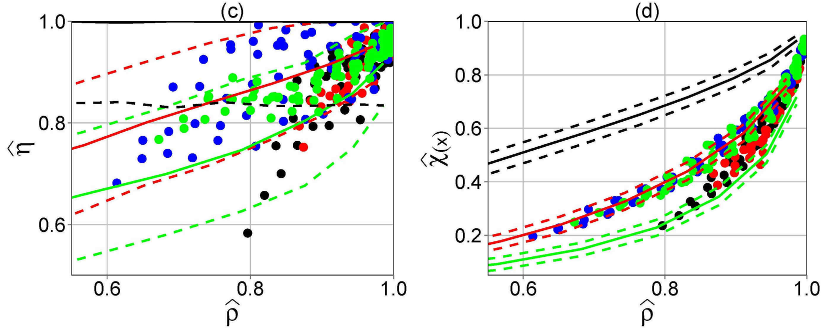

- η for $q = 0.99$, $\chi(x)$ for $x = 100$; $n = 58585$ sea states
- AD: Smith (black)
- AI: Gaussian (red), inverted logistic (green)

Spatial: Diagnosing dependence - storm peaks

Estimates of (a) η with $q = 0.90$ and (b) $\chi(x)$ with $x = 10$, plotted against Spearman's ρ for storm-peak H_S sample of size $n = 916$. Points and lines as described in previous slide; Kereszturi et al. [2016]

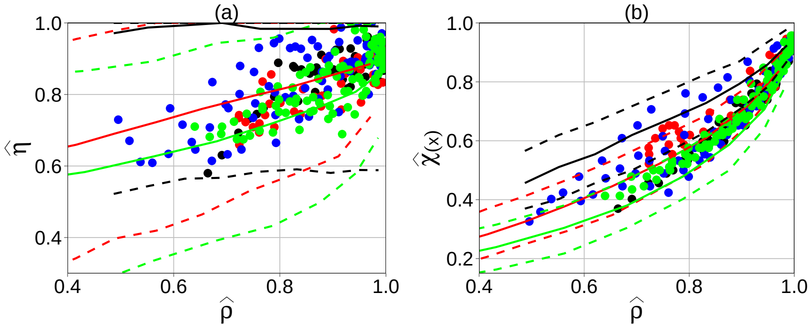

- η for $q = 0.9$, $\chi(x)$ for $x = 10$; $n = 916$ storm peak events
- AD: Smith (black)
- AI: Gaussian (red), inverted logistic (green)

Spatial: models and estimation

Marginal model

- Estimate non-stationary model
- Propagate uncertainty; this will be large in general

Max-stable process

- Smith, Schlather, Brown-Resnick, ...
- All AD; conservative estimates
- Wadsworth and Tawn [2012], Davison et al. [2012]

Composite likelihood

- Full likelihood unavailable; approximated by pairwise
- Padoan et al. [2010]

Censored likelihood

- Dependence structure required for peaks over threshold margins not block maxima
- Threshold selection required; confirmed choice not affecting main inferences; need to propagate uncertainty
- Huser and Davison [2014]

Spatial: Smith

Construction

- Max-stable process $Z(x) \sim \max_{i\geq 1} \xi_i f(x, U_i)$
- $f(x, U_i) \sim N(U_i, \Sigma)$
  $=$ storm profile
- $\xi_i =$ storm intensity
  from point process
- $U_i =$ storm centre
  uniform RV

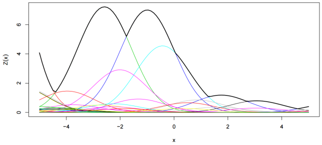

Estimation

- (Censored composite) likelihood available but messy
- 1D ("strip") : Estimate $\Sigma = \sigma^2$
- 2D (neighbourhood): Estimate $\Sigma = \begin{bmatrix} \sigma_1^2 & \sigma_{12} \\ \sigma_{12} & \sigma_2^2 \end{bmatrix}$

Spatial: Smith dependence anisotropy

Box plots: 1D parameter estimates; black lines: 2D parameter estimates; ML estimates with bootstrap 95% uncertainty bands accounting for uncertainty from marginal and dependence estimation; spatial dependence is higher WE than NS, consistent with large spatial events sweeping down from north

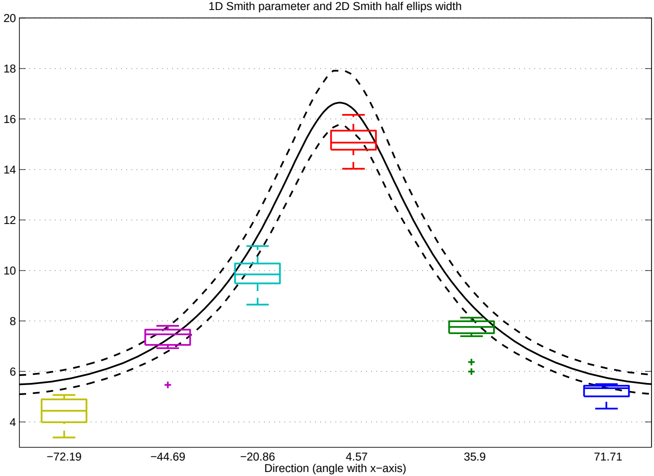

Spatial: Schlather & Brown-Resnick

Construction

- Max-stable process $Z(x) \sim \max_{i \geq 1} \xi_i Y_i(x)$
- Schlather
  - $Y_i(x) =$ standard normal Gaussian field, correlation $\rho(h)$
  - $\rho(h) = \exp(-0.5 h' \Sigma^{-1} h)$ here
- Brown-Resnick
  - $$Y_i(x) = \frac{\exp(W_i(x) - \gamma(x - U_i))}{n^{-1} \sum_{l=1}^{n} \exp(W_i(x_l) - \gamma(x_l - U_i))}$$
  - $W_i(x) =$ fractional Brownian motion, Hurst parameter $H \in [0, 1]$, variogram $2\gamma(h) = (h' \Sigma^{-1} h)^H$
  - Dieker and Mikosch [2014]

Estimation

- 1D: Estimate $\Sigma = \sigma^2$ (for range of H with BR)
- 2D: Estimate $\Sigma = \begin{bmatrix} \sigma_1^2 & \sigma_{12} \\ \sigma_{12} & \sigma_2^2 \end{bmatrix}$

Spatial: Dependence anisotropy

Smith, Schlather and Brown-Resnick consistent; confirmed that censored likelihood threshold not affecting relative size of dependence with direction

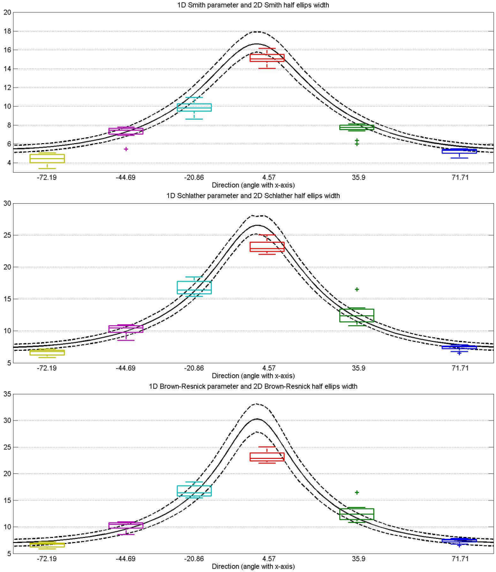

Spatial: Dependence location effect?

Box plots: 1D parameter estimates; black lines: 2D parameter estimates; ML estimates with bootstrap 95% uncertainty bands accounting for uncertainty from marginal and dependence estimation; spatial dependence is higher WE than NS, consistent with large spatial events sweeping down from north

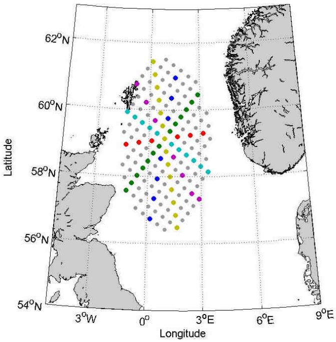

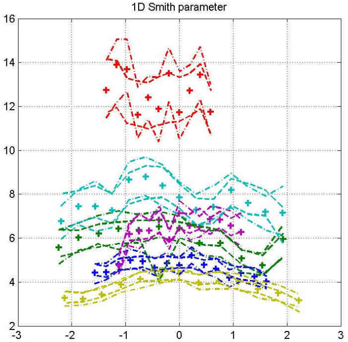

Spatial: Dependence location effect?

Box plots: 1D parameter estimates; black lines: 2D parameter estimates; ML estimates with bootstrap 95% uncertainty bands accounting for uncertainty from marginal and dependence estimation; spatial dependence is higher WE than NS, consistent with large spatial events sweeping down from north

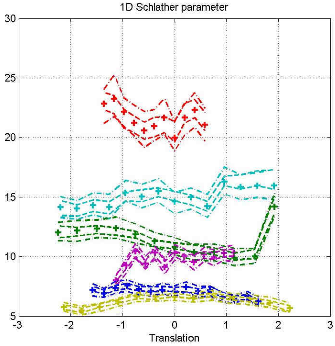

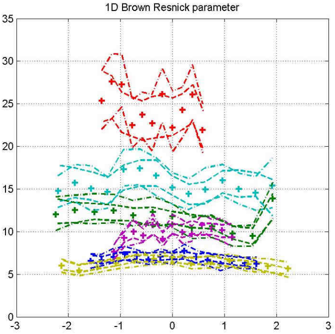

Summary

- Evidence for covariate effects in marginal, conditional and spatial extremes of ocean storms
  - Modelling non-stationarity essential for understanding extreme ocean storms, and estimating marine risk well
  - Non-parametric P-spline flexible basis for covariate description
  - Essential that non-stationary models are used for marginal, conditional and spatial extremes inference of ocean environment
  - Cradle-to-grave uncertainty quantification
- Further investigation of covariate effects in spatial ocean extremes needed
  - Anisotropy in North Sea hindcast, maybe absolute location (or fetch) effect?
  - Currently examining satellite altimeter measurements
  - Asymptotic independence?
- Goal : Bayesian inference for whole-basin spatial models with 4D covariates

References

C N Behrens, H F Lopes, and D Gamerman. Bayesian analysis of extreme events with threshold estimation. *Stat. Modelling*, 4:227–244, 2004.

A. Brezger and S. Lang. Generalized structured additive regression based on Bayesian P-splines. *Comput. Statist. Data Anal.*, 50:967–991, 2006.

V. Chavez-Demoulin and A.C. Davison. Generalized additive modelling of sample extremes. *J. Roy. Statist. Soc. Series C: Applied Statistics*, 54:207–222, 2005.

A. C. Davison, S. A. Padoan, and M. Ribatet. Statistical modelling of spatial extremes. *Statist. Sci.*, 27:161–186, 2012.

AB Dieker and T Mikosch. Exact simulation of Brown-Resnick random fields at a finite number of locations. *(arXiv preprint arXiv:1406.5624)*, 2014.

A. Frigessi, O. Haug, and H. Rue. A dynamic mixture model for unsupervised tail estimation without threshold selection. *Extremes*, 5: 219–235, 2002.

M. Girolami and B. Calderhead. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *J. Roy. Statist. Soc. B*, 73:123–214, 2011.

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, K. C. Ewans, and D. Randell. Non-stationary conditional extremes of northern North Sea storm characteristics. *Environmetrics*, 25:172–188, 2014.

M. Kereszturi, J. Tawn, and P. Jonathan. Assessing extremal dependence of north sea storm severity. *(Accepted by Ocean Engineering in April 2016, draft at www.lancs.ac.uk/∼jonathan)*, 2016.

A. MacDonald, C. J. Scarrott, D. Lee, B. Darlow, M. Reale, and G. Russell. A flexible extreme value mixture model. *Comput. Statist. Data Anal.*, 55:2137–2157, 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.

L. Raghupathi, D. Randell, E. Ross, K. Ewans, and P. Jonathan. Multi-dimensional predictive analytics for risk estimation of extreme events. *(Accepted for IEEE High-Performance Computing, Data and Analytics Conference (HIPC2016), draft at www.lancs.ac.uk/∼jonathan)*, 2016.

D. Randell, K. Turnbull, K. Ewans, and P. Jonathan. Bayesian inference for non-stationary marginal extremes. *(Accepted for publication in Environmetrics June 2016, draft at www.lancs.ac.uk/∼jonathan)*, 2016.

G. O. Roberts and O. Stramer. Langevin diffusions and Metropolis-Hastings algorithms. *Methodology and Computing in Applied Probability*, 4:337–358, 2002.

J.L. Wadsworth and J.A. Tawn. Dependence modelling for spatial extremes. *Biometrika*, 99:253–272, 2012.

T. Xifara, C. Sherlock, S. Livingstone, S. Byrne, and M Girolami. Langevin diffusions and the Metropolis-adjusted Langevin algorithm. *Stat. Probabil. Lett.*, 91(2002):14–19, 2014.