## Efficient adaptive covariate modelling for extremes

Slides at www.lancs.ac.uk/~jonathan

Matthew Jones, David Randell, Emma Ross, Elena Zanini, Philip Jonathan

## Structural damage



Ike, Gulf of Mexico, 2008 (Joe Richard)


North Sea, Winter 2015-16 (The Inertia)

## 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, spatial and temporal 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

The ocean environment is an amazing thing to study ... especially if you like to combine beautiful physics, measurement and statistical modelling!

## Fundamentals

- Environmental extremes vary smoothly with multidimensional covariates

■ Model parameters are non-stationary
■ Environmental extremes exhibit spatial and temporal dependence

- Characterise these appropriately

■ Uncertainty quantification for whole inference

- Data acquisition (simulator or measurement)
- Data pre-processing (storm peak identification)
- Hyper-parameters (extreme value threshold)
- Model form (marginal measurement scale effect, spatial extremal dependence)
- Statistical and computational efficiency
- Slick algorithms
- Parallel computation
- Bayesian inference


## A typical sample

Typical data for South China Sea location. Sea state (grey) and storm peak (black) $H_{S}$ on season and direction




## Outline

Directional-seasonal covariate models for $H_{S}^{s p}$

- Introductory example using P-splines
- Adaptive splines
- Partition models

■ South China Sea example as "connecting theme"

- Focus on the generalised Pareto (GP) inference


## Simple gamma-GP model



## Simple gamma-GP model

■ Sample of peaks over threshold $y$, with covariates $\theta$
■ $\theta$ is 1 D in motivating example : directional
■ $\theta$ is $n D$ later: e.g. $4 D$ spatio-directional-seasonal
■ Below threshold $\psi$

- y follows truncated gamma with shape $\alpha$, scale $1 / \beta$
- Hessian for gamma better behaved than Weibull
- Above $\psi$
- y follows generalised Pareto with shape $\xi$, scale $\sigma$

■ $\xi, \sigma, \alpha, \beta, \psi$ all functions of $\theta$

- $\psi$ for pre-specified threshold probability $\tau$

■ Generalise later to estimation of $\tau$
■ Frigessi et al. [2002], Behrens et al. [2004], MacDonald et al. [2011]

- Randell et al. [2016]


## Simple gamma-GP model

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

$$
= \begin{cases}\tau \times f_{T G}(y \mid \alpha, \beta, \psi) & \text { for } y \leq \psi \\ (1-\tau) \times f_{G P}(y \mid \xi, \sigma, \psi) & \text { for } y>\psi\end{cases}
$$

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

$$
\begin{aligned}
&= \prod_{i: y_{i} \leq \psi} f_{T G}\left(y_{i} \mid \alpha, \beta, \psi\right) \prod_{i: y_{i}>\psi} f_{G P}\left(y_{i} \mid \xi, \sigma, \psi\right) \\
& \times \tau^{n_{B}}(1-\tau)^{\left(1-n_{B}\right)} \text { where } n_{B}=\sum_{i: y_{i} \leq \psi} 1
\end{aligned}
$$

Estimate all parameters as functions of $\theta$

## Standard P-spline model

■ Physical considerations suggest $\alpha, \beta, \rho, \xi, \sigma, \psi$ and $\tau$ vary smoothly with covariates $\theta$
■ 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 $n D$ covariates, $\boldsymbol{B}$ takes the form of tensor product $\boldsymbol{B}_{\theta_{n}} \otimes \ldots \otimes \boldsymbol{B}_{\theta_{\kappa}} \otimes \ldots \otimes \boldsymbol{B}_{\theta_{2}} \otimes \boldsymbol{B}_{\theta_{1}}$
- Spline roughness with respect to each covariate dimension $\kappa$ given by quadratic form $\lambda_{\eta \kappa} \beta_{\eta \kappa}^{\prime} \boldsymbol{P}_{\eta \kappa} \boldsymbol{\beta}_{\eta \kappa}$
- $\boldsymbol{P}_{\eta \kappa}$ is a function of stochastic roughness penalties $\delta_{\eta \kappa}$

■ Brezger and Lang [2006]

## P-splines



Kronecker product


## Priors and conditional structure

Priors

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

Conditional structure

$$
\begin{aligned}
& f(\tau \mid \boldsymbol{y}, \Omega \backslash \tau) \propto f(\boldsymbol{y} \mid \tau, \Omega \backslash \tau) \times f(\tau) \\
& f\left(\boldsymbol{\beta}_{\eta} \mid \boldsymbol{y}, \Omega \backslash \boldsymbol{\beta}_{\eta}\right) \propto f\left(\boldsymbol{y} \mid \boldsymbol{\beta}_{\eta}, \Omega \backslash \boldsymbol{\beta}_{\eta}\right) \times f\left(\boldsymbol{\beta}_{\eta} \mid \boldsymbol{\delta}_{\eta}, \boldsymbol{\lambda}_{\eta}\right) \\
& f\left(\boldsymbol{\lambda}_{\eta} \mid \boldsymbol{y}, \Omega \backslash \boldsymbol{\lambda}_{\eta}\right) \propto f\left(\boldsymbol{\beta}_{\eta} \mid \boldsymbol{\delta}_{\eta}, \boldsymbol{\lambda}_{\eta}\right) \times f\left(\boldsymbol{\lambda}_{\eta}\right) \\
& \eta \in \Omega=\{\alpha, \beta, \rho, \xi, \sigma, \psi, \tau\}
\end{aligned}
$$

## 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]

## p-splines: GP parameter estimates




[^0]
## Inference with adaptive splines

- Advantages
- Arbitrary location of knots, and number of knots
- Estimate number, location, coefficient of knots
- Reversible-jump MCMC:
- Birth-death
- Split-combine (local birth-death)
- Detailed balance


■ Biller [2000], Zhou and Shen [2001], DiMatteo et al. [2001], Wallstrom et al. [2008]

## Inference with adaptive splines : e.g. birth-death



## Inference with adaptive bases: birth-death

Acceptance probability

$$
\alpha\left(m^{\prime} \mid m\right)=\min \left\{1, \frac{f\left(m^{\prime}\right)}{f(m)} \times \frac{f\left(y \mid m^{\prime}\right)}{f(y \mid m)} \times \frac{q\left(m \mid m^{\prime}\right)}{q\left(m^{\prime} \mid m\right)} \times\left|\frac{\partial m^{\prime}}{\partial m}\right|\right\}
$$

Dimension-jumping proposals: $\beta_{1}$ ( $p$-vector $) \rightarrow \beta_{2}((p+1)$-vector)

$$
\begin{aligned}
\eta & =B_{1} \beta_{1}=B_{2} \beta_{2}^{*} \\
\Rightarrow \hat{\beta}_{2}^{*} & =\left[\left(B_{2}^{\prime} B_{2}\right)^{-1} B_{2}^{\prime} B_{1}\right] \beta_{1}=G \beta_{1}
\end{aligned}
$$

$$
\begin{aligned}
\beta_{2} & =\left[\begin{array}{l|l}
G & \vdots \\
0 \\
1
\end{array}\right] \times\left[\begin{array}{c}
\beta_{1} \\
u
\end{array}\right] \\
u & \sim N(0, \bullet)
\end{aligned}
$$

## Adaptive splines: GP parameter estimates



[^1]
## Partition model

- Pros \& cons

■ Naturally local, nD

- Piecewise constant
- Estimate

■ Number of cells

- Centroid locations
- Cell coefficients

■ Reversible-jump MCMC

- Birth-death
- Detailed balance


■ Green [1995], Heikkinen and Arjas [1998], Denison et al. [2002], Costain [2008], Bodin and Sambridge [2009]

## Partition model: GP parameter estimates



[^2]
## Qualitative comparison of different estimates <br> P-splines: $n_{\xi}=6 \times 6, n_{\nu}=6 \times 6$ <br> Adaptive splines: $n_{\xi}^{m o}=3 \times 3, n_{\nu}^{m o}=4 \times 4$





Partition: $n_{\xi}^{m o}=1, n_{\nu}^{m o}=7$



## Summary

- Covariate effects important in environmental extremes

■ Need to tackle big problems $\Rightarrow$ need efficient models
■ Need to provide solutions as "end-user" software $\Rightarrow$ stable inference
■ P-splines: straightforward, global roughness per dimension

- Adaptive splines: optimally-placed knots

■ All splines: nD basis is tensor product of marginal bases

- Partition: piecewise constant, naturally nD
- Partition mixture model
- Combinations useful
- Conditional, spatial and temporal extremes


## References

C N Behrens, H F Lopes, and D Gamerman. Bayesian analysis of extreme events with threshold estimation. Stat. Modelling, 4:227-244, 2004.
C. Biller. Adaptive bayesian regression splines in semiparametric generalized linear models. J. Comput. Graph. Statist., 9:122-140, 2000.

Thomas Bodin and Malcolm Sambridge. Seismic tomography with the reversible jump algorithm. Geophysical Journal International, 178:1411-1436, 2009.
A. Brezger and S. Lang. Generalized structured additive regression based on Bayesian P-splines. Comput. Statist. Data Anal., 50:967-991, 2006.
D. A. Costain. Bayesian partitioning for modeling and mapping spatial case?control data. Biometrics, 65:1123-1132, 2008.
D. G. T. Denison, N. M. Adams, C. C. Holmes, and D. J. Hand. Bayesian partition modelling. Comput. Stat. Data Anal., 38:475-485, 2002.
I. DiMatteo, C. R. Genovese, and R. E. Kass. Bayesian curve-fitting with free-knot splines. Biometrika, 88:1055-1071, 2001.
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.
P.J. Green. Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. Biometrika, 82:711-732, 1995.

Juha Heikkinen and Elja Arjas. Non-parametric bayesian estimation of a spatial poisson intensity. Scand. J. Stat., 25:435-450, 1998.
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.
D. Randell, K. Turnbull, K. Ewans, and P. Jonathan. Bayesian inference for non-stationary marginal extremes. Environmetrics, 27:439-450, 2016.
G. O. Roberts and O. Stramer. Langevin diffusions and Metropolis-Hastings algorithms. Methodology and Computing in Applied Probability, 4:337-358, 2002.
G. Wallstrom, J. Liebner, and R. E. Kass. An implementation of Bayesian adaptive regression splines (BARS) in C with S and R wrappers. Journal of Statistical Software, $26,2008$.
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.
S. Zhou and X. Shen. Spatially adaptive regression splines and accurate knot selection schemes. J. Am. Statist. Soc., 96:247-259, 2001.

## Supporting material

## Partition model: $\psi$



## Partition model: $\xi$ and $\nu$ traces















[^0]:    Copyright of Shell

[^1]:    Copyright of Shell

[^2]:    Copyright of Shell

