Extremes for metocean: Multivariate

Philip Jonathan

Lancaster University

IOGP, London, September 2026
(slides at *www.lancs.ac.uk/∼jonathan*)

Overview of masterclass

- A: Fundamentals
- B: Covariates
- C: Multivariate
- D: covXtreme
- E: Recent developments & take-aways

Overview of C: Multivariate

- Motivating examples
- Understanding extremal dependence
- Conditional extremes
- What can go wrong?
- Further reading

All models are wrong, some models are useful

Motivating examples

- "Associated values" e.g. for $U_{10}|\{H_S^{sp} = 100\text{-year level}\}$
- "Weibull-log normal model" for the joint distribution (H_S, T_P)
- Copula models
- Temporal evolution of sea state H_S during a storm
- Spatial dependence of extreme H_S^{sp} within a hurricane or typhoon

- All these problems address the extremal dependence between two or more variables

Bivariate extremes

Basics

$$
\begin{aligned}
\text{CDF: } F(x,y) &= C(u,v) \quad \text{where } u = F_X(x) \text{ and } v = F_Y(y) \\
\text{Density: } f_{X,Y}(x,y) &= f_X(x) f_Y(y) \times c(u,v)
\end{aligned}
$$

- C and c are copula CDF and density, f_X and f_Y are marginal densities
- We are interested in what happens to C and c when u and v are large

Classic examples

- Bivariate Gaussian: $f(x,y) = \phi(x,y;\mathbf{0},[1,\rho;\rho,1])$, where ϕ is the bivariate Gaussian density with correlation $\rho \in [-1,1]$
- Bivariate logistic: $C(u,v) = \exp\left(-\left[(-\log u)^{1/\alpha} + (-\log v)^{1/\alpha}\right]^{\alpha}\right)$, where dependence $\alpha \in (0,1]$

Extremal dependence

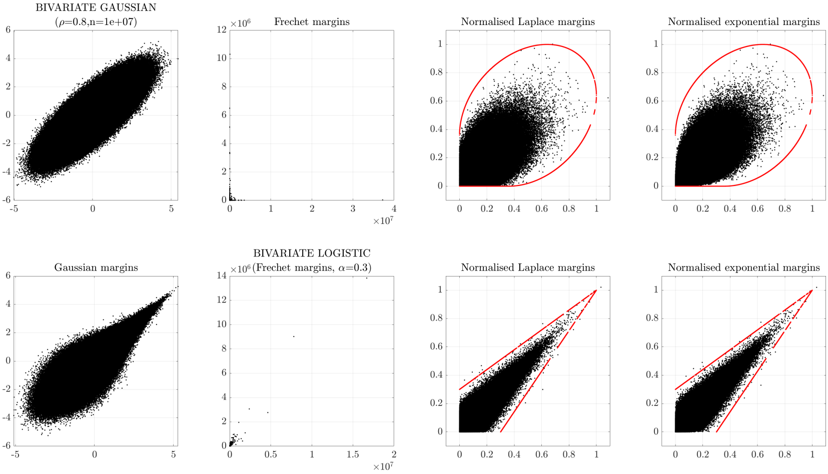

- Frechet $F(x) = \exp(-1/x)$, $x > 0$; Laplace $F(x) = \exp(x)/2$, $x < 0$ and $= 1 - \exp(-x)/2$, $x \geq 0$; exponential $F(x) = 1 - \exp(-x)$, $x \geq 0$
- Gaussian never touches (1,1) in normalised space, but the logistic can!
- What does this mean?

Extremal dependence

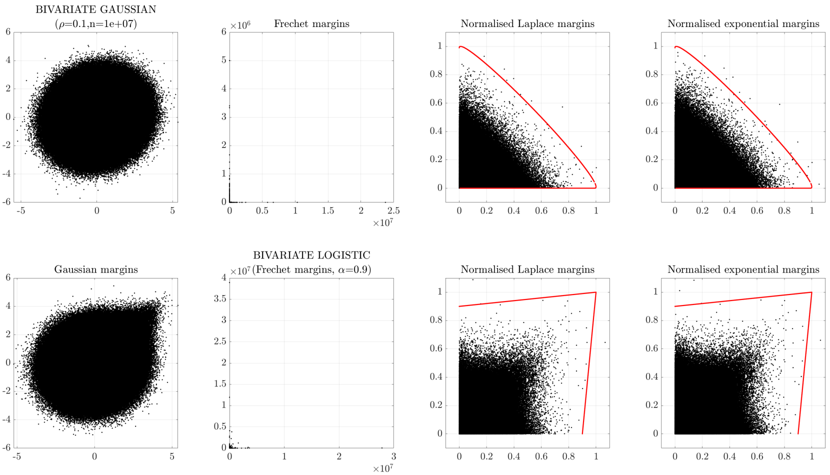

- Gaussian never touches (1,1) in normalised space, but the logistic can!

Asymptotic dependence (AD) and asymptotic independence (AI)

- Intuition
  - AD: both variables tend to get large together
  - AI: variables tend not to get large together

- Bivariate Gaussian exhibits AI for all correlations $\rho < 1$
- Bivariate logistic exhibits AD for all dependences $\alpha < 1$

- When we fit multivariate extreme value models, we must be able to distinguish AD and AI
  - Lots of copula models assume AD, and may be inappropriate for many environmental extreme value applications
  - Using an inappropriate dependence model causes bias
  - Classical multivariate extreme value theory (using "componentwise maxima") always leads to AD
- "Mea copula"
  - Financial crash 2007-8 (see MacKenzie and Spears 2014)
  - Proliferation of Gaussian copulas (assuming AI inappropriately)

Quantifying AD and AI

Variables $U = F_X(X)$ and $V = F_Y(Y)$ on uniform margins

χ quantifies AD

- $\chi(u) = \mathbb{P}(V > u \mid U > u) = \frac{\mathbb{P}(U>u, V>u)}{\mathbb{P}(U>u)} = 2 - \frac{\log C(u,u)}{\log(u)}$
- $\chi(u) \to \chi$ as $u \to \infty$
- $\chi = 1$: perfect dependence
- $\chi \in (0, 1)$: AD
- $\chi = 0$: AI or perfect independence

$\bar{\chi}$ quantifies AI

- $\bar{\chi}(u) = 2\frac{\log \mathbb{P}(U>u)}{\log \mathbb{P}(U>u, V>u)} - 1$
- $\bar{\chi}(u) \to \bar{\chi}$ as $u \to \infty$
- $\bar{\chi} = 1$: AD or perfect dependence
- $\bar{\chi} \in (0, 1)$: AI
- $\bar{\chi} = 0$: perfect independence

Practical implication

- Estimate $\chi(u)$ and $\bar{\chi}(u)$ from data to test for AD and AI

AI in practice

- Many environmental extremes problems involve AI, at least in part ... bivariate Gaussian is the obvious example
- But convergence of $\chi(u)$ and $\bar{\chi}(u)$ to χ and $\bar{\chi}$ is very slow!

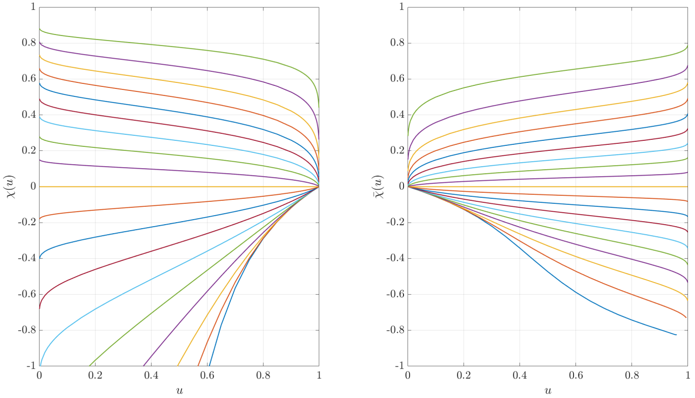

$\chi(u)$ and $\bar{\chi}(u)$ for bivariate Gaussian ($\Rightarrow \chi = 0, \bar{\chi} = \rho$); colours are correlations ρ on -0.9, -0.8, ..., 0.9 (Recreated from Coles et al. 1999)

Conditional extremes

- Conditional extremes gives a straightforward approach to estimating multivariate tails, admitting both AD and AI
- Used widely in flood modelling

- $\boldsymbol{X} = (X_1, ..., X_j, ..., X_p)$ consists of p random variables
- Each X and Y have standard Laplace margins ($f(x) = \exp(-|x|)/2$, $x \in \mathbb{R}$)
- Seek a model for $\boldsymbol{X}|(Y = y)$ for $y > u$

- Assume we can find p-dimensional scaling $\boldsymbol{a}, \boldsymbol{b} > \boldsymbol{0}$ such that

$$\begin{aligned} \mathbb{P}(\boldsymbol{Z} \leq \boldsymbol{z} | Y = y) \quad &\to \quad G(\boldsymbol{z}) \quad \text{as} \quad u \to \infty \\ \text{for} \quad \boldsymbol{Z} \quad &= \quad \frac{\boldsymbol{X} - \boldsymbol{a}(y)}{\boldsymbol{b}(y)} \end{aligned}$$

- Non-degenerate G is unknown, and estimated empirically
- Similar thinking to that used to stabilise distribution of maxima

Conditional extremes

- Typical scaling is $\boldsymbol{a} = \boldsymbol{\alpha} y$ and $\boldsymbol{b} = y^{\boldsymbol{\beta}}$, $\boldsymbol{\alpha} \in [-1, 1]^p$, $\boldsymbol{\beta} \in (-\infty, 1]^p$
- So simply fit regression model

$$\boldsymbol{X} | (Y = y) = \boldsymbol{\alpha} y + y^{\boldsymbol{\beta}} \boldsymbol{Z} \quad \text{for} \quad y > u$$

- $\alpha = 1$, $\beta = 0$: AD or perfect dependence
- $\alpha \in (0, 1)$: AI

- Heffernan and Tawn [2004] find choices for α and β for popular bivariate cases; e.g. bivariate Gaussian : $\alpha = \rho^2$, $\beta = 1/2$
- Simulate under the fitted model to predict e.g.
  - Conditional distribution of T_P^{sp} given than H_S^{sp} exceeds its 1,000-year value
  - Joint conditional distribution of H_S^{sp} and T_P^{sp} given than U_{10}^{sp} exceeds its 1,000-year value

Fitting a simple conditional extremes model

- Fit models to physical-scale $\dot{X}$ and $\dot{Y}$
  - We need full-sample marginal models; we could use gamma-GP
- Transform sample from physical-scale to Laplace-scale X and Y
  - Using $F_{\text{Laplace}}(x) = F_{\text{GmmGP}}(\dot{x})$ (called the probability integral transform)
- Fit the model $X|(Y = y) = \alpha y + y^{\beta}(\mu + \sigma W)$ for $y > u$
  - Explore stability of extimated α as a function of u
  - Since distribution of W is unknown, assume it is standard Gaussian for model fitting, then estimate it empirically from fit residuals for simulation purposes
- Simulate new values of Y and $X|(Y = y)$ using the fitted model on Laplace scale
- Transform the simulated data back to physical scale to estimate conditional return values etc.

Illustrative simulation

- Sample of H_S^{sp} and T_P^{ass}
- Want to estimate the distribution of T_P^{ass} for values of H_S^{sp} much larger than those seen in the data

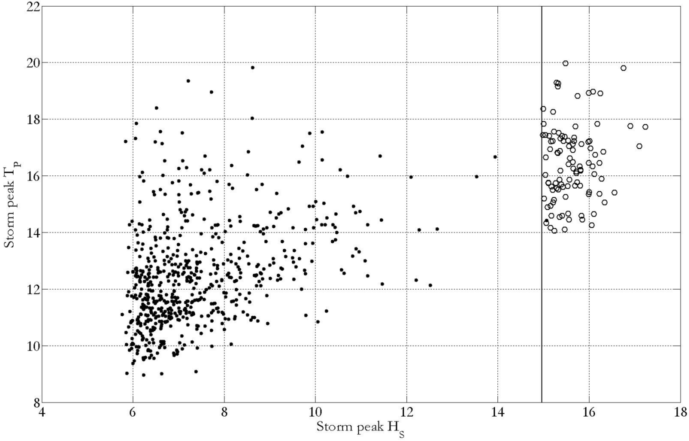

Fitting the simple conditional extremes model

- W assumed to follow a standard Gaussian
- Density for a single pair x, y on Laplace scale, with $y > u$

$$f_{\mathrm{CE}}(x|y,\alpha,\beta,\mu,\sigma) = \phi(x;\alpha y + \mu y^{\beta}, (\sigma y^{\beta})^2)$$

  with parameters $\alpha \in [-1,1]$, $\beta \in (-\infty, 1]$, $\mu \in \mathbb{R}$ and $\sigma > 0$
- Sample likelihood for pairs $\{x_i, y_i\}_{i=1}^{n}$ with $y_i > u$

$$\mathcal{L}_{\mathrm{CE}}(\alpha,\beta,\mu,\sigma|\boldsymbol{x},\boldsymbol{y}) = \prod_{i=1}^{n} f_{\mathrm{CE}}(x_i|y_i,\alpha,\beta,\mu,\sigma)$$

- Maximum likelihood optimisation

Non-stationary conditional extremes

- We can allow α, β, μ and σ to vary with direction (say)
- Example from northern North Sea
- Details in Jonathan et al. [2014]

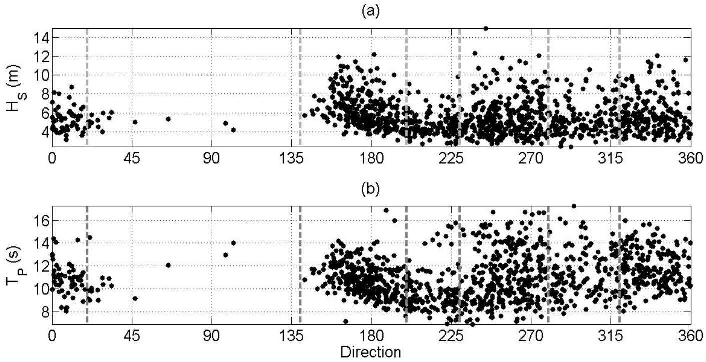

H_S^{sp} and T_P^{ass}

Non-stationary conditional extremes

- Dependence between H_S^{sp} and T_P^{ass} varies with direction

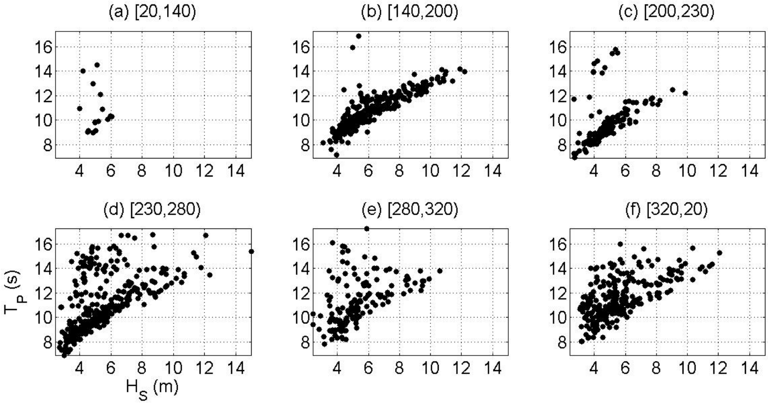

Fitting the non-stationary conditional extremes model

- α, β, μ and σ now vary with direction θ on $\mathcal{D}$
- Density for a single pair (x_i, y_i) with $y_i > u$

$$f_{\mathrm{CE}}(x_i|y_i, \alpha_{\mathcal{I}(i)}, \beta_{\mathcal{I}(i)}, \mu_{\mathcal{I}(i)}, \sigma_{\mathcal{I}(i)}) = \phi(x_i; \alpha_{\mathcal{I}(i)} y + \mu_{\mathcal{I}(i)} y^{\beta_{\mathcal{I}(i)}}, (\sigma_{\mathcal{I}(i)} y^{\beta_{\mathcal{I}(i)}})^2)$$

- Sample likelihood for pairs $\{x_i, y_i\}_{i=1}^{n}$ with $y_i > u$

$$\mathcal{L}_{\mathrm{CE}}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\mu}, \boldsymbol{\sigma}|\boldsymbol{x}, \boldsymbol{y}) = \prod_{i=1}^{n} f_{\mathrm{CE}}(x_i|y_i, \alpha_{\mathcal{I}(i)}, \beta_{\mathcal{I}(i)}, \mu_{\mathcal{I}(i)}, \sigma_{\mathcal{I}(i)})$$

  for parameter vectors $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$, with $\boldsymbol{\alpha} = \{\alpha_j\}_{j=1}^{m}$ etc.
- Penalised likelihood

$$\mathcal{L}^{*}_{\mathrm{CE}}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\mu}, \boldsymbol{\sigma}|\boldsymbol{x}, \boldsymbol{y}) = \mathcal{L}_{\mathrm{CE}}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\mu}, \boldsymbol{\sigma}|\boldsymbol{x}, \boldsymbol{y}) - \sum_{\eta \in \{\alpha,\beta,\mu,\sigma\}} \lambda_\eta \boldsymbol{\gamma}'_\eta \boldsymbol{D}' \boldsymbol{D} \boldsymbol{\gamma}_\eta$$

  for spline coefficients $\boldsymbol{\gamma}$ and roughness coefficients λ
- Maximise $\mathcal{L}^{*}_{\mathrm{CE}}$, using cross-validation to select λ_α, λ_β, λ_μ and λ_σ

Parameter estimates

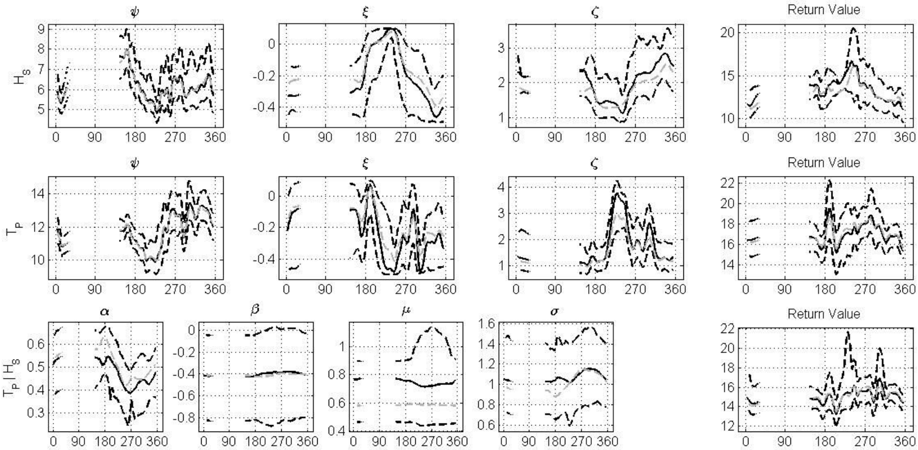

Summary of marginal and CE parameters with direction for fitted model

- Marginal threshold ψ, GP shape ξ and GP scale ζ
- Note difference between marginal and conditional return values for T_P^{ass}

Weibull-log normal model

- Developed explicitly for joint modelling of $Y(H_S)$ and $X(T_P)$ for large values of Y
  - See e.g. Haver [1987]
- Assumes that $X|\{Y = y\}$ follows a log-normal distribution:
  - $f_{X|Y}(x|y) = \frac{1}{\sqrt{2\pi}\sigma(y)x} \exp(-\frac{(\log x - \mu(y))^2}{2\sigma^2(y)})$
  - Estimates for parameters $\mu(y)$ and $\sigma(y)$ then regressed on y to extrapolate
  - Some explicit forms for $\mu(y)$ recommended
- Assumes Weibull model for Y over a suitably large threshold
  - We consider a modified approach in which the marginal model for Y is assumed GP

Extrapolating μ and σ in Haver model

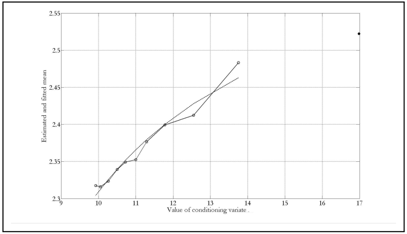

$\mu(y)$

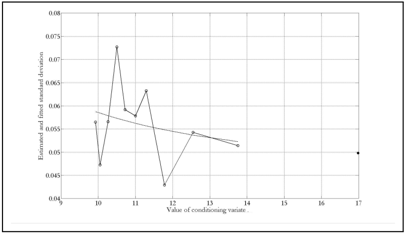

$\sigma(y)$

- Justification for functional forms of $\mu(y)$ and $\sigma(y)$ is empirical
  - Extrapolation is problematic
  - Yields more biased and variable estimates of extreme quantiles; see Jonathan et al. [2010]
- The conditional extremes model has a better theoretical underpinning

Conditional extremes: what can go wrong?

- Marginal transformation required $\Rightarrow$ all the issues associated with marginal modelling and covariates
- Selection of appropriate threshold(s) u the dependence modelling
  - Use diagnostics e.g. $\hat{\alpha}(u)$ on u
  - Model averaging probably better when choice of u not obvious
- Small sample size is a particular problem
- Complex tail structures requiring multiple covariates or very high u
- Slow convergence to asymptotic form

Implications

- Propagate and quantify uncertainty carefully
  - No excuses!
- Model averaging in place of hyper-parameter selection
- Use cross-validation and bootstrapping in frequentist analysis; otherwise use Bayesian inference
- See Towe et al. [2023] for quantification of uncertainties in estimation of associated values

All models are wrong, some models are useful

If you want to read more

Metocean-related extensions of conditional extremes

- Non-stationary conditional extremes: Jonathan et al. [2014]
- Spatial conditional extremes: Shooter et al. [2019], Shooter et al. [2021b], Shooter et al. [2021a], Shooter et al. [2022]
- Time-series conditional extremes: Winter and Tawn [2016], Winter and Tawn [2017], Tendijck et al. [2019], Tendijck et al. [2024]

Classic multivariate extreme value analysis

- A gentle introduction: Coles [2001]
- Copulas: Joe [2014]
- Spatial extremes: Padoan et al. [2010]

Metocean applications with angular-radial models

- Murphy-Barltrop et al. [2025], Mackay et al. [2026]
- More in final session, hopefully

References

S. Coles. *An introduction to statistical modelling of extreme values*. Springer, Berlin, 2001.

S Coles, J Heffernan, and J Tawn. Dependence measures for extreme value analyses. *Extremes*, 2:339–365, 1999.

S. Haver. On the joint distribution of heights and periods of sea waves. *Ocean Eng.*, 14:359–376, 1987.

J. E. Heffernan and J. A. Tawn. A conditional approach for multivariate extreme values. *J. Roy. Statist. Soc. B*, 66:497–546, 2004.

H. Joe. *Dependence modelling with copulas*. CRC Press, 2014.

P. Jonathan, J. Flynn, and K. C. Ewans. Joint modelling of wave spectral parameters for extreme sea states. *Ocean Eng.*, 37: 1070–1080, 2010.

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

E. Mackay, C. J. R. Murphy-Barltrop, J. Richards, and P. Jonathan. Deep learning joint extremes of metocean variables using the SPAR model. *J. Offshore Mech. Arct. Eng.*, 148:021201, 2026.

D. MacKenzie and T. Spears. A device for being able to book P&L: the organizational embedding of the Gaussian copula. *Soc. Stud. Sci.*, 44:418–440, 2014.

C. J. R. Murphy-Barltrop, E. Mackay, and P. Jonathan. Inference for bivariate extremes via a semi-parametric angular-radial model. *Extremes*, 28:209–238, 2025.

S. A. Padoan, M. Ribatet, and S. A. Sisson. Likelihood - based inference for max - stable processes. *J. Am. Statist. Soc.*, 105:263–277, 2010.

R. Shooter, E. Ross, J. A. Tawn, and P. Jonathan. On spatial conditional extremes for ocean storm severity. *Environmetrics*, 30:e2562, 2019.

R. Shooter, E Ross, A. Ribal, I. R. Young, and P. Jonathan. Spatial conditional extremes for significant wave height from satellite altimetry. *Environmetrics*, 32:e2674, 2021a.

R Shooter, J A Tawn, E Ross, and P Jonathan. Basin-wide spatial conditional extremes for severe ocean storms. *Extremes*, 24: 241–265, 2021b.

R. Shooter, E Ross, A. Ribal, I. R. Young, and P. Jonathan. Multivariate spatial conditional extremes for extreme ocean environments. *Ocean Eng.*, 247:110647, 2022.

S. Tendijck, E. Ross, D. Randell, and P. Jonathan. A non-stationary statistical model for the evolution of extreme storm events. *Environmetrics*, 30:e2541, 2019.

S. Tendijck, P. Jonathan, D. Randell, and J. A. Tawn. Temporal evolution of the extreme excursions of multivariate kth order Markov processes with application to oceanographic data. *Environmetrics*, 35:e2834, 2024.

R Towe, D Randell, J Kensler, G Feld, and P Jonathan. Estimation of associated values from conditional extreme value models. *Ocean Eng.*, 272:113808, 2023.

H. C. Winter and J. A. Tawn. Modelling heatwaves in central France: a case-study in extremal dependence. *J. Roy. Statist. Soc. C*, 65:345–365, 2016.

H. C. Winter and J. A. Tawn. kth-order Markov extremal models for assessing heatwave risks. *Extremes*, 20:393–415, 2017.