

Highlights

Multivariate spatial conditional extremes for extreme ocean environments

Rob Shooter, Emma Ross, Agustinus Ribal, Ian R. Young, Philip Jonathan

- Characterise joint extremal spatial dependence between multiple metocean variables.
- Methodology motivated by model of Heffernan and Tawn (2004).
- Satellite scatterometer (Metop) and hindcast data (Reistad, Breivik, Haakenstad, Aarnes, Furevik and Bidlot 2011) incorporated.
- Marginal directional and seasonal covariate effects accommodated.
- Extremal spatial dependence decays over 600-800 km for conditioning on an extreme occurrence of wind speed.

Multivariate spatial conditional extremes for extreme ocean environments

Rob Shooter^a, Emma Ross^b, Agustinus Ribal^c, Ian R. Young^d and Philip Jonathan^{e,f,*}

^aMet Office Hadley Centre, FitzRoy Road, Exeter, EX1 3PB, UK.

^bShell Global Solutions International BV, 1031 HW Amsterdam, The Netherlands.

^cDepartment of Mathematics, Faculty of Mathematics and Natural Sciences, Hasanuddin University, Makassar 90245, Indonesia.

^dDepartment of Infrastructure Engineering, University of Melbourne, Melbourne, VIC 3010, Australia.

^eDepartment of Mathematics and Statistics, Lancaster University, LA1 4YW, UK.

^fShell Research Ltd., London SE1 7NA, UK.

ARTICLE INFO

Keywords:

Conditional extremes
Covariate effects
Extremal spatial dependence
Joint extremes
Metop

ABSTRACT


The joint extremal spatial dependence of wind speed and significant wave height in the North East Atlantic is quantified using Metop satellite scatterometer and hindcast observations for the period 2007-2018, and a multivariate spatial conditional extremes (MSCE) model, ultimately motivated by the work of Heffernan and Tawn (2004). The analysis involves (a) registering individual satellite swaths and corresponding hindcast data onto a template transect (running approximately north-east to south-west, between the British Isles and Iceland), (b) non-stationary directional-seasonal marginal extreme value analysis at a set of registration locations on the transect, (c) transformation from physical to standard Laplace scale using the fitted marginal model, (d) estimation of the MSCE model on the set of registration locations, and assessment of quality of model fit. A joint model is estimated for three spatial quantities: Metop wind speed, hindcast wind speed and hindcast significant wave height. Results suggest that, when conditioning on extreme Metop wind speed, extremal spatial dependence for all three quantities decays over approximately 600-800 km.

1. Introduction

Severe ocean events at a location often involve extreme values for more than one oceanographic variable, leading to interest in the development and tailoring of statistical methods for multivariate and conditional extremes in a met-ocean context; for example, extreme loads on an offshore structure often arise from a combination of wind, wave and current. Further, simultaneous occurrences of extreme ocean events at different locations often represent higher risk than extreme events at one location; for example, extreme storm events often impact multiple structures simultaneously, requiring integrated planning of activities such as unmanning. Storms load coastal defences over a spatial neighbourhood, rather than at a single location. Design and maintenance of a wind farms requires characterisation of the joint spatial field for metocean variables over a spatial neighbourhood. This leads to interest in the development of statistical models for the joint characteristics of multiple spatial fields of variables such as wind speed and significant wave height. The objective of this article is to develop a straightforward model for Multivariate Spatial Conditional Extremes (MSCE), and demonstrate its usefulness in met-ocean application. Specifically, we seek a model for the joint behaviour of multiple met-ocean variables in space, given an occurrence of an extreme value of one of the variables at some location. The idea underpinning the methodology developed is the conditional extremes model of Heffernan and Tawn (2004). As discussed in Section 3, the MSCE model can be seen as the latest in a sequence of extensions of the conditional extremes approach to address specific applications' requirements. The main use of the MSCE model is to provide a principled approach to characterising the joint structure of spatial fields, to be applied in any situation where environmental hazard results from extreme values of at least one metocean variable.

Like its conditional extremes predecessor, inference using MSCE amounts essentially to fitting a non-linear regression model for observations of random variables on a standard Laplace marginal scale. The model admits different types

*Corresponding author

 philip.jonathan@shell.com (P. Jonathan)

ORCID(s): 0000-0002-2529-0249 (R. Shooter); 0000-0002-0287-0611 (E. Ross); 0000-0002-3286-5565 (A. Ribal); 0000-0003-2233-9227 (I.R. Young); 0000-0001-7651-9181 (P. Jonathan)

of extremal dependence (including asymptotic dependence and asymptotic independence, Coles, Heffernan and Tawn 1999) and is computationally rather straightforward to estimate.

Other statistical approaches to spatial extremes are motivated by the theory of max-stable processes (MSPs; see e.g., Brown and Resnick 1977, Smith 1990, Schlather 2002, Davison, Padoan and Ribatet 2012, Ribatet 2013 and the recent review of Huser and Wadsworth 2020). However, typical MSP models require an assumption that the extremal spatial dependence takes a particular form. Some models in principle are able to permit different classes of extremal dependence (e.g., Wadsworth, Tawn, Davison and Elton 2017) but can be fairly restrictive or computationally unwieldy.

Here, we are interested not only in quantifying the extremal spatial dependence for quantities such as wind speed and significant wave height, but also the cross-dependence between different quantities. In this context, the hierarchical max-stable spatial model of Reich and Shaby (2012) has been extended (Vettori 2017, Reich and Shaby (2018), Vettori, Huser and Genton 2019) to the multivariate case. Further, Genton, Padoan and Sang (2015) introduced multivariate max-stable spatial processes. These approaches certainly have their merits, but also have at least some of the limitations outlined above pertaining to MSP models.

Estimation using the MSCE model for a sample of spatial data for a number of different quantities, first requires that the data are presented on a common marginal standard Laplace scale for all quantities at all locations. This transformation is achieved by estimating non-stationary directional-seasonal extreme value models for each quantity at each location (e.g. Zanini, Eastoe, Jones, Randell and Jonathan 2020), and then application of the probability integral transform.

The layout of the article is as follows. Section 2 provides a description of the motivating application to wind speed and significant wave height in the North East Atlantic, satellite and hindcast data sources, and data pre-processing steps. Section 3 then presents the MSCE methodology. Results of applying the MSCE model in the North East Atlantic are given in Section 4, and Section 5 provides a discussion. Supporting plots for the exploratory data analysis are given in the appendix.

2. Motivating application

We seek to characterise the joint spatial conditional structure of extreme values of wind speed and significant wave height in the North East Atlantic between the British Isles and Iceland, subject to systematic directional and seasonal variability. In this section, we describe data sources, and data pre-processing necessary prior to MSCE inference. Two sources of wind data are available, namely satellite-observed wind speed $St1Wnd$ and direction, and hindcast model wind speed $HndWnd$ and direction. Hindcast significant wave height $HndWav$ and wave direction are also used. These data sources are combined to yield a set of observations of $St1Wnd$, $HndWnd$ and $HndWav$ and corresponding directional and seasonal data, on a transect of equally-spaced registration locations lying between approximately $2^\circ W$, $67^\circ N$ and $22^\circ W$, $55^\circ N$.

The objective of the analysis is to evaluate our ability to quantify the joint structure of spatial fields corresponding to significant wave height and wind speed, based on hindcast and satellite data, conditional on the occurrence of an extreme wind speed or significant wave height event. For the North East Atlantic, one use case for the model would be joint assessment of extreme loads on multiple offshore facilities, requiring joint characterisation of wind and wave fields, critical e.g. in the planning and execution of operations such as unmanning (Towe, Zanini, Randell, Feld and Jonathan 2021).

The appeal of direct earth observation by satellite is that it may provide, in the medium-to-long term, a rich source of data to inform ocean-related human activities in real time. At present, satellite data alone is not sufficient for design purposes of course because of poor temporal coverage and length of historical time-series for a location; hindcast data is clearly a more useful data source at present. However, the quality of hindcast output for extreme values of waves and wind at a location is also uncertain, usually requiring calibration to local measurements, especially to accommodate directional and other covariate effects. Similarly, satellite output currently requires calibration. For all these reasons, it is important to quantify how inferences for extremes from satellite observations compare with those from hindcasts, and to consider inferences for extremes which combine hindcast and satellite input.

Metop and hindcast data sources

Metop-A, B and C are polar-orbiting meteorological satellites, forming the EUMETSAT Polar System (EPS) series. Metop-A (launched 19 October 2006), Metop-B (launched 17 September 2012) and Metop-C (launched 7 November 2018) are in a low polar orbit, at an altitude of approximately 800 km over the Earth's surface. Metop uses the Advanced

SCATterometer (ASCAT) to measure wind speed and direction over the oceans. ASCAT is a real aperture radar, operating at 5.255 GHz (C-band) and using vertically polarised antennas. It transmits long pulses of microwave energy with “chirp” linear frequency modulation towards the sea surface. Winds over the sea cause centimetre scale disturbances of the sea surface which modify radar back-scattering characteristics dependent on both wind speed and direction. Two sets of three antennas measure the characteristics of the back-scattered signal in two 500 km wide swaths to each side of the satellite ground track, make sequential observations of the back-scattering coefficient of each point of interest from three directions. The characteristics of the back-scattered signals allow estimation of surface wind speed and direction. Over the North East Atlantic, the daily pass of each Metop satellite is from the north-east to the south-west. All Metop satellite data for the current work were sourced from the Australian Ocean Data Network (AODN). These datasets have been calibrated and quality controlled as described by Ribal and Young (2020b) and Ribal and Young (2020a). At wind speeds above 30m/s scatterometers tend to saturate (Ribal, Tamizi and Young 2021). However, as shown by Ribal and Young (2020b), Metop-A and Metop-B yield unbiased data for wind speeds up to a minimum of 25m/s. Note that only Metop-A and Metop-B data are considered here.

The NORA10 hindcast (NORwegian ReAnalysis 10 km, Reistad et al. 2011) is a combined high-resolution atmospheric downscaling and wave hindcast for the Norwegian Sea, the North Sea and the Barents Sea, based on the European Reanalysis dataset (ERA-40), outputting 3-hourly wave fields at a resolution of 10 km for the period 1957–2018.

Registered data

Figure 1 gives the locations of 14 equally-spaced registration locations of the registration transect selected, lying in the North East Atlantic between the British Isles and Iceland. The end points of the registration transect were chosen so that the transect covers the greatest number of satellite passes, thus maximising the size of the Metop sample for MSCE analysis. For each available satellite pass, we find the nearest point on the satellite transect to each of the registration locations, and allocate the value of wind speed and direction from the “matched” satellite location to the registration location. If the maximum “matched” distance (calculated using the spherical law of cosines) for any registration location corresponding to a specific satellite pass exceeds 50 km, the pass is not registered. Hindcast wind speed and direction, significant wave height and wave direction are registered similarly, based on matching spatially between the hindcast grid and the registration transect at the times of already-registered satellite passes. In this way, a total of 1532 joint observations of $St1Wnd$, $HndWnd$ and $HndWav$ were isolated for analysis.

Referring to Figure 1, the south-west location is taken as the conditioning location (square), and all other locations (discs) are used as remote locations. The colour-coding scheme used indicates that the conditioning quantity is $St1Wnd$ (green) at the conditioning location, but that all of $St1Wnd$ (green), $HndWnd$ (orange) and $HndWav$ (blue) are present in the MSCE model as conditioned variates at remote locations.

Figure 2 shows the dependence between each of these quantities at selected remote locations and the conditioning variate $St1Wnd$ at the conditioning location r_j , $j = 0$. The dependence between $St1Wnd$ at a given remote location (r_j , $j = 1, 2, \dots, 13$), and $St1Wnd$ at location r_0 , is similar to that between $HndWnd$ at location r_j and $St1Wnd$ at location r_0 . There is evidence for curvature in the relationship between $HndWav$ at location r_j and $St1Wnd$ at location r_0 , reflecting the typical drag-type squared relationship between wind forcing and resulting significant wave height. There is also evidence that the dependence between $HndWav$ at location r_j and $St1Wnd$ at location r_0 decays more slowly with increasing distance (or increasing r_j) than that between wind speed at r_j and $St1Wnd$ at location r_0 . Again, this is indicative of greater spatial dependence for significant wave height in general compared to wind speed. Figures A1–A3 in the appendix provide supporting scatter plots illustrating the spatial dependence for each of $St1Wnd$, $HndWnd$ and $HndWav$ fields separately.

Marginal transformation to standard Laplace scale

The registered data also show systematic variation with direction and season, as illustrated in Figure 3. Seasonal effects on all three quantities are clear in columns 4–6 of the figure, for all locations considered. The directional effect is more interesting. For locations r_0 , r_3 and r_6 , the prevailing direction for wind and H_S is from the west as might be expected in the north Atlantic. However, at locations r_9 and r_{13} , the directional dependence is more confused; here for significant wave height, there is evidence for large values emanating from the Norwegian Sea to the north or from the North Sea to the south.

Figure 3 makes clear that, to transform the registered data from physical to standard Laplace scale ready for MSCE inference requires fitting a non-stationary marginal model to the data for each of $St1Wnd$, $HndWnd$ and $HndWav$ at each registration location, so that the effects of directional and seasonal variation are captured adequately. Non-stationary

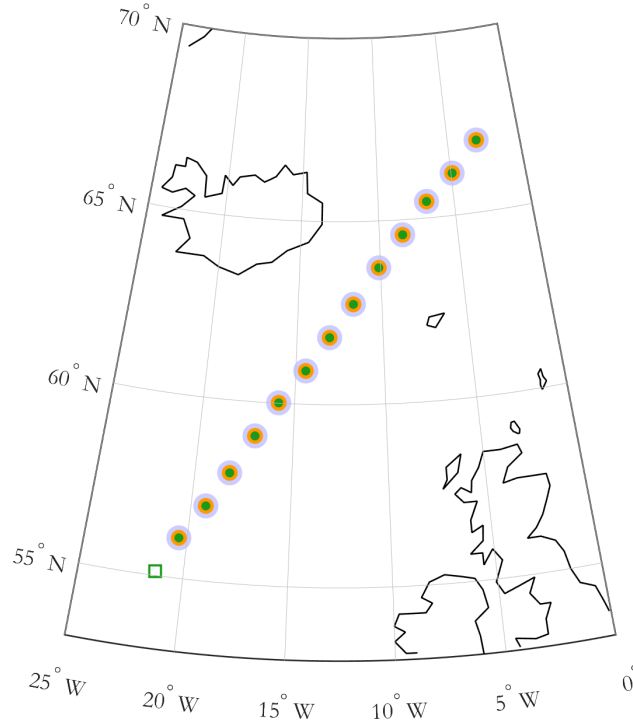


Figure 1: Map of the registration locations. Squares indicate the conditioning location r_j , $j = 0$, and discs indicate other (remote) locations r_j , $j = 1, 2, \dots, 13$ with increasing distance from location r_0 . Colour coding indicates that the conditioning quantity is St1Wnd (green square, at location r_0), and that all of St1Wnd (green), HndWnd (orange) and HndWav (blue) are used as conditioned quantities (discs) at remote locations. The Metop satellite trajectories are from north-east to south-west. (The green-orange-blue colour scheme appears as three shades of grey with decreasing intensity when viewed in black-and-white.)

directional-seasonal marginal extreme value models are estimated using the PPC (penalised piecewise constant) extreme value methodology, described previously in Ross, Sam, Randell, Feld and Jonathan (2018). Details of the analysis are withheld for brevity, in favour of the brief description here. The directional covariate domain is partitioned into eight octants centred on cardinal and semi-cardinal directions, and the seasonal covariate domain partitioned into “summer” and “winter” intervals centred on “seasonal degrees” 0° and 180° . This yields a partition of the full directional-seasonal covariate domain into 16 directional-seasonal “bins”. Then a piecewise constant generalised Pareto extreme value model for peaks over threshold is estimated simultaneously for all bins, such that the generalised Pareto shape parameter is constant everywhere on the covariate domain, and the generalised Pareto scale parameter is assumed constant within each bin. Further, the variation of the generalised Pareto scale parameter between bins is penalised. The extent of penalisation is regulated using cross-validation, to give the best “out-of-sample” predictive performance using the model. Moreover, a bootstrap scheme admitting a range of different plausible extreme value thresholds is incorporated, such that model uncertainty can be quantified reasonably. Software for the analysis is available at ESCADES-GITHub (2019).

Models are estimated independently for each of St1Wnd, HndWnd and HndWav at each registration location. Using simulation under the fitted model, the 100-year maximum value of St1Wnd is around 35 ms^{-1} , with 95% uncertainty interval of approximately $(28,45) \text{ ms}^{-1}$ at each of the registration locations. The corresponding values for HndWnd are similar. For HndWav, there is evidence that the 100-year maximum value reduces from around 22 m, (16,30) m at south-western locations (corresponding to small r_j) to around 17 m, (13,26) m at north-eastern locations. These values are in general agreement with expectations and previous estimates at these locations.

Using the fitted models with bootstrap median parameter estimates, and the probability integral transform, the registered data are transformed to standard Laplace scale, independently for each quantity at each registration location. The resulting Laplace-scale data are shown in Figure 4, and form the input for the MSCE inference discussed in Sections 3

Multivariate spatial conditional extremes

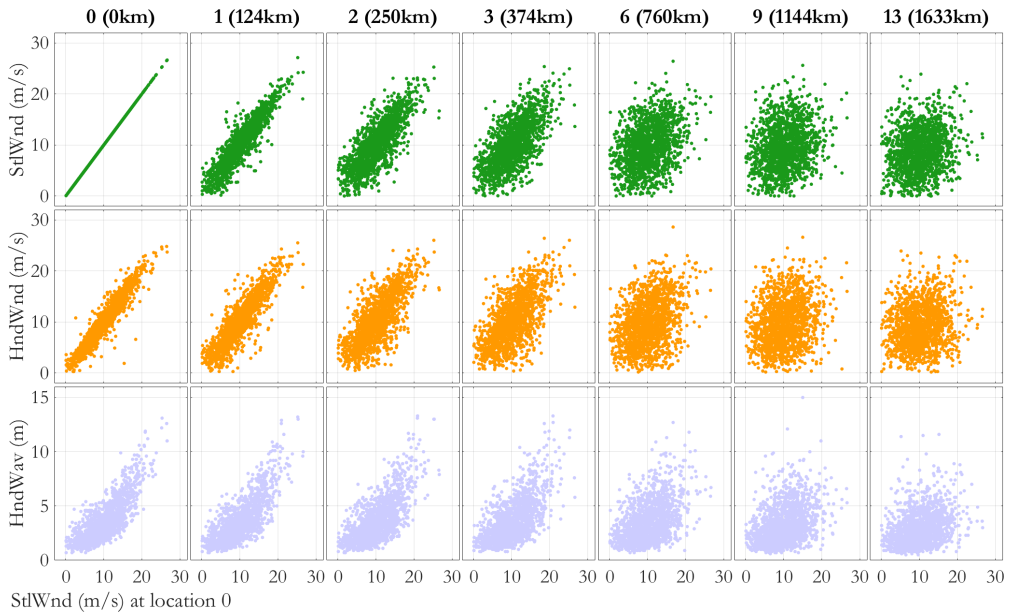


Figure 2: Scatter plots of StlWnd (green), HndWnd (orange) and HndWav (blue; all on physical scale) at selected locations r_j , $j = 0, 1, 2, 3, 6, 9, 13$ against StlWnd at the conditioning location r_j , $j = 0$.




Figure 3: Directional and seasonal covariate effects in registered data (on physical scale) for selected locations. The first three columns give directional variation for each of StlWnd (green), HndWnd (orange) and HndWav (blue) at locations r_j , $j = 0, 3, 6, 9, 13$. The corresponding seasonal variation is illustrated in the fourth to sixth columns. Direction is defined as the direction from which fluid flows, measured clockwise from north. Season is defined as the day of the year (1, 2, ..., 365 or 1, 2, ..., 366 for leap years) mapped linearly on to the interval (0, 360].

and 4 below. The characteristics of Figure 4 are discussed further in motivating the results in Sections 4.

The conditional extremes model of Heffernan and Tawn (2004) and derivatives is motivated by asymptotic arguments. Such arguments in multivariate extreme value theory generally assume that variables have common marginal distributions. The conditional extremes model form is most simply represented when that marginal distribution is the standard Laplace (Keef, Papastathopoulos and Tawn 2013a). From an applications perspective, the marginal transformation from physical to Laplace scales, achieved using the PPC model, also allows the effects of covariates to be captured.

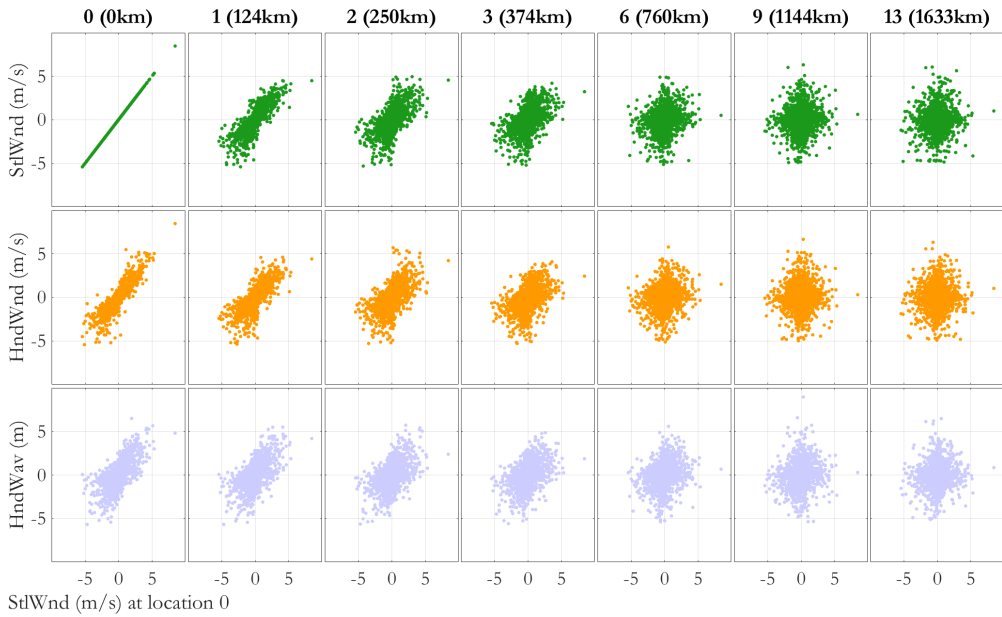


Figure 4: Scatter plots of StlWnd (green), HndWnd (orange) and HndWav (blue; all on standard Laplace scale) at selected locations r_j , $j = 0, 1, 2, 3, 6, 9, 13$ against StlWnd at the conditioning location r_j , $j = 0$. See Figure 2 for the corresponding plot on physical scale.

Hence, on Laplace scale following transformation using a non-stationary marginal extreme value model, the data are likely to exhibit little or no covariate dependence, and inference is likely to be more straightforward.

3. Methodology

3.1. Introduction

The multivariate spatial conditional extremes model is an extension of a sequence of earlier models for conditional extremes (Heffernan and Tawn 2004) and spatial conditional extremes (Tawn, Shooter, Towe and Lamb 2018, Shooter, Ross, Tawn and Jonathan 2019, Wadsworth and Tawn 2019, Shooter, Tawn, Ross and Jonathan 2021c, Shooter, Ross, Ribal, Young and Jonathan 2021b).

The underlying idea for all these models is a description of the conditional distribution of one or more conditioned variates Y given a large value $x > u$ of a conditioning variate X from Heffernan and Tawn (2004), motivated by asymptotic arguments as $u \rightarrow \infty$. For variables expressed on standard Laplace scale with positive association (as explained in Keef, Tawn and Lamb 2013b), this relationship can be written

$$Y|X = x = \alpha x + x^\beta Z, \quad x > u$$

for $\alpha \in (0, 1]$, $\beta \in (-\infty, 1]$, where the residual Z follows some unknown non-degenerate distribution G . For the purposes of parameter estimation, it is commonly assumed that $Z \sim N(\mu, \sigma^2)$ for $\mu \in \mathbb{R}$ and $\sigma > 0$, a normal distribution. Inference therefore involves estimation of α , β , μ and σ .

The methodology can be extended to multivariate conditional extremes by considering a vector $\mathbf{Y} = (Y_1, Y_2, \dots, Y_m)$ of m conditioned variates, with the model form becoming $\mathbf{Y}|X = x = \alpha x + x^\beta \mathbf{Z}$ for $x > u$. Vectors α , β now have elements $\alpha_k \in (0, 1]$, $\beta_k \in (-\infty, 1]$, and $\mathbf{Z} \sim G$ now represents the joint residual over m conditioned quantities. We might approximate this using a multivariate normal distribution $N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ with mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. Inference then involves estimation of α , β , $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$.

The methodology can also be extended to a spatial context, as described by Shooter et al. (2019) and Wadsworth and Tawn (2019), by considering a stationary spatial process $X(\cdot)$ on domain \mathbb{R} with standard Laplace marginal distributions. Then for a conditioning location r_0 and remote location r with separation $d = |r - r_0|$, we might write $X(r)|X(r_0) = x = \alpha(d)x + x^{\beta(d)}Z(d)$ for $x > u$. Inference for a finite set of locations r_0, r_1, \dots, r_p reduces to multivariate conditional extremes with a particular choice of parametric form for residual process \mathbf{Z} , to encode the fact

that as d increases, the dependence between locations will typically decrease to zero in environmental applications. In particular, for locations r_0, r corresponding to large d , we expect that $\alpha(d) \rightarrow 0$ and $\beta(d) \rightarrow 0$ so that $Z(d)$ must follow the standard Laplace distribution marginally. However, for small d , a multivariate normal choice for G might still be more appropriate. For this reason, a choice for G with marginal delta-Laplace (or generalised Gaussian) distribution is useful, admitting marginal standard normal and standard Laplace forms (Wadsworth and Tawn, 2019). The residual dependence structure of G might be represented by a (conditional) Gaussian field, as explained further in Section 3.2.

3.2. Outline of model

The MSCE model incorporates aspects of both multivariate conditional extremes and spatial conditional extremes. From the perspective of inference for observations at a finite set of locations r_0, r_1, \dots, r_p , the MSCE model reduces to the multivariate conditional extremes model with a specific parametric choice for the distribution G of the residual process Z .

Suppose that random variable $X_{j,k}$ represents quantity $k = 1, 2, \dots, m$ at location $r_j, j = 0, 1, \dots, p$ on standard Laplace scale, and that \mathbf{X} represents the set of ‘‘remote’’ random variables $X_{1,1}, X_{2,1}, \dots, X_{p,1}, X_{1,2}, X_{2,2}, \dots, X_{p,2}, \dots, X_{1,m}, \dots, X_{p,m}$. For brevity, we also write \mathbf{X} as $\{X_{j,k}\}$ for $(j, k) \in \mathcal{I}_{\text{Rmt}}$, for the ordered set

$$\mathcal{I}_{\text{Rmt}} = \{(1, 1), (2, 1), \dots, (p, 1), (1, 2), \dots, (p, 2), \dots, (1, m), \dots, (p, m)\}$$

and define the function $\mathcal{A}(j, k) = p(k - 1) + j$ which returns the location of the pair (j, k) in the ordered set \mathcal{I}_{Rmt} . Then, for a large value x of the conditioning random variable $X_{0,1}$ corresponding to quantity $k = 1$ at location r_0 , we assume we can write

$$\mathbf{X} | \{X_{0,1} = x\} = \boldsymbol{\alpha}x + x^\beta \mathbf{Z}, \quad x > u \quad (1)$$

where the vector of parameters $\boldsymbol{\alpha}$ has elements $\alpha_{\mathcal{A}(j,k)} \in (0, 1]$ for $(j, k) \in \mathcal{I}_{\text{Rmt}}$, the exponent vector $\boldsymbol{\beta}$ has elements $\beta_{\mathcal{A}(j,k)} \in (-\infty, 1]$, and element-wise multiplication of terms is assumed. We assume further that the residual process \mathbf{Z} follows a delta-Laplace distribution with conditional Gaussian covariance structure, parameterised as $\mathbf{Z} \sim \text{DL}(\boldsymbol{\mu}, \boldsymbol{\sigma}^2, \boldsymbol{\delta}; \boldsymbol{\Sigma})$ for mean vector $\boldsymbol{\mu}$ with elements $\mu_{\mathcal{A}(j,k)} \in \mathbb{R}$, marginal variance vector $\boldsymbol{\sigma}^2$ with elements $\sigma_{\mathcal{A}(j,k)}^2 > 0$, delta-Laplace parameter vector $\boldsymbol{\delta}$ with elements $\delta_{\mathcal{A}(j,k)} > 0$. The marginal density function for any $(j, k) \in \mathcal{I}_{\text{Rmt}}$ is given by

$$f_{Z_{j,k}}(z_{j,k}) = \frac{\delta_{j,k}}{2\kappa_{j,k}\sigma_{j,k}\Gamma\left(\frac{1}{\delta_{j,k}}\right)} \exp\left\{-\left|\frac{z - \mu_{j,k}}{\kappa_{j,k}\sigma_{j,k}}\right|^{\delta_{j,k}}\right\} \quad (2)$$

where $\kappa_{j,k}^2 = \Gamma(1/\delta_{j,k})/\Gamma(3/\delta_{j,k})$ and $\Gamma(\cdot)$ represents the gamma function. The mean and variance of this distribution are respectively $\mu_{j,k}$ and $\sigma_{j,k}^2$, regardless of the choice of $\delta_{j,k}$. The dependence structure of \mathbf{Z} is described on Gaussian scale, with correlation matrix

$$\boldsymbol{\Sigma}_{\mathcal{A}(j,k)\mathcal{A}(j',k')} = \frac{\boldsymbol{\Sigma}_{\mathcal{A}^*(j,k)\mathcal{A}^*(j',k')} - \boldsymbol{\Sigma}_{\mathcal{A}^*(j,k)\mathcal{A}^*(0,1)}\boldsymbol{\Sigma}_{\mathcal{A}^*(0,1)\mathcal{A}^*(j',k')}}{\left(1 - \boldsymbol{\Sigma}_{\mathcal{A}^*(j,k)\mathcal{A}^*(0,1)}^2\right)^{1/2}\left(1 - \boldsymbol{\Sigma}_{\mathcal{A}^*(0,1)\mathcal{A}^*(j',k')}^2\right)^{1/2}}$$

for (j, k) and $(j', k') \in \mathcal{I}_{\text{Rmt}}$. The correlation matrix is simply that of a standard Gaussian field evaluated at the p locations $r_j, j = 1, \dots, p$ for all m quantities, conditioned on the value of one quantity ($k = 1$) at an external location r_0 . The correlation matrix for the corresponding unconditioned standard Gaussian field is

$$\boldsymbol{\Sigma}_{\mathcal{A}^*(j,k)\mathcal{A}^*(j',k')}^* = \lambda_{k,k'}^{|k-k'|} \exp\left(-\left(\frac{\text{dist}(r_j, r_{j'})}{\rho_{k,k'}}\right)^{\kappa_{k,k'}}\right) \quad (3)$$

for pairs (j, k) and (j', k') in the extended ordered set $\mathcal{I}_{\text{Rmt}}^* = \mathcal{I}_{\text{Rmt}} \cup \{(0, 1)\}$ (i.e. including the conditioning quantity $k = 1$ at conditioning location r_0), where function $\mathcal{A}^*(j, k)$ now identifies the location of pair (j, k) in $\mathcal{I}_{\text{Rmt}}^*$. Parameters $\lambda_{k,k'} \in [0, 1]$ are the assumed common correlations between quantities at any one location for $k \neq k'$, and $\lambda_{k,k'} = 1$ when $k = k'$. Further $\rho_{k,k'} > 0$ and $\kappa_{k,k'} > 0$ are the scale and exponent parameters of the assumed powered exponential dependence of the standard Gaussian field for quantities k and k' . We choose to write the unique set of

parameters $\{\lambda_{k,k'}\}_{k'>k}$ for estimation as vector λ , and the unique sets $\{\rho_{k,k'}\}_{k' \geq k}$ and $\{\kappa_{k,k'}\}_{k' \geq k}$ as vectors ρ and κ respectively. $\text{dist}(\mathbf{r}_j, \mathbf{r}_{j'})$ is the distance between locations \mathbf{r}_j and $\mathbf{r}_{j'}$, calculated in metres using spherical distance on the Earth's surface.

Thus the joint distribution G of \mathbf{Z} can be written

$$G(\mathbf{z}) = \Phi_{mp} \left(\left(\Phi^{-1}(F_{Z_{1,1}}(z_{1,1})), \Phi^{-1}(F_{Z_{2,1}}(z_{2,1})), \dots, \Phi^{-1}(F_{Z_{j,k}}(z_{j,k})), \dots, \Phi^{-1}(F_{Z_{p,m}}(z_{p,m})) \right); \mathbf{0}, \mathbf{\Sigma} \right)$$

where Φ is the cumulative distribution function of a standard Gaussian distribution, and $\Phi_{mp}(\mathbf{z}; \mathbf{0}, \mathbf{\Sigma})$ is the cumulative distribution function of a mp -dimensional Gaussian distribution with mean $\mathbf{0}$ and covariance matrix $\mathbf{\Sigma}$ evaluated at \mathbf{z} . $F_{j,k}(z)$ is the delta-Laplace marginal cumulative distribution function for quantity k at location j with density given by Equation 2. By differentiating the expression for $G(\mathbf{z})$, an expression for the log-density of \mathbf{Z} can be found, and hence an expression for the sample log-likelihood required for inference (see Shooter et al. 2021c).

3.3. Inference

Inference involves estimation of parameter vectors α , β , μ , σ , δ , and λ , ρ and κ . The variation of any one of α , β , μ , σ and δ with distance is described using a piecewise linear representation (with parameters estimated at each of n_{Nod} nodes) for each of the m quantities of interest. Thus for generic parameter η (i.e. any of α , β , μ , σ , δ for a specific quantity, varying with distance d), with node values $\{\eta_{\ell}^N\}_{\ell=1}^{n_{\text{Nod}}}$, the assumed piecewise linear representation is $\eta(d) = (h_L \eta_{\ell^*}^N + h_U \eta_{\ell^*+1}^N) / h$, where $h_L = d - d_{k^*}$, $h_U = d_{k^*+1} - d$ and $h = d_{k^*+1} - d_{k^*}$, and $k^* = \underset{k}{\text{argmax}} (d_k : d_k < d)$.

The full inference therefore requires the estimation of parameter set Ω given by

$$\Omega = \left(\{ \alpha_{\ell,k}^N, \beta_{\ell,k}^N, \mu_{\ell,k}^N, \sigma_{\ell,k}^N, \delta_{\ell,k}^N \}, \{ \lambda_{k,k'} \}_{k'>k}, \{ \rho_{k,k'}, \kappa_{k,k'} \}_{k' \geq k} \right), \quad \ell = 1, 2, \dots, n_{\text{Nod}}, \quad k, k' = 1, 2, \dots, m.$$

Noting that the dimensions of λ , ρ and κ are respectively $m(m-1)/2$, $m(m+1)/2$ and $m(m+1)/2$, Ω therefore contains $m(5n_{\text{Nod}} + (3m+1)/2)$ parameters.

We use Bayesian inference to estimate the joint posterior distribution of MSCE model parameters. An adaptive MCMC algorithm based on Roberts and Rosenthal (2009) is used for parameter inference, described in Shooter et al. (2019), Shooter et al. (2021c) and Shooter et al. (2021b). Briefly, random search is used to find a reasonable starting solution. Then a Metropolis-within-Gibbs algorithm is used iteratively to sample each of the individual parameters in turn for a total of $n_1 = 250$ iterations. Subsequently we use the adaptive MCMC algorithm to update all parameters jointly for a further $n_2 = 19750$ iterations.

Uniform prior distributions were assumed for each parameter; it was confirmed that posterior densities were not obviously restricted by prior specification. In particular, we allow the node values of α to be > 1 (as discussed by Tendijck, Eastoe, Tawn, Randell and Jonathan 2021 for sub-asymptotic levels), and hence did not adopt the conditional quantile constraints of Keef et al. (2013a).

Data and prototype MATLAB code for the analysis discussed in this article are available at Shooter, Ross and Jonathan (2021a).

4. Results

Results of applying the MSCE model to the Laplace-scale sample illustrated in Figure 4 are now discussed. Inspection of Figure 4 provides some guidance regarding the decay of α with distance we might expect to infer. Consider the scatter plot of St1Wnd at location r_1 on St1Wnd at conditioning location r_0 in the first row and second column of the figure. For large values (say ≥ 4) of St1Wnd at r_0 , the values of St1Wnd at r_1 are also relatively large, suggesting that the corresponding value of α estimated (see Equation 1) should be near unity. As we move across the first row to larger distances r_j , it is clear that typical values of St1Wnd at r_j (for large values of St1Wnd at r_0) are centred around zero, suggesting that α at these locations will be near zero. Indeed, for all of St1Wnd, HndWnd and HndWav, it appears that the value of α for all quantities will be near zero for locations r_j with $j \geq 6$, or $r_j > 760$ km. At a smaller distance r_j (with $j \leq 3$), the joint characteristics of the three quantities of interest appear rather similar on Laplace scale.

The corresponding posterior estimates for α , together with those for β , μ , σ and δ are shown in Figure 5. As anticipated α decays from a value > 0 at r_1 to around zero for distances over 600 km. Estimates for α are generally somewhat larger for St1Wnd than for HndWnd than for HndWav as might be expected, but the differences are small given the widths of credible intervals. It is interesting that the profiles for μ with distance are also similar for the three quantities, reducing

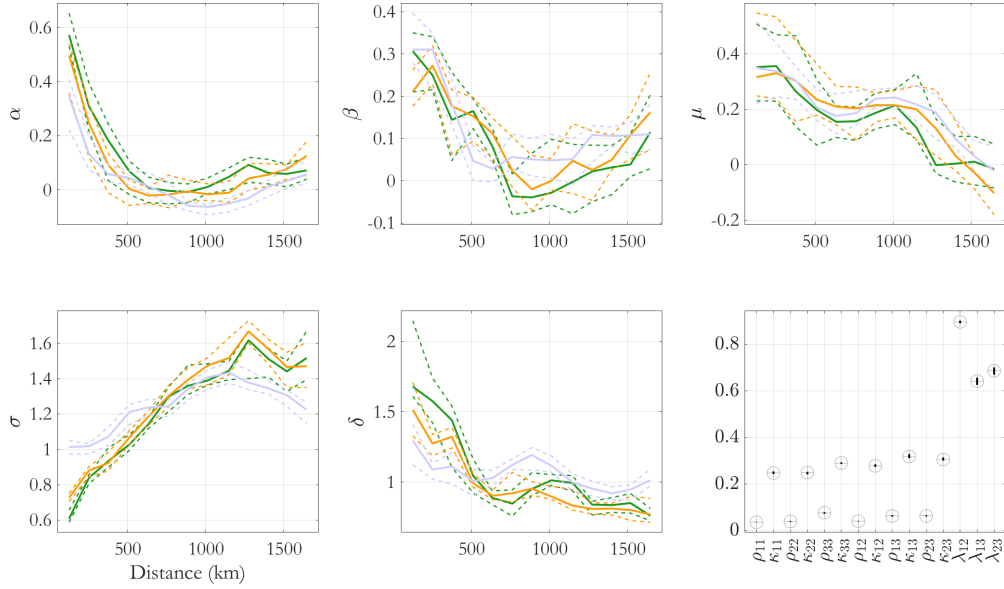


Figure 5: Posterior parameter estimates (shown as mean (solid lines) and 95% credible intervals (dashed lines)) for marginal MSCE model parameters α , β , μ , σ , δ , and residual dependence parameters ρ , κ and λ . Conditioning on St1Wnd at location r_0 with conditioning value equal to the 0.75 quantile of the standard Laplace distribution. Colour coding indicates the conditioned quantity: St1Wnd (green), HndWnd (orange), and HndWav (blue). Estimates for residual dependence parameters are given by large circles, centred on the posterior mean estimate, and vertical black lines reflecting posterior 95% credible intervals. For further interpretation of residual dependence parameters, see Equation 3.

from around 0.4 for small distances to around zero for large distances. Estimates for β reduce from approximately 0.3 at small distances to zero or 0.1 with increasing distance. Estimates for σ increase towards approximately $\sqrt{2}$ with increasing distance, as suggested by the model formulation: for large distances, the effect of conditioning on $X_{0,1}$ is negligible, so that $X_{j,k}|\{X_{0,1} = x\}$ is similar to the unconditioned $X_{j,k}$ which is standard Laplace distributed, with variance equal to 2. The behaviour of σ for HndWav is somewhat different, reflecting the difference between H_S and the other two wind speed variates. The decay of δ with increasing distance to around unity, suggests that $X_{j,k}|\{X_{0,1} = x\}$ is Laplace-distributed (with $\delta = 1$) for large distances, but more Gaussian (with $\delta = 2$) for small distances, as expected.

Parameter estimates for the residual dependence structure (see Equation 3) in the bottom right panel indicate rather similar shape and scale estimates for distance decay of pairs of residuals for St1Wnd (ρ_{11}, κ_{11}), pairs of residuals for HndWnd (ρ_{22}, κ_{22}), and pairs of residuals for HndWav (ρ_{33}, κ_{33}). Shape and scale parameters for cross-dependence between different quantities take similar values, although it is noteworthy that residual parameters involving HndWav tend to be somewhat larger. We note that the values of ρ and κ have been scaled so that they fall comfortably in the interval $[0,1]$ for convenience during MCMC inference. The actual values of these parameters (e.g. appropriate for input to Equation 3) are given by 100ρ and 5κ respectively. The large value for $\lambda_{1,2} \approx 0.9$ indicates high zero-distance residual correlation between St1Wnd and HndWnd. Similar smaller values for $\lambda_{1,3}$ and $\lambda_{2,3}$ around 0.7 indicate lower zero-distance residual correlation between a wind speed variate and the H_S variate HndWnd, again as might be expected from physical considerations.

Figure 6 shows conditional mean profiles for St1Wnd (green), HndWnd (orange) and HndWav (blue) estimated under the fitted MSCE model, for conditioning on $X_{0,1}$ with conditioning value corresponding to the 0.95 quantile of the standard Laplace distribution, at approximately 2.3. The characteristics of the conditional mean are rather similar to those of α and μ in Figure 5. At small distances, the conditional mean of St1Wnd is clearly larger than that for HndWnd, which is itself larger than for HndWav. However, after approximately 600 km, all conditional mean profiles have decayed to a value of around 0.2. The profiles of conditional standard deviation with distance reflect the characteristics of σ in Figure 5.

Figures 7 and 8 explore the quality of fit of the MSCE model to the sample. Figure 7 shows 0.025, 0.25, 0.5, 0.75 and 0.975 quantiles from marginal simulations for St1Wnd (green), HndWnd (orange) and HndWav (blue) with distance




Figure 6: Conditional mean ($\alpha x + x^\beta \mu$) and standard deviation (σx^β , both shown as mean and 95% credible intervals) from estimated MSCE model for conditioning on St1Wnd at location r_0 with conditioning value equal to the 0.95 quantile of the standard Laplace distribution. Colour coding indicates the conditioned quantity: St1Wnd (green), HndWnd (orange), and HndWav (blue).

for conditioning on St1Wnd at location r_0 with conditioning value equal to the 0.75 quantile of the standard Laplace distribution. Also shown in black are the corresponding quantiles calculated directly from the sample. There is good agreement, indicating that marginally at least the model is able to capture the distance-dependent features of the data.

Figure 8 compares the characteristics of observed residuals (black) and corresponding residuals simulated under the fitted model (red). Diagonal panels compare histograms for selected (location,quantity) pairs, and off-diagonal panels show scatter plots for different (location,quantity) pairs. Conditioning is again on St1Wnd at location r_0 with conditioning value equal to the 0.75 quantile of the standard Laplace distribution. The quality of agreement between the empirical distributions of observed and simulated residuals was also quantified using bootstrapping to estimate a null distribution for the Kullback-Leibler divergence between the distributions of bootstrap resamples of the observed residuals. Then the tail probability corresponding to the Kullback-Leibler divergence between the distributions of observed and simulated residuals in the null distribution is estimated. We found that approximately 20% of the values exceeded the 95% percentile of the null distribution. This indicates that there is reasonable if not excellent agreement between the samples of observed and simulated residuals, and that the residual dependence model is able to capture the sample characteristics.

5. Discussion

This article outlines a multivariate spatial conditional extremes (MSCE) model to describe the dependence between extremes of multiple different spatial fields. The model is the latest extension of the conditional extremes model of Heffernan and Tawn (2004). The MSCE model quantifies the conditional distribution of a number of quantities measured on a common set of locations, conditional on the occurrence of a large value of one of the quantities at a conditioning location. The key characteristic of the model is smooth variation of marginal and residual dependence with increasing distance between spatial locations. Bayesian inference is used for parameter estimation. The model is applied to observations of wind speed from satellite observation (St1Wnd) and hindcast (HndWnd), and hindcast significant wave height (HndWav) on a spatial transect lying between the British Isles and Iceland, conditioning on a large value of St1Wnd at the most south-western location on the transect.

There is evidence that the dependence between St1Wnd at neighbouring locations is greater than that between St1Wnd and HndWnd at the same neighbouring locations, and that these are both greater than the dependence between St1Wnd and HndWav at the neighbouring locations. However, the overall trends of MSCE model parameters with distance

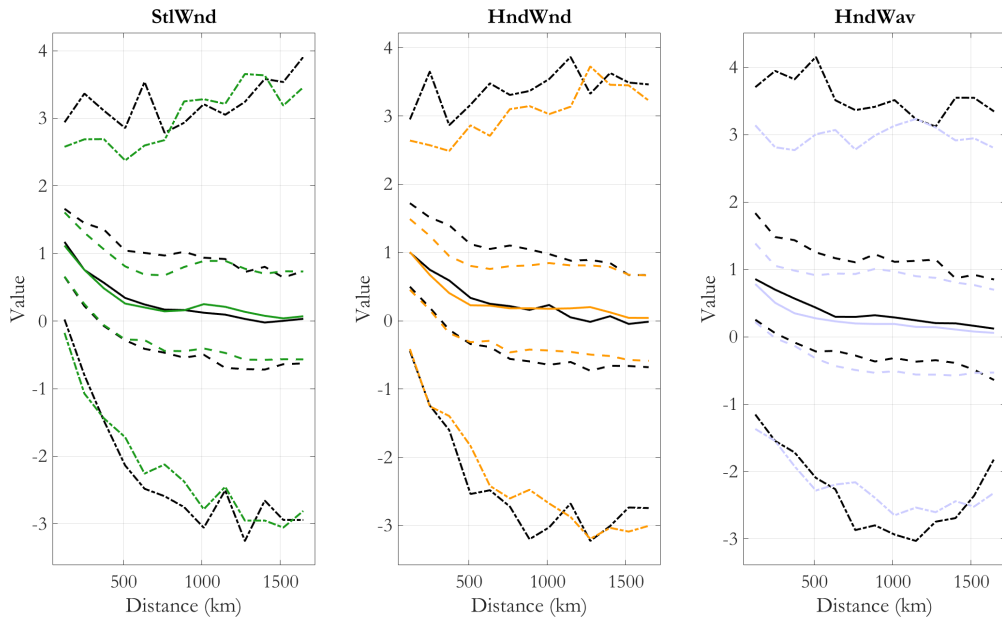


Figure 7: Model validation. Comparison of observed (Laplace-scale) data with distance with simulations under the fitted MSCE model. Conditioning on $StlWnd$ at location r_0 with conditioning value equal to the 0.75 quantile of the standard Laplace distribution. Shown are 0.025, 0.25, 0.5, 0.75 and 0.975 quantiles of the observations (black) and simulation. Colour coding indicates the conditioned quantity for simulation: $StlWnd$ (green), $HndWnd$ (orange), and $HndWav$ (blue).

between locations are similar for each of $StlWnd$, $HndWnd$ and $HndWnd$. The conditional mean for all quantities decays to a baseline level at a distance of approximately 600 km. We note however that figures in the appendix, considering dependence for extremes of each of $StlWnd$, $HndWnd$ and $HndWnd$ separately, suggest that the extent of spatial dependence for $HndWnd$ (Figure A3) is somewhat larger than for the other quantities (Figure A1 and Figure A2). We speculate that a distance of approximately 600-800km is indicative of the spatial extent of coherence for wind systems in the North East Atlantic. It would be interesting to estimate the corresponding distance for other ocean basins.

The effect of choice of conditioning quantity, conditioning location and conditioning value were examined for a number of cases, although not exhaustively. Results with similar general characteristics to those reported here were obtained.

The original intention for this work was to combine all of (a) satellite scatterometer measurements for wind speed and direction from Metop, (b) satellite altimeter measurements for significant wave height and wave direction from JASON (e.g. Shooter et al. 2021b), and (c) corresponding hindcast data for all the variables in (a) and (b) in one MSCE model. However, it quickly became apparent that the number of approximately joint measurements (given space and time) from Metop and JASON available is small, and so far insufficient for joint modelling from scatterometry and altimetry.

The current analysis uses satellite observations of average wind speed and direction corresponding to a relatively short period of time on the daily satellite swath over the North Atlantic, and corresponding spatially- and temporally-matched data from the NORA10 hindcast. These observations are not guaranteed to be representative of wind speeds and directions at the locations of interest, since the satellites pass over the North Atlantic at approximately the same time each day, a source of measurement bias. Moreover, the temporal extent of a severe storm in the North Atlantic is of the order of days, meaning that multiple satellite observations from the same storm event are possible from the daily pass for each satellite. These observations are therefore likely to be correlated in time, at least to some extent. Generally, it would be preferable to perform extreme value analysis on “storm peak” wind speed and direction, which can reasonably be assumed to be temporally independent; however, these data are not available from the daily passes per satellite.

Inference for the MSCE model is straightforward using the adaptive MCMC algorithm of Roberts and Rosenthal (2009), and convergence of MCMC chains is relatively rapid; in practice, 10000 MCMC iterations is more than sufficient. We believe that the MSCE methodology is an interesting extension to the statistician’s and met-ocean engineer’s tool kits, providing a practically applicable yet statistically principled approach to quantification of conditional ex-

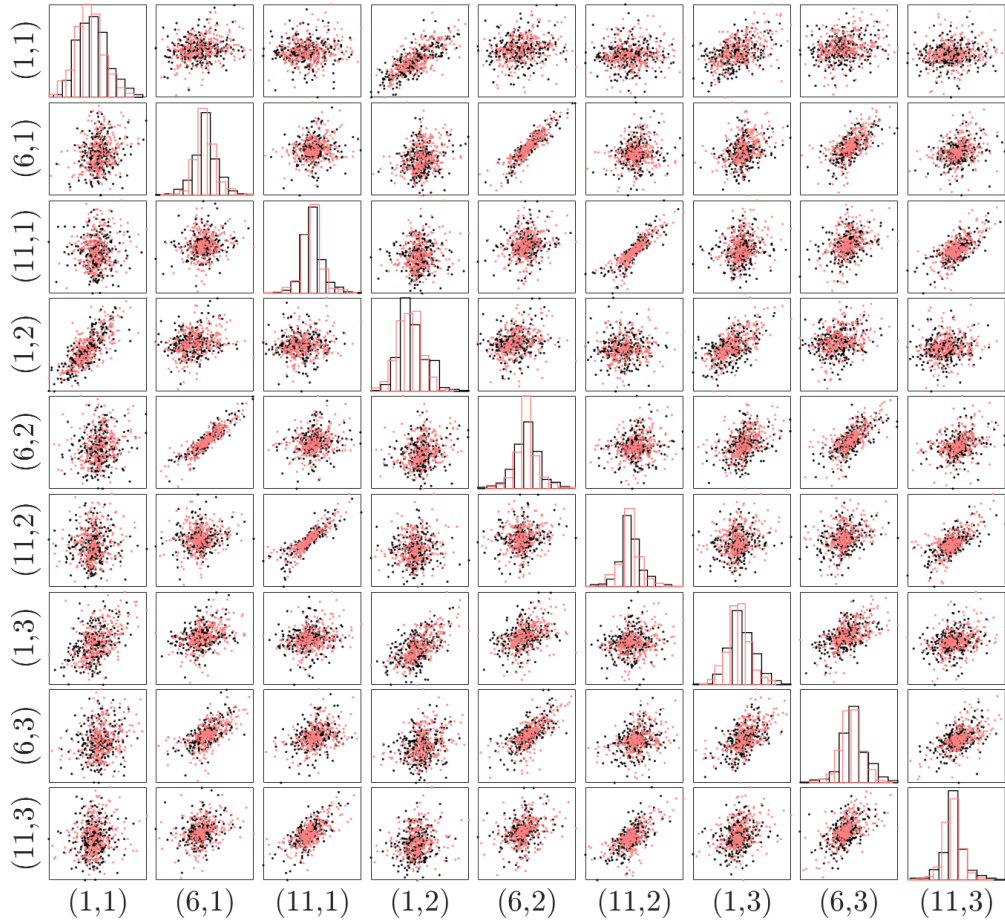


Figure 8: Model validation. Comparison of observed residuals (black) and residuals simulated under the fitted MSCE model (red). Diagonal panels show histograms of residuals for selected location-quantity pairs $(j, k) \in \mathcal{I}_{\text{Rmt}}$. Off-diagonal panels show scatter plots of residuals for selected pairs $(j, k), (j', k') \in \mathcal{I}_{\text{Rmt}}$. Conditioning on StlWnd at location r_0 with conditioning value equal to the 0.75 quantile of the standard Laplace distribution.

tremes behaviour over multiple spatial fields.

6. Acknowledgement

We thank David Randell from Shell for assistance with data preparation. I. R. Young acknowledges ongoing financial support from the Integrated Marine Observing System (IMOS) and the Victorian Government through the Department of Environment, Land, Water and Planning, Australia. Data and prototype MATLAB code for the analysis discussed in this article are available at Shooter et al. (2021a).

References

- Brown, B.M., Resnick, S.I., 1977. Extreme values of independent stochastic processes. *J. Appl. Probab.* 14, 732–739.
- Coles, S., Heffernan, J., Tawn, J., 1999. Dependence measures for extreme value analyses. *Extremes* 2, 339–365.
- Davison, A.C., Padoan, S.A., Ribatet, M., 2012. Statistical modelling of spatial extremes. *Statist. Sci.* 27, 161–186.
- Genton, M.G., Padoan, S.A., Sang, H., 2015. Multivariate max-stable spatial processes. *Biometrika* 102, 215–230.
- Heffernan, J.E., Tawn, J.A., 2004. A conditional approach for multivariate extreme values. *J. R. Statist. Soc. B* 66, 497–546.
- Huser, R., Wadsworth, J.L., 2020. Advances in statistical modelling of spatial extremes. *Wiley Interdisciplinary Reviews: Computational Statistics* doi:10.1002/wics.1537.
- Keef, C., Papastathopoulos, I., Tawn, J.A., 2013a. Estimation of the conditional distribution of a vector variable given that one of its components is large: additional constraints for the Heffernan and Tawn model. *J. Mult. Anal.* 115, 396–404.
- Keef, C., Tawn, J.A., Lamb, R., 2013b. Estimating the probability of widespread flood events. *Environmetrics* 24, 13–21.
- Reich, B., Shaby, B., 2018. Modeling of multivariate spatial extremes. *Researches.One* , 18.09.00013v1.
- Reich, B.J., Shaby, B.A., 2012. A hierarchical max-stable spatial model for extreme precipitation. *Ann. Appl. Stat.* 6, 1430–1451.
- Reistad, M., Breivik, O., Haakenstad, H., Aarnes, O.J., Furevik, B.R., Bidlot, J.R., 2011. A high-resolution hindcast of wind and waves for the North Sea, the Norwegian Sea, and the Barents Sea. *J. Geophys. Res.* 116, 1–18.
- Ribal, A., Tamizi, A., Young, I.R., 2021. Calibration of scatterometer wind speed under hurricane conditions. *Journal of Atmospheric and Oceanic Technology* 38, 1859 – 1870.
- Ribal, A., Young, I.R., 2020a. Calibration and cross validation of global ocean wind speed based on scatterometer observations. *J. Atmos. Ocean. Technol.* 37, 279 – 297.
- Ribal, A., Young, I.R., 2020b. Global calibration and error estimation of altimeter, scatterometer, and radiometer wind speed using triple collocation. *Remote Sens.* 12, 1997.
- Ribatet, M., 2013. Spatial extremes: max-stable processes at work. *J. Soc. Fr. Statistique* 154, 156–177.
- Roberts, G.O., Rosenthal, J.S., 2009. Examples of adaptive MCMC. *J. Comp. Graph. Stat.* 18, 349–367.
- Ross, E., Sam, S., Randell, D., Feld, G., Jonathan, P., 2018. Estimating surge in extreme North Sea storms. *Ocean Eng.* 154, 430–444.
- Schlather, M., 2002. Models for stationary max-stable random fields. *Extremes* 5, 33–44.
- Shooter, R., Ross, E., Jonathan, P., 2021a. Multivariate spatial conditional extremes. <https://github.com/ygraigarw/MultivariateSpatialConditionalExtremesPublic>.
- Shooter, R., Ross, E., Ribal, A., Young, I.R., Jonathan, P., 2021b. Spatial conditional extremes for significant wave height from satellite altimetry. *Environmetrics* 32, e2674.
- Shooter, R., Ross, E., Tawn, J.A., Jonathan, P., 2019. On spatial conditional extremes for ocean storm severity. *Environmetrics* 30, e2562.
- Shooter, R., Tawn, J.A., Ross, E., Jonathan, P., 2021c. Basin-wide spatial conditional extremes for severe ocean storms. *Extremes* 24, 241–265.
- Smith, R.L., 1990. Max-stable processes and spatial extremes. Unpublished article, available electronically from www.stat.unc.edu/postscript/rs/spatex.pdf.
- Tawn, J., Shooter, R., Towe, R., Lamb, R., 2018. Modelling spatial extreme events with environmental applications. *Spatial Statistics* 28, 39 – 58.
- Tendijck, S., Eastoe, E., Tawn, J., Randell, D., Jonathan, P., 2021. Modelling the extremes of bivariate mixture distributions with application to oceanographic data. Submitted to *JASA*, draft at www.lancs.ac.uk/~jonathan.
- ESCADES-GITHUB, 2019. Environmental contours for safe design of ships and other marine structures. URL: github.com/ECSADES.
- Towe, R., Zanini, E., Randell, D., Feld, G., Jonathan, P., 2021. Efficient estimation of distributional properties of extreme seas from a hierarchical description applied to calculation of un-manning and other weather-related operational windows. *Ocean Eng.* 238, 109642.
- Vettori, S., 2017. Models and Inference for Multivariate Spatial Extremes. KAUST Research Repository. URL: <https://doi.org/10.25781/KAUST-G2PH8>.
- Vettori, S., Huser, R., Genton, M.G., 2019. Bayesian modeling of air pollution extremes using nested multivariate max-stable processes. *Biometrics* 75, 831–841.
- Wadsworth, J.L., Tawn, J.A., 2019. Higher-dimensional spatial extremes via single-site conditioning. In submission, arxiv.org/abs/1912.06560.
- Wadsworth, J.L., Tawn, J.A., Davison, A.C., Elton, D.M., 2017. Modelling across extremal dependence classes. *J. Roy. Statist. Soc. C* 79, 149–175.
- Zanini, E., Eastoe, E., Jones, M., Randell, D., Jonathan, P., 2020. Covariate representations for non-stationary extremes. *Environmetrics* 31, e2624.

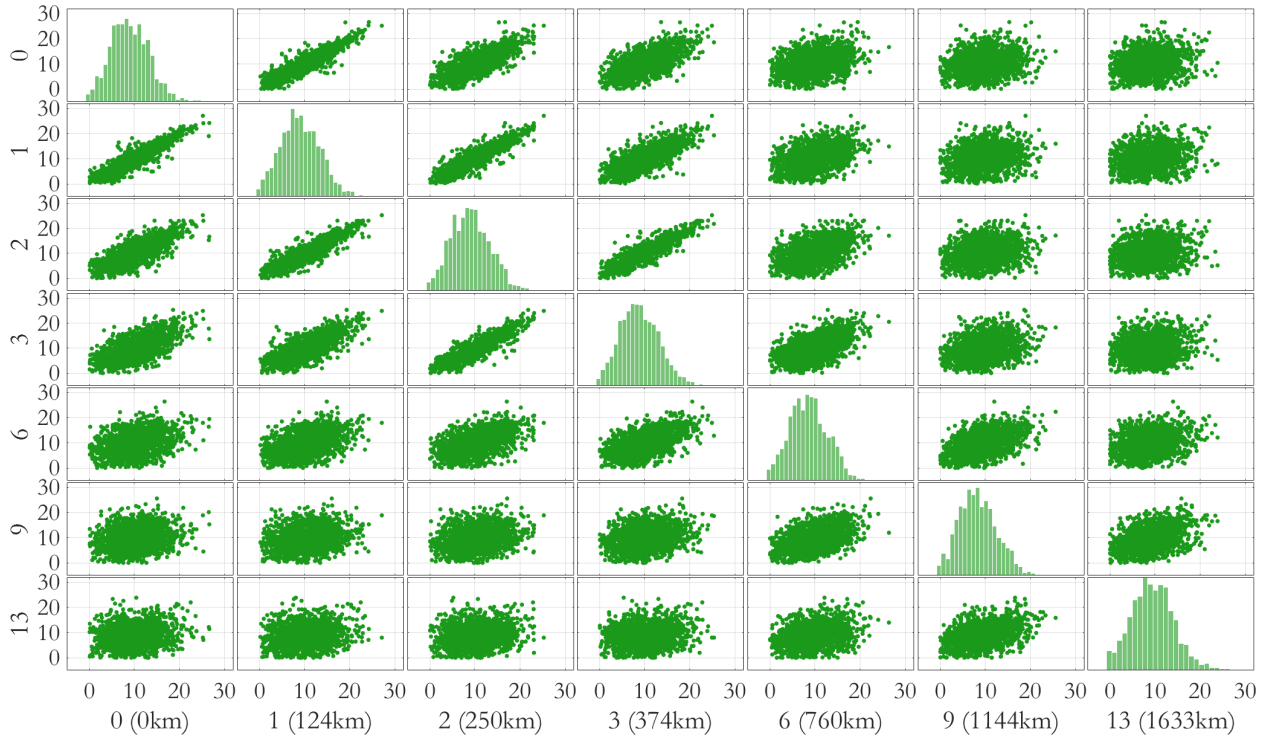


Figure A1: Off-diagonal entries give scatter plots of $St1Wnd$ on physical scale for selected pairs of registration locations r_j , $j = 0, 1, 2, 3, 6, 9, 13$ (see Figure 1). Diagonal elements give empirical densities for $St1Wnd$ at the locations.

A. Scatter plots of original data

Figures A1-A3 provide scatter plots and histograms for $St1Wnd$, $HndWnd$ and $HndWav$ on their original physical scales, for selected representative registration locations. The figures reveal a number of interesting features. For example, comparison of Figures A1 and A2 suggests that dependence with distance is rather similar for $St1Wnd$ and $HndWnd$. In contrast, the decay of dependence with distance is more gradual for $HndWav$ as might be expected from physical considerations.

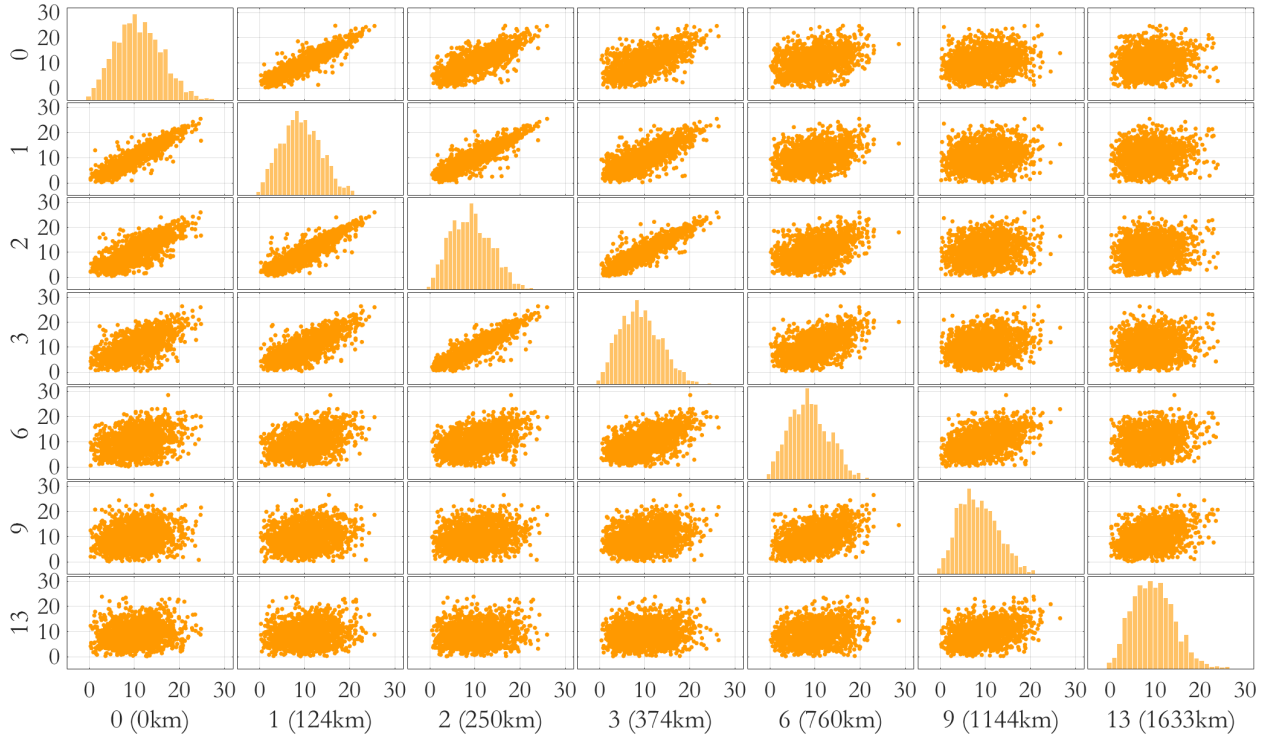


Figure A2: Off-diagonal entries give scatter plots of HndWnd on physical scale for selected pairs of registration locations r_j , $j = 0, 1, 2, 3, 6, 9, 13$. Diagonal elements give empirical densities for HndWnd at the locations.




Figure A3: Off-diagonal entries give scatter plots of HndWav on physical scale for selected pairs of registration locations r_j , $j = 0, 1, 2, 3, 6, 9, 13$. Diagonal elements give empirical densities for HndWav at the locations.