Ocean Engineering 273 (2023) 113959

Contents lists available at ScienceDirect

Ocean Engineering

journal homepage: www.elsevier.com/locate/oceaneng

Model-free environmental contours in higher dimensions

Ed Mackay [a,*], Guillaume de Hauteclocque [b]

[a] University of Exeter, Penryn, Cornwall, TR10 9FE, UK
[b] Bureau Veritas, Paris La Defense, France

ARTICLE INFO

Keywords:
Joint extremes
Multivariate statistics
Serial correlation
Extreme response
Open-source code

ABSTRACT

This paper presents a method for estimating environmental contours in an arbitrary number of dimensions. The method, referred to as Direct-IFORM, does not require a model for the joint distribution of the variables. It can therefore be applied in higher dimensions without degradation in performance. The method involves multiple univariate analyses under rotations of the axes to estimate return values in various directions, which are used to form a contour. The method accounts for serial correlation in the observations, which removes the positive bias that occurs when this is neglected. An efficient open-source code is provided for estimating Direct-IFORM contours. A four-dimensional example is presented for contours of wind speed, wave height, wave period and wind-wave misalignment direction. The application of the method to the design of offshore wind turbines is discussed.

1. Introduction

The environmental contour method is widely used in offshore design to obtain approximations for long-term extreme responses. Many approaches have been proposed for estimating environmental contours (Manuel et al., 2018; Ross et al., 2020). Most methods for estimating contours require the estimation of the joint distribution of the observations as a prior step to constructing the contour (Haselsteiner et al., 2021a). The requirement to fit a joint model to the data limits the application of most methods to low dimensions, due to the so-called 'curse of dimensionality', whereby observations become increasingly sparse in higher dimensions due to the exponential increase in the volume of the variable space.

Recent work has shown that fitting a joint model to the data is a major source of uncertainty, with different approaches yielding contours which result in design loads that vary by ±50% in some cases (de Hauteclocque et al., 2022). A wide range of methods has been proposed for estimating the joint density function. A relatively recent review is presented by Ross et al. (2020). Most commonly, two approaches are used, either hierarchical conditional models or copula models. In the conditional modelling approach, the joint density function $f_{X,Y}$ of variables X, Y is written in conditional form, as $f_{X,Y}(x,y) = f_X(x) f_{Y|X}(y|x)$, where $f_X(x)$ is the marginal density of X and $f_{Y|X}(y|x)$ is the density of Y conditional on X. The approach extends to higher dimensions in the obvious way. Usually, a parametric form is assumed for both f_X and $f_{Y|X}$, with popular choices being Weibull and log-normal models (Haver, 1985; Mathisen and Bitner-Gregersen, 1990). The parameters of the conditional distribution are then modelled as a function of the independent variable, X. There are several drawbacks to this approach. Assuming a model for the bulk of the observations does not guarantee a good fit to the tail of the distribution, which is the region of most interest when estimating extreme responses. Secondly, the model for the conditional dependence structure is based on ad-hoc assumptions, that are not based on any physical or mathematical principles, and give no rationale for extrapolating outside the range of observations. These limitations mean that the fitted models are often in poor agreement with observations (de Hauteclocque et al., 2022).

In copula-based approaches, the joint density function is expressed as $f_{X,Y}(x,y) = f_X(x)\, f_Y(y)\, c(F_X(x), F_Y(y))$, where f_X, f_Y, F_X, F_Y are the marginal density and distribution functions, and c is the copula density function (see e.g. Nelsen (2006) and Joe (2015)). As well as assuming a form for the margins, a parametric form for the copula is also assumed. Common choices for the copula include Gaussian (sometimes referred to as the Nataf model — see Silva-González et al. (2013)), Frank, Gumbel, and Farlie–Gumbel–Morgenstern (Montes-Iturrizaga and Heredia-Zavoni, 2015, 2017). Different choices of copula can lead to large differences in the probabilities in the joint tail regions. Moreover, there is no a-priori reason to suppose that the dependence structure of the observations follows a particular parametric form.

A further limitation of most environmental contour methods is that they do not account for serial correlation in the observations. As explained in detail in Mackay et al. (2021), this can lead to an inherent positive bias in the estimates of return values (see also Vanem,

---

* Corresponding author.
E-mail addresses: e.mackay@exeter.ac.uk (E. Mackay), guillaume.de-hauteclocque@bureauveritas.com (G.d. Hauteclocque).

https://doi.org/10.1016/j.oceaneng.2023.113959
Received 26 November 2022; Received in revised form 12 January 2023; Accepted 12 February 2023

2022; Coe et al., 2022). This is because serial correlation reduces the effective number of independent observations per year, so return values calculated under the assumption that all observations are independent will overestimate the probability of a large value occurring in any given year.

This work describes a method for estimating environmental contours which addresses these limitations, by removing the need to fit a joint model and accounting for serial correlation. The method, referred to as Direct Inverse First Order Reliability Method (D-IFORM), was initially presented by Derbanne and de Hauteclocque (2019). D-IFORM provides a mathematically-justified way of extrapolating, based on extreme value theory. The D-IFORM method was originally presented for two-dimensional applications. An example application of the method in three dimensions was presented by Mackay and Hardwick (2022), who estimated contours of significant wave height, current speed and wave–current misalignment direction. The aim of the present work is to show how the D-IFORM method can be applied in higher dimensions, obtaining contours in lower dimensions as special cases. As explained in further detail in the next section, because there is no requirement to fit a joint probability model, the D-IFORM method can be applied in higher dimensions, without suffering from the curse of dimensionality. The method for calculating D-IFORM contours has been implemented in open-source code in MATLAB and Python, available from https://github.com/edmackay/Direct-IFORM, together with the data used in this study.

The remainder of the paper is organised as follows. Section 2 presents the motivation for the D-IFORM method, and compares this to alternative methods for deriving contours. Section 3 describes the implementation of the D-IFORM method in an arbitrary number of dimensions. An example application is presented in Section 4, where we consider four-dimensional contours of wind and wave parameters, and the potential application for the design of offshore wind turbines. Finally, a discussion and conclusions are presented in Section 5.

2. Motivation for D-IFORM method

The objective of the environmental contour (EC) method is to provide a means of estimating the long-term extreme response of a structure, using a small number of response simulations. To achieve this, the EC method aims to find a region of the environmental parameter space (referred to as the 'design region'), such that a structure that can withstand all environmental conditions within that region has a probability of failure of less than or equal to a specified value. The environmental contour represents the boundary of the design region. It is assumed that if the structure can withstand all conditions along the contour, then it can withstand any environmental condition inside the contour. For simplicity, we will present the motivation for the D-IFORM method in two dimensions. The generalisation to higher dimensions is considered in the next section.

The first environmental contour method, known as the Inverse First Order Reliability Method (IFORM), was presented by Winterstein et al. (1993). The method is based on the First-Order Reliability Method (FORM) (Madsen et al., 2006). In the FORM method, a failure surface for a structure is estimated as a function of environmental variables X and Y. The failure surface is transformed to independent standard normal space using the Rosenblatt transformation (Rosenblatt, 1952). The probability of failure of the structure is then estimated by linearising the failure surface at the highest probability point in standard normal space, and calculating the probability contained in the corresponding half-space region (see Fig. 1). If the distance from the linearised failure surface to the origin is β, then the probability of an observation falling in the exceedance region is $\alpha = 1 - \Phi(\beta)$, where Φ is the standard normal cumulative distribution function (CDF). Winterstein et al. (1993) observed that all half-plane regions which contain probability α are at a fixed radius of β from the origin, and the intersection of all such regions forms a circle. Therefore, if the desired probability of failure of the structure is α, then the FORM approach can be made more efficient by calculating the response of the structure for combinations of environmental variables which lie on a circle of radius β in standard normal space. Transforming this circle back to the original variable space gives the IFORM environmental contour.

Huseby et al. (2013) observed that if the motivation is to define half-space regions which contain a given probability, then the use of the Rosenblatt transformation is unnecessary. Instead, environmental contours can be defined as the intersection of all half-space regions in the original variable space, that contain fixed probability level α. These half-space regions are interpreted as the linearised boundaries of failure regions of possible structures (see Fig. 1). To estimate these regions, Huseby et al. (2013) proposed to first fit a joint probability model to the observations, then generate a large Monte Carlo sample from the fitted model. The distance from the origin of the half-space region at angle θ is then calculated by projecting the Monte Carlo sample onto a line at angle θ, and calculating the empirical quantile at exceedance probability α. The resulting method is referred to as the Direct Sampling (DS) method. As well as the uncertainty in the fitted joint distribution model, this procedure introduces a further sampling uncertainty, due to the random Monte Carlo simulation procedure — see Mackay and Jonathan (2021) for a discussion of the sampling uncertainties in extreme quantiles.

Derbanne and de Hauteclocque (2019) noted that the procedure could be simplified further. Since IFORM and DS environmental contours are defined as the intersection of half space regions, there is no need to fit a joint probability model. Instead, the half-space regions can be viewed as univariate exceedance regions under rotations of the axes. For each given rotation angle, a standard univariate analysis can be applied. Univariate extreme value theory and inference is well-established and widely used (Coles, 2001; Beirlant et al., 2004). Derbanne and de Hauteclocque (2019) proposed that a standard peaks-over-threshold (POT) analysis is conducted on the observations projected onto lines at various angles to the origin, whereby a generalised Pareto (GP) distribution is fitted to exceedances of a pre-determined threshold. The return values at each angle (and hence the desired half-space regions) can then be calculated in terms of the estimated GP parameters. The boundary of the region formed by taking the intersection of the half-spaces defined by the directional return values gives the Direct IFORM (D-IFORM) contour. Since extreme value theory tells us that threshold exceedances from any distribution (that satisfies certain regularity conditions) will converge to a GP distribution (Coles, 2001), the D-IFORM method gives a mathematically-justified and general means of forming environmental contours, which does not rely on ad hoc models.

In the POT method, serial correlation in the observations is accounted for by declustering. This is the process of modelling only peak observations which are sufficiently separated in time that they can be considered independent. Various procedures exist for identifying peaks which meet this condition (see e.g. Coles (2001) and Beirlant et al. (2004)), and an efficient method for declustering is described in Section 3.4. Accounting for the serial correlation in the observations removes the positive bias in estimates of return values, caused by assuming serially-correlated observations are independent (Mackay et al., 2021).

Finally, it is worth mentioning that while the IFORM method and its variants make a linear approximation to the failure surface, other methods have been proposed for defining environmental contours which involve different approximations for the shape of the failure region. Haselsteiner et al. (2017) and Chai and Leira (2018) noted that the linear approximation to the failure region is non-conservative if the failure region is convex. To account for this, they make the conservative assumption that structural failure occurs anywhere outside the environmental contour. The corresponding exceedance regions are defined in terms of the total probability outside the contour, either in standard normal space (Chai and Leira, 2018), or in the original variable space (Haselsteiner et al., 2017). In both cases, the

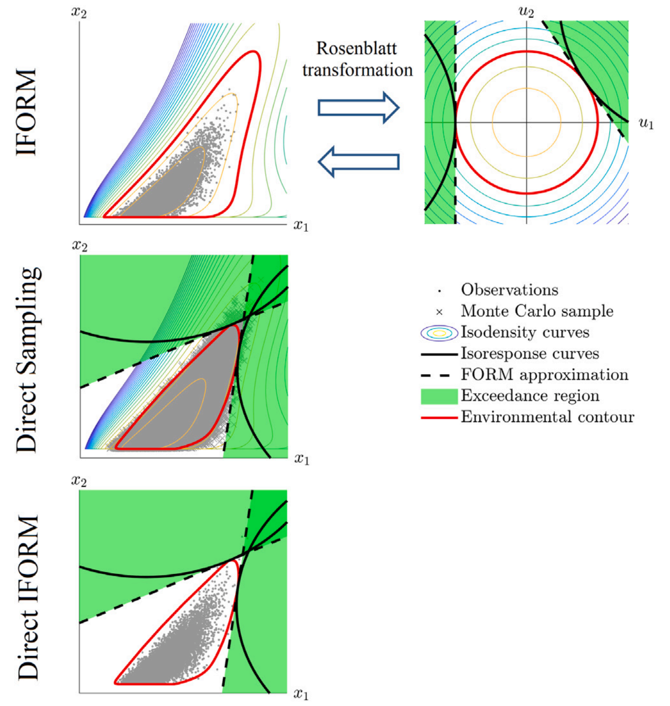

Fig. 1. Illustration of IFORM, Direct Sampling and Direct IFORM methods. Direct Sampling simplifies the IFORM method by removing the need to apply the Rosenblatt transformation. Direct IFORM simplifies the Direct Sampling method by removing the need to fit a joint probability model. The two response curves shown represent different response quantities, with a single response level shown for each.

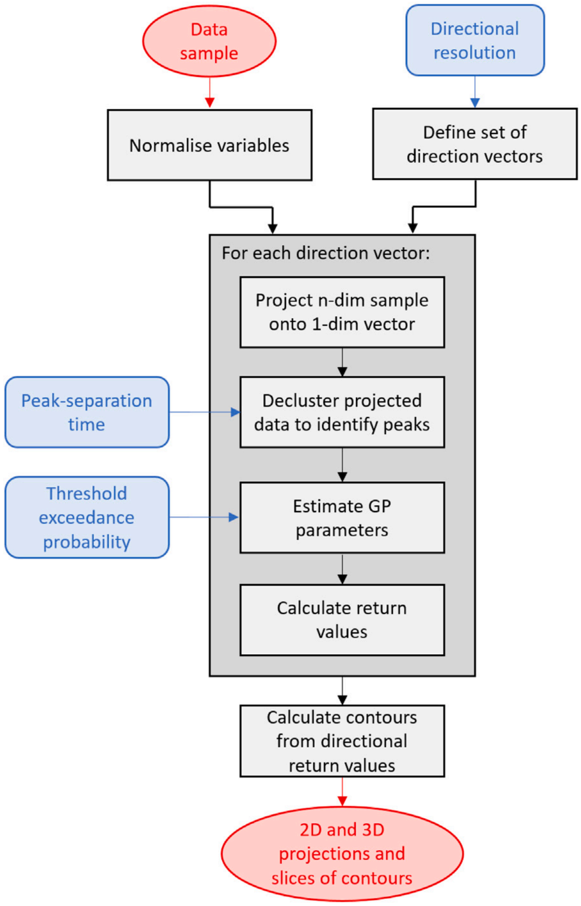

Fig. 2. Flowchart of steps for estimating D-IFORM environmental contours. Red: Inputs/Outputs. Blue: User-defined hyper-parameters. Black: Computational steps.

derived contours typically result in a more conservative set of design conditions and correspondingly more severe responses (see e.g. Mackay and Haselsteiner (2021), de Hauteclocque et al. (2022) and Haselsteiner et al. (2021b, 2022)). These methods both require a model for the joint probability of the observations, and are therefore subject to the same uncertainties as the IFORM and DS methods. They are therefore not considered further in this work.

3. Implementation of D-IFORM in higher dimensions

Throughout this section, we assume that we have a sample of n observations of a d-dimensional random vector $\mathbf{X} = (X_1, \ldots, X_d)$, denoted $\{\mathbf{x}_j\}_{j=1:n}$, where $\mathbf{x}_j = (x_{1,j}, \ldots, x_{d,j})$. A flowchart of the steps involved in estimating D-IFORM contours is shown in Fig. 2. In the following subsections, we describe each of these steps in detail.

3.1. Normalisation

The scales of environmental variables can be very different. For example, wind speeds are typically in the range 0–50 m/s, whereas values of wave steepness are typically less than 0.1. Whilst this does not directly affect the inference, if contours were constructed using the raw variables, then a higher resolution of direction vectors would be required in certain direction ranges, to account for the differing scales. To avoid this issue, we normalise each variable to give roughly equal scales in each dimension. In the present implementation, we have normalised each variable by its standard deviation. We also subtract the median of each variable, so that data is centred. This step is not strictly necessary, as the location of the origin does not affect the inference. However, it slightly simplifies the presentation later on. We denote the normalised sample as $\{\mathbf{y}_j\}_{j=1:n}$, where

$$\mathbf{y}_j = \left(\frac{x_{1,j} - \mathrm{med}(X_1)}{\mathrm{std}(X_1)}, \ldots, \frac{x_{d,j} - \mathrm{med}(X_d)}{\mathrm{std}(X_d)} \right). \tag{1}$$

In practice, the sample median and standard deviation are used in place of the population median and standard deviation. Contours are calculated in the normalised space, then converted back to the original variable space at the end.

3.2. Selection of direction vectors

When working in two or three dimensions, it is natural to use polar or spherical coordinates to specify angles. However, the use of spherical coordinates becomes cumbersome in higher dimensions, so vector notation is used here instead. To apply the D-IFORM method, we need to select a number of direction vectors, $\mathbf{u}$, for analysis. Preferably, we would like the directions to be evenly-spaced. This is equivalent to choosing evenly-spaced points on the surface of the d-dimensional hypersphere. This is only possible when $d = 2$, and the 2-dimensional hypersphere is a circle, although various approximate algorithms exist for $d > 2$.

In this work, we have opted to address the problem by choosing evenly-spaced points on the surface of the unit hypersphere with respect to the L^1 norm, given by $\{\mathbf{w} : \|\mathbf{w}\|_1 = 1\}$, where $\|\mathbf{w}\|_1 = \sum_{i=1}^{d} |w_i|$. In contrast to the standard hypersphere with respect to the L^2 (Euclidean) norm, it is straightforward to create evenly spaced points on the surface of the L^1 hypersphere. We can then normalise so that $\mathbf{u} = \mathbf{w}/\|\mathbf{w}\|_2$ is a unit vector with respect to the L^2 norm, where $\|\mathbf{w}\|_2 = \left(\sum_{i=1}^{d} w_i^2\right)^{1/2}$. If we restrict our attention to the points in the orthant $\{\mathbf{w} : w_i \geq 0\}$, then the surface of the L^1 hypersphere in $\mathbb{R}^d$ is a regular $d-1$ simplex, with vertices given by the standard basis vectors

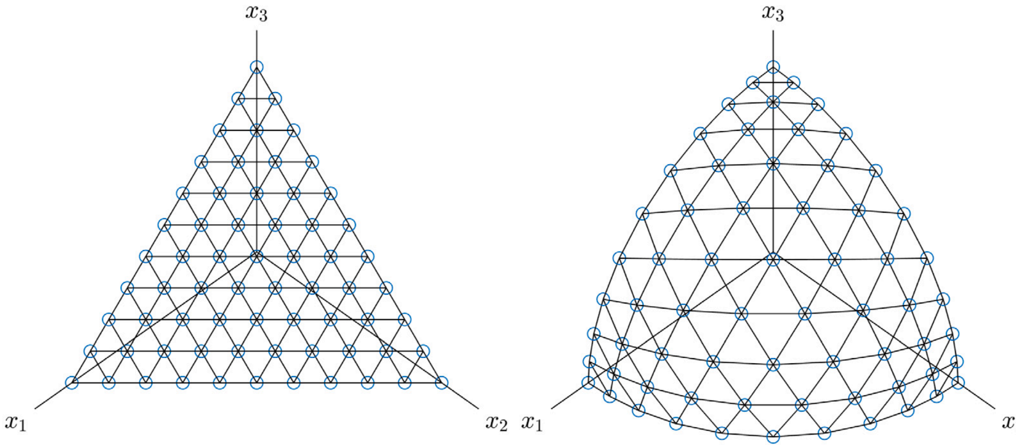

Fig. 3. Left: Evenly-spaced points on surface of unit L^1-sphere in 3 dimensions. Right: Equivalent points after normalisation to lie on surface of unit L^2-sphere in 3 dimensions.

of $\mathbb{R}^d$. Regularly-spaced points, with $m+1$ points along each edge of the simplex, can be written

$$\begin{aligned} &w_1 \in \{i_1/m : i_1 = 0,\dots,m\} \\ &w_2 \in \{i_2/m : (i_2 = 0,\dots,m) \wedge (w_2 \le 1 - w_1)\} \\ &w_3 \in \{i_3/m : (i_3 = 0,\dots,m) \wedge (w_3 \le 1 - w_1 - w_2)\} \\ &\;\vdots \qquad\quad \vdots \\ &w_d = 1 - \sum_{i=1}^{d-1} w_i. \end{aligned} \tag{2}$$

This gives $\binom{d+m-1}{d-1} = (m+1)(m+2)\cdots(m+d-1)/(d-1)!$ points on the surface of the L^1-hypersphere, within a single orthant of $\mathbb{R}^d$. These points all lie on a regular Cartesian grid with spacing $1/m$. Therefore, a simple algorithm to generate these points is to first generate a dense grid with spacing $1/m$ in each dimension, then keep only the points whose coordinates sum to one. Although this is a somewhat inefficient way to generate the direction vectors, the computational cost is very small (typically fractions of a second) in comparison to other steps in the method. Direction vectors in other orthants of $\mathbb{R}^d$ can be obtain by reflection.

The resulting points on the first octant of the L^1-sphere in 3 dimensions are shown in Fig. 3, for $m = 10$. The normalised points lying on the surface of L^2-sphere are also shown for comparison. The rescaled points on the L^2-sphere have a slightly larger spacing in the region closer to the direction $(1,1,1)$, due to the stretching of the surface. However, the non-equal spacing of direction vectors is not critical for the analysis. The important factor is to have a sufficient angular resolution to produce a reasonably smooth contour (see Section 3.6). Moreover, the simplicity of the algorithm for generating the direction vectors in higher dimensions makes it attractive.

3.3. Projection of data

We denote the projection of the observations onto the line in the direction of the unit vector $\mathbf{u} = (u_1,\dots,u_d)$, as $\{r_{\mathbf{u},j}\}_{j=1:n}$, where

$$r_{\mathbf{u},j} = \mathbf{u} \cdot \mathbf{y}_j = \sum_{i=1}^{d} u_i\, y_{i,j}. \tag{3}$$

An illustration of the projection of a sample onto three different direction vectors is shown in Fig. 4. The projected sample is one-dimensional, and can be thought of as a sample of observations of the random variable $R_{\mathbf{u}} = \mathbf{u} \cdot \mathbf{Y}$, corresponding to the shortest distance from the origin to the hyperplane containing $\mathbf{Y}$, with normal vector $\mathbf{u}$. Alternatively, the projected variable, $R_{\mathbf{u}}$, can be interpreted as the first coordinate of the data under a rotation of the coordinate axes.

For a given direction $\mathbf{u}$, the problem of finding return values of $R_{\mathbf{u}}$ is a univariate problem, where a standard univariate POT method can be applied. Since all observations are used in the projection at each direction, increasing the number of dimensions or resolution of direction vectors only affects computational time, but does not decrease the available data. This is the key feature of the D-IFORM method that enables its use for higher-dimensional problems.

3.4. Declustering

To account for the serial correlation in the data, the return values of the random vector Y in direction $\mathbf{u}$ are estimated using a POT analysis. However, the points in the time series $\{r_{\mathbf{u},j}\}_{j=1:n}$ which correspond to local peaks, will depend on the projection direction $\mathbf{u}$. This is illustrated in Fig. 5. This means we need to identify peaks for each projection direction. In the present implementation, peaks are defined as local maxima within a moving window of $\pm h$ hours, where the value of h is set by the user. A moving maxima envelope over a window of $\pm h$ hours is defined as

$$r_{\mathbf{u},\max}(t,h) = \max\{r_{\mathbf{u}}(\tau) : \tau \in [t-h, t+h]\}. \tag{4}$$

Peaks are defined as points where original time series is equal to moving-maximum envelope, i.e. values of t such that $r_{\mathbf{u}}(t) = r_{\mathbf{u},\max}(t,h)$ (see Fig. 6). A check is then performed to see if there are any peaks which are separated by less than h hours, which can happen if there are two equal maxima in the same window. In this case, we discard the second equal maxima in the window. For most metocean variables this is unlikely to occur, so this final check is quick to make. This algorithm also works with time series with gaps (often the case for measured datasets), as gaps do not affect the moving maxima envelope.

This declustering algorithm has the advantage that it does not require definition of a threshold or a definition of peak prominence. It only requires a separation time, specified in terms of the length of the moving maxima window. A rigorous selection of peak separation time can be determined by looking at decorrelation times, using visualisations such as the extremogram (Davis and Mikosch, 2009). This is defined as the limiting conditional probability $\lim_{\mu\to\infty} \Pr(X_{t+\tau} > \mu | X_t > \mu)$, which estimates the extent of asymptotic dependence in the time series at lag τ. For stationary series, the quantity $\Pr(X_{t+\tau} > \mu | X_t > \mu) - \Pr(X_t > \mu)$ will tend to zero when X_t and $X_{t+\tau}$ are independent. This quantity can be plotted as a function of τ for various thresholds μ, to select de-correlation timescales used in declustering routines (see Mackay and Johanning (2018)). Typical choices for peak separation times are in the range of 2–5 days.

The efficiency of the declustering algorithm is strongly dependent on the implementation. Looping over the time series, moving the window and calculating maxima can be very slow. MATLAB has an

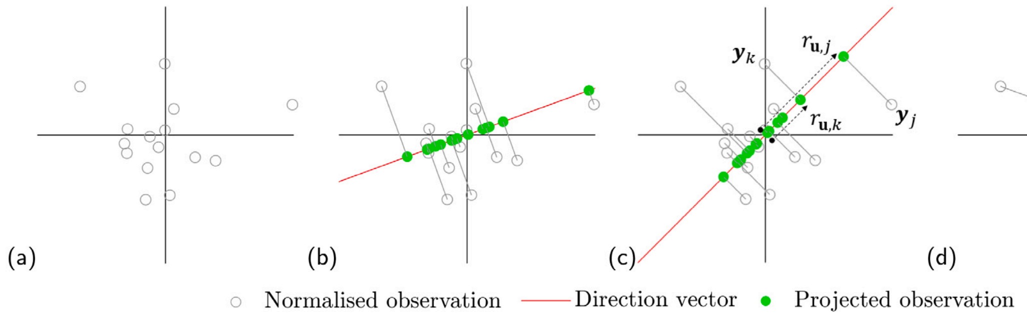

Fig. 4. Illustration of projection of observations onto lines at various angles to the origin. (a) Original normalised sample. (b)–(d) Original normalised sample, plus projected sample at various angles. The distances of the projected observations from the origin are shown in (c) for two observations. Note that all observations are used to create the projected sample at each angle.

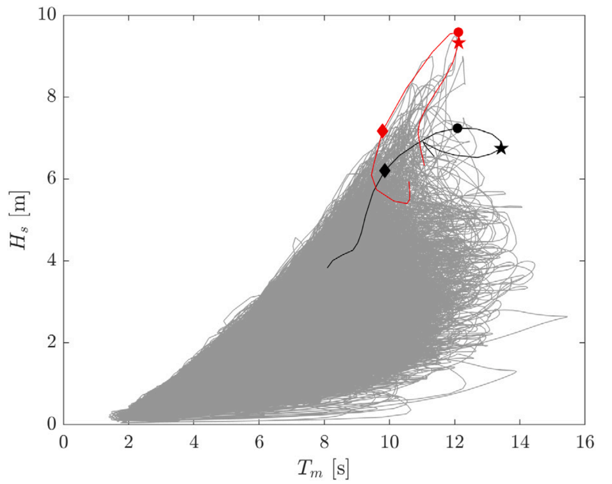

Fig. 5. Illustration of why declustering must be conducted for each projection direction. Grey lines show time series of H_s and T_m. Red and black lines show 24-h sections of time series, centred on a local maximum of H_s, indicated by a circle. The corresponding maximum period and steepness in these sections of time series are indicated by stars and diamonds, respectively. The points corresponding to 'peaks' in the time series therefore depend on the projection direction.

efficient native function for calculating a moving maximum of an array, `movmax`, which is several orders of magnitude faster than looping through the time series.

An example of the effect of neglecting serial correlation on estimated return periods is shown in Fig. 7, using significant wave height (H_s) and wind speed (U_{10}) data for the example presented in Section 4. Return periods calculated under the incorrect assumption that all hourly observations are independent are shown in red. Return periods for the declustered peaks are shown in blue. It is clear that neglecting the serial correlation in the observations leads to a positive bias in the return value at a given return period. However, since the largest hourly observation and largest observed peak must coincide in any dataset, the empirical return periods for the largest hourly value and largest peak are equal. This is an effect of finite sample size. To mitigate for this, return periods inferred from fitted GP models (see following subsection) are also shown. These indicate that the positive bias caused by neglecting serial correlation persists outside the range of observations. Further discussion of these effects can be found in Mackay et al. (2021).

3.5. Inference

For a given projection direction, the return values of $R_{\mathbf{u}}$ are calculated using a POT analysis, where exceedances of a threshold, μ, are modelled using a generalised Pareto (GP) distribution. The cumulative distribution function (CDF) of the GP distribution is

$$F(y) = \begin{cases} 1 - \left(1 + \xi \dfrac{y-\mu}{\sigma}\right)^{-1/\xi}, & \xi \neq 0, \\ 1 - \exp\left(-\dfrac{y-\mu}{\sigma}\right), & \xi = 0, \end{cases} \tag{5}$$

where the support is $\mu \leq y < \infty$ for $\xi \geq 0$ and $\mu \leq y \leq \mu - \sigma/\xi$ for $\xi < 0$. Here, ξ and σ are shape and scale parameters. For a given set of GP parameters, the return value of $R_{\mathbf{u}}$ at return period T, is then given by

$$R_{\mathbf{u},T} = \begin{cases} \mu + \dfrac{\sigma}{\xi}\left((\lambda T)^{\xi} - 1\right), & \xi \neq 0, \\ \mu + \sigma \log(\lambda T), & \xi = 0, \end{cases} \tag{6}$$

where λ is the average number of threshold exceedances per year.

The GP parameters are estimated independently for each direction, $\mathbf{u}$. We set the threshold estimate, $\hat{\mu}$, as an empirically-estimated quantile of $R_{\mathbf{u}}$ at a fixed exceedance probability, ζ. The choice of ζ is a bias–variance trade-off between the asymptotic GP model being appropriate for the exceedances (bias) and having a sufficient number of observations to estimate the model (variance). In practice, we found that setting $\zeta = 0.1$ gave reasonable results for our datasets. The fit of the model for a given threshold can be assessed using standard diagnostic plots (see e.g. Coles (2001)).

The scale and shape parameters are estimated using a constrained maximum likelihood method. When the shape parameter, ξ, is non-negative, the distribution is unbounded from above. For most metocean variables, this is physically unrealistic, and we would expect the distribution to have some finite upper bound. We therefore constrain the estimate $\hat{\xi} < 0$. A GP distribution with $\xi = -1$ is a uniform distribution. When $\xi < -1$ the likelihood of observations increases toward the upper end point of the distribution. This is also physically unlikely for most metocean variables, so we also set a lower bound for the estimate as $\hat{\xi} \geq -1$. We also place an additional constraint that $\max\{r_{\mathbf{u},j}\} \leq \hat{\mu} - \hat{\sigma}/\hat{\xi}$, so that the parameter estimates are consistent with the observations. The parameter values that minimise the negative log-likelihood function are then found using a simplex search method (Lagarias et al., 1998), using the moment estimators as a first guess. The moment estimators are given by $\hat{\xi}_{mom} = \frac{1}{2}\left(1 - (\bar{r}/s_r)^2\right)$, $\hat{\sigma}_{mom} = \bar{r}(1 - \hat{\xi}_{mom})$, where $\bar{r}$ and s_r are the sample mean and standard deviation.

It is reasonable to expect that the GP parameters are smoothly-varying functions of direction $\mathbf{u}$. This could be accounted for directly in the inference, by fitting a non-stationary model, where the smooth variation of the parameters is modelled using e.g. splines (Randell et al., 2015; Zanini et al., 2020) or piecewise-linear functions (Barlow et al., 2023). An alternative option is to first estimate the GP parameters, then smooth the estimated return values, as suggested by Derbanne and de Hauteclocque (2019). For the non-stationary models, an optimal 'roughness' for the parameters can be estimated as part of the inference. For the post-inference smoothing of return values, the choice of

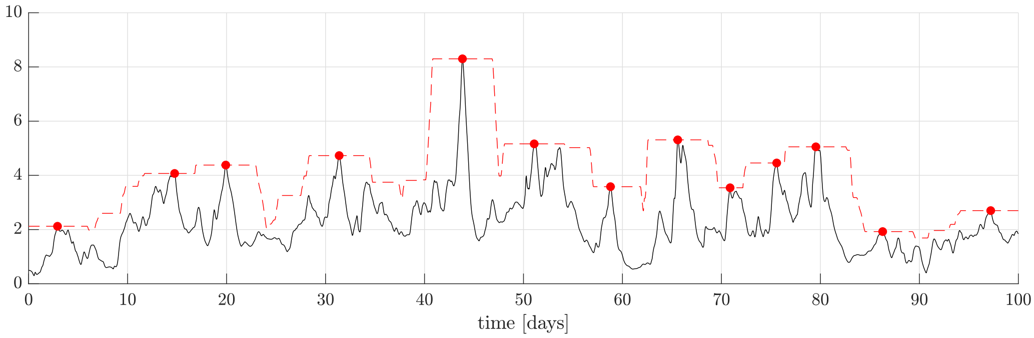

Fig. 6. Illustration of criteria used for declustering. Solid line: time series of variable to be declustered. Dashed line: Maximum value of variable within ±3 days of each time step. Circles: Declustered peaks, defined as local maxima within a window of ±3 days — these correspond to points where the time series meets the local maxima envelope.

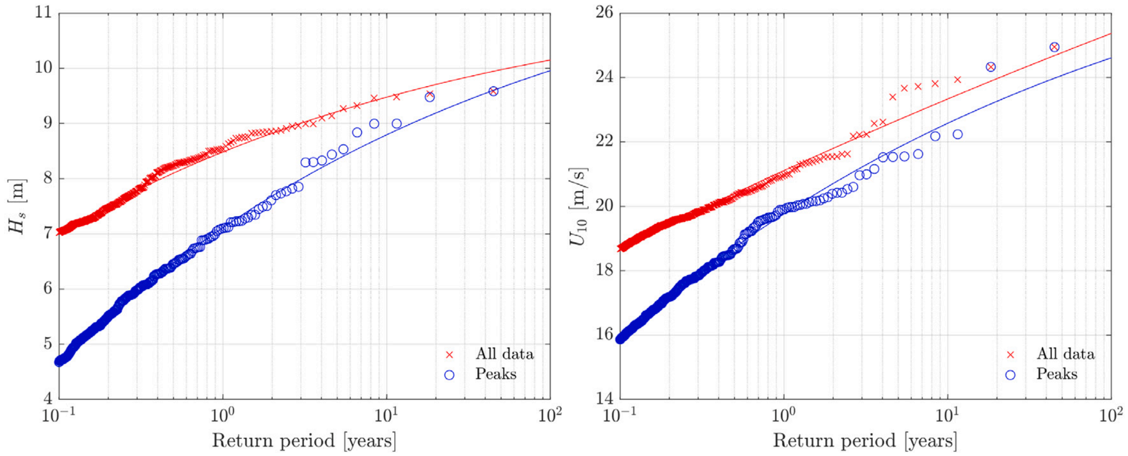

Fig. 7. Effect of neglecting serial correlation on estimated return periods for H_s (left) and U_{10} (right). Crosses denote return periods calculated under the incorrect assumption that all hourly are independent. Circles denote return periods for declustered peaks. Solid lines are fitted generalised Pareto models.

optimal smoothing is more ad hoc. However, both approaches become difficult to implement and slow to compute in higher dimensions. In the examples presented in Section 4, the GP parameter estimates and return values from the independent analyses were reasonably smooth, meaning that there would be little benefit to using a more complex inference scheme. Also, given the approximate nature of extreme response estimates from environmental contours, using an efficient but less-than-optimal inference scheme seems appropriate.

3.6. Calculation of contours

Finding the intersection of a large number of half-spaces in higher dimensions is a difficult computational problem. An efficient solution to this problem was proposed by Hafver et al. (2022), who showed that environmental contours can be interpreted as the boundary of Voronoi cells. This geometric interpretation leads to a simple algorithm for constructing contours, utilising efficient code for calculating Voronoi cells in higher dimensions. Here, we present a sketch of the argument for the Voronoi interpretation of environmental contours, and refer to Hafver et al. (2022) for further details.

The Voronoi cell of a point $\mathbf{o} \in \mathbb{R}^d$, with respect to a set of points $S \subset \mathbb{R}^d$ is denoted $\mathrm{Vor}(\mathbf{o}, S)$. It consists of the set of points that are at least as close to $\mathbf{o}$ as any point in S. That is

$$\mathrm{Vor}(\mathbf{o}, S) = \left\{ \mathbf{x} \in \mathbb{R}^d \,:\, \|\mathbf{x} - \mathbf{o}\|_2 \leq \min_{\mathbf{s} \in S} \|\mathbf{x} - \mathbf{s}\|_2 \right\}. \quad (7)$$

In the original IFORM and Direct Sampling definitions, environmental contours are the boundary of a region defined as the intersection of half-spaces. The half-spaces represent non-exceedance regions at various angles to the origin (see left hand plot of Fig. 8). Consider, the set of reflections, S, of the origin, $\mathbf{o}$, in the boundaries of the exceedance regions at each direction, illustrated in the right hand plot of Fig. 8. For a given set of direction vectors, U, and return period, T, the set of reflections of the origin, $S(U,T)$, is given by

$$S(U,T) = \{2R_{\mathbf{u},T}\,\mathbf{u} \,:\, \mathbf{u} \in U\}. \quad (8)$$

For a given direction, $\mathbf{u}$, the half-space defined by the non-exceedance region $\{\mathbf{y} \,:\, \mathbf{y} \cdot \mathbf{u} \leq R_{\mathbf{u},T}\}$ is equal to the set of points that are closer to the origin, than to the reflection point $2R_{\mathbf{u},T}\,\mathbf{u}$. Following this reasoning, we see that the intersection of all such non-exceedance regions is the set of points that is closer to the origin than to the set of reflections, S. That this, the Voronoi cell of the origin with respect to S, given in (7). Fig. 8 provides an illustration of the equivalence of the Voronoi definition of a contour, with the 'standard' definition in terms of intersections of non-exceedance regions. For the MATLAB implementation, the native function `voronoin` is used to calculate Voronoi cells in arbitrary numbers of dimensions, which utilises the Quickhull algorithms[1] (Barber et al., 1996).

In the present work, we have set the origin as the joint median of the sample. However, Hafver et al. (2022) showed that contours defined in terms of Voronoi cells are invariant to the location of the origin, provided that it is within the contour. Moreover, Hafver et al. (2022) also showed that if $U_1 \subseteq U_2$, then $\mathrm{Vor}(\mathbf{o}, S(U_2,T)) \subseteq \mathrm{Vor}(\mathbf{o}, S(U_1,T))$. That is, estimating contours using fewer direction vectors is conservative, and adding more direction vectors to the analysis will reduce the volume contained within the contour. An illustration of this is shown in Fig. 9, where contours are shown using 4, 8, and 16 direction vectors for a

---

[1] http://www.qhull.org/.

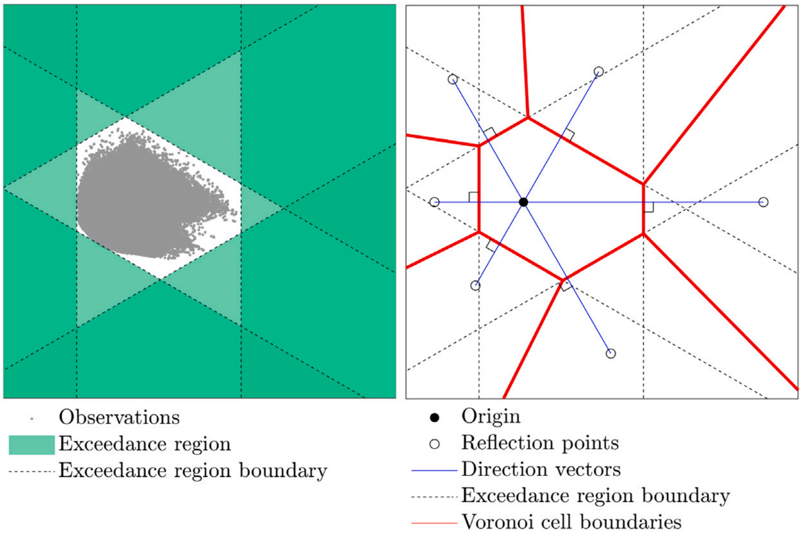

Fig. 8. Illustration of equivalent definitions of environmental contours. Left: 'Standard' definition in terms of intersection of non-exceedance regions. Right: Definition in terms of the Voronoi cell for the origin, with respect to the set of reflections of the origin in the boundaries of the exceedance regions.

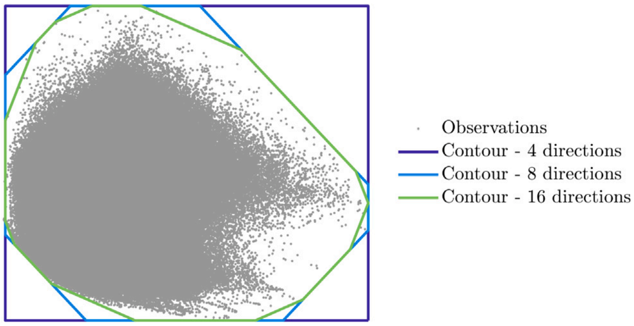

Fig. 9. Illustration of effect of calculating contours using different numbers of direction vectors.

2D case. In the case that $U = \{(\pm 1, 0), (0, \pm 1)\}$, the contour is defined by the intersection of the half-spaces representing the marginal return values. Clearly, such a contour would lead to conservative designs, if responses increase monotonically in each variable, as is often (but not always) the case for responses dependent on wave heights, wind speeds and current speeds.

The Voronoi method for calculating contours from half-spaces avoids the occurrence of 'loops' in the contour, which can occur when using the Monte Carlo method proposed by Huseby et al. (2013). This problem is neatly avoided using the Voronoi interpretation.

3.7. *Visualisation*

The contours that are calculated are output as lists of vertices of a hyper-surface, that is, a set of m points $C = \{\mathbf{c}_1, \ldots, \mathbf{c}_m\}$, where each point $\mathbf{c}_i = (c_{i,1}, \ldots, c_{i,d})$ is a d-dimensional vector. However, for selecting design conditions, it is useful to visualise the contour in 2 or 3 dimensions, either by projecting the contour into lower dimensions, or by taking slices through the contour for fixed values of certain variables.

Projection of the contour into lower dimensions is straightforward. Suppose that we wish to visualise the contour in three dimensions, corresponding to variables X_j, X_k and X_l. We define the set $C_{jkl} = \{\mathbf{c}'_1, \ldots, \mathbf{c}'_m\}$ as the restriction of C to dimensions X_j, X_k and X_l, that is, we simply ignore the values of vertices of the hyper-surface in the dimensions that we are not interested in, so that $\mathbf{c}'_i = (c_{i,j}, c_{i,k}, c_{i,l})$, $i = 1, \ldots, m$. We then note that the closed convex hull of a set is the intersection of all closed half-spaces containing it (Rockafellar, 1970). So the projected contour in dimensions X_j, X_k and X_l can be defined as the convex hull of C_{jkl}. In MATLAB, this is implemented in the native function `convhulln`, which also uses the Quickhull algorithms (Barber et al., 1996). This function outputs a triangulation of the vertices of C_{jkl}, which form the convex hull, so that a mesh of the 3D surface can be plotted.

A 2D slice through the contour for fixed values of certain variables can be calculated as follows. First, we calculate a triangulation of the hyper-surface into $(d-1)$-simplices, using the function `convhulln`. For each simplex which intersects with the values of the $d-2$ fixed variables, we find coordinates of the intersection. This can be calculated by interpolating along each edge of the simplex which intersects the fixed values. The contour of the 2D slice is then defined as the convex hull of the intersection points.

4. Example: environmental contours of wind and wave parameters

We consider the particular problem of defining design conditions for the design of offshore wind turbines. The design standard IEC 61400-3-1 (2019) requires designers to estimate extreme responses using an environmental contour. In design load case (DLC) 1.6 it is required that the design is checked for combinations of wind speed and significant wave height along a 50-year environmental contour, and that the spectral peak period should be chosen as the period that causes the highest loads at the particular combination of wind speed and wave height. Using the period that leads to the highest loads is conservative. However, this combination of probabilistic and deterministic choices of variables cannot be interpreted consistently in terms of failure probability and implied reliability (see the discussion in Haselsteiner et al. (2022)). The standard also requires that the misalignment in wind and wave directions should be considered, although it does not provide guidance on probabilistic means to do this.

In this section, we show how the D-IFORM method can be used to calculate environmental contours of wind speed (U_{10}), significant wave height (H_s), mean wave period (T_m), and wind-wave misalignment direction (θ_{rel}). These contours provide a rigorous means of selecting design conditions for offshore wind turbines that have a specified return period. The example uses data for a site in the Celtic Sea, off the south west coast of the UK, which has been identified for the development of floating offshore wind farms. The data comes from a

Table 1
Values of hyperparameters used in the example.

Spacing of L^1 directional vectors (see Eq. (2))	0.1
Peak separation time	2 days
Threshold exceedance probability	0.1

31-year hindcast of wind, wave and current conditions, with data at hourly intervals (Mackay and Hardwick, 2022).

As environmental contours are the boundary to a convex set, it is advantageous to work with combinations of variables for which the isodensity lines of the joint probability density function also define a convex region. Although the joint density is not estimated as part of the D-IFORM method, a visual inspection of scatter plots can indicate whether the isodensity lines are approximately convex. Wave breaking is governed by steepness. In deep water, a spectrally-averaged steepness can be defined as $s_m = 2\pi H_s/gT_m^2$. So, for a fixed steepness, we have H_s proportional to T_m^2. Moreover, wind drag on the sea surface is approximately quadratic, and for a Pierson–Moskowitz spectrum, the value of H_s is proportional to U_{10}^2 (Tucker and Pitt, 2001, §5.5.1.1). It is therefore useful to work with $\sqrt{H_s}$, rather than H_s, so that the region of observations is more convex. Scatter plots of H_s against U_{10} and T_m are shown in Fig. 10, showing the sample cloud is concave on the upper left boundary. Scatter plots of the normalised variables (see Section 3.1) are also shown, where the normalised wave height is calculated from $\sqrt{H_s}$. In these cases, the sample cloud of the normalised variables are more convex in shape.

As directions are periodic, it does not make sense to estimate extreme directions. Instead, we work with the components of wave height that are in line (longitudinal) and transverse to the wind direction. We therefore work with the following variables:

- U_{10}, hourly-mean wind speed at 10 m above sea level
- $H_L = \sqrt{H_s}\cos(\theta_{rel})$, the longitudinal component of $\sqrt{H_s}$, in line with the wind direction.
- $H_T = \sqrt{H_s}\sin(\theta_{rel})$, the transverse component of $\sqrt{H_s}$, at right angles to the wind direction.
- T_m, mean wave period

Note that H_L and H_T are not intended to represent physical transport of energy in particular directions. Rather, they simply intended to be convenient quantities to work with for the calculation of contours. If return values of combinations of H_L and H_T are known, then the corresponding return values of H_s and θ_{rel} can be retrieved using the inverse transformation:

$$H_s = H_L^2 + H_T^2, \tag{9}$$

$$\theta_{rel} = \mathrm{atan2}(H_T, H_L), \tag{10}$$

where atan2 is the four-quadrant inverse tan function.

The application of the D-IFORM method requires the specification of three hyper-parameters, listed in Table 1. For the analysis we have used L^1 direction vectors with a spacing of 0.1 in each dimension (see Eq. (2)). From Section 3.2, we see that using this spacing gives 286 direction vectors in each orthant of $\mathbb{R}^4$. Using reflection to obtain direction vectors in each of the 16 orthants and removing duplicates (i.e. those vectors where the distance from the reflection hyperplane is zero), results in 2720 direction vectors. The declustering procedure used a peak-separation time of 48 hours, and identified between approximately 2100 and 2600 peaks at each direction. The threshold exceedance probability was fixed at $\zeta = 0.1$, which resulted in between 210 and 260 exceedances used to fit the GP model at each direction. For the fitted GP distributions, the upper bound for the shape parameter estimate, of $\hat{\xi} = 0$, was reached for 5 out of 2720 directions, and the lower bound of $\hat{\xi} = -1$ was reached once. This relatively low occurrence is likely due to the use of a reasonably long dataset of 31 years. If a shorter dataset had been used, then the upper bound would likely have been reached more often, due to sampling uncertainty in estimating the GP parameters (see e.g. Mackay et al. (2011)). The total computational time for this example was approximately 30 s on a laptop with an Intel® Core™ i7-7500u processor. I.e. the time required to compute the POT analysis for each direction vector is less than $30/2720 = 0.011$ s, since the total computational time also includes the other steps in the analysis.

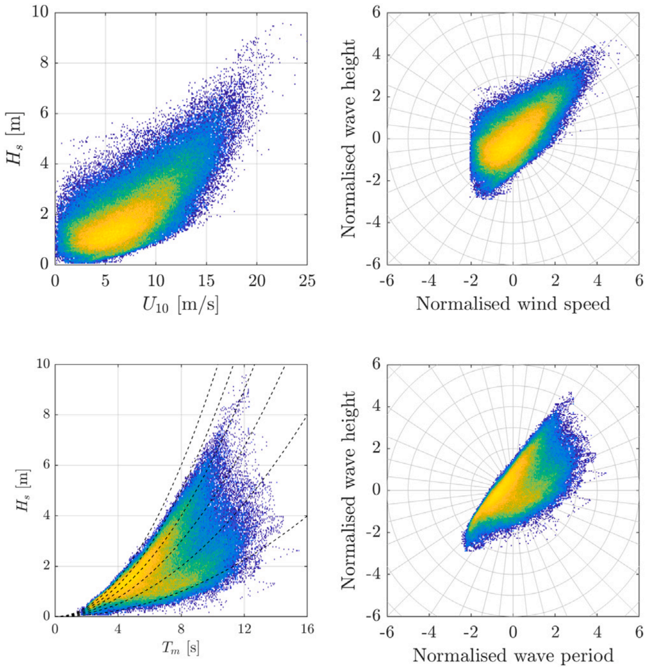

Fig. 10. Scatter plots for original and normalised variables, colour denotes empirical density. Normalised wave height uses $\sqrt{H_s}$, so that resulting sample cloud is more convex. Radial lines in right hand plots indicate direction vectors used in the analysis. The non-constant spacing of direction vectors is due to the method used to create an approximately regular spacing of direction vectors in higher dimensions. Dashed lines on lower left plot indicate constant significant steepness at levels 0.01, 0.02, …, 0.06.

Examples of estimated return values as a function of projection angle are shown in Figs. 11 and 12, together with the declustered peaks at each angle. In Fig. 11 the projection angle is the Cartesian angle calculated in the plane where the first dimension is normalised wind speed and the second dimension is normalised wave period. In Fig. 12 the plane has first dimension of normalised H_L and the second dimension of normalised H_T. In both cases, there is a smooth variation of the threshold and return values (up to the 50-year level) with projection angle, indicating that using independent analyses at each angle gives reasonable results and using more sophisticated inference is unlikely to result in large changes. Examples of the fit of the GP model to the threshold exceedances are shown in Fig. 13 for various projection directions. A 95% confidence interval for the empirical exceedance probabilities is also shown (see Mackay and Jonathan (2021) for details). In all cases shown, the GP model provides a good fit to the observations, with deviations within the range expected due to sampling uncertainties. Similarly, good fits were observed for other directions (not shown here).

The comparison of the tails of the empirical and fitted distributions, shown in Fig. 13, provides the most direct diagnostic tool for assessing the fit of the models to the observation. Projections of the contours into 3D and 2D provide a more heuristic indication of adequacy of the estimated contours (see Figs. 14 and 15). Broadly speaking, we are looking to check that the contours provide a reasonable description of the range of observations. The question of how many points we would expect to observe outside the contours is complex to answer, as multiple effects are present. Firstly, exceedances tend to occur in clusters due to serial

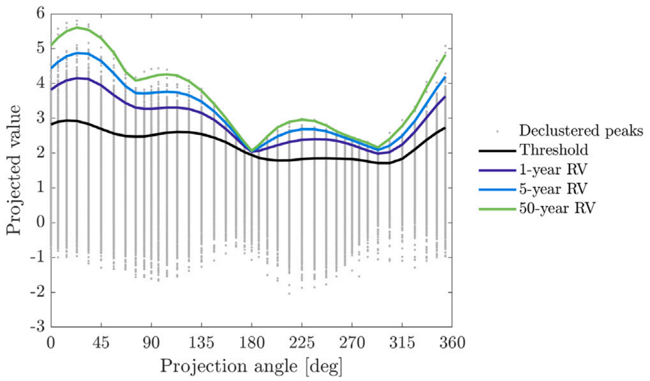

Fig. 11. Declustered peaks against projection angle, with GP threshold and estimated return values. The projection angle is the Cartesian angle calculated in the plane where the first dimension is normalised wind speed and the second dimension is normalised wave period.

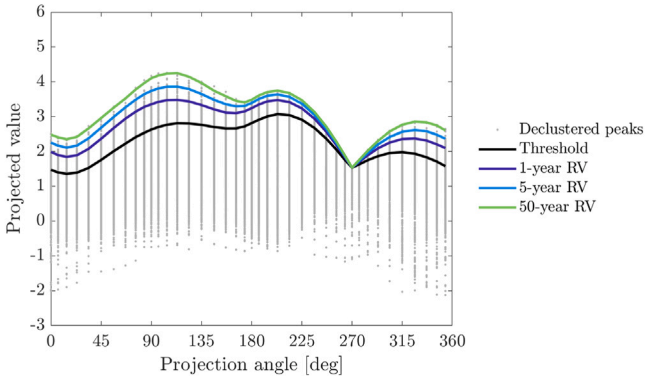

Fig. 12. As previous figure, but with first dimension being normalised H_L and the second dimension being normalised H_T.

correlation of the variables. Secondly, the contour is defined in terms of an intersection of multiple non-exceedance regions, so it difficult to say how many observations would be expected to fall outside the contour in a given period. The difference between the probability of an observation falling in a particular half-space exceedance region and the probability of an observation falling anywhere outside the contour was discussed in detail by Mackay and Haselsteiner (2021).

For a particular projection direction, $\mathbf{u}$, the problem of calculating the probability that the maximum observation in N years does not exceed the T-year return value, is straightforward. First, define $M_{\mathbf{u},N}$ to be the random variable representing the maximum value of $R_{\mathbf{u}}$ in an N-year period. Then the T-year return value, $R_{\mathbf{u},T}$, is defined as the solution of

$$P(M_{\mathbf{u},1} \leq R_{\mathbf{u},T}) = 1 - \frac{1}{T}. \tag{11}$$

As the maximum value in each year is independent, the probability of not exceeding $R_{\mathbf{u},T}$ in N years is

$$P(M_{\mathbf{u},N} \leq R_{\mathbf{u},T}) = \left(P(M_{\mathbf{u},1} \leq R_{\mathbf{u},T})\right)^N = \left(1 - \frac{1}{T}\right)^N. \tag{12}$$

In this particular example, $N = 31$, so $P(M_{\mathbf{u},31} \leq R_{\mathbf{u},50}) \approx 0.53$. Therefore, there is a roughly 50–50 chance that we will have at least one observation falling outside the 50-year contour in any particular direction. It is difficult to generalise to points falling anywhere outside the contour, as the peaks are correlated between directions. Consider a 2D example, and denote the coordinates of an observation, not necessarily a peak, as $(x_i, y_i) = \rho_i(\sin(\theta_i), \cos(\theta_i))$. The projection of this observation onto direction vector $\mathbf{u} = (\cos(\alpha), \sin(\alpha))$ is $r_{\mathbf{u},i} = x_i \cos(\alpha) + y_i \sin(\alpha) = \rho_i \cos(\alpha - \theta_i)$. So, the projected distance from the origin is periodic in the projection direction, α. However, as noted in Section 3.3, the points which correspond to local peaks in the time series, will depend on the projection direction. It is therefore it is difficult to make more general statements about how exceedances vary with projection direction.

Examples of 3D projections of the 4D contours are shown in Fig. 14. The left had plot shows the 50-year contour of U_{10}, H_L and H_T, and the right hand plot shows the 50-year contour of U_{10}, T_m and H_s. In the latter case, the contour vertices for H_s have been calculated from those for H_L and H_T using (9). In both cases, the 3D contours (surfaces) provide a good description of the extent of the observations. Fig. 15 shows 2D projections for contours of H_s vs. U_{10}, T_m vs. U_{10}, H_s vs. T_m, and $H_{s,long}$ vs. $H_{s,trans}$. In the latter case, the variables $H_{s,long} = H_s \cos(\theta_{rel})$ and $H_{s,trans} = H_s \sin(\theta_{rel})$ are shown in preference to H_L and H_T, so that the dimensions are in metres. In all cases, the contours provide a good description of the observations. In the case of the contours for (H_s, U_{10}) and (H_s, T_m), the lower values of observations are not tracked well by the contours. This is a result of how the contours are calculated in terms of H_L and H_T. The reason for this is discussed further below. However, it should be noted that when estimating extreme responses, it is only the largest values of H_s at a given wave period, wind speed and misalignment direction that are of interest. The 2D projected contours therefore provide an adequate description in this respect.

It is interesting to consider the variation in the 50-year H_s as a function of θ_{rel}, indicated in the lower right plot of Fig. 15. As would be expected, the largest values of H_s occur when the wind and wave direction are aligned, with a 50-year value of around 9.5 m. When wind and wave directions are perpendicular (i.e. $H_{s,long} = 0$), the 50-year H_s is just under 6 m. Finally, when winds and waves are in opposing directions, the 50-year H_s is around 4 m.

These 2D and 3D plots are useful for assessing the fit of the contours to the observations. However, for specifying design conditions, we need to know the joint extreme conditions of all values. Fig. 16 shows 2D slices through the 4D contours for $U_{10} = 11$ m/s and various values of T_m. This corresponds to somewhere around the rated speed of an offshore wind turbine, where the largest loads on some components can occur. For example, the IEA 15 MW reference turbine has a rated wind speed of 10.6 m/s (Gaertner et al., 2020). The slices show that for lower periods, the largest H_s is limited by wave breaking. This is evident by only small differences between 1-year and 50-year return values. For longer periods, the breaking limit is not always reached, and there are differences in return values with return period and with misalignment direction. The plots also illustrate that the joint distribution of H_L and H_T is 'hollow'. That is, for a given wind speed and wave period, there is only a narrow range of H_s that is observed, with low values of H_s at larger wind speed being unlikely due to the physical connection between winds and waves. As the contours are derived in terms of H_L and H_T, the contours track the outer range of these variables, rather than the lower values of H_s at a given wind speed or wave period, explaining the trends shown in Fig. 15.

5. Discussion and conclusions

This work has presented a method for estimating environmental contours in arbitrary numbers of dimensions. The key insight utilised in the D-IFORM approach is that the half-space exceedance regions, which are used to define IFORM and Direct Sampling contours, correspond to univariate exceedance regions under rotations of the axes. This removes the need to fit a joint probability model, removing a significant source of uncertainty associated with other environmental contour methods. It also allows all observations to be used in the analysis at each direction, so that increasing the number of variables does not increase uncertainties, as happens when a joint distribution model is fitted in higher dimensions.

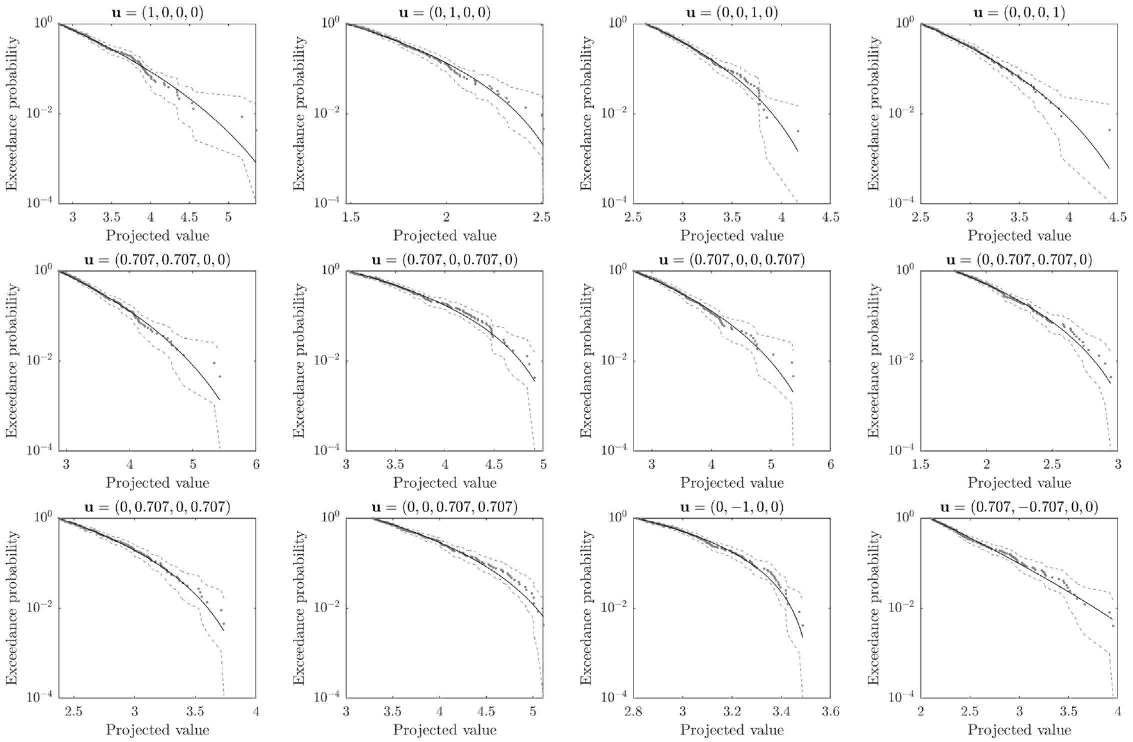

Fig. 13. Fit of GP model for various projection directions. Dots: peaks of projected data. Dashed lines: 95% CI for empirical exceedance probability. Solid lines: GP model.

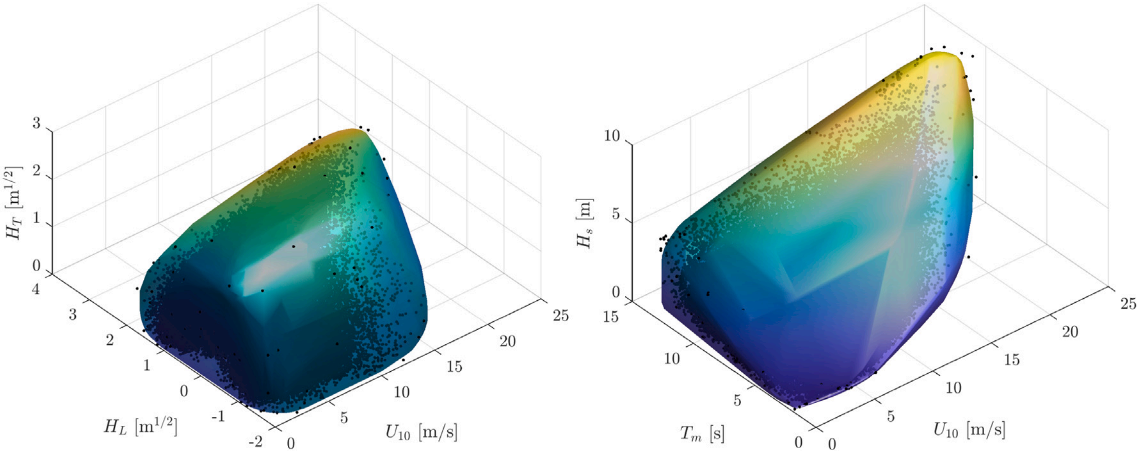

Fig. 14. Examples of 3D projections of 4D contours compared with observations (dots). Left: 50-year contour of U_{10}, H_L and H_T. Right: 50-year contour of U_{10}, T_m and H_s.

The application of the method requires only three parameters to be specified by the user: (1) the spacing of the directional vectors, (2) the peak separation time, and (3) the threshold exceedance probability for fitting the GP model. The GP model itself is asymptotically justified from extreme value theory. In comparison, methods which involve fitting a model for the joint distribution require much stronger assumptions about the form of the marginal distributions and dependence structure.

For the four-dimensional example presented here, the computational effort is very low, requiring around 30 s. It is interesting to consider what would happen if other variables which may influence the response of an offshore structure were included in the analysis. For example, current speed, current direction and water level may also have an effect on the response. Including these variables would result in a seven-dimensional problem. Using a spacing of 0.1 for direction vectors, as in the example in Section 4, would give 209762 directions for the analysis (see Section 3.2). Compared to the example presented, this would still result in reasonable computational times of less than an hour.

The use of environmental contours provides a reduction of one dimension for the problem of calculating extreme responses, compared to ‘full long-term analysis’ methods, where the response must be estimated for the full variable space. This reduction in dimensionality is useful, but in higher dimensions the number of simulations required to fully characterise a contour can be very large. High dimensional contours can be used as a way of choosing plausible combinations of parameters for assessing the response. It is not a requirement to conduct response simulations for all points on a contour. Instead, sensitivity studies can be used to indicate regions likely to lead to high response. Engineering judgement can be applied in the same way as for lower dimensional contours, when values of other variables must be decided for the response analysis.

Finally, it is important to note that extreme responses calculated using the D-IFORM method will be subject to the same limitations as

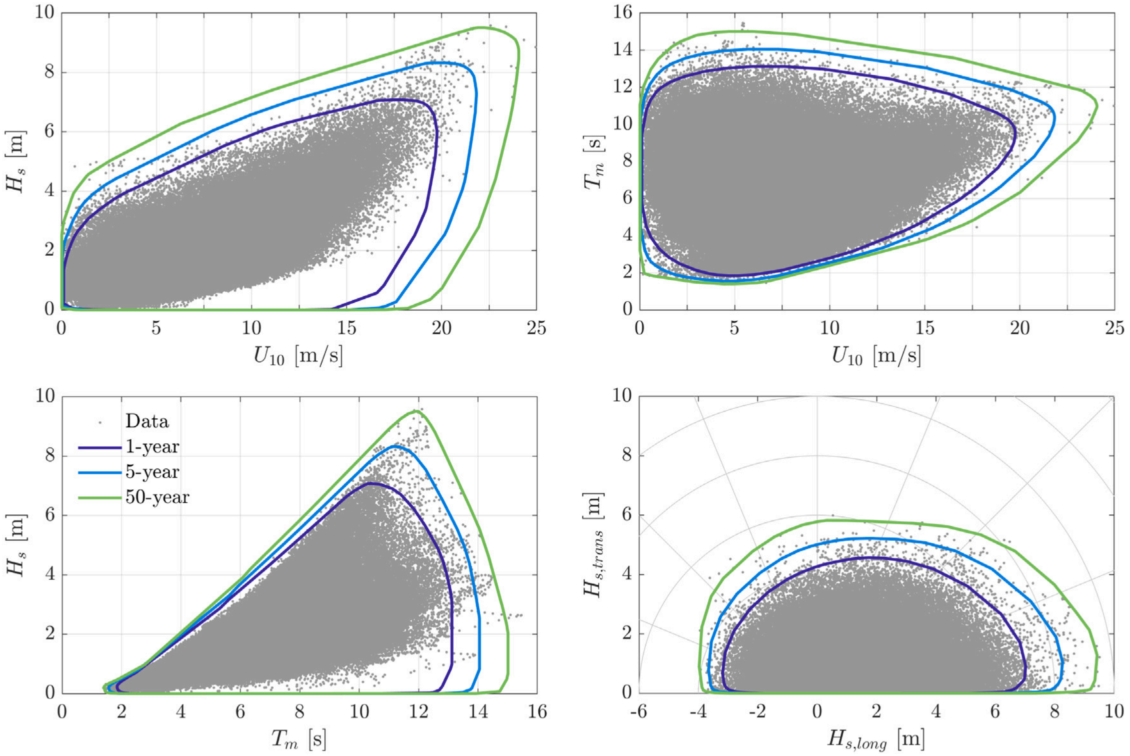

Fig. 15. Examples of 2D projections of 4D contours compared with observations (dots). Angular-radial grid on lower right plot indicates wind-wave misalignment angle and H_s (radius).

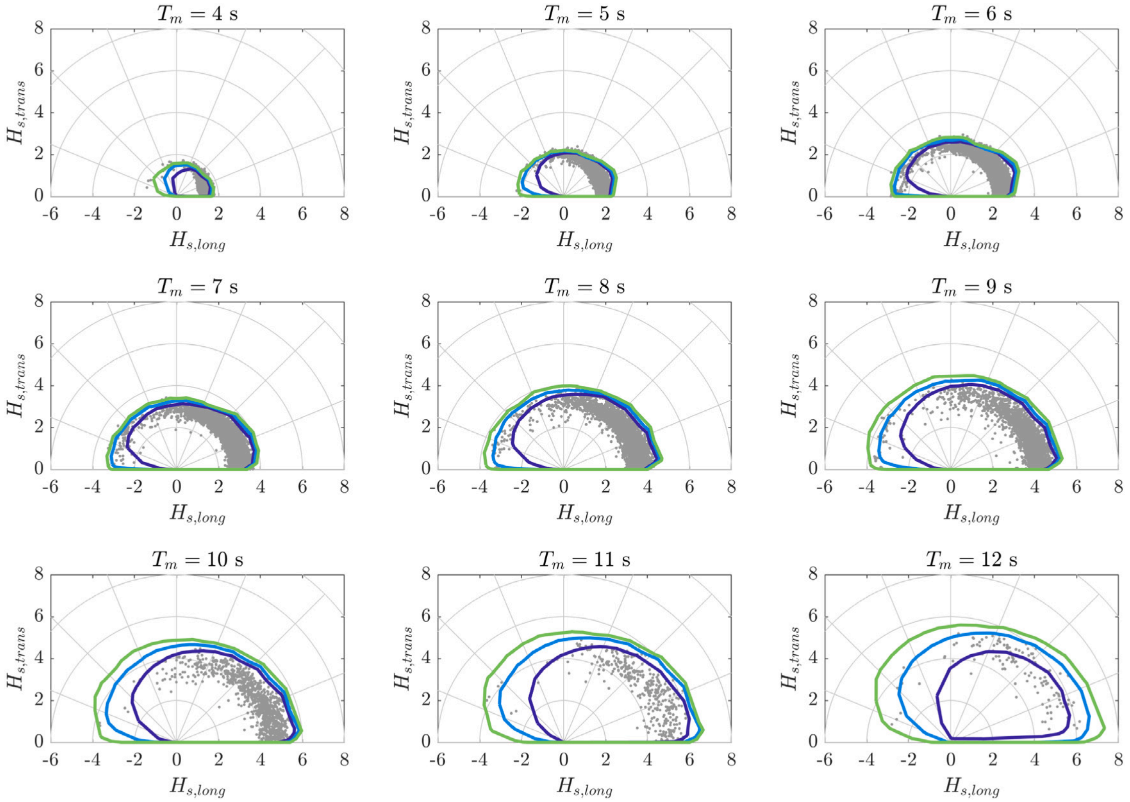

Fig. 16. Examples of 2D slices through 4D contours for $U_{10} = 11$ m/s and various values of T_m, compared with observations (dots). Contour return periods are 1-, 5- and 50-years. Bin size for observations is ± 1 m/s for U_{10} and ± 0.5 s for T_m.

the standard IFORM and Direct Sampling methods, in terms of the assumption of a linearised failure surface and neglecting short-term variability in the response. The effect of these assumptions has been discussed elsewhere in the literature (see e.g. de Hauteclocque et al. (2022) and Haselsteiner et al. (2022) for recent examples). Moreover, the extreme responses are sensitive to the variable space in which the contours are constructed. This is because a linear approximation to a failure surface in one variable space will not necessarily be linear when transformed into another variable space (see e.g. Huseby et al. (2013) and Vanem (2017)). This highlights the uncertainty introduced by using the FORM approximation of a linear failure surface. Nevertheless, despite these uncertainties, extreme responses inferred using D-IFORM environmental contours have been shown to be in good agreement with extreme responses calculated from response-based analyses, for a wide range of structures (de Hauteclocque et al., 2022).

CRediT authorship contribution statement

Ed Mackay: Conceptualization, Methodology, Software, Writing – original draft. **Guillaume de Hauteclocque:** Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgements

EM was funded under the Tidal Stream Industry Energiser Project (TIGER), France, a European Union INTERREG V A France (Channel) England Research and Innovation Programme, which is cofinanced by the European Regional Development Fund (ERDF), United Kingdom.

References

Barber, C.B., Dobkin, D.P., Huhdanpaa, H., 1996. The Quickhull algorithm for convex hulls. ACM Trans. Math. Softw. 22 (4), 469–483.

Barlow, A.M., Mackay, E., Eastoe, E., Jonathan, P., 2023. A penalised piecewise-linear model for non-stationary extreme value analysis of peaks over threshold. Ocean Engineering 267, 113265. http://dx.doi.org/10.1016/j.oceaneng.2022.113265.

Beirlant, J., Goegebeur, Y., Segers, J., Teugels, J., 2004. Statistics of Extremes: Theory and Applications. Wiley, Chichester, UK.

Chai, W., Leira, B.J., 2018. Environmental contours based on inverse SORM. Mar. Struct. 60, 34–51. http://dx.doi.org/10.1016/j.marstruc.2018.03.007.

Coe, R.G., Manuel, L., Haselsteiner, A.F., 2022. On limiting the influence of serial correlation in metocean data for prediction of extreme return levels and environmental contours. Ocean Eng. 266, 113032. http://dx.doi.org/10.1016/j.oceaneng.2022.113032.

Coles, S., 2001. An Introduction to Statistical Modeling of Extreme Values. Springer, http://dx.doi.org/10.1198/tech.2002.s73.

Davis, R., Mikosch, T., 2009. The extremogram: A correlogram for extreme events. Bernoulli 15, 977–1009. http://dx.doi.org/10.3150/09-BEJ213.

Derbanne, Q., de Hauteclocque, G., 2019. A new approach for environmental contour and multivariate de-clustering. In: 38th International Conference on Ocean, Offshore and Arctic Engineering. Glasgow, p. OMAE2019/95993. http://dx.doi.org/10.1115/OMAE2019-95993.

Gaertner, E., J., R., Sethuraman, L., Zahle, F., Anderson, B., Barter, G., Abbas, N., Meng, F., Bortolotti, P., Skrzypinski, W., Scott, G., Feil, R., Bredmose, H., Dykes, K., Shields, M., Allen, C., Viselli, A., 2020. Definition of the IEA 15-Megawatt Offshore Reference Wind Turbine. NREL/TP-5000-75698, National Renewable Energy Laboratory, Golden, CO, URL: https://www.nrel.gov/docs/fy20osti/75698.pdf.

Hafver, A., Agrell, C., Vanem, E., 2022. Environmental contours as Voronoi cells. Extremes 25, 1–36. http://dx.doi.org/10.1007/S10687-022-00437-7.

Haselsteiner, A.F., Coe, R.G., Manuel, L., Chai, W., Leira, B., Clarindo, G., Guedes Soares, C., Hannesdóttir, Á., Dimitrov, N., Sander, A., Ohlendorf, J.H., Thoben, K.D., de Hauteclocque, G., Mackay, E., Jonathan, P., Qiao, C., Myers, A., Rode, A., Hildebrandt, A., Schmidt, B., Vanem, E., Huseby, A.B., 2021a. A benchmarking exercise for environmental contours. Ocean Eng. 236, 109504. http://dx.doi.org/10.1016/J.OCEANENG.2021.109504.

Haselsteiner, A.F., Frieling, M., Mackay, E., Sander, A., Thoben, K.D., 2022. Long-term extreme response of an offshore turbine: How accurate are contour-based estimates? Renew. Energy 181, 945–965. http://dx.doi.org/10.1016/j.renene.2021.09.077.

Haselsteiner, A.F., Mackay, E., Thoben, K.-D., 2021b. Reducing conservatism in highest density environmental contours. Appl. Ocean Res. 117, 102936. http://dx.doi.org/10.1016/j.apor.2021.102936.

Haselsteiner, A.F., Ohlendorf, J.-H., Wosniok, W., Thoben, K.-D., 2017. Deriving environmental contours from highest density regions. Coast. Eng. 123, 42–51. http://dx.doi.org/10.1016/j.coastaleng.2017.03.002.

de Hauteclocque, G., Mackay, E., Vanem, E., 2022. Quantitative assessment of environmental contour approaches. Ocean Eng. 245, 110374. http://dx.doi.org/10.1016/j.oceaneng.2021.110374.

Haver, S., 1985. Wave climate off Northern Norway. Appl. Ocean Res. 7, 85–92. http://dx.doi.org/10.1016/0141-1187(85)90038-0.

Huseby, A.B., Vanem, E., Natvig, B., 2013. A new approach to environmental contours for ocean engineering applications based on direct Monte Carlo simulations. Ocean Eng. 60, 125–135. http://dx.doi.org/10.1016/j.oceaneng.2012.12.034.

International Electrotechnical Commission, 2019. Wind Energy Generation Systems - Part 3-1: Design Requirements for Fixed Offshore Wind Turbines. IEC 61400-3-1.

Joe, H., 2015. Dependence Modeling with Copulas. CRC Press, pp. 1–457.

Lagarias, J.C., Reeds, J.A., Wright, M.H., Wright, P.E., 1998. Convergence properties of the Nelder–Mead simplex method in low dimensions. SIAM J. Optim. 9 (1), 112–147. http://dx.doi.org/10.1137/S1052623496303470.

Mackay, E.B.L., Challenor, P.G., Bahaj, A.S., 2011. A comparison of estimators for the generalised Pareto distribution. Ocean Eng. 38, 1338–1346. http://dx.doi.org/10.1016/j.oceaneng.2011.06.005.

Mackay, E.B., Hardwick, J.P., 2022. Joint extremes of waves and currents at tidal energy sites in the English Channel. In: 42nd International Conference on Ocean, Offshore and Arctic Engineering. Hamburgh, p. OMAE2022/79348. http://dx.doi.org/10.1115/OMAE2022-79348.

Mackay, E., Haselsteiner, A.F., 2021. Marginal and total exceedance probabilities of environmental contours. Mar. Struct. 75, 102863. http://dx.doi.org/10.1016/j.marstruc.2020.102863.

Mackay, E., de Hauteclocque, G., Vanem, E., Jonathan, P., 2021. The effect of serial correlation in environmental conditions on estimates of extreme events. Ocean Eng. 242, 110092. http://dx.doi.org/10.1016/j.oceaneng.2021.110092.

Mackay, E., Johanning, L., 2018. Long-term distributions of individual wave and crest heights. Ocean Eng. 165, 164–183. http://dx.doi.org/10.1016/j.oceaneng.2018.07.047.

Mackay, E., Jonathan, P., 2021. Sampling properties and empirical estimates of extreme events. Ocean Eng. 239, 109791. http://dx.doi.org/10.1016/j.oceaneng.2021.109791.

Madsen, H.O., Krenk, S., Lind, N., 2006. Methods of Structural Safety. Dover Publications, Mineola, New York.

Manuel, L., Nguyen, P.T., Canning, J., Coe, R.G., Eckert-Gallup, A.C., Martin, N., 2018. Alternative approaches to develop environmental contours from metocean data. J. Ocean Eng. Mar. Energy 4, 293–310. http://dx.doi.org/10.1007/s40722-018-0123-0.

Mathisen, J., Bitner-Gregersen, E., 1990. Joint distributions for significant wave height and wave zero-up-crossing period. Appl. Ocean Res. 12 (2), 93–103. http://dx.doi.org/10.1016/S0141-1187(05)80033-1.

Montes-Iturrizaga, R., Heredia-Zavoni, E., 2015. Environmental contours using copulas. Appl. Ocean Res. 52, 125–139. http://dx.doi.org/10.1016/j.apor.2015.05.007.

Montes-Iturrizaga, R., Heredia-Zavoni, E., 2017. Assessment of uncertainty in environmental contours due to parametric uncertainty in models of the dependence structure between metocean variables. Appl. Ocean Res. 64, 86–104. http://dx.doi.org/10.1016/j.apor.2017.02.006.

Nelsen, R.B., 2006. An Introduction to Copulas. Springer.

Randell, D., Feld, G., Ewans, K., Jonathan, P., 2015. Distributions of return values for ocean wave characteristics in the South China Sea using directional-seasonal extreme value analysis. Environmetrics 26, 442–450. http://dx.doi.org/10.1002/env.2350.

Rockafellar, R.T., 1970. Convex Analysis. Princeton University Press.

Rosenblatt, M., 1952. Remarks on a multivariate transformation. Ann. Math. Stat. 23 (3), 470–472. http://dx.doi.org/10.1214/aoms/1177729394.

Ross, E., Astrup, O.C., Bitner-Gregersen, E., Bunn, N., Feld, G., Gouldby, B., Huseby, A., Liu, Y., Randell, D., Vanem, E., Jonathan, P., 2020. On environmental contours for marine and coastal design. Ocean Eng. 195, 106194. http://dx.doi.org/10.1016/j.oceaneng.2019.106194.

Silva-González, F., Heredia-Zavoni, E., Montes-Iturrizaga, R., 2013. Development of environmental contours using Nataf distribution model. Ocean Eng. 58, 27–34. http://dx.doi.org/10.1016/j.oceaneng.2012.08.008.

Tucker, M.J., Pitt, E.G., 2001. Waves in Ocean Engineering, second ed. Elsevier.

Vanem, E., 2017. A comparison study on the estimation of extreme structural response from different environmental contour methods. Mar. Struct. 56, 137–162. http://dx.doi.org/10.1016/j.marstruc.2017.07.002.

Vanem, E., 2022. Analysing multivariate extreme conditions using environmental contours and accounting for serial dependence. Renew. Energy http://dx.doi.org/10.1016/j.renene.2022.11.033.

Winterstein, S.R., Ude, T.C., Cornell, C.A., Bjerager, P., Haver, S., 1993. Environmental parameters for extreme response: Inverse FORM with omission factors. In: 6th International Conference on Structural Safety & Reliability (ICOSSAR).

Zanini, E., Eastoe, E., Jones, M., Randell, D., Jonathan, P., 2020. Flexible covariate representations for extremes. Environmetrics 31, e2624. http://dx.doi.org/10.1002/env.2624.