



*Highly advanced Probabilistic design and Enhanced Reliability methods
for high-value, cost-efficient offshore WIND*

Title: Methods for adaptive calculation of FLS loads and reliability, and their impact on probabilistic fatigue design

Deliverable no: D4.3

Delivery date: 20 June 2024

Lead beneficiary: EDF R&D

Dissemination level: Public



*This project has received funding from the European Union's Horizon
2020 Research and Innovation Program under Grant Agreement No.
101006689*

Author information:

Name	Organization	Address	Email
E. Fekhari	EDF R&D	<i>6 quai Watier, Chatou, France</i>	elias.fekhari@edf.fr
V. Chabridon	EDF R&D	<i>6 quai Watier, Chatou, France</i>	vincent.chabridon@edf.fr
S. Marelli	ETH Zürich	<i>Stefano-Francini-Platz, 5, Zürich, Switzerland</i>	marelli@ibk.baug.ethz.ch
E. Ardillon	EDF R&D	<i>6 quai Watier, Chatou, France</i>	emmanuel.ardillon@edf.fr

Document information:

Version	Date	Description	Prepared by	Reviewed by	Approved by
1.0	20-06-2024	Official	Authors listed above	N. Dimitrov	N. Dimitrov

Acknowledgements:

The authors are grateful to N. Dimitrov for reviewing this document, for numerous technical discussions, and for coordinating the HIPERWIND project. The authors would also like to thank Professor B. Sudret for the fruitful discussions on probabilistic fatigue modeling and reliability analysis. Finally, the authors would like to thank several members of the HIPERWIND project such as A. Lovera, E. Vanem, M. Capaldo, M. Guiton and C. Peyrard for their technical assistance.

Abstract:

Wind energy companies constantly deploy cutting-edge wind turbines to reach unexploited wind resources. In these highly uncertain environments, numerical models simulating such systems can be used to assess the performance and reliability of wind turbines. This numerical procedure is often computationally expensive. However, promising new kernel-based techniques aim at simplifying such a computational burden. This report first demonstrates the use of a fast sampling technique on the numerical model simulating a real offshore wind turbine operating in Teesside, UK. This sampling method conveniently extracts a small number of relevant input points (i.e., environmental conditions) from a measured data set to be used as simulation inputs for the estimation of cumulative lifetime damage. Secondly, a reliability analysis is conducted on additional sources of uncertainties related to the system itself (e.g., yaw misalignment, soil stiffness, and material properties). As the risk measure estimated in this reliability analysis presented an important sensitivity to the critical damage threshold, a perturbation-based robustness analysis was also applied on the threshold. Finally, the impact of the wake on the reliability and robustness related to fatigue damage was studied, using the wake modeling results from Work Package 3.

Contents

1	Introduction	6
1.1	Industrial context and motivations	6
1.2	Uncertainty quantification framework	6
1.3	Deliverable objectives and structure of the report	7
2	Teesside multi-physics wind turbine model	8
2.1	Introduction	8
2.2	Numerical simulation model	8
2.3	Fatigue assessment	9
2.4	Synthesis	11
3	Quantification and modeling of input uncertainties	12
3.1	Introduction	12
3.2	Definition of the environmental random inputs	12
3.3	Definition of the system random inputs	13
3.4	Probabilistic fatigue assessment	14
3.5	Synthesis	14
4	A first uncertainty propagation approach for cumulated fatigue damage estimation	15
4.1	Introduction	15
4.2	A first integration approach based on kernels	17
4.3	Numerical experiments	24
4.4	Synthesis	28
5	A second integration approach based on sparse polynomial chaos expansions	31
5.1	Introduction	31
5.2	Polynomial chaos expansions	31
5.3	Numerical experiments	32
5.4	Synthesis	35
6	Reliability and robustness analyses in the context of fatigue	35
6.1	Introduction	35
6.2	Surrogate modeling for reliability analysis	36
6.3	Reliability and robustness analysis	39
6.4	Synthesis	42
7	Reliability and robustness analyses including the wake	43
7.1	Introduction	43
7.2	Surrogate model for reliability analysis	44
7.3	Nominal reliability and robustness analysis	46
7.4	Synthesis	46
8	Conclusion	47

List of Figures

1.1	General uncertainty quantification framework (De Rocquigny et al., 2008, adapted by Ajenjo, 2023).	7
2.1	Diagram of the chained OWT simulation model.	8
2.2	Teesside wind farm layout (left). Monopile OWT diagram from (Chen et al., 2018) (right).	10
2.3	Angular distribution of the wind and waves with a horizontal cross-section of the OWT structure and the mudline. Red crosses represent the discretized azimuths for which the fatigue is computed.	11
3.1	Copulogram of the Teesside measured data ($N = 10^4$ in gray). Marginals are represented by univariate KDE plots (diagonal), the dependence structure with scatter plots in the rank space (upper triangle). Scatter plots on the bottom triangle are set in the physical space.	13
3.2	Illustration of a probabilistic S-N curve according to the model defined in Gu���� et al. (2007).	15
4.1	Histogram of the log-damage, at mudline, azimuth 45 deg. (Monte Carlo reference sample).	17
4.2	Kernel mean embedding of a continuous and discrete probability distribution	18
4.3	Kernel mean embedding of continuous and discrete distributions.	20
4.4	Greedy KH algorithm.	21
4.5	Kernel illustrations (left to right: energy-distance, squared exponential, and Mat���� 5/2).	22
4.6	Sequential KH for increasing design sizes ($n \in \{10, 20, 40\}$) built on a candidate set of $N = 8196$ points drawn from a complex Gaussian mixture π .	22
4.7	Bayesian quadrature on a one-dimensional case.	23
4.8	Analytical benchmark results on the test case #1.	26
4.9	Analytical benchmark results on the test case #2.	27
4.10	Mean damage estimation workflows for the industrial Teesside use case. The orange parts represent optional alterations to the workflow: the first one is an alternative to input data subsampling where the underlying distribution is sampled from, the second one improves mean damage calculation by using optimal weights over the output data.	28
4.11	Copulogram of the KH design of experiments with corresponding outputs in color (log-scale) on the Teesside case ($n = 10^3$). The color scale ranges from blue for the lowest values to red for the largest.	29
4.12	Mean estimation convergence (at the mudline, azimuth $\theta = 45$ deg.) on the Teesside case. Monte Carlo confidence intervals are all computed by bootstrap.	30
5.1	Finite elements model of a 23-bar truss structure used in the first numerical example, Section 5.3.1. The horizontal elements are highlighted in blue, while the diagonal ones are highlighted in magenta.	33
5.2	Comparison of different integration strategies for the finite-element truss structure in Section 5.3.1. Left panel: convergence of the integral value (mean) as a function of the number of model evaluations in the experimental design. Right panel: convergence of the relative mean absolute error for the same configurations.	34
5.3	Comparison of different integration strategies for the finite-element truss structure in Section 5.3.2. Left panel: convergence of the integral value (mean) as a function of the number of model evaluations in the experimental design. Right panel: convergence of the relative mean absolute error for the same configurations.	34
5.4	Comparison of different integration strategies for the finite-element truss structure in Section 5.3.2. Left panel: convergence of the integral value (mean) as a function of the number of model evaluations in the experimental design. Right panel: convergence of the relative mean absolute error for the same configurations.	35
6.1	Learning set of the mean damage surrogate model. A Halton sequence is first built (in blue) and completed by a set KH-generated points (in orange) in a subdomain defined a priori (in gray).	37
6.2	Mean damage evaluated on the mixed design illustrated in Fig. 6.1.b.	38
6.3	Three-dimensional plot of the surrogate model $\tilde{D}_{\varepsilon_0}$ (in blue) and learning set (in black).	38
6.4	Cross-section of the surrogate model $\tilde{D}_{\varepsilon_0}$ (in shades of blue) for given values of K_{soil} and Θ_{yaw} . The darker the shade, the closer to the cross-section.	39
6.5	Leave-one-out validation results of the surrogate model $\tilde{D}_{\varepsilon_0}$.	39
6.6	Perturbations in terms of standard deviation of a lognormal distribution (left) and a truncated normal distribution (right).	41
6.7	Perturbed-law based indices for relative perturbations of the standard deviations of $(K_{\text{soil}}, \Theta_{\text{yaw}}, \varepsilon)$. The failure probabilities studied are each estimated by FORM-IS method on the surrogate model (with sample size $N = 5 \times 10^4$).	41

6.8	Perturbed-law based indices for relative perturbations of the standard deviation of D_{cr} . The failure probabilities studied are each estimated by FORM-IS method on the surrogate model (with sample size $N = 5 \times 10^4$).	42
7.1	Dissimilarity measure between Teesside's ambient metocean distribution and the wake-perturbed distributions at each wind turbine (see Lovera et al., 2023 for further details)	43
7.2	Mean damage evaluated on the mixed design illustrated in Fig. 6.1.b.	44
7.3	Three-dimensional plot of the surrogate model $\tilde{D}_{\varepsilon_0}$ (in blue) and learning set (in black).	45
7.4	Cross-section of the surrogate model $\tilde{D}_{\varepsilon_0}$ (in shades of blue) for given values of K_{soil} and Θ_{yaw} . The darker the shade, the closer to the cross-section.	45
7.5	Leave-one-out validation results of the surrogate model $\tilde{D}_{\varepsilon_0}$.	45
7.6	Perturbed-law based indices for relative perturbations of the standard deviations of $(K_{soil}, \Theta_{yaw}, \varepsilon)$. The failure probabilities studied are each estimated by FORM-IS method on the surrogate model (with sample size $N = 5 \times 10^4$).	46
7.7	Perturbed-law based indices for relative perturbations of the standard deviation of D_{cr} . The failure probabilities studied are each estimated by FORM-IS method on the surrogate model (with sample size $N = 5 \times 10^4$).	47

List of Tables

2.1	Teesside OWT datasheet.	9
2.2	S-N curve numerical values of welded tubular joints in different environmental conditions (source: DNV-RP-C203, 2016).	11
3.1	Description of the environmental data.	12
3.2	Marginal distributions of the system random variables.	14
4.1	Kernels considered in the following numerical experiments.	21
4.2	Analytical test cases	25
5.1	Joint distribution of the input parameters for the 23-bar truss example in Section 5.3.1	33
6.1	Nominal reliability analysis (size $N = 5 \times 10^4$ for IS and SS. $p_0 = 0.1$ for SS).	40
7.1	Nominal reliability analysis (size $N = 5 \times 10^4$ for IS and SS. $p_0 = 0.1$ for SS).	46

1 Introduction

1.1 Industrial context and motivations

Offshore wind turbine (OWT) new technologies tend to reach for more difficult and uncertain environmental conditions. This industry needs probabilistic tools to manage risks associated with OWT operation and maintenance. To do so, industrial and academic players put much effort in order to develop simulators such as costly multi-physics computational tools (such as, e.g., the DIEGO tool developed by EDF R&D, see [Milano et al., 2019](#)). Such computational models often require to be deployed on high-performance computing (HPC) facility in order to be used intensively in industrial studies.

In order to simulate realistic OWT cases, several sources of uncertainties have to be taken into account. If it seems obvious, in the OWT field, that environmental conditions (also known as “metocean” conditions) play a major role as an uncertainty source, one must also admit that other variables are tainted with uncertainty, either arising from a kind of variability (e.g., the soil mechanical properties) or describing the lack of knowledge about some system parameters or, finally, trying to encompass some modeling approximations.

To properly take all these uncertainties into account and to provide efficient predictive tools for a more robust decision-making, a global methodology has been developed from the last 30 years in the Uncertainty Quantification (UQ) community ([De Rocquigny et al., 2008](#); [Smith, 2014](#); [Sullivan, 2015](#)). This methodology is briefly recalled in the next subsection.

1.2 Uncertainty quantification framework

The field of UQ in computer simulation emerged mostly from, or at least, in parallel with, the field of the design and analysis of computer experiments ([Santner et al., 2003](#); [Fang et al., 2006](#)). The UQ methodology is made of several steps which enable to identify and model the various sources of uncertainties, then to propagate them through the computer model, and finally to estimate a quantity of interest (QoI). In addition to this “forward” UQ, a “backward” analysis can be realized in order to perform calibration or sensitivity analysis. In the following paragraph, the overall methodology is briefly presented. However, the reader is invited to refer to the textbooks by [De Rocquigny et al. \(2008\)](#); [Smith \(2014\)](#); [Sullivan \(2015\)](#) or to the recent tutorial by [Blanchard et al. \(2023\)](#) for a concise and updated overview.

Figure 1.1 illustrates the main steps of the generic uncertainty quantification method, which are briefly summarized hereafter:

- **Step A – Problem specification:** This step involves identifying the system under study and constructing a numerical model capable of accurately simulating its behavior. Specifying the problem also involves defining a set of parameters inherent to the numerical model. These parameters include both the input variables and the output variables produced by the simulation. In this document, the numerical model is considered a black box, as opposed to approaches that are integrated into the numerical solution schemes for the system’s behavioral equations (referred to as “intrusive approaches” [Le Maître and Knio, 2010](#)). In general, numerical models are first calibrated against measured data and undergo a process of validation and verification to reduce modeling errors (see, e.g., [Oberkampf and Roy, 2010](#); [Damblin, 2015](#); [Carmassi, 2018](#)).
- **Step B – Uncertainty modeling:** The goal of this step is to identify and model all sources of uncertainty associated with the input variables. In most cases, uncertainty is modeled in the probabilistic framework ([Sullivan, 2015](#)), but other approaches could be considered ([Ajenjo, 2023](#)).
- **Step C – Uncertainty propagation:** This step propagates the uncertain inputs through the computer model. As a result, the output of the numerical model (usually scalar) also becomes uncertain. The goal is to estimate a quantity of interest, which is a statistic related to the random output variable under study. The uncertainty propagation method may differ depending on the quantity of interest targeted (e.g., a moment, a quantile, a rare event probability) ([Morio and Balesdent, 2015](#)).
- **Step C’ – Inverse analysis:** In this step, either calibration or sensitivity analysis can be performed to study the role allocated to each uncertain input leading to the uncertain output.
- **Metamodeling:** Given the high computational cost associated with some simulations, statistical techniques can be used at the step A to emulate these expensive simulators with a limited number of simulations. UQ can then be performed using a “surrogate model” (or “metamodel”) to reduce the computational cost. This

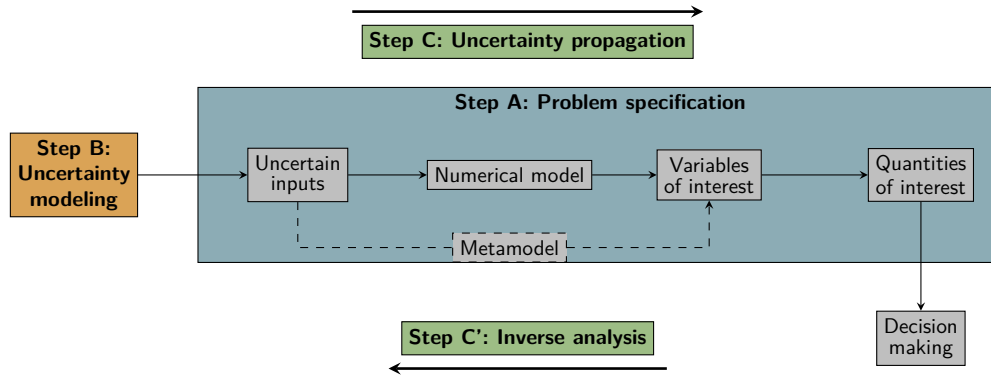


Figure 1.1: General uncertainty quantification framework (De Rocquigny et al., 2008, adapted by Ajenjo, 2023).

step is often essential in practice to perform UQ.

All in all, this step-by-step methodology serves only as a set of guidelines for the analyst to better identify what is needed to achieve a complete UQ analysis. For a particular use case, it is obvious that no generic approach can be used as a “plug-and-play” approach. Usually, a lot of effort has to be put into tuning existing methods and adapting techniques with respect to the intrinsic difficulties imposed by the case. As will be shown in the rest of this report, the overall goal is to apply this methodology (at least a part of it) to a challenging offshore wind turbine case for which performing UQ should help to assess the reliability and robustness of a given design.

1.3 Deliverable objectives and structure of the report

The present report is the SMART Deliverable N° 4.3 of the HIPERWIND¹ Project, and in particular the WP4, whose main objectives are to develop and validate mathematical models for structural reliability. Therefore, the proposed developments should support probabilistic design in practice in line with the Specific Objective #3 (SO3) of the project² while maintaining consistency with the design load case approach currently used in the international design standards (IEC-61400-1, 2019). In particular, this work is dedicated to the treatment of the Task 4.3 entitled “Reliability for fatigue limit states” and is thus limited to the consideration of fatigue damage. Note that most of the contributions were adapted from the PhD thesis of Fekhari (2024).

To sum up, the scientific objectives tackled in this report as listed below:

Objectives

- ▷ **(O1)** – Provide an updated overview of the main state-of-the-art techniques in the field of UQ for OWT, especially in the areas of damage calculation and fatigue reliability.
- ▷ **(O2)** – Identify the limitations of the methods available in the literature and to propose new innovative techniques for OWT mean damage calculation and reliability assessment. These methods will be challenged and, if possible, validated through the realistic use case.
- ▷ **(O3)** – Propose methods to evaluate the robustness of reliability estimates with respect to different uncertainty modeling hypotheses, as well as scale effects that occur when moving from a single turbine to a wind farm (e.g., considering the wake effect).

These three objectives are clearly challenging and necessitate, in order to be solved, relying on a large panel of methods, implying their own formalism, their own hypotheses as well as their own limitations. For these reasons, this report not only provides reminders of theoretical elements when necessary, but also illustrates the proposed methods on toy cases for the sake of clarity and pedagogy. Finally, numerical results are always discussed and synthesis subsections are provided to help the reader to get the most useful information at the end of each section.

The structure of this report is as follows. After this introductory section, Section 2 is devoted to the presentation of the main industrial use case studied in this report: the Teesside case and its numerical modeling and

¹“Highly advanced Probabilistic design and Enhanced Reliability methods for high-value, cost-efficient offshore WIND”

²SO3: Develop an efficient reliability calculation framework, that treats uncertainties in a consistent manner and makes a full probabilistic design approach feasible (WP4).

simulation. Then, Section 3 deals with the modeling and quantification of input uncertainties (both the environmental and the system uncertain inputs), and states the problem of probabilistic fatigue assessment, i.e., damage calculation. Sections 4 and 5 follow with contributions on mean damage estimation: first with a kernel-based approach, then with a spectral approach. Section 6 goes a step further by addressing the problems of reliability assessment, i.e., computing a failure probability in the context of a fatigue limit-state, as well as providing tools to quantify the robustness of such an estimate with respect to the input probabilistic modeling of both resistance and stress variables. As for Section 7, it proposes a very similar study (reliability and robustness assessment), but including the wake effect and the wind farm scale effect. Finally, Section 8 provides a general conclusion and draws some perspectives for future research.

2 Teesside multi-physics wind turbine model

2.1 Introduction

Wind energy is a highly competitive industry with increasing regulations regarding its impact on ecosystems, land and sea use, landscapes, or air traffic management (Beauregard et al., 2022). During the long process of winning calls for tenders, obtaining construction permits, or through wind farm exploitation, an advanced technical understanding of such systems might offer a competitive advantage.

The operation of offshore wind turbines (OWTs) is driven by multiple physics coupled. This behavior results from different external effects which are highly turbulent and uncertain. Among them, the *metocean* (abbreviation of “meteorology” and “oceanography”) environmental conditions play a primary role. Yet, many other types of external factors affect the exploitation of offshore wind turbines such as the corrosion of the structure, global scour, marine growth, stress concentration factor induced by the manufacturing process, etc.

In this context, numerical models have been developed to certify the structural integrity of OWTs w.r.t. external stresses. A wind farm project planned at a given location should pass different validation procedures (regarding fatigue, ultimate loads, buckling, serviceability, etc.) established by international standards such as the International Electrotechnical Commission (IEC-61400-1, 2019). As wind turbine structures face a large number of stress cycles in their lifetime (up to 10^8 for 25 years of operation), this section will particularly focus on modelling fatigue damage assessment. The present work applies different steps of the UQ methodology on an industrial use case using elements from the PhD thesis of Fekhari (2024).

2.2 Numerical simulation model

An OWT is a complex system interacting with its environment. However, multi-physics numerical models allow to simulate its response against a set of stresses. The use case considered here is a model of a Siemens SWT 2.3MW bottom-fixed turbine on a monopile foundation (see the datasheet in Table 2.1), operating since 2014 in Teesside, UK (see the wind farm layout and wind turbine diagram in Fig. 2.2). This simulation model consists of a chain of three numerical codes executed sequentially. As illustrated in Fig. 2.1, a simulation over a time period is the sequence of, first, a turbulent wind speed field generation, then an OWT simulation (computing various outputs including mechanical stress), and finally, a post-processing phase to assess the fatigue damage of the structure.




Figure 2.1: Diagram of the chained OWT simulation model.

The first block of the chain consists of a turbulent wind field simulator called “TurbSim” (developed by Jonkman, 2009 from the National Renewable Energy Laboratory, USA) that uses, as a turbulence model, called the Kaimal spectrum (Kaimal et al., 1972), as recommended by the IEC-61400-1 (2019) standards. Moreover, to extrapolate the wind speed vertically, the shear is modeled by a power law. Since the wind field generation shows inherent stochasticity, each 10-minute long simulation is repeated with different pseudorandom seeds and one averages the estimated damage over these repetitions. This question has been widely studied by some authors,

Table 2.1: Teesside OWT datasheet.

Siemens SWT-2.3-93	
Rated power	2.3 MW
Rotor diameter	93 m
Hub height	83 m
Cut-in, cut-out wind speed	4 m/s, 25 m/s

(e.g., Slot et al., 2020), who concluded that the six repetitions recommended by the IEC-61400-1 (2019) may be insufficient to properly average the effect of stochasticity. Thus, in the following, the simulations are repeated eleven times (picking an odd number also directly provides the median value over the repetitions). This number of repetitions was chosen to suit the maximum number of simulations and the storage capacity of the generated simulations.

In the second block, “DIEGO” software (for “Dynamique Intégrée des Éoliennes et Génératrices Offshore”³) is a numerical model developed by EDF R&D to simulate the aero-hydro-servo-elastic behavior of OWTs in the time domain. Different extensive code-to-code comparisons between DIEGO and other aero-hydro-servo-elastic models showed comparable results. For instance considering bottom-fixed OWT (Popko et al., 2021) or floating OWT (Robertson et al., 2020; Kim et al., 2022), DIEGO was compared to “OpenFAST” (developed by NREL), “HAWC2” (developed by DTU), “BLADED” (developed by DNV), and “DeepLines Wind” (developed by Principia).

This code relies on a flexible multibody approach to model the nonlinear structural dynamics and uses Morison’s equation together with a first-order potential solution to simulate the hydrodynamics (Milano, 2021). It takes the turbulent wind speed field generated by TurbSim as input and computes the dynamical behavior of the system (including the multiaxial mechanical stress at different nodes of the structure). For the application of interest, the control system is modeled by the open source DTU controller (Hansen and Henriksen, 2013). As for the waves, they are modeled by the parametric wave spectrum is named “JONSWAP”, for the “JOint North Sea WAve Project” (Hasselmann et al., 1973). Further details regarding wind turbine model and the software used are provided in Capaldo et al. (2021). To avoid numerical perturbations and reach the stability of the dynamical system, our simulation period is extended to 1000 seconds and the first 400 seconds are cropped in the post-processing step. This chained OWT numerical simulation model has been deployed on an EDF R&D HPC facility to benefit from parallel computing speed up (a single simulation on one CPU takes around 45 minutes).

2.3 Fatigue assessment

As described in Fig. 2.1, a typical DIEGO simulation returns a 10-minute multiaxial stress time series at each node $i \in \mathbb{N}$ of the one-dimensional meshed structure. Since classical fatigue laws are established for uniaxial stresses, the first step is to compute one equivalent Von Mises stress time series at each structural node.

The foundation and the tower of an OWT are a succession of tubes with various sections connected by bolted or welded joints. The proposed work focuses on the welded joints at the mudline level, identified as a critical area for the structure. This hypothesis is confirmed in the literature by different contributions, see e.g., the results of Müller and Cheng (2018, Fig. 12), or Katsikogiannis et al. (2021). At the top of the turbine, the fatigue is commonly studied at the blade root, which was not studied here since the blades in composite material have different properties (see e.g., Dimitrov, 2013). Note that the OWT simulations provide outputs allowing us to similarly study any node along the structure (without any additional computational effort).

To compute fatigue in a welded joint, the external circle of the welding ring is discretized for a few azimuth angles $\theta \in \mathbb{R}_+$ (see the red points in the monopile cross-section on the right in Fig. 2.3). The equivalent Von Mises stress time series is then reported on the external welding ring for an azimuth θ . According to Li and Zhang (2020) and our own experience, the most critical azimuth angles are roughly aligned with the main wind and wave directions (whose distributions are illustrated in Fig. 2.3). Looking at these illustrations, the wind and wave conditions have a very dominant orientation, which is explained by the closeness of the wind farm to the shore. Then, it is assumed that azimuth angles in these directions will be more stressed, leading to higher fatigue

³In English, “Integrated Dynamics of Wind Turbines and Offshore Generators”.

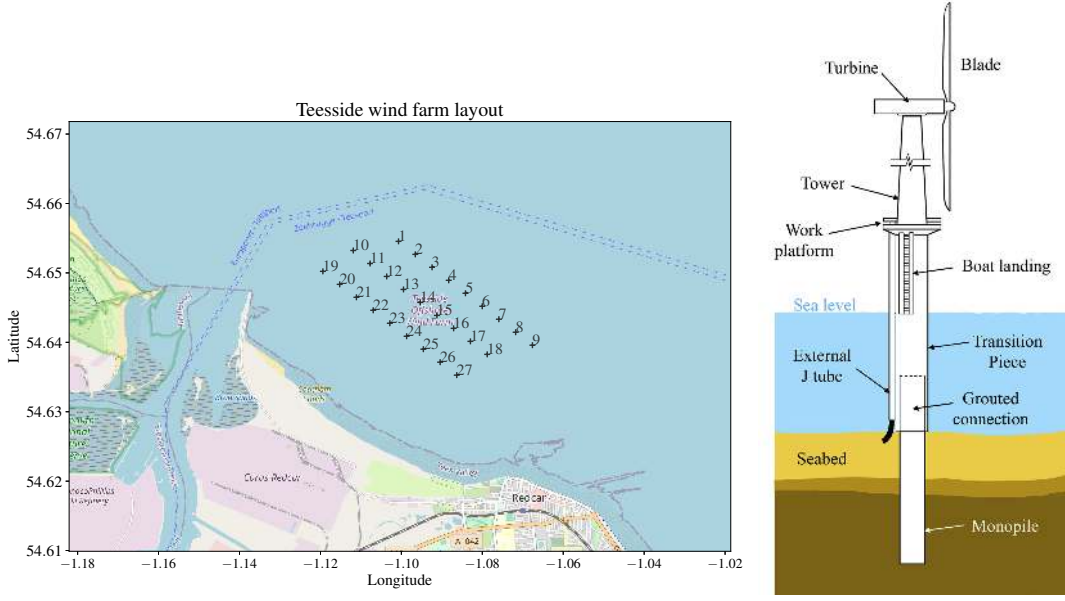


Figure 2.2: Teesside wind farm layout (left). Monopile OWT diagram from (Chen et al., 2018) (right).

damage.

To assess fatigue damage, rainflow counting (Dowling, 1972) first identifies the stress cycles and their respective amplitudes (using the implementation of the ASTM E1049-85 rainflow cycle counting algorithm from the Python package `rainflow`⁴). Then, for each identified stress cycle of amplitude, $s \in \mathbb{R}_+$, the “Stress vs. Number of cycles” curve (also called the “S-N curve” or “Wöhler curve”) allows one to estimate the number N_c of similar stress cycles necessary to reach fatigue crack initiation. A well-admitted simplification of the S-N curve is to consider it as log-linear on two segments with slopes $(-m_1, -m_2) \in \mathbb{R}_+^2$ and intercepts $(\log(a_1), \log(a_2)) \in \mathbb{R}_+^2$:

$$\log(N_c(s)) = \begin{cases} \log(a_1) - m_1 \log(s), & \text{for } s \in [s_{\min}, s_e] \\ \log(a_2) - m_2 \log(s), & \text{for } s \in [s_e, s_{\max}], \end{cases} \quad (2.1)$$

where N_c is the predicted number of cycles to failure for stress range s , s_{\min} is the minimal (resp. maximal) stress range identified by the rainflow counting, and s_e is the stress range axis of the intersection of the two log-lines formed by the S-N curve.

The expression of this curve in two linear segments arises from the concept of endurance limit of a material, denoted by s_e , under which the effect of fatigue on a material should be considerably smaller. According to DNV-RP-C203 (2016), the S-N curve is altered for welded tubular joints by taking into account the tube’s thickness:

$$N_c(s) = \begin{cases} a_1 \left(s \left(\frac{t}{t_{\text{ref}}} \right)^h \right)^{-m_1}, & \text{for } s \in [s_{\min}, s_e] \\ a_2 \left(s \left(\frac{t}{t_{\text{ref}}} \right)^h \right)^{-m_2}, & \text{for } s \in [s_e, s_{\max}], \end{cases} \quad (2.2)$$

with t_{ref} the reference thickness (for tubular welded joints $t_{\text{ref}} = 25$ mm), t the plate thickness, and h the thickness exponent. The numerical values considered in the present work derive from the standards DNV-RP-C203 (2016, Sec. 2.4.6), and are reproduced in Table 2.2.

Most S-N curves are built over zero mean stress cycles, however, other empirical models such as the “Goodman line” or the “Gerber parabola” were developed to consider various stress mean (Suresh, 1998). In these cases, the S-N curve becomes a three-dimensional envelope depending on the number of cycles, the stress amplitude, and the mean stress. In this work, the nonzero correction is not considered as the values of mean stress were found to be negligible compared to the yield stress of the steel material studied.

⁴<https://github.com/iamlikeme/rainflow>

Table 2.2: S-N curve numerical values of welded tubular joints in different environmental conditions (source: DNV-RP-C203, 2016).

Environment	m_1	$\log(a_1)$	m_2	$\log(a_2)$	h
Air	3.0	12.48	5.0	16.13	0.25
Seawater with cathodic protection	3.0	12.18	5.0	16.13	0.25
Seawater free corrosion	3.0	12.03	3.0	12.03	0.25

Finally, a popular approach to compute the damage is to linearly aggregate the fatigue contribution of each stress cycle using Palmgren-Miner's rule. This rule defines a *cumulative damage* d_c by summing the fatigue contributions of each stress cycle $s^{(j)}$, $j \in \{1, \dots, k\}$, regardless of their order of appearance:

$$d_c = \sum_{j=1}^k \frac{1}{N_c(s^{(j)})}. \quad (2.3)$$

In this theory, the material reaches ultimate fatigue ruin when $d_c \geq 1$. A common practice when using Palmgren-Miner's rule is to gather the stress cycles in a set of bins. This practice induces an integration error, which becomes significant as the number of bins is reduced. In the following, the cumulative damage is computed without binning, as defined in Eq. (2.3).

As alternative to Palmgren-Miner's rule, spectral methods were also introduced to quantify fatigue damage in the late 1980s. The main idea is to infer a probability density function (PDF) over the amplitudes of the stress cycles identified by rainflow counting, typically using a mixture of parametric distributions. From this PDF, one can derive the fatigue endurance and the cumulated damage (see further details in the review of Dirlik and Benasciutti, 2021). Spectral approaches showed to be sufficient at early design phases but not ideal for some tasks (e.g., for blades' fatigue Ragan and Manuel, 2007). Finally, nonlinear fatigue models were developed in the 1990s (Fatemi and Yang, 1998) to take into account the order in which the loading cycles were applied to the structure. For OWT applications, the recent work of Rocher et al. (2020) studied a probabilistic version of nonlinear fatigue models. Unfortunately, this refined approach requires a larger computational effort and a calibration step over experimental tests of various parameters (Freysinet et al., 2023).




Figure 2.3: Angular distribution of the wind and waves with a horizontal cross-section of the OWT structure and the mudline. Red crosses represent the discretized azimuths for which the fatigue is computed.

2.4 Synthesis

This section introduced the chain of numerical models used to assess the fatigue damage cumulated in the wind turbine tower over time. Overall, one evaluation of this chain is the sequence of a turbulent wind generation, a multi-physics wind turbine simulation in the time domain, and a post-processing step to assess the fatigue damage. In the following, the uncertainties associated to the input variables should be defined.

Variable	Notation	Unit	Description
Mean wind speed	U	m.s^{-1}	10-min. average horizontal wind speed
Wind turbulence	σ_U	m.s^{-1}	10-min. wind speed standard deviation
Wind direction ⁵	θ_{wind}	deg.	10-min. average wind direction
Significant wave height	H_s	m	Significant wave height
Peak wave period	T_p	s	Peak 1-hour spectral wave period
Wave direction	θ_{wave}	deg.	10-min. average wave direction

Table 3.1: Description of the environmental data.

3 Quantification and modeling of input uncertainties

3.1 Introduction

According to the general UQ framework described in Fig. 1.1, this section focuses on the definition of the uncertain inputs and their corresponding probabilistic model (step B). In the following, the input variables considered as uncertain are split into two groups (assumed independent), respectively called *environmental variables* and *system variables*. First, the environmental variables are a collection of variables characterizing the long-term metocean conditions near a wind farm. This source of uncertainty is empirically well-defined thanks to wind measurement campaigns, however, it presents a complex dependence structure. The second group of uncertain inputs is related to the wind turbine system. A wide range of uncertainties can be taken into account in such systems, such as material properties, manufacturing quality, soil conditions, control error, corrosion, marine growth, aerodynamic damping, etc. Among them, a restricted list of four independent variables is considered according to sensitivity analysis results from the literature and expert knowledge (Dimitrov et al., 2018; Teixeira et al., 2019b; Velarde et al., 2019; Petrovska, 2022).

3.2 Definition of the environmental random inputs

In the lifespan of a wind farm project, environmental data is collected at different phases. In order to decide on the construction of a wind farm, meteorological masts with wind speed cup anemometers, and wave buoys are usually installed on a potential site for a few years. After its construction, each wind turbine is equipped with monitoring instruments (e.g., cup anemometers). In total, five years of wind data have been collected on the turbines which are not affected by the wake on this site. Their acquisition system (usually called “SCADA” for “Supervisory Control And Data Acquisition”) has a sampling period of ten minutes. The wave data arise from a buoy placed in the middle of the farm. These data describe the physical features listed in Table 3.1. A limitation of the present study is that its controller-induced uncertainty (like wind misalignment) is not considered.

The Teesside farm is located close to the coast, making the environmental conditions very different depending on the direction (see the wind farm layout in Fig. 2.2). Since measures are also subject to uncertainties, a few checks were made to ensure that the data were physically consistent. Truncation bounds were applied since this study is focused on central tendency estimation (i.e., mean behavior) rather than extreme values. In practice, this truncation only removes extreme data points (associated with storm events). In addition, a simple trigonometric transform is applied to each directional feature to take into account their cyclic structure. Finally, the remaining features are rescaled (i.e., using a min-max normalization).

Teesside’s environmental data is illustrated by its copulogram in Fig. 3.1, a graphical tool to visualize multivariate data. The copulogram exhibits the marginals with univariate KDE plots (in the diagonal), and the dependence structure with scatter plots in the normalized rank space (in the upper triangle). The scatter plots of normalized ranks are actually a representation of the empirical copula density. In the lower triangular matrix, the scatter plots are set in the physical space, merging the effects of the marginals and the dependencies (as in the usual visualization offered by the matrix plot). All together, the copulogram gives a synthetic decomposition of the dataset and was implemented in a new open source Python package named `copulogram`⁶.

⁵Note that the two directional variables could be replaced by a wind-wave misalignment variable for a bottom-fixed technology, however, our framework can be directly transposed to floating models.

⁶GitHub repository: <https://github.com/efekhari27/copulogram>

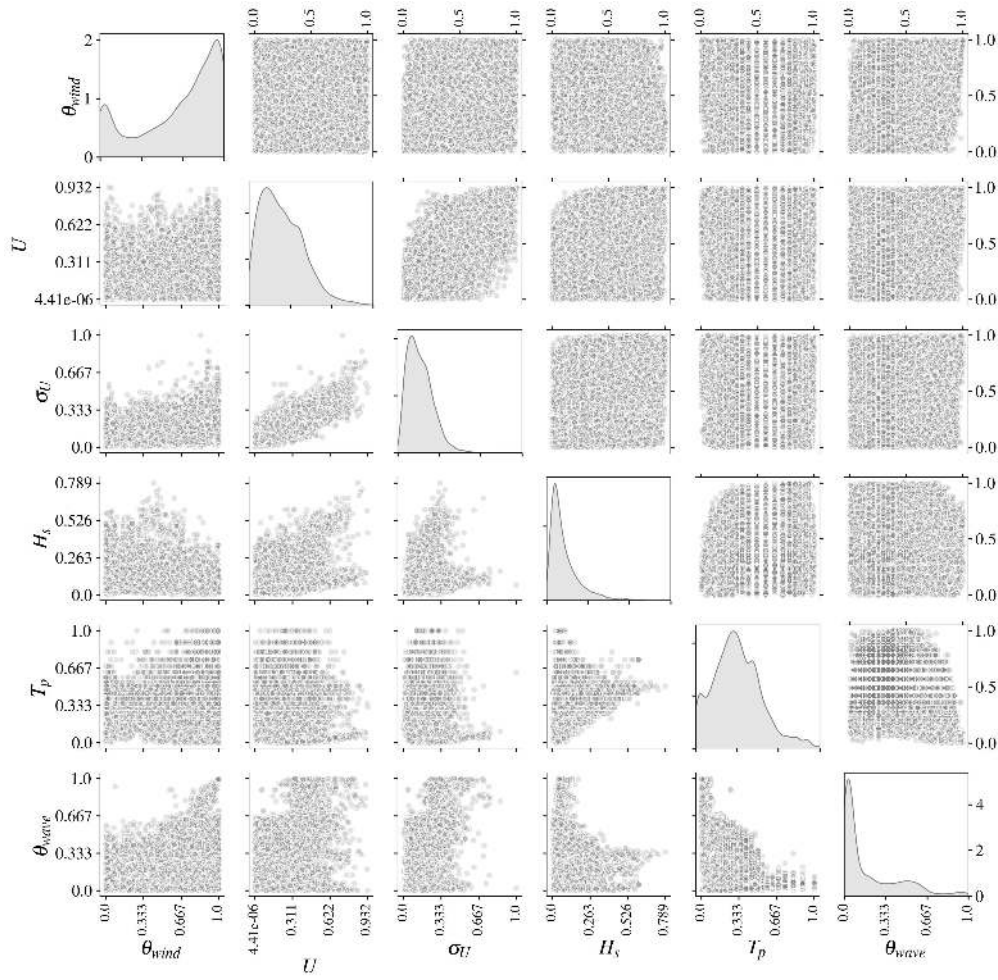


Figure 3.1: Copulogram of the Teesside measured data ($N = 10^4$ in gray). Marginals are represented by univariate KDE plots (diagonal), the dependence structure with scatter plots in the rank space (upper triangle). Scatter plots on the bottom triangle are set in the physical space.

The copulogram in Fig. 3.1, shows the dependence structure underlying the environmental input data studied. In general, building an probabilistic model to accurately represent the environmental random vector $\mathbf{X} = (U, \sigma_U, \theta_{\text{wind}}, H_s, T_p, \theta_{\text{wave}})^T \in \mathbb{R}^6$ can be challenging. Various approaches were proposed for this purpose, either using conditional distributions (e.g., [Bitner-Gregersen, 2015](#); [Vanem et al., 2024](#)), or vine copulas (e.g., [Vanem, 2016](#); [Montes-Iturrizaga and Heredia-Zavoni, 2016](#); [Lin and Dong, 2019](#)). Nonparametric methods like KDE were also applied in this context (e.g., [Han et al., 2018](#)). Nonparametric techniques generally struggle to model the distributions' tails, even if the tails are essential to qualify structures for ultimate events. However, they are highly flexible and often easier to implement than parametric methods. A hybrid inference strategy was also proposed in [Fekhari \(2024, Chapter 3\)](#), composing some well-known parametric models for the marginals (e.g., Weibull distribution for the wind speed), with a flexible dependence modeling by the empirical Bernstein copula ([Sancetta and Satchell, 2004](#)).

3.3 Definition of the system random inputs

As mentioned earlier, multiple parameters in a wind turbine system can be considered as uncertain. Our study focuses on the effects of uncertainties on fatigue damage over the structure. Therefore, the literature review of the sensitivity analysis on offshore wind turbine fatigue helped us narrow down a few system variables. [Petrovska \(2022\)](#) explored the sensitivity analysis of many variables on the fatigue of a wind turbine in Teesside. Even if the use of the Morris method is questionable (see e.g., [Da Veiga et al., 2021](#)), the results allowed us to screen out some variables. For example, the uncertainties related to the corrosion, the wind shear exponent, or the nacelle

mass showed a limited impact on the fatigue. By crossing the conclusions of various research with the expert knowledge among partners from the HIPERWIND European project, the system variables considered uncertain in the following are described in Table 3.2 and gathered by the random vector $\mathbf{Z} = (K_{\text{soil}}, \Theta_{\text{yaw}}, \varepsilon)^\top \in \mathbb{R}^3$. Each of them is assumed independent, with a marginal probabilistic model arising from the literature (see e.g., JCSS, 2011; Drexler and Muskulus, 2021). Note that the critical damage, above which fatigue failure should occur is also assumed random with the random variable $D_{\text{cr}} \in \mathbb{R}$ (also called “Miner’s damage sum” at failure in JCSS, 2011).

Name	Notation	Marginal model	Remark
Soil stiffness coeff.	K_{soil}	$\mathcal{N}(\mu = 1, \sigma^2 = 0.2^2)$	Truncation: $[0.5, 2]$
Yaw misalignment	Θ_{yaw}	$\mathcal{N}(\mu = 0, \sigma^2 = 5^2)$	Truncation: $[-10, 10]$
S-N curve coeff.	ε	$\mathcal{LN}(\mu_\varepsilon = 1, \sigma_\varepsilon^2 = 0.2^2)$	–
Critical damage	D_{cr}	$\mathcal{LN}(\mu_{D_{\text{cr}}} = 1, \sigma_{D_{\text{cr}}}^2 = 0.3^2)$	–

Table 3.2: Marginal distributions of the system random variables.

3.4 Probabilistic fatigue assessment

The definition of a fatigue endurance model has a main impact on fatigue damage assessment. However, the S-N curves usually describing the endurance of a material are built on repeated laboratory experiments. Even if the need for random S-N curves has long been expressed in the field of fatigue experiments (Lieudarde, 1982), their probabilistic description was better formalized in (Guédé et al., 2007; Sudret, 2013).

The models proposed in Guédé et al. (2007) are based on the experimental procedure used to build the S-N curves. For identical steel specimens, a cyclic loading with fixed amplitude is repeated until fatigue crack initiation. Because of variations in the material’s microstructure, the fatigue endurance for the same cyclic solicitation is random. This variation is commonly assumed to follow a lognormal distribution in the literature. A probabilistic model of the S-N curve naturally comes:

$$\ln(N_c(s, \varepsilon)) = \ln(a) - m \ln(s) + \ln(\varepsilon) \quad (3.1a)$$

$$\Rightarrow N_c(s, \varepsilon) = a s^{-m} \varepsilon, \quad (3.1b)$$

where $\varepsilon \sim \mathcal{LN}(\mu_\varepsilon = 1, \sigma_\varepsilon^2 = 0.2^2)$ is assumed when no measurement is available (according to DNV-RP-C203, 2016, Appendix F.5). This uncertainty can be injected into the Miner-Palmgren rule defined in Eq. (2.3), leading to:

$$d_c(\varepsilon) = \sum_{j=1}^k \frac{1}{N_c(s^{(j)}, \varepsilon)} = \sum_{j=1}^k \frac{1}{a (s^{(j)})^{-m} \varepsilon} = \frac{1}{\varepsilon} \sum_{j=1}^k \frac{1}{a (s^{(j)})^{-m}}, \quad (3.2)$$

which can be assessed as a pure post-processing of fatigue damage results computed with a single deterministic S-N curve.

In wind energy standards, S-N curves for design are actually a conservative envelope of the measured fatigue endurance. DNV-RP-C203 (2016, Appendix F.7) describes how to define a design S-N curve from fatigue measures. Assuming that the fatigue endurance follows a Gaussian distribution on a logarithmic scale, such that $\ln(\varepsilon) \sim \mathcal{N}(\mu_{\ln(\varepsilon)}, \sigma_{\ln(\varepsilon)}^2)$, the design S-N curve $N_c^{\text{design}}(s)$ is the curve at two standard deviations $\sigma_{\ln(\varepsilon)}$ below the median curve. Using the design S-N curve given in DNV-RP-C203 (2016, Sec. 2.4.6) and the normality assumption, one can reconstruct the median S-N curve by taking:

$$q_{50\%}[\ln(N_c(s, \varepsilon))] = \ln(N_c^{\text{design}}(s)) + 2 \sigma_{\ln(\varepsilon)}. \quad (3.3)$$

Fig. 3.2 illustrates the design curve defined by DNV for tubular joints DNV-RP-C203 (2016, Sec. 2.4.6) and the reconstructed probabilistic model according to the previous assumptions.

3.5 Synthesis

This section defines two sources of uncertainties considered in this problem: the environmental variables and the system ones, following the typology introduced in the PhD thesis of Fekhari (2024). Among the system

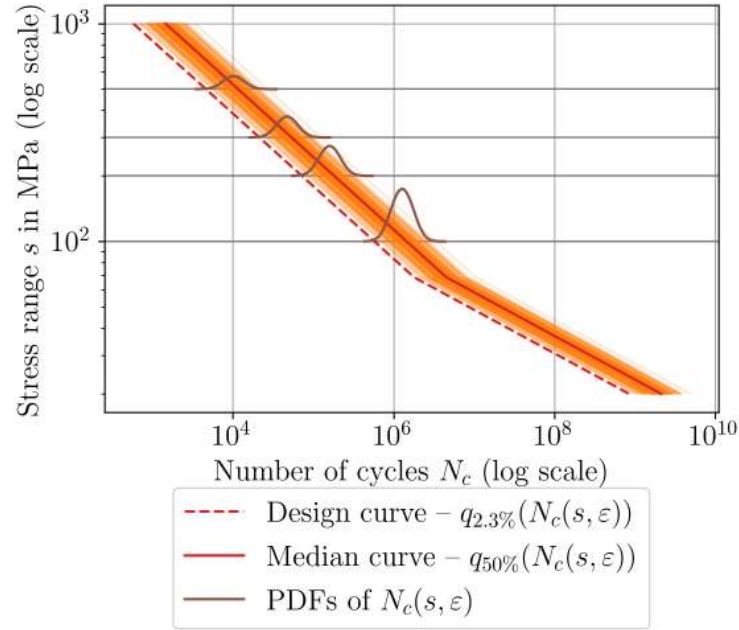


Figure 3.2: Illustration of a probabilistic S-N curve according to the model defined in Gu  d   et al. (2007).

uncertainties, a focus is done on the construction of a probabilistic S-N curve according to the standards. In the following, the two sources of uncertainties will differently impact the output quantities of interest studied.

4 A first uncertainty propagation approach for cumulated fatigue damage estimation

4.1 Introduction

In the wind energy community, the definition of fatigue loads (or load cases) considering metocean uncertainties is an important topic. Following the uncertainty quantification framework (see Fig. 1.1), it is the propagation of the metocean uncertainties introduced in Section 3 through the simulation model described in Section 2. To validate the structural integrity of an OWT design regarding the fatigue limit states, international standards such as IEC-61400-1 (2019) and DNV-RP-C203 (2016) recommend to evaluate the OWT simulation models over a representative set of metocean configurations. In this case, the quantity of interest studied is the lifetime damage at some specific hot spots of the structure. In fact, one can demonstrate that assessing the cumulative damage over the lifetime is equivalent to estimating the mean of a cumulative damage over 10 minute w.r.t. the metocean uncertainties. However, international standards suggest using regular grid to explore the domain induced by the environmental variables, which requires over 2×10^5 simulations of the costly OWT simulation model (see Graf et al., 2016). The goal of this work is to consider the uncertainties related to site-specific environmental conditions (with the random vector $\mathbf{X} \in \mathcal{D}_{\mathbf{X}} \subset \mathbb{R}^d$), and efficiently propagate them through an OWT simulation model (denoted by $d_c : \mathcal{D}_{\mathbf{X}} \rightarrow \mathbb{R}$) to estimate a quantity of interest $\psi(Y) = \psi(d_c(\mathbf{X}))$ on the cumulated fatigue damage.

Regarding the probabilistic modeling of the metocean conditions, two cases may occur: either measured data are directly available (i.e., the “given-data” context) or a theoretical parametric form for the joint input probability distribution can be postulated. Such existing parametric joint distributions often rely on prior data fitting combined with expert knowledge. When a considerable amount of environmental data is available, nonparametric approaches such as the empirical Bernstein copula have been studied in Fekhari (2024, Chapter 3) to capture complex dependence structures. Alternatively, an idea is to directly use the data as an empirical representation of input uncertainties in order to avoid an additional inference error.

In the UQ literature for wind energy, authors have proposed numerous propagation approaches to estimate the quantity of interest. Depending on the nature of $\psi(Y)$, one often distinguishes between two types of uncertainty propagation: a central tendency estimation (e.g., focusing on the output mean value or the variance) and a

tail estimation (e.g., focusing on a high-order quantile or a failure probability). When uncertainty propagation aims at central tendency estimation, the usual methods can be split into two groups. First, those relying on sampling, i.e., mainly Monte Carlo sampling (Graf et al., 2016), quasi-Monte Carlo sampling (Müller and Cheng, 2018), geometrical subsampling (Kanner et al., 2018), or deterministic quadrature rules (Van den Bos, 2020). All these methods estimate the quantity directly on the numerical simulator's outputs. Second, those that rely on the use of surrogate models to emulate the costly numerical model by a statistical model. Among a large panel of surrogates, one can mention, regarding wind energy applications, the use of polynomial chaos expansions (Dimitrov et al., 2018; Murcia et al., 2018), GP regression (Huchet, 2018; Teixeira et al., 2019a; Slot et al., 2020; Wilkie and Galasso, 2021), or artificial neural networks (Bai et al., 2023). When uncertainty propagation aims at studying the tail of the output distribution such as in risk or reliability assessment, one usually desires to estimate a quantile or a failure probability. In the wind energy literature, failure probability estimation has been studied, e.g., in time-independent reliability assessment (Zwick and Muskulus, 2015; Slot et al., 2020; Wilkie and Galasso, 2021) or regarding time-dependent problems (Abdallah et al., 2019; Lataniotis, 2019).

In uncertainty quantification, modelers and analysts often need to determine whether inputs are influential or not in order to prioritize their effort (in terms of experimental data collecting, simulation budget, or expert elicitation). Sometimes, they want to get a better understanding of the OWT numerical models' behavior or to enhance the input uncertainty modeling. All these questions are intimately related to the topic of sensitivity analysis (Saltelli et al., 2008; Da Veiga et al., 2021). In the wind energy literature, one can mention, among others, some works related to Spearman's rank correlation analysis and the use of the Morris method in Velarde et al. (2019) and Petrovska (2022). Going to variance-based analysis, the direct calculation of Sobol' indices after fitting a polynomial chaos surrogate model has been applied in many works (e.g., in Murcia et al., 2018) while the use of distributional indices (e.g., based on the Kullback–Leibler divergence) has been investigated by Teixeira et al. (2019b).

The present section is based on Fekhari (2024, Chapter 4) to focus on the problem of uncertainty propagation, and more specifically, on the mean fatigue damage estimation (i.e., $\psi(Y) = \mathbb{E}[d_c(\mathbf{X})]$). Such a problem is usually encountered, by engineers, during the design phase. Most of the time, current standards as well as common engineering practices make them use regular grids, as recalled in Graf et al., 2016. Altogether, one can describe three alternative strategies: (i) direct sampling on the numerical model (e.g., using Monte Carlo), (ii) sampling on a static surrogate model (e.g., using GP regression), or (iii) using an “active learning” strategy (i.e., progressively adding evaluations of the numerical model to enhance the surrogate model fitting process). In practice, fitting a surrogate model in the context of OWT fatigue damage can be challenging due to the nonlinearity of the code. Moreover, the surrogate model validation procedure complexifies the process. Finally, active learning strategies restrict the potential number of parallel simulations, which limits the use of HPC facilities. Thus, the proposed work consists of evaluating the advantages and limits of a kernel-based method, called “kernel herding” (Chen et al., 2010), as a tool for given-data, fast, and fully-distributable uncertainty propagation in OWT simulators. Additionally, this sampling method is highly flexible, allowing one to complete an existing design of experiments. Such a property can be crucial in practice when the analyst is asked to include some specific points to the design (e.g., characteristic points describing the system's behavior required by experts or by standards, as explained in Huchet, 2018).

4.1.1 Definition of the quantity of interest

For a realization $\mathbf{x} \in \mathcal{D}_{\mathbf{X}}$ of environmental conditions, at a structural node i , an azimuth angle θ ; $k \in \mathbb{N}$ stress cycles of respective amplitude $\{s_{i,\theta}^{(j)}(\mathbf{x})\}_{j=1}^k$ are identified. Then, Miner's rule defines the damage function $g_{i,\theta}(\mathbf{x}) : \mathcal{D}_{\mathbf{X}} \rightarrow \mathbb{R}_+$ by:

$$d_c^{i,\theta}(\mathbf{x}) = \sum_{j=1}^k \frac{1}{N_c(s_{i,\theta}^{(j)}(\mathbf{x}))}. \quad (4.1)$$

As defined by the DNV standards for OWT fatigue design (DNV-RP-C203, 2016), the quantity of interest in the present section is the “mean damage” $D^{i,\theta}$ (also called “mean cumulative damage”), computed at a node i , for an azimuth angle θ :

$$D^{i,\theta} = \mathbb{E}[d_c^{i,\theta}(\mathbf{X})] = \int_{\mathcal{D}_{\mathbf{X}}} d_c^{i,\theta}(\mathbf{x}) f_{\mathbf{X}}(\mathbf{x}) d\mathbf{x}. \quad (4.2)$$

To get a preview of the distribution of this output random variable $d_c^{i,\theta}(\mathbf{X})$, a histogram of a large Monte Carlo

simulation ($N_{\text{ref}} = 2000$) is represented in Fig. 4.1 (with a log scale). In this case, the log-damage histogram presents a little asymmetry but it is frequently modeled by a normal distribution (see e.g., Teixeira et al., 2019b).




Figure 4.1: Histogram of the log-damage, at mudline, azimuth 45 deg. (Monte Carlo reference sample).

The present section will first introduce various kernel-based methods for central tendency estimation. Then, numerical experiments obtained on both analytical and industrial cases will be analyzed. Finally, it will present some discussions and draw some conclusions.

4.2 A first integration approach based on kernels

Various methods aim at approximating the integral of a function against a probability measure. In the case of OWT mean damage estimation, these methods can be used for defining efficient design load cases. This problem is equivalent to the central tendency estimation of $\mathbf{Y} = d_c(\mathbf{X})$, the image of the environmental random variable \mathbf{X} by the damage function $d_c(\cdot)$. Considering a measurable space $\mathcal{D}_{\mathbf{X}} \subset \mathbb{R}^d, d \in \mathbb{N}^*$, associated with a known probability measure π , this section studies the approximation of integrals of the form $I_{\pi}(g) = \int_{\mathcal{D}_{\mathbf{X}}} g(\mathbf{x})d\pi(\mathbf{x})$, for a measurable function $g : \mathcal{D}_{\mathbf{X}} \rightarrow \mathbb{R}$.

4.2.1 Quadrature rules and quasi-Monte Carlo methods

Numerical integration authors may call this generic problem *probabilistic integration* (Briol et al., 2019). In practice, this quantity is estimated on an n -sized set of damage realizations $\mathbf{y}_n = \{g(\mathbf{x}^{(1)}), \dots, g(\mathbf{x}^{(n)})\}$ of an input sample $\mathbf{X}_n = \{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)}\}$. A weighted arithmetic mean of the realizations $\{g(\mathbf{x}^{(1)}), \dots, g(\mathbf{x}^{(n)})\}$ is called a *quadrature rule* with a set of unconstrained weights $\mathbf{w}_n = \{w_1, \dots, w_n\} \in \mathbb{R}^n$, such that:

$$I_{\pi}(g) = \int_{\mathcal{D}_{\mathbf{X}}} g(\mathbf{x})d\pi(\mathbf{x}) \approx \sum_{i=1}^n w_i g(\mathbf{x}^{(i)}). \quad (4.3)$$

The numerical experiment framework often implies that the function g is costly to evaluate, making the realization number limited. For a given sample size n , the goal is to find a set of tuples of nodes and weights $\{\mathbf{x}^{(i)}, w_i\}_{i=1}^n$ (i.e., a quadrature rule), giving the best approximation of the studied quantity. In the literature, a large panel of numerical integration methods has been proposed to tackle this problem. For example, Van den Bos (2020) recently developed a quadrature rule based on arbitrary sample sets which was applied to a similar industrial OWT use case.

Alternatively, sampling methods rely on generating a set of points \mathbf{X}_n drawn from the input distribution to compute the arithmetic mean of their realizations (i.e., uniform weights $\{w_i = \frac{1}{n}\}_{i=1}^n$). Among them, low-discrepancy sequences, also called “quasi-Monte Carlo” sampling (e.g., Sobol’, Halton, Faure sequences Leobacher and Pillichshammer, 2014) are known to improve the standard Monte Carlo convergence rate and will be used as a deterministic reference method in the following numerical experiments.

4.2.2 Kernel discrepancy

Quasi-Monte Carlo sampling methods widely rely on a uniformity metric, called *discrepancy*. This section first presents the link between discrepancy and numerical integration. Then it introduces a kernel-based discrepancy,

generalizing the concept to non-uniform measures. This tool is eventually used to build a sequential quadrature rule by subsampling from a finite dataset.

Quantization of probability measures and quadrature When dealing with probabilistic integration such as Eq. (4.3), a quadrature rule is a finite representation of a continuous measure π by a discrete measure $\zeta_n = \sum_{i=1}^n w_i \delta(\mathbf{x}^{(i)})$ (weighted sum of Dirac distributions at the design points \mathbf{X}_n). In the literature, this procedure is also called *quantization* of a continuous measure π . Overall, numerical integration is a particular case of probabilistic integration against a uniform input measure. For uniform measures, the Koksma-Hlawka inequality (Morokoff and Caflisch, 1995) provides an upper bound on the absolute error of a quadrature rule:

$$\left| \int_{[0,1]^p} g(\mathbf{x}) d\mathbf{x} - \frac{1}{n} \sum_{i=1}^n g(\mathbf{x}^{(i)}) \right| \leq V(g) D_n^*(\mathbf{X}_n). \quad (4.4)$$

As presented in Oates (2021), $V(g) = \sum_{u \subseteq \{1, \dots, p\}} \int_{[0,1]^u} \left| \frac{\partial^u g}{\partial \mathbf{x}_u}(\mathbf{x}_u, 1) \right| d\mathbf{x}_u$, quantifies the complexity of the integrand, while $D_n^*(\mathbf{X}_n)$ evaluates the discrepancy to uniformity of the design \mathbf{X}_n . Therefore, the Koksma-Hlawka inequality shows that the quadrature rule's accuracy relies on the good quantization of π by \mathbf{X}_n . For a uniform measure π , the star discrepancy $D_n^*(\mathbf{X}_n)$ is a metric assessing how far from uniformity a sample \mathbf{X}_n is. When generalizing to a non-uniform measure, a good quantization of π should also lead to a good approximation of the quantity.

Reproducing kernel Hilbert space and kernel mean embedding To generalize the Koksma-Hlawka inequality to any probability measure, let us assume that the integrand g lives in a specific function space $\mathcal{H}(k)$. $\mathcal{H}(k)$ is a *reproducing kernel Hilbert space* (RKHS), which is an inner product space of functions $g : \mathcal{D}_{\mathbf{X}} \rightarrow \mathbb{R}$. Considering a symmetric and positive definite function $k : \mathcal{D}_{\mathbf{X}} \times \mathcal{D}_{\mathbf{X}} \rightarrow \mathbb{R}$, later called a “reproducing kernel” or simply a “kernel”, an RKHS verifies the following axioms:

- The “feature map” $\phi : \mathcal{D}_{\mathbf{X}} \rightarrow \mathcal{H}(k)$; $\phi(\mathbf{x}) = k(\cdot, \mathbf{x})$ belongs to the RKHS: $k(\cdot, \mathbf{x}) \in \mathcal{H}(k), \forall \mathbf{x} \in \mathcal{D}_{\mathbf{X}}$;
- The “reproducing property”: $\langle g, k(\cdot, \mathbf{x}) \rangle_{\mathcal{H}(k)} = g(\mathbf{x}), \quad \forall \mathbf{x} \in \mathcal{D}_{\mathbf{X}}, \forall g \in \mathcal{H}(k)$.

Note that it can be shown that every positive semi-definite kernel defines a unique RKHS (and vice versa) with a feature map ϕ , such that $k(\mathbf{x}, \mathbf{x}') = \langle \phi(\mathbf{x}), \phi(\mathbf{x}') \rangle_{\mathcal{H}(k)}$. This framework allows us to embed a continuous or discrete probability measure in an RKHS, as illustrated in Fig. 4.2. For any measure π , let us define its *kernel mean embedding* (Sejdinovic et al., 2013), also called “potential” $P_\pi(\mathbf{x})$ in Pronzato and Zhigljavsky (2020), associated with the kernel k as:

$$P_\pi(\mathbf{x}) := \int_{\mathcal{D}_{\mathbf{X}}} k(\mathbf{x}, \mathbf{x}') d\pi(\mathbf{x}'). \quad (4.5)$$

Figure 4.2: Kernel mean embedding of a continuous and discrete probability distribution

Respectively, the potential $P_{\zeta_n}(\mathbf{x})$ of a discrete distribution $\zeta_n = \sum_{i=1}^n w_i \delta(\mathbf{x}^{(i)})$, $w_i \in \mathbb{R}_+$ associated with the kernel k can be written as:

$$P_{\zeta_n}(\mathbf{x}) = \int_{\mathcal{D}_{\mathbf{X}}} k(\mathbf{x}, \mathbf{x}') d\zeta_n(\mathbf{x}') = \sum_{i=1}^n w_i k(\mathbf{x}, \mathbf{x}^{(i)}). \quad (4.6)$$

The potential $P_\pi(\mathbf{x})$ of the targeted measure π will be referred to as “target potential” and the potential $P_{\zeta_n}(\mathbf{x})$ associated with the discrete distribution ζ_n called “current potential” when its support is the current design \mathbf{X}_n . When $P_{\zeta_n}(\mathbf{x})$ is close to $P_\pi(\mathbf{x})$, it can be interpreted as ζ_n being an adequate quantization or representation of π (which leads to a good estimation of a quantity such as $I_\pi(g)$ from Eq. (4.3)). Potentials can be computed in closed forms for specific pairs of distribution and associated kernel. Summary tables of some of these cases are detailed in Briol (2019) (section 3.4), Pronzato and Zhigljavsky (2020) (section 4), and extended in Fekhari et al. (2023). However, in most cases, the target potentials must be estimated on a large and representative sample, typically a large quasi-Monte Carlo sample of π .

The *energy* of a measure π is defined as the integral of the potential P_π against the measure π , which leads to the following scalar quantity:

$$\varepsilon_\pi := \int_{\mathcal{D}_\mathbf{x}} P_\pi(\mathbf{x}) d\pi(\mathbf{x}) = \iint_{\mathcal{D}_\mathbf{x}^2} k(\mathbf{x}, \mathbf{x}') d\pi(\mathbf{x}) d\pi(\mathbf{x}'). \quad (4.7)$$

Finally, using the reproducing property and writing the Cauchy-Schwarz inequality on the absolute quadrature error leads to the following inequality, similar to the Koksma-Hlawka inequality Eq. (4.4) (see Briol et al. (2019)):

$$\left| \sum_{i=1}^n w_i g(\mathbf{x}^{(i)}) - \int_{\mathcal{D}_\mathbf{x}} g(\mathbf{x}) d\pi(\mathbf{x}) \right| = \left| \langle g, P_{\zeta_n}(\mathbf{x}) \rangle_{\mathcal{H}(k)} - \langle g, P_\pi(\mathbf{x}) \rangle_{\mathcal{H}(k)} \right| \quad (4.8a)$$

$$= \left| \langle g, (P_{\zeta_n}(\mathbf{x}) - P_\pi(\mathbf{x})) \rangle_{\mathcal{H}(k)} \right| \quad (4.8b)$$

$$\leq \|g\|_{\mathcal{H}(k)} \|P_\pi(\mathbf{x}) - P_{\zeta_n}(\mathbf{x})\|_{\mathcal{H}(k)}. \quad (4.8c)$$

Maximum mean discrepancy A metric of discrepancy and quadrature error is offered by the *maximum mean discrepancy* (MMD). This distance between two probability distributions π and ζ is given by the worst-case error for any function within a unit ball of the Hilbert space $\mathcal{H}(k)$, associated with the kernel k :

$$\text{MMD}_k(\pi, \zeta) := \sup_{\|g\|_{\mathcal{H}(k)} \leq 1} \left| \int_{\mathcal{D}_\mathbf{x}} g(\mathbf{x}) d\pi(\mathbf{x}) - \int_{\mathcal{D}_\mathbf{x}} g(\mathbf{x}) d\zeta(\mathbf{x}) \right|. \quad (4.9)$$

According to the inequality in Eq. (4.8c), $\text{MMD}_k(\pi, \zeta) = \|P_\pi - P_\zeta\|_{\mathcal{H}(k)}$, meaning that the MMD fully relies on the difference of potentials. Moreover, Sriperumbudur et al. (2010) defines a kernel as “characteristic kernel” when the following equivalence is true: $\text{MMD}_k(\pi, \zeta) = 0 \Leftrightarrow \pi = \zeta$. This property makes the MMD a metric on $\mathcal{D}_\mathbf{x}$. The diagram in Fig. 4.3 illustrates the popularity of this metric, which relies on the so-called “kernel-trick”, allowing to implicitly compute dot-products in the RKHS between two objects without computing their embedding. The squared MMD has been used for other purposes than numerical integration: e.g., statistical testing (Gretton et al., 2006), global sensitivity analysis (Da Veiga, 2015), or clustering distributions (Lovera et al., 2023). It can be developed as follows:

$$\text{MMD}^2(\pi, \zeta) = \|P_\pi(\mathbf{x}) - P_\zeta(\mathbf{x})\|_{\mathcal{H}(k)}^2 \quad (4.10a)$$

$$= \langle (P_\pi(\mathbf{x}) - P_\zeta(\mathbf{x})), (P_\pi(\mathbf{x}) - P_\zeta(\mathbf{x})) \rangle_{\mathcal{H}(k)} \quad (4.10b)$$

$$= \langle P_\pi(\mathbf{x}), P_\pi(\mathbf{x}) \rangle_{\mathcal{H}(k)} - 2 \langle P_\pi(\mathbf{x}), P_\zeta(\mathbf{x}) \rangle_{\mathcal{H}(k)} + \langle P_\zeta(\mathbf{x}), P_\zeta(\mathbf{x}) \rangle_{\mathcal{H}(k)} \quad (4.10c)$$

$$= \iint_{\mathcal{D}_\mathbf{x}^2} k(\mathbf{x}, \mathbf{x}') d\pi(\mathbf{x}) d\pi(\mathbf{x}') - 2 \iint_{\mathcal{D}_\mathbf{x}^2} k(\mathbf{x}, \mathbf{x}') d\pi(\mathbf{x}) d\zeta(\mathbf{x}') + \iint_{\mathcal{D}_\mathbf{x}^2} k(\mathbf{x}, \mathbf{x}') d\zeta(\mathbf{x}) d\zeta(\mathbf{x}'). \quad (4.10d)$$

Taking a discrete distribution with uniform weights $\zeta_n = \frac{1}{n} \sum_{i=1}^n \delta(\mathbf{x}^{(i)})$, the squared MMD reduces to:

$$\text{MMD}_k(\pi, \zeta_n)^2 = \varepsilon_\pi - \frac{2}{n} \sum_{i=1}^n P_\pi(\mathbf{x}^{(i)}) + \frac{1}{n^2} \sum_{i,j=1}^n k(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}). \quad (4.11)$$

4.2.3 Kernel herding sampling

The MMD is used below to build a quadrature rule by sampling from a known measure. In other words, to quantize a known target measure π by a design sample \mathbf{X}_n . For practical reasons, the design construction is done sequentially. Moreover, since each realization is supposed to be obtained at the same unitary cost, the quadrature weights are first fixed as uniform during the construction of the design \mathbf{X}_n .

Kernel herding (KH), proposed by Chen et al. (2010), is a sampling method that offers a quantization of the measure π by minimizing a squared MMD when adding points iteratively. With a current design \mathbf{X}_n and its




Figure 4.3: Kernel mean embedding of continuous and discrete distributions.

corresponding discrete distribution with uniform weights $\zeta_n = \frac{1}{n} \sum_{i=1}^n \delta(\mathbf{x}^{(i)})$, a KH iteration consists in solving an optimization of the following criterion, selecting the point $\mathbf{x}^{(n+1)} \in \mathcal{D}_{\mathbf{X}}$ as follows:

$$\mathbf{x}^{(n+1)} \in \arg \min_{\mathbf{x} \in \mathcal{D}_{\mathbf{X}}} \left\{ \text{MMD}^2 \left(\pi, \frac{1}{n+1} \left(\delta(\mathbf{x}) + \sum_{i=1}^n \delta(\mathbf{x}^{(i)}) \right) \right) \right\}. \quad (4.12)$$

In the literature, two formulations of this optimization problem can be found. The first one uses the Frank-Wolfe algorithm (or “conditional gradient algorithm”) to compute a linearization of the problem under the convexity hypothesis (see [Lacoste-Julien et al., 2015](#) and [Briol et al., 2015](#) for more details). The second one is a straightforward greedy optimization. Due to the combinatorial complexity, the greedy formulation is tractable for sequential construction only. Let us develop the MMD estimator introduced in Eq. (4.11):

$$\text{MMD}^2 \left(\pi, \frac{1}{n+1} \left(\delta(\mathbf{x}) + \sum_{i=1}^n \delta(\mathbf{x}^{(i)}) \right) \right) = \varepsilon_{\pi} - \frac{2}{n+1} \sum_{i=1}^{n+1} P_{\pi}(\mathbf{x}^{(i)}) + \frac{1}{(n+1)^2} \sum_{i,j=1}^{n+1} k(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}) \quad (4.13a)$$

$$= \varepsilon_{\pi} - \frac{2}{n+1} \left(P_{\pi}(\mathbf{x}) + \sum_{i=1}^n P_{\pi}(\mathbf{x}^{(i)}) \right) \quad (4.13b)$$

$$+ \frac{1}{(n+1)^2} \left(\sum_{i,j=1}^n k(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}) + 2 \sum_{i=1}^n k(\mathbf{x}^{(i)}, \mathbf{x}) - k(\mathbf{x}, \mathbf{x}) \right), \quad (4.13c)$$

where the “potential”, also called kernel mean embedding, of a measure π is given by $P_{\pi}(\mathbf{x}) = \int_{\mathcal{D}_{\mathbf{X}}} k(\mathbf{x}, \mathbf{x}') d\pi(\mathbf{x}')$, and the energy of the measure π is defined as $\varepsilon_{\pi} = \int \int_{\mathcal{D}_{\mathbf{X}}^2} k(\mathbf{x}, \mathbf{x}') d\pi(\mathbf{x}) d\pi(\mathbf{x}')$. In the previously developed expression, only a few terms actually depend on the next optimal point $\mathbf{x}^{(n+1)}$ since the target energy, denoted by ε_{π} , and $k(\mathbf{x}, \mathbf{x}) = \sigma^2$ are constant assuming a stationary kernel. Therefore, the greedy minimization of the MMD can be equivalently written as:

$$\mathbf{x}^{(n+1)} \in \arg \min_{\mathbf{x} \in \mathcal{D}_{\mathbf{X}}} \left\{ \frac{1}{n+1} \sum_{i=1}^n k(\mathbf{x}^{(i)}, \mathbf{x}) - P_{\pi}(\mathbf{x}) \right\} = \arg \min_{\mathbf{x} \in \mathcal{D}_{\mathbf{X}}} \left\{ \frac{n}{n+1} P_{\zeta_n}(\mathbf{x}) - P_{\pi}(\mathbf{x}) \right\}. \quad (4.14)$$

Remark 1 For the sequential and uniformly weighted case, the formulation in Eq. (4.14) is almost similar to the Frank-Wolfe formulation. Based on a few numerical experiments, these two versions generate very close designs, especially as n becomes large. [Pronzato and Rendas \(2023\)](#) express the Frank-Wolfe formulation in the sequential and uniformly weighted case as follows:

$$\mathbf{x}^{(n+1)} \in \arg \min_{\mathbf{x} \in \mathcal{D}_{\mathbf{X}}} \{ P_{\zeta_n}(\mathbf{x}) - P_{\pi}(\mathbf{x}) \}. \quad (4.15)$$

Remark 2 In practice, the optimization problem is solved by a brute-force approach on a fairly dense finite subset $S \subseteq \mathcal{D}_{\mathbf{X}}$ of candidate points with size $N \gg n$ that emulates the target distribution, also called the “candidate set”. This sample is also used to estimate the target potential $P_{\pi}(\mathbf{x}) \approx \frac{1}{N} \sum_{i=1}^N k(\mathbf{x}^{(i)}, \mathbf{x})$.

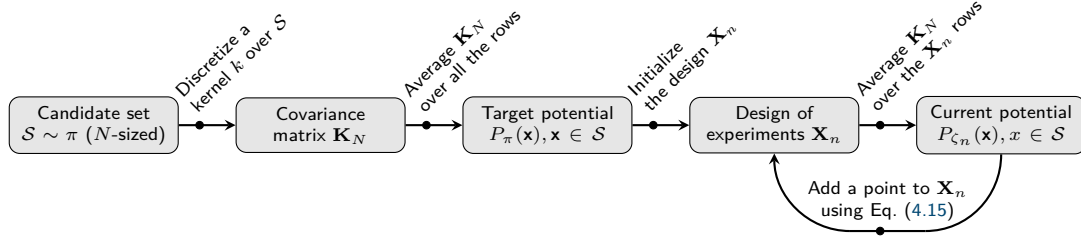


Figure 4.4: Greedy KH algorithm.

Energy-distance	$k_E(\mathbf{x}, \mathbf{x}') = \frac{1}{2} (\ \mathbf{x}\ + \ \mathbf{x}'\ - \ \mathbf{x} - \mathbf{x}'\)$	
Squared exponential	$k_G(\mathbf{x}, \mathbf{x}') = \prod_{i=1}^p k_{\theta_i}(x_i - x'_i)$	$k_{\theta}(x - x') = \exp\left(-\frac{(x-x')^2}{2\theta^2}\right)$
Matérn ($\nu = 5/2$)	$k_M(\mathbf{x}, \mathbf{x}') = \prod_{i=1}^p k_{5/2, \theta_i}(x_i - x'_i)$	$k_{5/2, \theta}(x - x') = \left(1 + \frac{\sqrt{5}}{\theta} x - x' \right) + \frac{5}{3\theta^2}(x - x')^2 \exp\left(-\frac{\sqrt{5}}{\theta} x - x' \right)$

Table 4.1: Kernels considered in the following numerical experiments.

The diagram illustrated in Fig. 4.4 summarizes the main steps of a KH sampling algorithm. One can notice that the initialization can either be done using a median point (maximizing the target potential) or from any existing design of experiments. This second configuration is practical when the analyst must include some characteristic points in the design (e.g., points with a physical interpretation).

As explained previously, choosing the kernel defines the function space on which the worst-case function is found (see Eq. (4.9)). Therefore, this sampling method is sensitive to the kernel's choice. A kernel is defined, both by the choice of its parametric family (e.g., Matérn, squared exponential) and the choice of its tuning. The so-called “support points” method developed by Mak and Joseph (2018) is a special case of KH that uses the characteristic and parameter-free “energy-distance” kernel (introduced by Székely and Rizzo, 2013). In the following numerical experiments, the energy-distance kernel will be compared with an isotropic tensor product of a Matérn kernel (with regularity parameter $\nu = 5/2$ and correlation lengths θ_i), or a squared exponential kernel (with correlation lengths θ_i) defined in Table 4.1. Since the Matérn and squared exponential kernels are widely used for GP regression (Rasmussen and Williams, 2006), they were naturally picked to challenge the energy-distance kernel. The correlation lengths for the squared exponential and Matérn kernels are set using the heuristic given in Fekhari et al. (2023), $\theta_i = n^{-1/d}$, $i \in \{1, \dots, d\}$.

Fig. 4.5 represents the covariance structure of the three kernels. One can notice that the squared exponential and Matérn $\nu = 5/2$ kernels are closer to one another than they are to the energy-distance. In fact, as ν tends to infinity, the Matérn kernel tends toward the squared exponential kernel (which has infinitely differentiable sample paths, see Rasmussen and Williams, 2006). For these two stationary kernels, the correlation length controls how fast the correlation between two points decreases as their distance from one another increases.

Meanwhile, the energy-distance is not stationary (but still positive and semi-definite). Therefore, its value does not only depend on the distance between two points but also on the norm of each of the points. Interestingly, the energy-distance kernel is almost similar to the kernel used by Hickernell (1998) to define a widely-used space-filling metric called the centered L^2 -discrepancy. A presentation of these kernel-based discrepancies from the design of experiment point of view is also provided in Chapter Two from Fang et al. (2018).

To illustrate the KH sampling of a complex distribution, Fig. 4.6 shows three nested samples (orange crosses for various sizes $n \in \{10, 20, 40\}$) of a mixture of Gaussian distributions with complex nonlinear dependencies (with density represented by the black isoprobability contours). In this example, the method seems to build a parsimonious design between each mode of the distribution (by subsampling directly without any transformation). The candidate set (in light gray) was generated by a large quasi-Monte sample of the underlying Gaussian mixture. In this two-dimensional case, this candidate set is sufficient to estimate the target potential P_π . However, the main bottleneck of KH is the estimation of the potentials, which becomes costly in high dimensions.

Other approaches take advantage of the progressive knowledge acquired sequentially from the outputs to select




Figure 4.5: Kernel illustrations (left to right: energy-distance, squared exponential, and Matérn 5/2).




Figure 4.6: Sequential KH for increasing design sizes ($n \in \{10, 20, 40\}$) built on a candidate set of $N = 8196$ points drawn from a complex Gaussian mixture π .

the following points in the design. These methods are sometimes called “active learning” or “adaptive strategies” (Fuhg. et al., 2021). Many of them rely on a sequentially updated Gaussian process (or Kriging) metamodel. To solve a probabilistic integration problem, the concept of Bayesian quadrature is introduced in the following.

4.2.4 Bayesian quadrature.

Gaussian processes for Bayesian quadrature Kernel methods and GP present a lot of connections and equivalences, thoroughly reviewed by Kanagawa et al. (2018). In numerical integration, GP have been used to build quadrature rules in the seminal paper of O’Hagan (1991), introducing the concept of *Bayesian quadrature* (BQ). Let us recall the probabilistic integration problem $I_\pi(g) = \int_{\mathcal{D}_x} g(\mathbf{x}) d\pi(\mathbf{x})$ (stated in Eq. (4.3)). From a general point of view, this quantity could be generalized by composing g with another function ψ (e.g., other moments, quantiles, exceedance probabilities). The quantity of interest then becomes, $I_\pi(\psi(g))$, for example when ψ is a monomial, it gives a moment of the output distribution.

Let us assume, adopting a Bayesian point of view, that G is a stochastic process describing the uncertainty affecting the knowledge about the true function g . Let G be a Gaussian process (GP) prior with a zero trend (denoted by $\mathbf{0}$) to ease the calculation, and a stationary covariance kernel (denoted by $k(\cdot, \cdot)$). The conditional posterior $G_n = (G|\mathbf{y}_n) \sim \text{GP}(\eta_n, s_n^2)$ has been conditioned on the function observations $\mathbf{y}_n = [g(\mathbf{x}^{(1)}), \dots, g(\mathbf{x}^{(n)})]^\top$ computed from the input design \mathbf{X}_n and is fully defined by the well-known “Kriging equations” (see e.g., Rasmussen and Williams, 2006):

$$\begin{cases} \eta_n(\mathbf{x}) &= \mathbf{k}_n^\top(\mathbf{x}) \mathbf{K}_n^{-1} \mathbf{y}_n \\ s_n^2(\mathbf{x}) &= k_n(\mathbf{x}, \mathbf{x}) - \mathbf{k}_n^\top(\mathbf{x}) \mathbf{K}_n^{-1} \mathbf{k}_n(\mathbf{x}) \end{cases} \quad (4.16)$$

where $\mathbf{k}_n(\mathbf{x})$ is the column vector of the covariance kernel evaluations $[k_n(\mathbf{x}, \mathbf{x}^{(1)}), \dots, k_n(\mathbf{x}, \mathbf{x}^{(n)})]$ and \mathbf{K}_n is the $(n \times n)$ variance-covariance matrix such that the (i, j) -element is $\{\mathbf{K}_n\}_{i,j} = k_n(\mathbf{x}^{(i)}, \mathbf{x}^{(j)})$.

In BQ, the main object is the random variable $I_\pi(G_n)$. Following the denomination of Briol et al. (2019),

its distribution on \mathbb{R} is the pushforward of G_n through the integration operator $I_\pi(\cdot)$, sometimes called *posterior distribution*:

$$I_\pi(G_n) = \int_{\mathcal{D}_x} (G(\mathbf{x})|\mathbf{y}_n) d\pi(\mathbf{x}) = \int_{\mathcal{D}_x} G_n(\mathbf{x}) d\pi(\mathbf{x}). \quad (4.17)$$

Fig. 4.7 provides a one-dimensional illustration of the BQ of an unknown function (dashed black curve) against a given input measure π (with corresponding gray distribution at the bottom). For an arbitrary design, one can fit a Gaussian process model, interpolating the function observations (black crosses). Then, multiple trajectories of this conditioned Gaussian process G_n are drawn (orange curves) whilst its mean function, also called “predictor”, is represented by the red curve. Therefore, the input measure π is propagated through the conditioned Gaussian process to obtain the random variable $I_\pi(G_n)$, with distribution represented on the right plot (brown curve). Again on the right plot, one can remark how the mean of this posterior distribution (brown line) is closer to the reference output expected value (dashed black line) than the arithmetic mean of the observations (black line). This plot was inspired by the paper of [Huszár and Duvenaud \(2012\)](#).

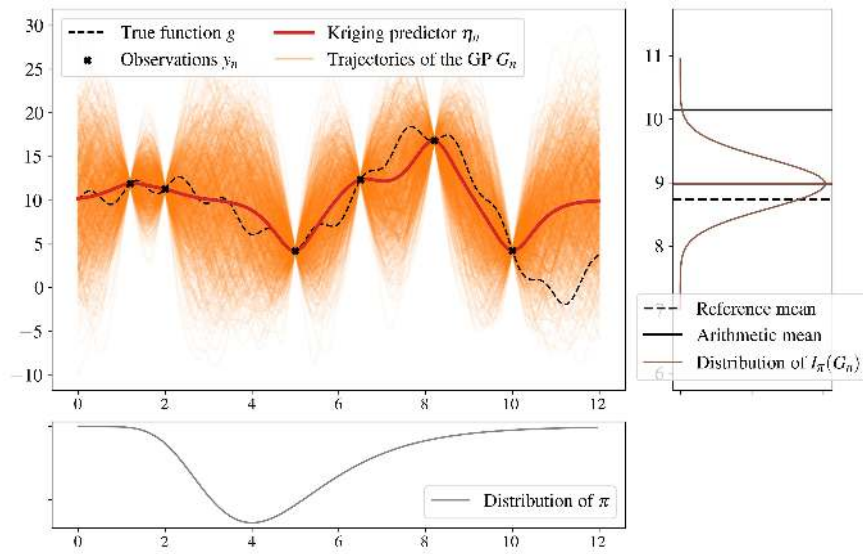


Figure 4.7: Bayesian quadrature on a one-dimensional case.

Optimal weights computed by Bayesian quadrature Taking the random process G_n as Gaussian conveniently implies that its posterior distribution $a_\pi(G_n)$ is also Gaussian. This comes from the linearity of the infinite sum of realizations of a Gaussian process. The posterior distribution is described in a closed form through its mean and variance by applying Fubini’s theorem (see the supplementary materials from [Briol et al., 2019](#) for the proof regarding the variance):

$$\bar{y}_n^{\text{BQ}} = \mathbb{E}[I_\pi(G_n)|\mathbf{y}_n] = \int_{\mathcal{D}_x} \eta_n(\mathbf{x}) d\pi(\mathbf{x}) = \left[\int_{\mathcal{D}_x} \mathbf{k}_n^\top(\mathbf{x}) d\pi(\mathbf{x}) \right] \mathbf{K}_n^{-1} \mathbf{y}_n = P_\pi(\mathbf{X}_n) \mathbf{K}_n^{-1} \mathbf{y}_n, \quad (4.18)$$

$$(\sigma_n^{\text{BQ}})^2 = \text{Var}(I_\pi(G_n)) = \iint_{\mathcal{D}_x^2} k_n(\mathbf{x}, \mathbf{x}') d\pi(\mathbf{x}) d\pi(\mathbf{x}') = \varepsilon_\pi - P_\pi(\mathbf{X}_n) \mathbf{K}_n^{-1} P_\pi(\mathbf{X}_n)^\top, \quad (4.19)$$

where $P_\pi(\mathbf{X}_n)$ is the row vector of potentials $[\int k_n(\mathbf{x}, \mathbf{x}^{(1)}) d\pi(\mathbf{x}), \dots, \int k_n(\mathbf{x}, \mathbf{x}^{(n)}) d\pi(\mathbf{x})]$, and ε_π is given in Eq. (4.7). As in the one-dimensional example presented in Fig. 4.7, the expected value of $I_\pi(G_n)$ expressed in Eq. (4.18) is a direct estimator of the quantity of interest expressed in Eq. (4.3). The so-called “Bayesian quadrature estimator” appears to be a simple linear combination of the observations by taking the row vector of “optimal weights” as:

$$\mathbf{w}_{\text{BQ}} = P_\pi(\mathbf{X}_n) \mathbf{K}_n^{-1} \quad (4.20)$$

For any given sample, an optimal set of weights can be computed, leading to the mean of the posterior distribution. Remark here that this enhancement depends on the evaluation of the inverse variance-covariance matrix \mathbf{K}_n^{-1} , which can present numerical difficulties, either when design points are too close, leading to poor conditioning of

the matrix. Moreover, a prediction interval on the BQ estimator can be computed since the posterior distribution is Gaussian, with a variance expressed in closed-form in Eq. (4.19). The expressions in Eq. (4.18) and Eq. (4.19) were extended to GPs in the case of constant and linear trends in Pronzato and Zhigljavsky (2020). In the following numerical experiments, the expression with a hypothesis of constant trend β_n is used, which leads to:

$$\mathbb{E}[I_\pi(G_n)] = \beta_n + P_\pi(\mathbf{X}_n)\mathbf{K}_n^{-1}(\mathbf{y}_n - \beta_n \mathbf{1}_n). \quad (4.21)$$

Then, an a posteriori 95% prediction interval around the mean Bayesian estimator is directly given by:

$$\bar{y}_n^{\text{BQ}} \in [\bar{y}_n^{\text{BQ}} - 2\sigma_n^{\text{BQ}}, \bar{y}_n^{\text{BQ}} + 2\sigma_n^{\text{BQ}}]. \quad (4.22)$$

Variance-based Bayesian quadrature rule. The link between the posterior variance and the squared MMD has been first made by (Huszár and Duvenaud, 2012, Proposition 1): the expected variance in the BQ $\text{Var}(I_\pi(G_n))$ is the MMD between the target distribution π and $\zeta_n = \sum_{i=1}^n \mathbf{w}_{\text{BQ}}^{(i)} \delta(\mathbf{x}^{(i)})$. The proof is reproduced below (as well as in Proposition 6.1 from Kanagawa et al., 2018):

$$\text{Var}(I_\pi(G_n)) = \mathbb{E}[(I_\pi(G_n) - I_{\zeta_n}(G_n))^2] \quad (4.23a)$$

$$= \mathbb{E}\left[\left(\langle G_n, P_\pi \rangle_{\mathcal{H}(k)} - \langle G_n, P_{\zeta_n} \rangle_{\mathcal{H}(k)}\right)^2\right] \quad (4.23b)$$

$$= \mathbb{E}\left[\langle G_n, P_\pi - P_{\zeta_n} \rangle_{\mathcal{H}(k)}^2\right] \quad (4.23c)$$

$$= \|P_\pi - P_{\zeta_n}\|_{\mathcal{H}(k)}^2 \quad (4.23d)$$

$$= \text{MMD}^2(\pi, \zeta_n). \quad (4.23e)$$

Note that the transition from Eq. (4.23c) to Eq. (4.23d) relies on the property stating that if G is a standard Gaussian process then $\forall g \in \mathcal{H}(k) : \langle G, g \rangle_{\mathcal{H}(k)} \sim \mathcal{N}(0, \|g\|_{\mathcal{H}(k)}^2)$. The method that sequentially builds a quadrature rule by minimizing this variance is called by the authors “sequential Bayesian quadrature”. According to the previous proof, this criterion can be seen as an optimally-weighted version of the KH criterion, as stated in the title of the paper from Huszár and Duvenaud (2012). Later, Briol et al. (2015) proved the weak convergence of $I_\pi(G_n)$ towards the target integral. Closer to wind turbine applications, Huchet (2018) and Huchet et al. (2019) introduced the “Adaptive Kriging Damage Assessment” method: a Kriging-based method for mean damage estimation that is very close to sequential Bayesian quadrature. However, this type of method inherits the limits from both KH and BQ since it searches for optimal design points among a candidate set and computes an inverse variance-covariance matrix. These numerical operations both scale hardly in high dimension.

Remark 3 Every quadrature method introduced in this section has been built without any observation of the possibly costly function g . Therefore, they cannot be categorized as active learning approaches. Contrarily, Kanagawa and Hennig (2019) presents a set of methods for BQ with transformations (i.e., adding a positivity constraint on the function g), which are truly active learning methods.

4.3 Numerical experiments

This section presents numerical results computed on two different analytical test cases, respectively in dimension 2 (test case #1) and dimension 10 (test case #2), with easy-to-evaluate functions $g(\cdot)$ and associated input distributions π . Therefore, reference values can easily be computed with great precision. For each test case, a large reference Monte Carlo sample ($N_{\text{ref}} = 10^8$) is taken. This first benchmark compares the mean estimation of test cases given by a quasi-Monte Carlo technique (abbreviated by QMC in the next figures) which consists herein using a Sobol’ sequence, and KH with the three kernels defined in Table 4.1. Notice that the quasi-Monte Carlo designs are first generated on the unit hypercube and then, transformed using the generalized Nataf transformation to follow the target distribution (Lebrun and Dutfoy, 2009). Additionally, the performances of KH for both uniform and optimally-weighted Eq. (4.21) estimators are compared.

All the following results and methods (i.e., the KH sampling and BQ methods) have been implemented in a new open source Python package named `otkerneldesign`⁷, which relies on `OpenTURNS` (Baudin et al., 2017). Finally, note that the numerical experiments for the test cases are available in the Git repository named `ctbenchmark`⁸.

⁷<https://efekhari27.github.io/otkerneldesign/master/index.html>

⁸<https://github.com/efekhari27/ctbenchmark>

Table 4.2: Analytical test cases

Test case #1	$d = 2$	$g_1(\mathbf{x}) = x_1 + x_2$	Gaussian mixture from Fig. 4.6
Test case #2	$d = 10$	$g_2(\mathbf{x}) = \prod_{i=1}^{10} \frac{ 4x_i - 2 + a_i}{1 + a_i}, \{a_i = 2\}_{i=1}^{10}$	Gaussian $\mathcal{N}(\mathbf{0.5}, \mathbf{I}_{10})$

4.3.1 Benchmark results on analytical test cases

The test cases summarized in Table 4.2 were chosen to cover a panel of complex probabilistic integration problems, completing the ones from Fekhari et al. (2022). To assess the complexity of numerical integration problems, Owen (2003) introduced the concept of the “effective dimension” of an integrand function (number of the variables that actually impact the integral). The author showed that functions built on sums yield a low effective dimension (unlike functions built on products). In the same vein, Kucherenko et al. (2011) build three classes of integrand sorted by difficulty depending on their effective dimension:

- *class A problems*: problem with a few dominant variables;
- *class B problems*: problem without unimportant variables, and important low-order interaction terms;
- *class C problems*: problems without unimportant variables, and important high-order interaction terms.

In the case of the two-dimensional Gaussian mixture problem (test case #1), the complexity is carried by the mixture of Gaussian distributions with highly nonlinear dependencies. Then, the 10-dimensional “GSobol function” (test case #2) with a set of coefficient $\{a_i = 2\}_{i=1}^{10}$ has an effective dimension equal to 10 and belongs to the hardest class C from Kucherenko et al. (2011). Probabilistic integration results are presented in Fig. 4.8 (test case #1) and Fig. 4.9 (test case #2). Kernel herding samples based on the squared exponential kernel are in green, those using the Matérn kernel are in orange, those using the energy-distance kernel are in red, while quasi-Monte Carlo samples built from Sobol’ sequences are in gray. Convergences of the arithmetic means are plotted on the left and MMDs on the right. The respective BQ estimators of the means are plotted in dashed lines.

Remark 4 *Different kernels are used in these numerical experiments. First, the generation kernel is used by the KH algorithm to generate designs (with the heuristic tuning defined in Section 4.2.3). Second, the BQ kernel allows computation of the optimal weights (arbitrarily set up as a Matérn 5/2 with the heuristic tuning). Third, the evaluation kernel must be common to allow a fair comparison of the computed MMD results (same as the BQ kernel).*

Regarding the results for the test case #1, the convergence plots are provided in Fig. 4.8. KH consistently converges faster than quasi-Monte Carlo in this case, especially for small sizes in terms of MMD. BQ weights tend to reduce the fluctuations in the mean convergence, which ensures better performance for any size. Overall, applying the weights enhances the convergence rate.

Then, the convergence plots corresponding to the test case #2 are provided in Fig. 4.9. Although quasi-Monte Carlo is known to suffer the “curse of dimensionality”, KH does not outperform it drastically in this example. In fact, KH with uniform weights performs worse than quasi-Monte Carlo while optimally-weighted KH does slightly better. Moreover, the results confirm that $\text{MMD}_{\text{BQ}} < \text{MMD}_{\text{unif}}$ for all the proposed experiments. The application of optimal weights to the quasi-Monte Carlo sample slightly improves the estimation in this case. Note that the prediction interval around the BQ estimator is not plotted for the sake of readability. Finally, one can note that the prediction interval around the BQ estimator is not plotted for the sake of readability.

In these two test cases, the MMD is shown to quantify numerical integration convergence well, which illustrates the validity of the inequality given in Eq. (4.8c), similar to the Koksma-Hlawka inequality. Between the three kernels employed for KH, the energy-distance kernel performs slightly better and presents the advantage of not needing tuning.

4.3.2 Application to the Teesside wind turbine fatigue estimation

In order to summarize the mean damage estimation strategies studied in this section, the diagram represented in Fig. 4.10 describes the different workflows computed. The simplest workflow is represented by the gray horizontal sequence. It directly subsamples a design of experiments from a large and representative dataset (previously referred to as candidate set). This workflow simply estimates the mean damage by computing an arithmetic average of the outputs.

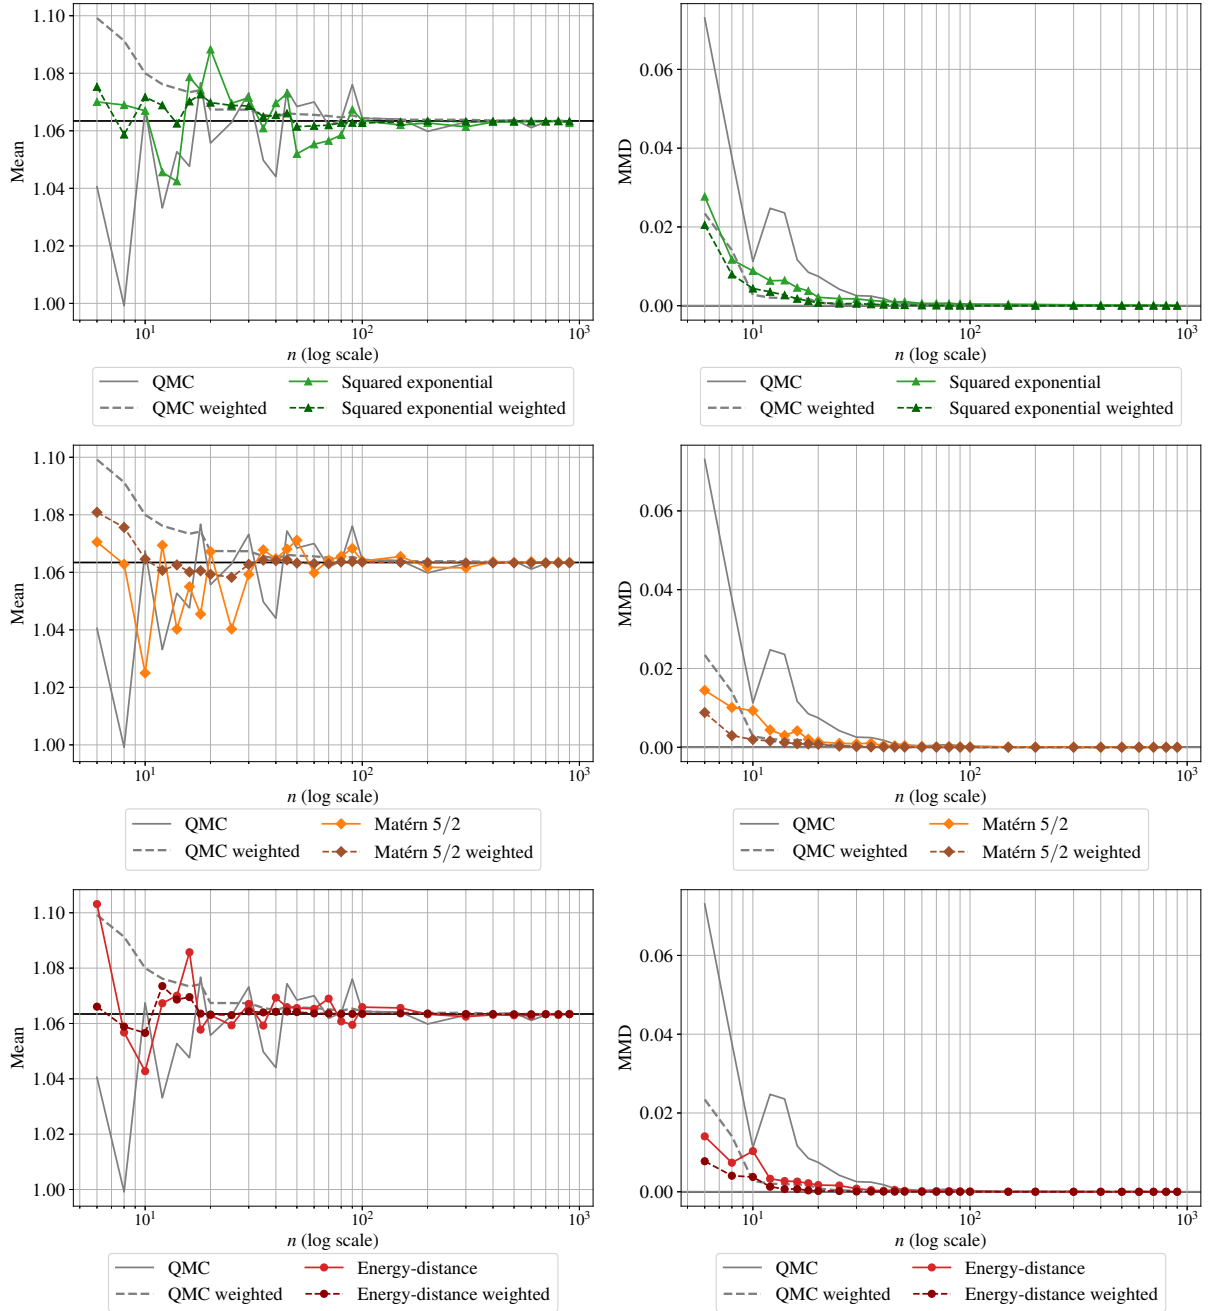


Figure 4.8: Analytical benchmark results on the test case #1.

Alternatively, one can respectively fit a joint distribution and sample from it. Here, this distribution is only known empirically via the candidate set. Since its dependence structure is complex (see Fig. 3.1), a parametric method might fit the distribution poorly (and therefore lead to a poor estimation of the quantity). Then, a nonparametric fit using the empirical Bernstein copula (introduced in Vanem et al., 2024) coupled with a KDE on each marginal is applied to the candidate set (with the EBC parameter $m = 100 > m_{\text{AMISE}}$ to avoid bias, see Lasserre, 2022, p.117). The sampling on this hybrid joint distribution is realized with a quasi-Monte Carlo method, here a Sobol' low-discrepancy sequence generates a uniform sample in the unit hypercube, then transformed according to the target distribution. Remember that quasi-Monte Carlo sampling is also sensitive to the choice of a low-discrepancy sequence, each presenting different properties (Sobol', Halton, Faure, etc.). Finally, the estimation by an arithmetic mean can be replaced by an optimally weighted mean. To do so, optimal weights must be computed, using the formulas introduced in Eq. (4.20).

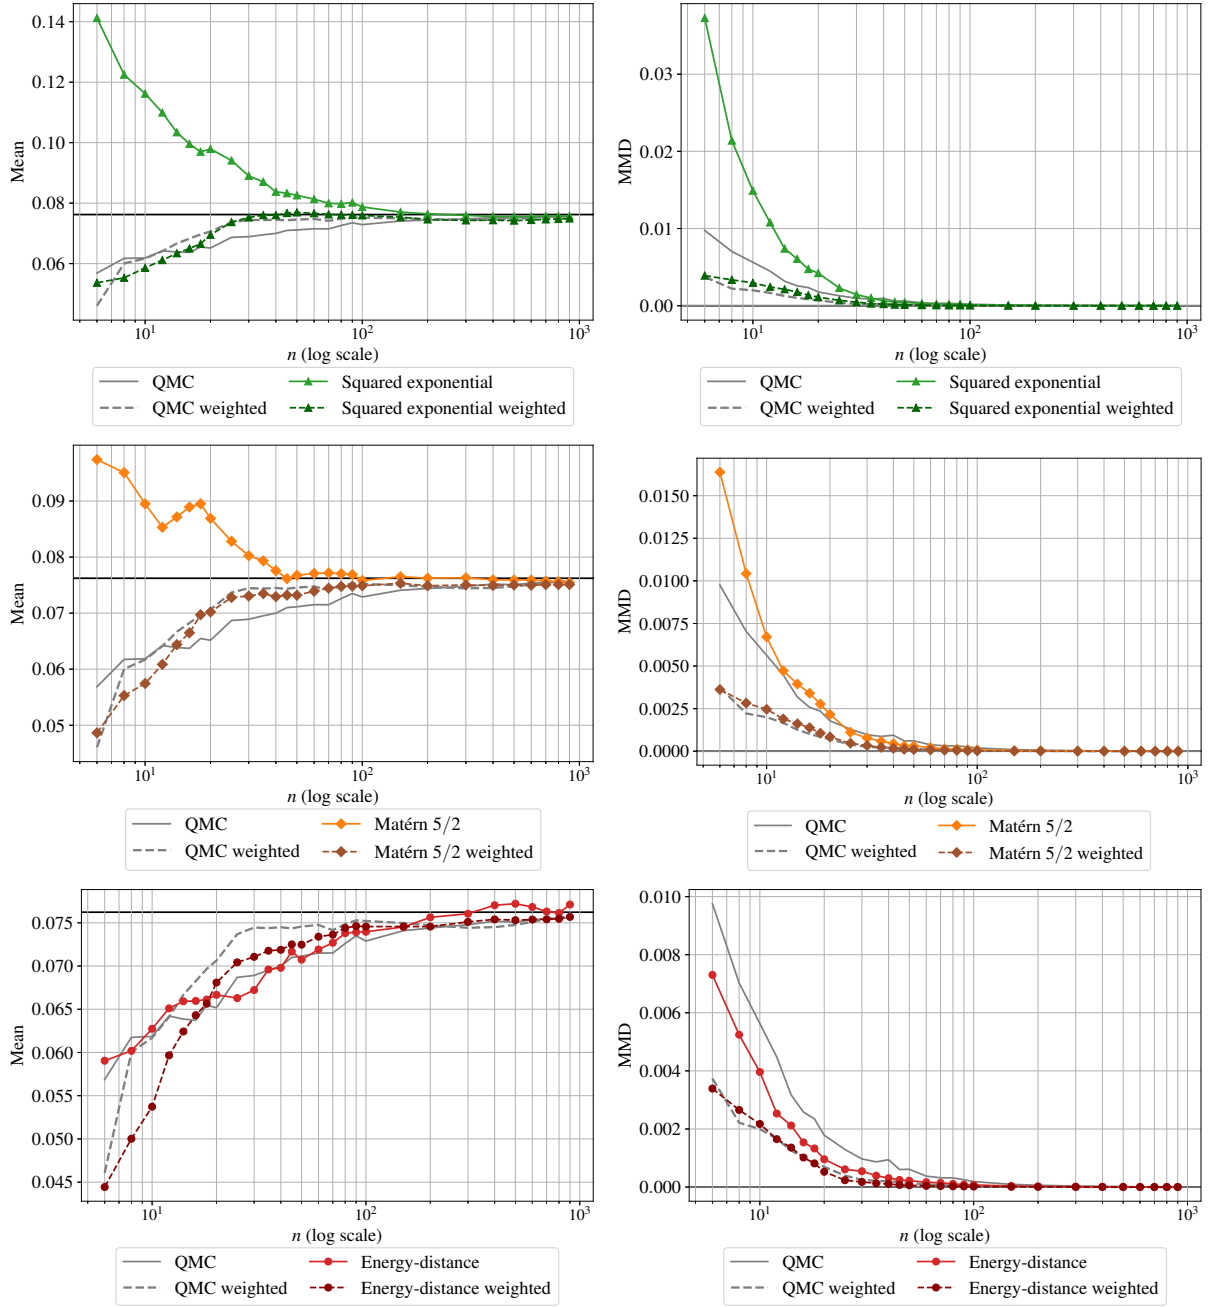


Figure 4.9: Analytical benchmark results on the test case #2.

The copulogram in Fig. 4.11 illustrates the intensity of the computed damages, proportionally to the color scale. Note that the numerical values of the damage scale are kept confidential since it models the state of an operating asset. Before analyzing the performance of the KH on this industrial application, one can notice that the copulogram Fig. 4.11 seems to be in line with the global sensitivity analysis presented in Murcia et al. (2018) and Li and Zhang (2020). In particular, the fact that the scatter plot of mean wind speed vs. turbulence wind speed is the main factor explaining the variance of the output $Y = g(\mathbf{X})$. Judging from these references, the numerical model does not seem to have a high effective dimension. However, the input dependence structure is still challenging, and the damage assessment induces strong nonlinearities (see Eq. (2.2)).

The results presented are compared in the following to a reference Monte Carlo sample of size $n_{\text{ref}} = 2000$, with a confidence interval computed by bootstrap (see Fig. 4.12). This reference is represented by a horizontal line intersecting with the most converged Monte Carlo estimation. Once again, the mean damage scale is hidden

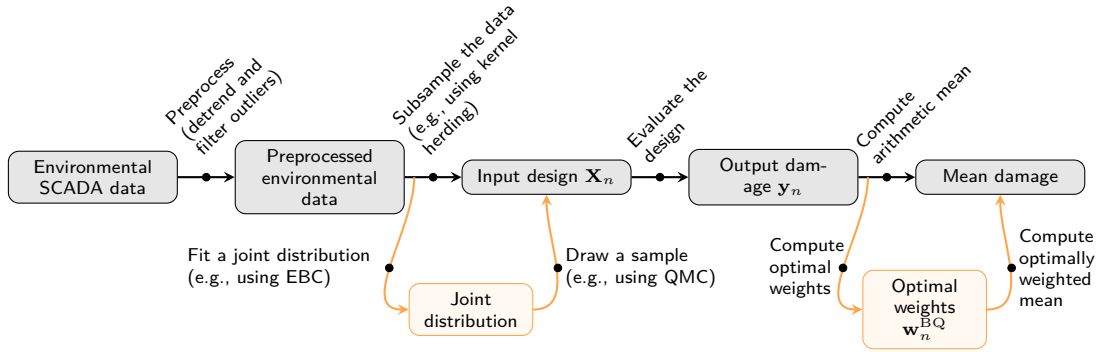


Figure 4.10: Mean damage estimation workflows for the industrial Teesside use case. The orange parts represent optional alterations to the workflow: the first one is an alternative to input data subsampling where the underlying distribution is sampled from, the second one improves mean damage calculation by using optimal weights over the output data.

for confidentiality reasons, but all the plots are represented for the same vertical scale. However, the KH quickly converges towards the confidence interval of the Monte Carlo obtained with the reference sample. In addition, the BQ estimator also offers a posteriori prediction interval, which can reassure the analyst. Finally, one can see that the BQ prediction intervals are smaller than the ones obtained by bootstrap on the reference Monte Carlo sample.

To provide more representative results, note that a set of scale parameters is computed with a Kriging procedure to define the kernel used to compute BQ intervals. Since other methods do not generate independent samples, bootstrapping them is not legitimate. Contrarily to the other kernels, one can notice that the energy-distance kernel presents a small bias with the Monte Carlo reference for most of the azimuth angles computed in this experiment. Meanwhile, combining nonparametric fitting with quasi-Monte Carlo sampling also delivers accurate results as long as the fitting step does not introduce a bias. In the present case, any potential bias due to poor fitting would be the result of a poorly tuned EBC. Fortunately, [Nagler et al. \(2017\)](#) formulated recommendations regarding how to tune EBCs. We follow these recommendations in the present work.

4.4 Synthesis

Wind energy assets are subject to highly uncertain environmental conditions. Uncertainty propagation through numerical models is performed to ensure their structural integrity (and energy production). For this case, the method recommended by the standards (regular grid sampling) is intractable for even moderate-fidelity simulators. In practice, such an approach can lead to poor uncertainty propagation, especially when facing simulation budget constraints.

In the present section, a real industrial wind turbine fatigue estimation use case is investigated, considering site-specific data. As a perspective, other sites with different environmental conditions could be studied. This use case induces two practical constraints: first, usual active learning methods are hard to set up on such a model (mainly due to the nonlinearity of the variable of interest), and they restrict the use of high-performance computing facilities; second, the input distribution of the environmental conditions presents a complex dependence structure which is hard to infer with common parametric approaches.

In this work, the association of KH sampling with BQ for central tendency estimation is explored theoretically and numerically. This method fits with the practical constraints induced by the industrial use case. To be more specific, the KH method easily subsamples the relevant points directly from a given dataset (here, from the measured environmental data). Moreover, the method is fully compatible with the intensive use of HPC facilities. Moreover, the present work outlined an upper bound based on the maximum mean discrepancy (MMD) on numerical integration absolute error. Kernel herding and BQ both aim at finding the optimal quadrature rule minimizing the MMD, and therefore the absolute integration error. The numerical experiments confirm that the MMD is an appropriate criterion since it leads to results being better or equivalent to quasi-Monte Carlo sampling. Finally, the proposed numerical benchmark relies on a Python package, called `otkerneldesign`, which implements the methods and allows anyone to reproduce the results.

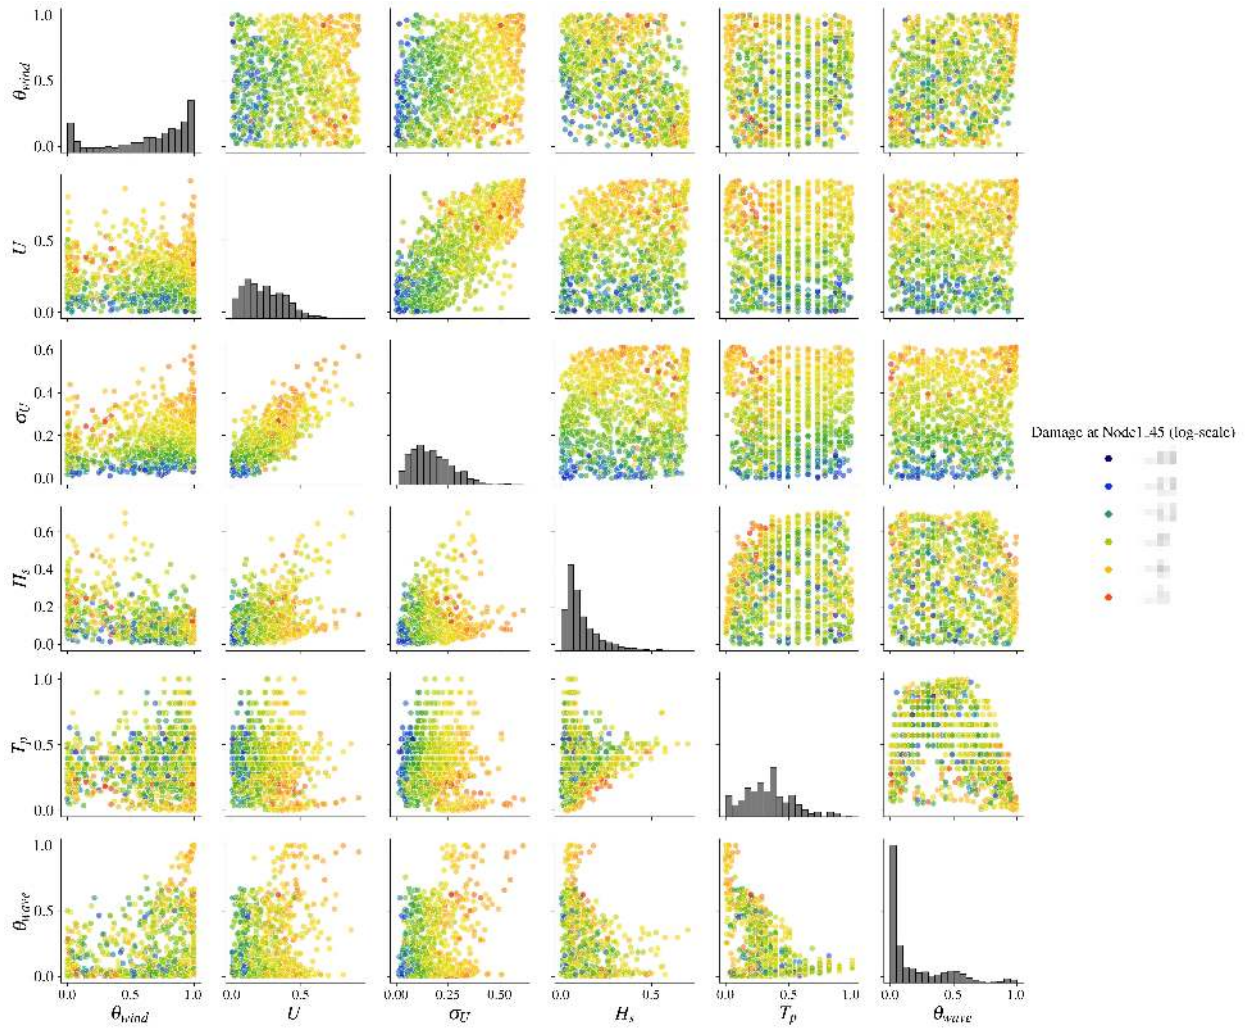


Figure 4.11: Copulogram of the KH design of experiments with corresponding outputs in color (log-scale) on the Teesside case ($n = 10^3$). The color scale ranges from blue for the lowest values to red for the largest.

The limits of the proposed method are reached when the input dimension of the problem increases, requiring a larger candidate set and therefore a larger covariance matrix. Moreover, the numerical experiments show that the method can be sensitive to the choice of the kernel and its tuning (although good practices can still be derived). From a methodological viewpoint, further interpretation of the impact of the different kernels could be explored. Besides, extensions of KH sampling for quantile estimation could be investigated, in a similar fashion as the work on randomized quasi-Monte Carlo for quantiles proposed by [Kaplan et al. \(2019\)](#). Kernel herding could also be used to quantize conditional distributions, using the so-called “conditional kernel mean embedding” concept reviewed by [Klebanov et al. \(2020\)](#). Finally, regarding the industrial use case, the next step will be to perform a reliability analysis by considering another group of random variables (related to the OWT). Further perspectives could be to explore the possibilities offered by reliability-oriented sensitivity analysis in the context of kernel-based indices, as studied in [Marrel and Chabridon \(2021\)](#).

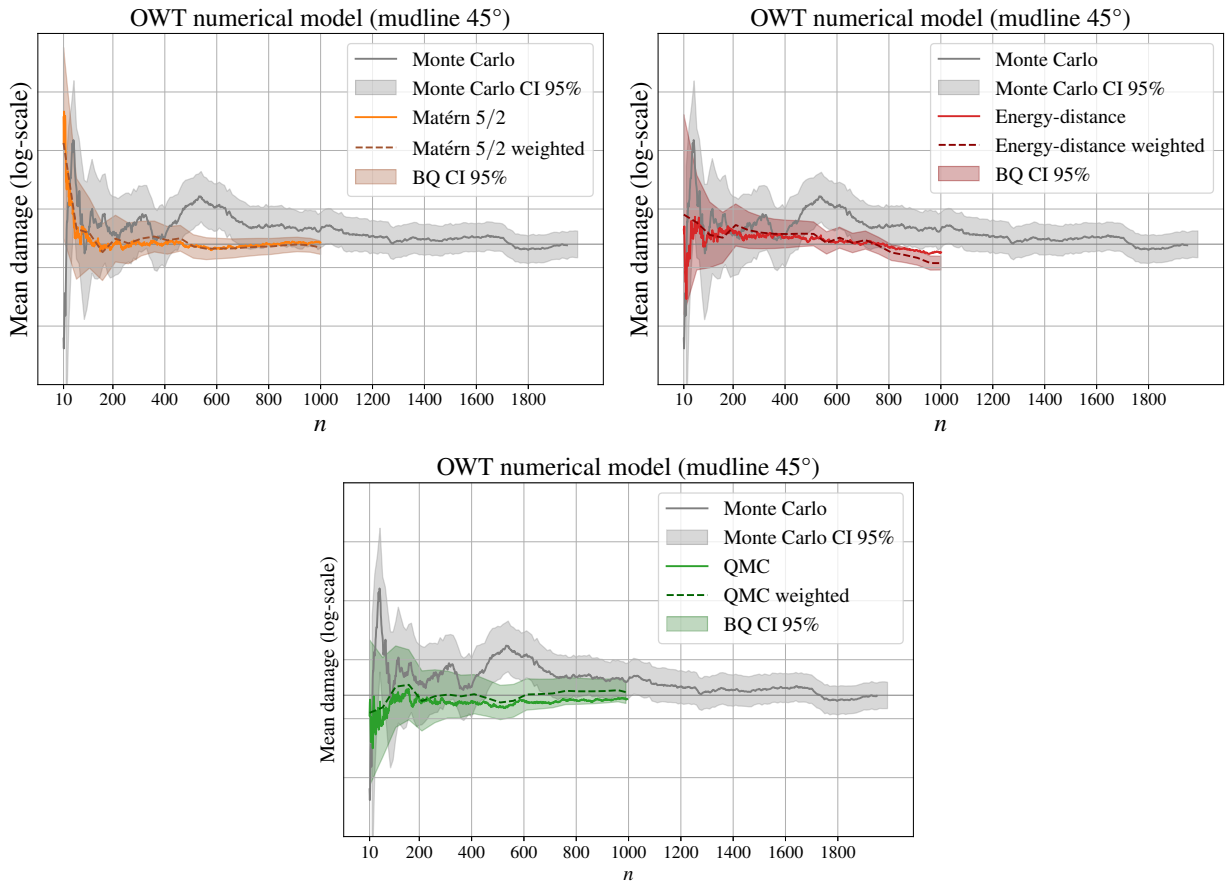


Figure 4.12: Mean estimation convergence (at the mudline, azimuth $\theta = 45$ deg.) on the Teesside case. Monte Carlo confidence intervals are all computed by bootstrap.

5 A second integration approach based on sparse polynomial chaos expansions

As an alternative estimation method for the integral in Eq. (4.2), we explore here a global approach based on the properties of spectral expansions in general, and data-driven polynomial chaos expansions (Xiu and Karniadakis, 2002) in particular.

We follow the formalism and overall computational framework from Torre et al. (2019a,b), adopting a combination of non-parametric marginals & vine-copula inference, and sparse polynomial chaos expansions.

5.1 Introduction

According to Eq. (4.2), two ingredients contribute to the expected damage estimation: the joint distribution of environmental parameters $f_{\mathbf{X}}(vx)$, $\mathbf{x} \in \mathbb{R}^M$, and the computational model of the damage $d_c^{i,\theta}$. Similarly to Section 4.2, the focus of this section will be on the integration methodology, rather than on its specific application to the wind-turbine case.

We write the integral of interest as (see also Eq (4.3)):

$$\mathbb{E}[h(\mathbf{x})] = \int_{\mathcal{D}_{\mathbf{X}}} h(\mathbf{x}) f_{\mathbf{X}}(\mathbf{x}) d\mathbf{x}, \quad (5.1)$$

where $h(\mathbf{x})$ is a generally expensive-to-evaluate computational model $y = h(\mathbf{x})$ that calculates a scalar quantity of interest (QoI) y .

The focus here is to integrate Eq. (5.1) while minimizing the number of calls to the expensive $h(\mathbf{x})$. We therefore assumed that the joint distribution $f_{\mathbf{X}}(\mathbf{x})$ is either known, or enough information is available to properly reconstruct an approximation from data. In the latter case, we use the data-driven vine-copula fitting approach from Aas et al. 2009; Torre et al. 2019a, as it is implemented in the UQLab software (Marelli and Sudret, 2014; Torre et al., 2024).

We also assume that the integrand $h(\mathbf{x})$ is a finite-variance function w.r.t. probability distribution $f_{\mathbf{X}}(\mathbf{X})$:

$$\int_{\mathcal{D}_{\mathbf{X}}} h(\mathbf{x})^2 f_{\mathbf{X}}(\mathbf{x}) d\mathbf{x} < \infty, \quad (5.2)$$

which is a relatively weak requirement for most physics-based models that calculate physically-interpretable quantities of interest, such as stresses or accelerations.

5.2 Polynomial chaos expansions

Consider the Hilbert space \mathcal{H} of finite variance functions w.r.t. to the probability distribution $f_{\mathbf{X}}(\mathbf{x})$ (Eq. (5.2)), endowed with inner product:

$$\langle h, g \rangle \stackrel{\text{def}}{=} \mathbb{E}_{\mathbf{X}}[h(\mathbf{x})g(\mathbf{x})] = \int_{\mathcal{D}_{\mathbf{X}}} h(\mathbf{x})g(\mathbf{x}) f_{\mathbf{X}}(\mathbf{x}) d\mathbf{x}. \quad (5.3)$$

Then any function $Y = h(\mathbf{X})$, $h \in \mathcal{H}$ admits a spectral decomposition known as *polynomial chaos expansion* (PCE, Xiu and Karniadakis (2002)) of the form:

$$h(\mathbf{x}) = \sum_{\alpha \in \mathbb{N}^M} a_{\alpha} \Psi_{\alpha}(\mathbf{x}), \quad (5.4)$$

where $a_{\alpha} = \langle h(\mathbf{x}), \Psi_{\alpha}(\mathbf{x}) \rangle \in \mathbb{R}$ are scalar coefficients (coordinates), $\Psi_{\alpha}(\mathbf{x})$ are multivariate polynomials of total degree $d = |\alpha|$, and $\alpha \in \mathbb{N}^M$ is a multi-index that represents the polynomial degree in each of the input variables.

In the special case of independent input variables, i.e. $f_{\mathbf{X}}(\mathbf{x}) = \prod_{i=1}^M f_{X_i}(x_i)$, the multivariate polynomials in

Eq. (5.4) can be obtained as tensor products of the univariate orthonormal polynomials $\psi_i^{(i)}(x_i)$, with $\langle \phi_i, \phi_j \rangle = \delta_{ij}$. With this construction, it immediately follows that also their multivariate counterparts $\Psi_{\alpha}(\mathbf{x})$ are also orthonormal:

$$\langle \Psi_{\alpha}, \Psi_{\beta} \rangle = \delta_{\alpha\beta}, \quad (5.5)$$

where the explicit dependence on \mathbf{x} was dropped for notational simplicity.

As with most spectral representations, an interesting property of PCE is that its coefficient encapsulate interesting information about the model behavior. Of particular relevance to the case under investigation, we note that $\Psi_0 \stackrel{\text{def}}{=} 1$, and therefore:

$$a_0 = \langle 1, h \rangle \stackrel{\text{def}}{=} \mathbb{E}_{\mathbf{X}} [h(\mathbf{x})] = \int_{\mathcal{D}_{\mathbf{X}}} h(\mathbf{x}) f_{\mathbf{X}}(\mathbf{x}) d\mathbf{x}. \quad (5.6)$$

In other words, thanks to the spectral properties of PCE, the integral in Eq. (5.1) can be directly estimated as the coefficient of the constant term in the polynomial chaos expansion of h . In case of non-independent input distributions, isoprobabilistic transforms (e.g. Rosenblatt transform (Rosenblatt, 1952)) can be used to map the original problem into the independent space.

Alternatively, Eq. (5.4) can be directly evaluated on any value of \mathbf{x} , providing a very efficient and accurate global surrogate model, which can run millions of times per second on common consumer hardware. This particular methodology does not require isoprobabilistic transforms in the presence of dependence between the input variables, as the coefficients are used only to calculate the response of the surrogate model. Once the surrogate PCE is available, the integral in Eq. (5.1) can be directly estimated as a sample average over a large enough sample of $f_{\mathbf{X}}(\mathbf{x})$:

$$\int_{\mathcal{D}_{\mathbf{X}}} h(\mathbf{x}) f_{\mathbf{X}}(\mathbf{x}) d\mathbf{x} = \mathbb{E} [h(\mathbf{x})] \approx \frac{1}{N} \sum_{i=1}^N h(\mathbf{x}^{(i)}), \quad (5.7)$$

where $\mathbf{x}^{(i)} \in \mathcal{X} = \{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}\}$ is a set of I.I.D. realizations of the input random vector \mathbf{X}

Advantages and disadvantages of both approaches in conjunction with PCE in the context of uncertainty quantification are explored in detail in Torre et al. (2019a), and are outside the scope of this report.

Calculation of the PCE coefficients

Three main methodologies are available for the calculation of PCE coefficients, namely intrusive Galerkin methods (Xiu and Karniadakis, 2002), stochastic collocation and sparse grids (Eldred, 2009), and sparse regression (Lüthen et al., 2021). Due to their flexibility and fast convergence with the number of runs of the computational model, we adopt the sparse-adaptive PCE framework from Blatman and Sudret (2011); Lüthen et al. (2022), as implemented in the UQLab software (Marelli and Sudret, 2014; Marelli et al., 2024).

5.3 Numerical experiments

We present in this section numerical results on three numerical and analytical examples meant to represent the following scenarios:

- An FEM model of a truss structure under random load, with both independent and correlated inputs, showcasing the role of correlation in the integration performance.
- The same 10D analytical model of *Test case #2* in Section 4.3.1, which showcases a problem without unimportant variables and strong interactions.
- A variation of the latter, but in a different configuration with reduced interactions and unimportant variables.

Every numerical experiment is replicated $N_R = 20$ times with different random seed to assess the robustness of the algorithm, and the results are presented as box plots for readability.

5.3.1 FEM model of a truss structure

Our first application is a FEM model of the truss structure under uncertain loads depicted in Figure 5.1 (modified from Dubourg (2011)). This problem includes $M = 10$ input variables, representing the uncertainty in the geometrical and material properties of the horizontal (A_1 and E_1) and diagonal (A_2 and E_2) elements, as well as the six loads P_1, \dots, P_6 , with lognormal marginal distributions and coupled by a Gaussian copula with $\rho_{ij} = 0.6$, $i \neq j$. For reference, the full list of marginal distributions in this numerical experiment are reported in Table 5.1.

Table 5.1: Joint distribution of the input parameters for the 23-bar truss example in Section 5.3.1

Name	Distribution	Moments	Notes
$E_{1,2}$	Lognormal	$\mu = 2.1 \times 10^{11}$, $\sigma = 2.1 \times 10^{10}$	
A_1	Lognormal	$\mu = 2. \times 10^{-3}$, $\sigma = 2.0 \times 10^{-4}$	Independent
A_2	Lognormal	$\mu = 1. \times 10^{-3}$, $\sigma = 1.0 \times 10^{-4}$	
$P_{1,\dots,6}$	Gumbel	$\mu = 5.0 \times 10^4$, $\sigma = 7.5 \times 10^3$	Gaussian copula with $\rho_{ij} = 0.6$




Figure 5.1: Finite elements model of a 23-bar truss structure used in the first numerical example, Section 5.3.1. The horizontal elements are highlighted in blue, while the diagonal ones are highlighted in magenta.

This experiment is intended to showcase and compare the performance of the different integration approaches described in section 5.1 on an actual engineering model with relatively high input dimension $M = 10$, some unimportant variables and a complex joint distribution of the input parameters.

We compare here three different integration approaches, for an increasing number of model evaluations:

- **LHS:** Direct Monte-Carlo integration from a Latin hypercube sample (McKay et al., 1979) of the input distributions (Eq. (5.7));
- **Spectral PCE:** Spectral integration according to Eq. (5.6), where a PCE is constructed on a reduced set of input variables after decorrelating the inputs with Rosenblatt transform. The experimental design is an LHS sample from the joint input distribution $f_{\mathbf{X}}(\mathbf{x})$;
- **Independent PCE:** Surrogate-based integration, by using a PCE constructed on the same experimental design as the previous case, but assuming the input distributions independent (following the recommendations in Torre et al. (2019a)). The integration is performed directly by substituting the computational model $h(\mathbf{x})$ in Eq. 5.7 with its PCE surrogate from Eq. 5.4.

All estimates are compared to a reference MCS mean estimator $\mu_{ref} = \frac{1}{N_{ref}} \sum_{i=1}^{N_{ref}} h(\mathbf{x}^{(i)})$, based on a large sample of size $N_{ref} = 10^8$.

The results of the three integration methods are reported in Figure 5.2. The box-plots highlight the robustness of the approach to the natural variability of the experimental design, as every calculation is repeated a total of $N_{rep} = 20$ times. The left panel of Figure 5.2 shows the convergence of the three methods to the mean value, while the right panel shows the relative absolute error of the integral, as given by:

$$\epsilon_R = \left| \frac{\hat{\mu}}{\mu_{ref}} - 1 \right|, \quad (5.8)$$

where $\hat{\mu}$ is the value estimated by the method under examination.

All methods clearly converge to the reference value when increasing number of model evaluations. However, direct MCS integration (LHS) converges distinctly slower than both PCE-based methods, and is particularly sensitive to random seed variability, as most of the box-plots show a much larger scatter around the reference value. Consistently with the findings in Torre et al. (2019a), the use of independent input variables followed by

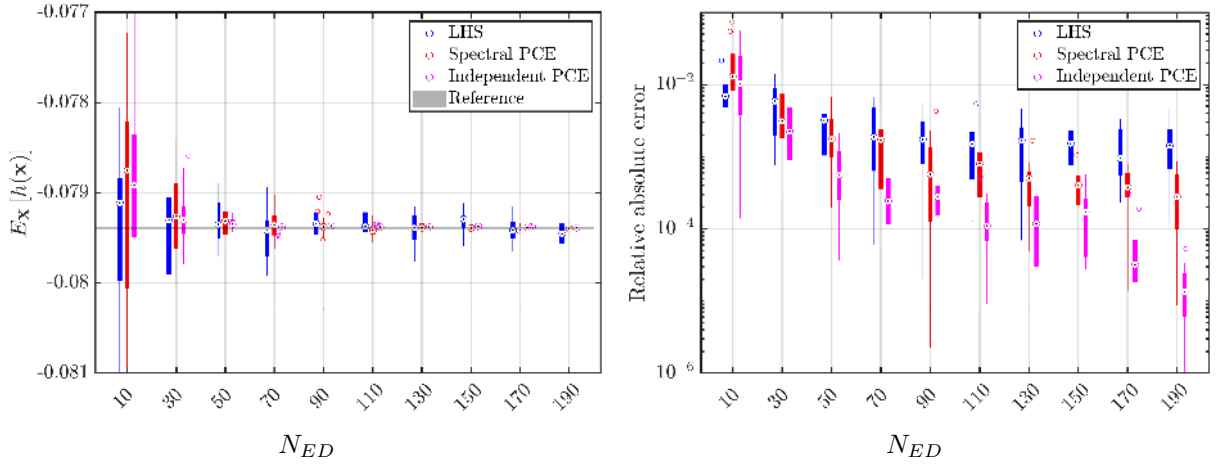


Figure 5.2: Comparison of different integration strategies for the finite-element truss structure in Section 5.3.1. Left panel: convergence of the integral value (mean) as a function of the number of model evaluations in the experimental design. Right panel: convergence of the relative mean absolute error for the same configurations.

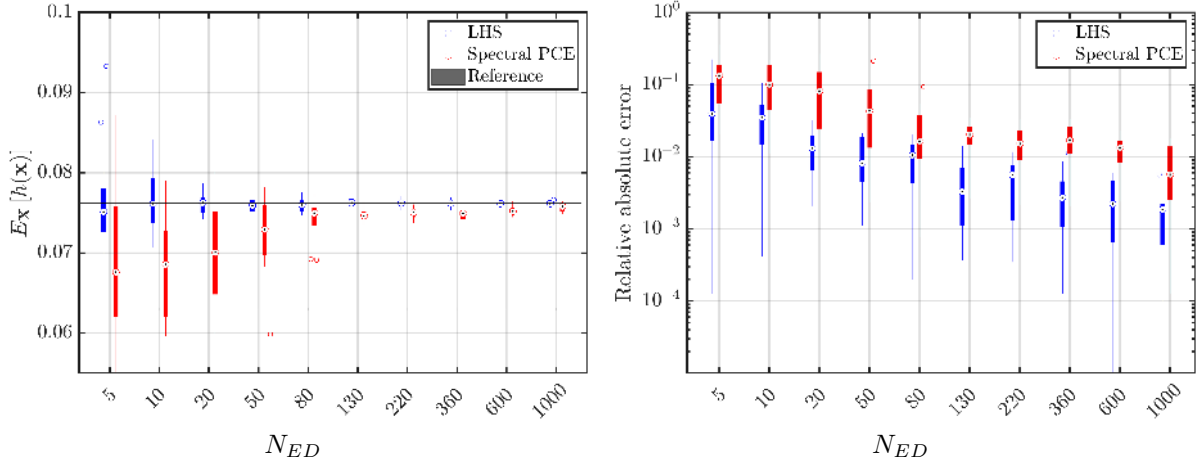


Figure 5.3: Comparison of different integration strategies for the finite-element truss structure in Section 5.3.2. Left panel: convergence of the integral value (mean) as a function of the number of model evaluations in the experimental design. Right panel: convergence of the relative mean absolute error for the same configurations.

numerical integration (*Independent PCE*), consistently outperforms both direct LHS sampling and the *Spectral PCE* approach. The difference appears to be more marked with higher numbers of model evaluations, up to multiple orders of magnitude for $N_{ED} = 190$.

5.3.2 High dimensional function with interactions

The second numerical experiment is a direct follow-up of Example #2 in Section 4.3.1, a high-dimensional function without unimportant variables, and with important interaction terms. In this case the input variables are independent, and therefore only results from the *spectral PCE* approach are reported. Figure 5.3 once again shows that both direct integration and spectral PCE-based integration converge to the reference mean value (based on a MCS integration with $N_{ref} = 10^8$). Nonetheless, PCE shows a biased behavior towards lower values of the integral, and an overall significantly worse performance in terms of relative absolute error. PCE integration also shows significantly higher inter-seed variability w.r.t. direct LHS integration.

As a final example, we slightly modify the previous case, by changing the weight of the input variables inside




Figure 5.4: Comparison of different integration strategies for the finite-element truss structure in Section 5.3.2. Left panel: convergence of the integral value (mean) as a function of the number of model evaluations in the experimental design. Right panel: convergence of the relative mean absolute error for the same configurations.

the reference function $g_2(\mathbf{x})$ (Table 4.2):

$$\tilde{g}_2(\mathbf{x}) = \prod_{i=1}^{10} \frac{|4x_i - 2| + a_i}{1 + a_i}, \quad \mathbf{a} = \{2, 5, 10, 20, 50, 100, 500, 500, 500, 500\}. \quad (5.9)$$

Additionally, the input distributions are considered independent standard normals. This function has fewer interactions and unimportant variables w.r.t. the previous version (although the interactive terms still account for more than 30% of the total variance of the model).

Figure 5.4 showcase the performance of spectral PCE integration. Overall, the performance is similar to that of the previous case, with direct integration performing significantly better than spectral integration. The difference in performance between direct and spectral integration, however, is significantly reduced in this experiment w.r.t. the previous case.

5.4 Synthesis

In this section we have introduced the use of a spectral expansion, namely sparse polynomial chaos expansions, to efficiently integrate expensive computational models over complex input distributions. While the methodology performs well in several realistic conditions typical of uncertainty quantification studies, its main limitation lies in the presence of strong interactions and high effective dimensionality, as it capitalizes on sparsity to achieve its efficiency.

6 Reliability and robustness analyses in the context of fatigue

6.1 Introduction

One of the main goals of this work is to evaluate the reliability of an OWT's monopile foundation w.r.t. fatigue solicitations. Let us recall the approach to assess fatigue damage over the structure's lifetime (see IEC-61400-1, 2019, Appendix H). First, the lifetime before decommissioning time $t_d \in \mathbb{R}_+$ forms the interval $T = [0, t_d]$, which can be used to define a probability space $(T, \mathcal{B}(T), \mathcal{U}(T) = \lambda/t_d)$. Then, for all $t \in T$ and assuming the probability space $(\Omega, \mathcal{A}, \mathbb{P})$, the random vector of the environmental conditions $\mathbf{X}(t, \cdot)$ is a measurable function on $(\Omega, \mathcal{A}) \rightarrow (\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$. Therefore, one can define the product probability space $(T \times \Omega, \mathcal{B}(T) \otimes \mathcal{A}, \mathcal{U}(T) \otimes \mathbb{P})$ and the random vector \mathbf{X} , which is a measurable function on $(T \times \Omega, \mathcal{B}(T) \otimes \mathcal{A}) \rightarrow (\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$. For a realization of \mathbf{X} , denoted by $\mathbf{X}(t^{(i)}, \omega^{(r)})$, and a given set of parameters $\mathbf{z} = (k_{\text{soil}}, \theta_{\text{yaw}}, \varepsilon)$ related to the system (defined in Table 3.2), one can perform a 10-minute Turbsim-DIEGO simulation and compute the corresponding cumulative damage (using Eq. (2.3)):

$$d_c^{10\text{min}} \left(\mathbf{X} \left(t^{(i)}, \omega^{(r)} \right) \mid \mathbf{Z} = \mathbf{z} \right). \quad (6.1)$$

In practice, T is often discretized into $N_T \in \mathbb{N}$ 10-minute intervals $\{t^{(i)}\}_{i=1}^{N_T}$ (IEC-61400-1, 2019, Appendix H). The 10-minute duration results from the typical wind spectral density (see Burton et al., 2021), presenting a “short-term” behavior (for turbulent wind with a return period below 10 minutes) and “long-term” behavior (otherwise, which is defined in Table 3.1). To cumulate the damage over the lifetime, one can write the sum of 10-minute damages, each averaged over $n_{\text{rep}} \in \mathbb{N}$ pseudorandom seed repetitions:

$$D(\mathbf{z}) = N_T \mathbb{E} [d_c^{10\text{min}}(\mathbf{X} | \mathbf{Z} = \mathbf{z})] \quad (6.2a)$$

$$\approx N_T \frac{1}{N_T n_{\text{rep}}} \sum_{i=1}^{N_T} \sum_{r=1}^{n_{\text{rep}}} d_c^{10\text{min}} \left(\mathbf{X} \left(t^{(i)}, \omega^{(r)} \right) | \mathbf{Z} = \mathbf{z} \right). \quad (6.2b)$$

In Section 4, KH was proposed as a method for given-data subsampling to propagate the uncertain environmental conditions on Teesside’s OWT model. This method showed equivalent performances to QMC for estimating the lifetime damage in Eq. (6.2b), while being more flexible than QMC. Considering the linear cumulative damage model typically used by the community (Miner’s rule), a damage value higher than one leads to fatigue crack initiation. The present section assesses the probability of such a rare event, considering both the environmental uncertainties (aggregated according to Eq. (6.2b)) and the uncertainties related to the system itself (described by the random vector $\mathbf{Z} \in \mathcal{D}_{\mathbf{Z}}$ with joint PDF $f_{\mathbf{Z}}$). Assuming that the critical damage D_{cr} is a random variable centered around one (equivalent to a “resistance” variable in the well-known “resistance-solicitation” paradigm in reliability analysis Lemaire et al., 2009), this failure probability is written as:

$$p_f = \int_{\mathcal{D}_{\mathbf{Z}} \times \mathbb{R}^+} \mathbb{1}_{\{D(\mathbf{z}) \geq D_{\text{cr}}\}} f_{\mathbf{Z}}(\mathbf{z}) f_{D_{\text{cr}}}(d_{\text{cr}}) d\mathbf{z} dd_{\text{cr}}. \quad (6.3)$$

Less information is available to define the probabilistic model of the system uncertainties than the environmental ones. Therefore, the robustness analysis of the failure probability w.r.t. the probabilistic model \mathbf{Z} should be studied. To do so, a perturbation-based approach using the *perturbed-law based sensitivity indices* (PLI), originally introduced by (Lemaître et al., 2015), is used in this section. However, the reliability and robustness analysis requires a number of evaluations of the Turbsim-DIEGO simulator $D(\cdot)$ that would be prohibitive without the use of surrogate models.

The present section is structured as follows: Subsection 6.2 presents the construction of a surrogate model of $D(\cdot)$, then Subsection 6.3.1 analyses the reliability of the monopile foundation of Teesside’s turbine for a nominal distribution of \mathbf{Z} , and finally Subsection 6.3.2 proposes a robustness analysis of p_f by perturbing the laws of both \mathbf{Z} and D_{cr} . Note that the contributions from this section were adapted from the PhD thesis of Fekhari (2024).

6.2 Surrogate modeling for reliability analysis

The prohibitive computational cost of the function $D(\cdot)$ requires fitting a surrogate model. This section presents the specific high-performance computer (HPC) wrapper developed for this application and how it is used to build a learning set for a GP regression model.

6.2.1 High-performance computer evaluation

The wrapper of the numerical chain including Turbsim and Diego (illustrated in Fig. 2.1) for reliability analysis has a nested double loop structure. The outer loop pilots the realizations of \mathbf{Z} while the inner loop concerns the environmental conditions and their repetitions for n_{rep} different pseudorandom seeds. Since ε is a coefficient applied to the damage as post-processing (see Subsection 3.4), one can fix it at a nominal value $\varepsilon = \varepsilon_0 = 1$. At this stage, the goal is to approximate $D(\cdot)$ for any value of $\mathbf{z} = (k_{\text{soil}}, \theta_{\text{yaw}}, \varepsilon_0)$. Using a KH design to explore the environmental conditions as discussed in Section 4, this approximation is written as:

$$D(\mathbf{z}) \approx D^{\text{KH}}(\mathbf{z}) = N_T \frac{1}{n_{\mathbf{X}} n_{\text{rep}}} \sum_{i=1}^{n_{\mathbf{X}}} \sum_{r=1}^{n_{\text{rep}}} d_c^{10\text{min}} \left(\mathbf{x}^{(i)}, \omega^{(r)} | \mathbf{Z} = \mathbf{z} \right), \quad (6.4)$$

where a KH design with size $n_{\mathbf{X}} \in \mathbb{N}$ is denoted by $\{\mathbf{x}^{(i)}\}_{i=1}^{n_{\mathbf{X}}}$. According to the convergence results obtained in Section 4, the KH size is fixed at $n_{\mathbf{X}} = 200$ and the repetitions at $n_{\text{rep}} = 11$, which implies a total of 2200 Turbsim-DIEGO simulations per evaluation of the function $D^{\text{KH}}(\cdot)$. In this setup, the CRONOS HPC from EDF

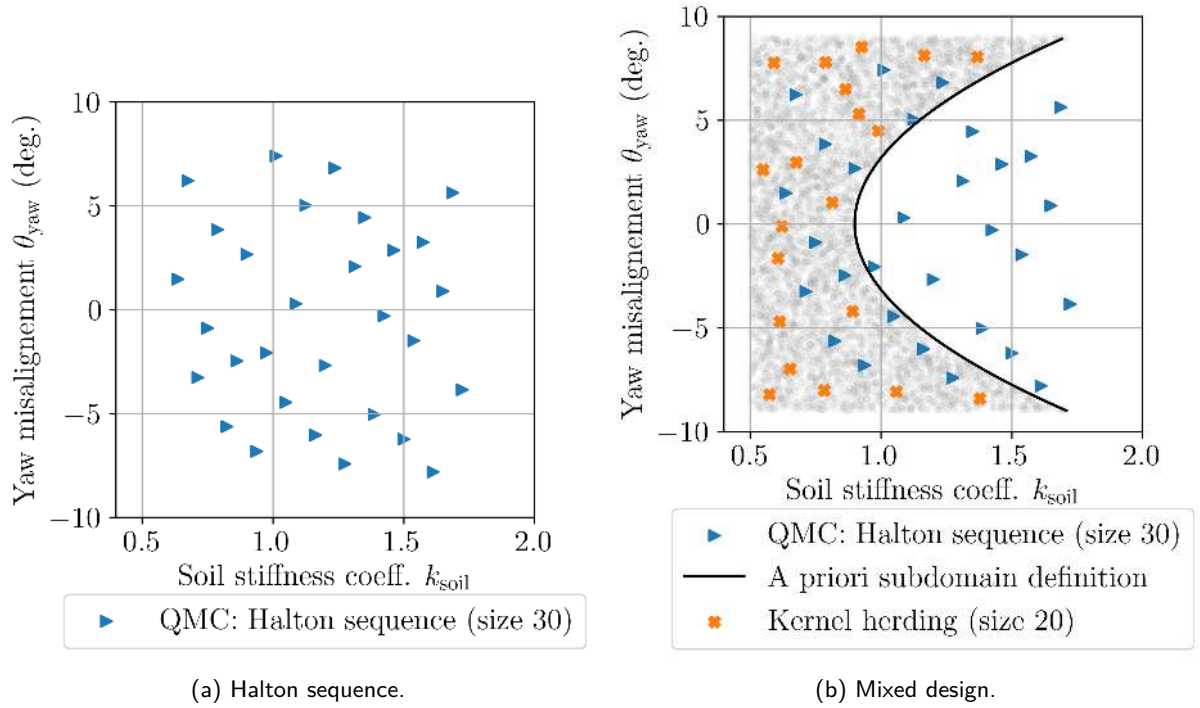


Figure 6.1: Learning set of the mean damage surrogate model. A Halton sequence is first built (in blue) and completed by a set KH-generated points (in orange) in a subdomain defined a priori (in gray).

R&D allows us to simultaneously perform those 2200 simulations in parallel. The random variable associated with the S-N curve uncertainty is introduced later as a product factor of Eq. (6.4). Note that the cumulative damage studied is actually the maximum value of $D^{KH}(\cdot)$ over the discretized azimuth angles (illustrated in Fig. 2.3), at the mudline level.

6.2.2 Design of experiments

To build a learning set, a space-filling design of experiments is created on the joint domain of (K_{soil}, Θ_{yaw}) . This design was first composed of 30 points generated by a Halton sequence (illustrated in Fig. 6.1.a) which was the most space-filling sequential method (compared to other QMC sequences and KH). Its evaluation and analysis showed that the highest damage values are the result of high absolute values of Θ_{yaw} or low values of K_{soil} . Therefore, the design was completed in a second phase by 20 points targeting these areas by applying kernel herding to a subdomain defined a priori (see the candidate set represented by the gray points in Fig. 6.1.b). Finally, the learning set is the union of the two complementary designs, later referred to as the “mixed design” (see Fig. 6.1.b). This mixed design, denoted by \mathbf{Z}_{n_Z} , has a size of $n_Z = 50$ points, which represents over 10^5 Turbsim-DIEGO simulations (each requiring around 45 minutes of CPU time). Fig. 6.2 shows the lifetime cumulated damage evaluated on the mixed design (with normalized values corresponding to the color scale).

6.2.3 Gaussian process regression

A GP regression with Matérn 5/2 and constant trend is fitted on the mixed design \mathbf{Z}_{n_Z} according to the Kriging equations, recalled in Eq. (4.16). The resulting surrogate model $\tilde{D}_{\varepsilon_0} : \mathbb{R}^2 \rightarrow \mathbb{R}$, which approximates $D^{KH}(\cdot, \cdot, \varepsilon_0)$, is represented by the blue three-dimensional surface in Fig. 6.3, and its learning set \mathbf{Z}_{n_Z} by the black crosses. A complementary visualization of this surrogate is proposed for cross-sections with fixed values $\theta_{yaw} = 0$ on Fig. 6.4.a, and $k_{soil} = 1$ on Fig. 6.4.b. On these two figures, learning points are plotted in a grayscale and the surrogate model in a blue scale. The darker the shade, the closer to the cross-section the points are.

One can first notice that the damage is not symmetric w.r.t. the yaw misalignment θ_{yaw} . When the nacelle perfectly aligns itself with the wind direction, this angle is equal to zero. For this bottom-fixed OWT, introducing a yaw error of the same amplitude has more impact in one direction than the other. In the meantime, the soil stiff k_{soil} has, as expected, a monotonic influence on the damage.

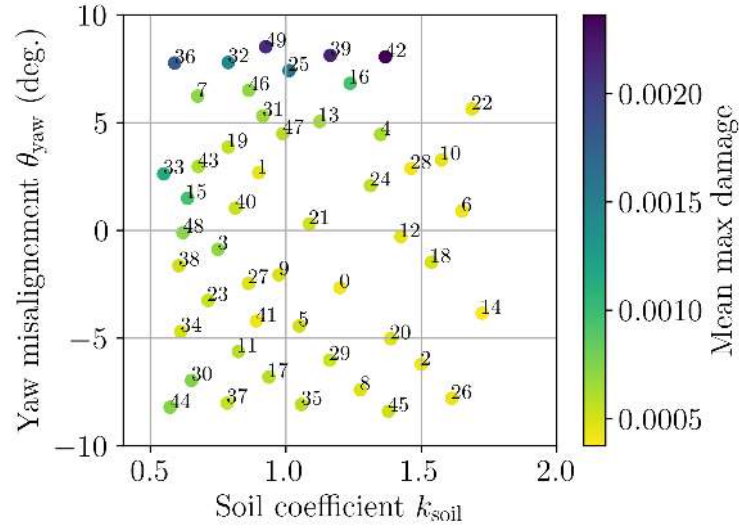


Figure 6.2: Mean damage evaluated on the mixed design illustrated in Fig. 6.1.b.

To validate this surrogate model, a leave-one-out (LOO) procedure is realized. Fig. 6.5.b represents the LOO squared-residuals at each point of the design (with values corresponding to the color scale). High residual values are mostly due to the strong nonlinearity of the code in some areas (as revealed by the cross-section in Fig. 6.4.b). Fig. 6.5.a shows the quantile-quantile plot comparing the LOO predictions with the lifetime damage evaluations on the learning set. The general coefficient of predictivity of $\hat{Q}_{LOO}^2 = 0.72$ is considered acceptable in this small data context, especially as the LOO procedure was shown to underestimate the true performance metric in Fekhari et al. (2023). However, the surrogate could be enhanced by adding points to the learning set in areas with high nonlinearities.

Remark 5 Active learning methods for reliability assessment could be a great option for a such costly computer model. However, the stochasticity of the function would disturb the learning criterion. Since the nonlinearities seem restricted to a small area, the present approach should be more robust. Another approach could be to fit a stochastic surrogate (Binois et al., 2019; Baker et al., 2022; Zhu, 2022) on a learning set before averaging on the pseudorandom seed repetitions.

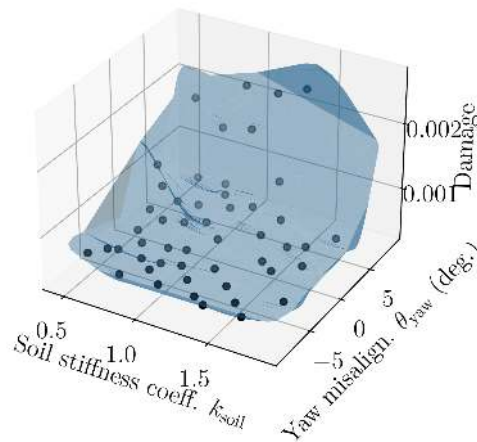


Figure 6.3: Three-dimensional plot of the surrogate model $\tilde{D}_{\varepsilon_0}$ (in blue) and learning set (in black).

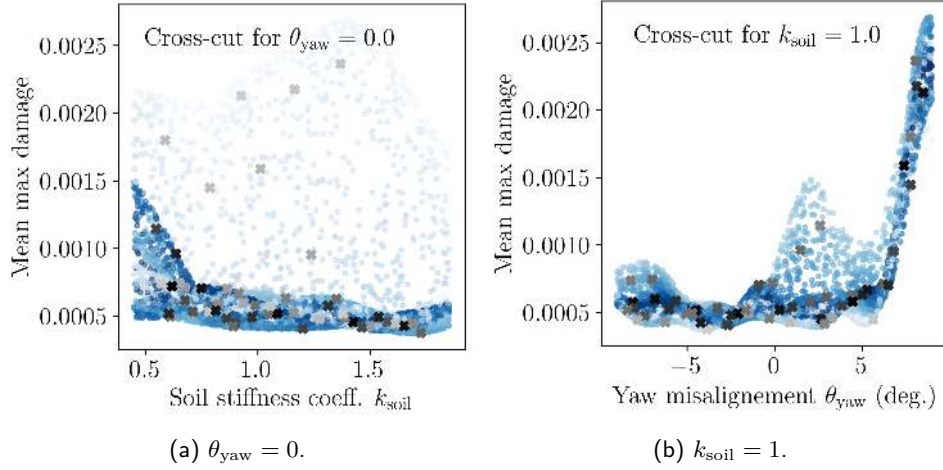


Figure 6.4: Cross-section of the surrogate model $\tilde{D}_{\varepsilon_0}$ (in shades of blue) for given values of K_{soil} and Θ_{yaw} . The darker the shade, the closer to the cross-section.

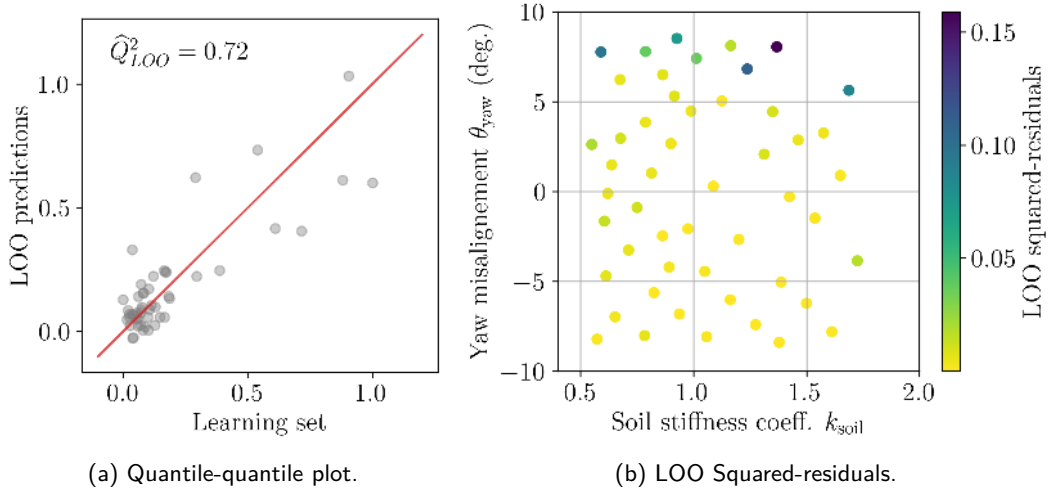


Figure 6.5: Leave-one-out validation results of the surrogate model $\tilde{D}_{\varepsilon_0}$.

6.3 Reliability and robustness analysis

In the wind energy industry, acceptable risk levels for fatigue are defined by the standards. The target failure probability of the order of magnitude of 10^{-4} over the last year of exploitation is recommended by the IEC-61400-1 (2019) while the recommendations of other standards are reviewed by Wang et al. (2022). Nielsen and Sørensen (2021) discussed the relevance of this risk level, defined from an economic point of view, and offered quantitative guidelines for lifetime extension. In this section, a reliability analysis is conducted using the surrogate model of the lifetime cumulative damage previously constructed in Subsection 6.2.3. Afterwards, the robustness of this reliability analysis is studied by applying perturbations to the input distributions and computing the perturbed-law based sensitivity indices (PLI) initially proposed by Lemaître et al. (2015).

6.3.1 Nominal reliability analysis

The S-N curve uncertainty can be expressed as a simple coefficient ε (see Subsection 3.4), which is applied to the lifetime cumulative damage as follows:

$$D(k_{soil}, \theta_{yaw}, \varepsilon) = \frac{1}{\varepsilon} D(k_{soil}, \theta_{yaw}, \varepsilon_0). \quad (6.5)$$

Therefore, the surrogate model of the lifetime cumulative damage $\tilde{D}_{\varepsilon_0}$ is modified without updating its learning such as:

$$\tilde{D}(\mathbf{z}) = \tilde{D}(k_{\text{soil}}, \theta_{\text{yaw}}, \varepsilon) = \frac{1}{\varepsilon} \tilde{D}_{\varepsilon_0}(k_{\text{soil}}, \theta_{\text{yaw}}). \quad (6.6)$$

The probability introduced in Eq. (6.3) then becomes:

$$p'_f = \int_{\mathcal{D}_{\mathbf{z}}} \mathbb{1}_{\{\tilde{D}(\mathbf{z}) \geq D_{\text{cr}}\}} f_{\mathbf{z}}(\mathbf{z}) d\mathbf{z}. \quad (6.7)$$

Table 6.1 presents the estimates of this quantity by various methods (FORM, FORM-IS and SS, see Morio and Balesdent, 2015) and for two different distributions of the resistance variable. D_{cr} either follows a lognormal distribution (which has a short left tail), or a truncated normal distribution on $[0, +\infty[$ (with a heavier left tail). All methods deliver similar values of p_f even if FORM-IS converges faster than SS when comparing the coefficients of variation (COV). The adequation of FORM with the simulation methods reveals that the LSF in this case is almost linear. This linearity is mostly due to the nature of the LSF, following a resistance-solicitation paradigm as introduced earlier. The probabilities are, as expected, much lower under the hypothesis of lognormal distribution for D_{cr} than for a normal distribution. However, this significant difference questions the robustness of this result w.r.t. the probabilistic model of both D_{cr} and \mathbf{Z} .

Table 6.1: Nominal reliability analysis (size $N = 5 \times 10^4$ for IS and SS. $p_0 = 0.1$ for SS).

Reliability method	$D_{\text{cr}} \sim \mathcal{LN}(\mu_{D_{\text{cr}}} = 1, \sigma_{D_{\text{cr}}}^2 = 0.3^2)$		$D_{\text{cr}} \sim \mathcal{N}(\mu = 1, \sigma^2 = 0.3^2)$	
	\hat{p}'_f	$\widehat{\text{COV}}$	\hat{p}'_f	$\widehat{\text{COV}}$
FORM	9.87×10^{-13}	–	3.35×10^{-6}	–
FORM-IS	9.84×10^{-13}	1%	3.36×10^{-6}	1%
SS	9.46×10^{-13}	7%	3.50×10^{-6}	4%

6.3.2 Robustness analysis using perturbed-law sensitivity indices

The method of Lemaître et al. (2015), later called *perturbed-law based sensitivity indices* (PLI) by Sueur et al. (2017) relies on perturbing input densities. The goal is to assess the robustness of a quantity of interest (here, a failure probability) w.r.t. these perturbations. This type of index is for example used in studies of thermal-hydraulic numerical models for nuclear safety (Iooss et al., 2022).

Assuming a random variable $Z_j \sim f_j \in \mathcal{D}_{Z_j}$ with mean $\mathbb{E}[Z_j] = \mu$ and variance $\text{Var}(Z_j) = \sigma^2$, the strategy is to find the “closest” distribution $f_{j\delta}$ under the constraint of moment perturbation of magnitude δ . The notion of proximity between distributions is quantified by Lemaître et al. (2015) in terms of Kullback–Leibler (KL) divergence. For example, a relative mean perturbation is defined as:

$$f_{j\delta} = \arg \min_{\substack{\pi \in \mathcal{P}, \\ \text{s.t.}, \mathbb{E}_{\pi}[Z_j] = \mathbb{E}_{f_j}[Z_j](1+\delta)}} \text{KL}(\pi || f_j), \quad (6.8)$$

where $\delta \in \mathbb{R}$ denotes the relative perturbation, and \mathcal{P} is a family of distributions.

Note that the perturbed distribution $f_{j\delta}$ might not belong to the parametric family of f_j . This is typically the case for bounded distributions (e.g., when perturbing the mean of a uniform distribution). Separate approaches exist to define perturbations for the problem studied, such as Lemaître et al. (2015), or Gauchy et al. (2022). However, to ease the computation in the following, the perturbations will conserve the initial parametric family (which seems reasonable for distributions in the exponential family).

The adapted expression of the PLI used hereafter (Gauchy et al., 2022) reflects the relative impact of a perturbation on the quantity of interest:

$$\text{PLI}(f_{j\delta}) = \frac{p_{f,j\delta} - p_f}{p_f}, \quad (6.9)$$

where $p_{f,j\delta}$ is the probability obtained when injecting $f_{j\delta}$ in Eq. (6.7). Although this PLI expression is simple to manipulate, it is not symmetric, i.e., when $p_{f,j\delta} \rightarrow 0$, $\text{PLI}(f_{j\delta}) \rightarrow -1$ and when $p_{f,j\delta} \rightarrow 1$, $\text{PLI}(f_{j\delta}) \rightarrow 1/p_f$.




Figure 6.6: Perturbations in terms of standard deviation of a lognormal distribution (left) and a truncated normal distribution (right).

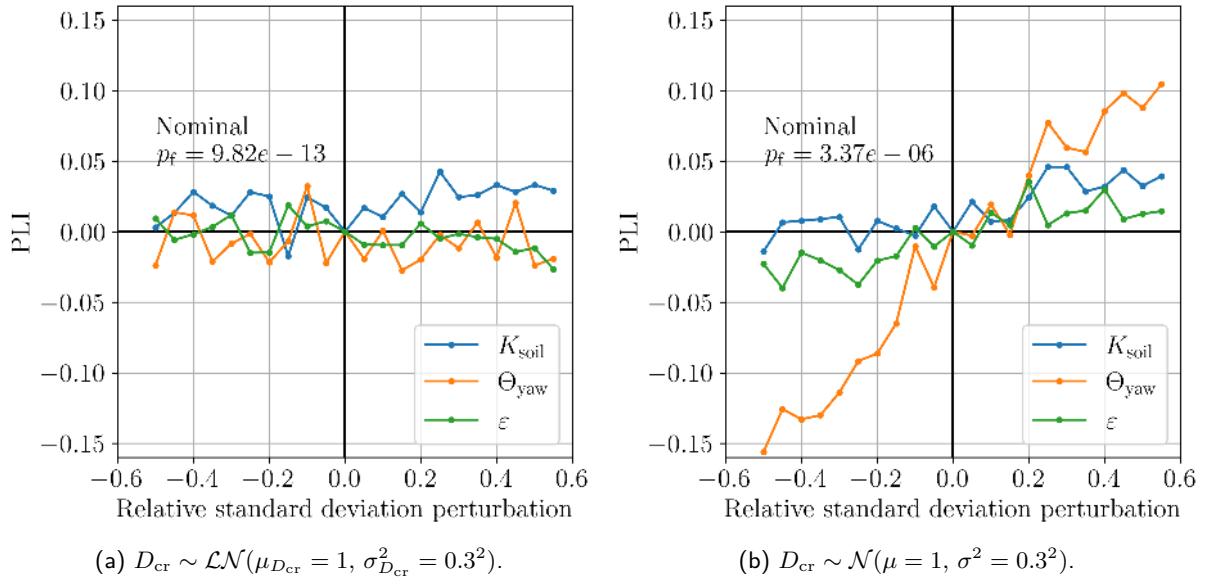


Figure 6.7: Perturbed-law based indices for relative perturbations of the standard deviations of $(K_{soil}, \Theta_{yaw}, \varepsilon)$. The failure probabilities studied are each estimated by FORM-IS method on the surrogate model (with sample size $N = 5 \times 10^4$).

In the following, each variable in \mathbf{Z} is perturbed one by one in terms of relative standard deviation, such that $\sigma_{j\delta} = \sigma_j(1 + \delta)$. The illustration of such perturbations is given in Fig. 6.6, for distributions of K_{soil} on the left, and of Θ_{yaw} on the right. This strategy assumes that the analyst has enough information to determine the mean of the variables Z_j .

The resulting PLI are presented in Fig. 6.7 for relative perturbations of the standard deviations of $(K_{soil}, \Theta_{yaw}, \varepsilon)$. Each failure probability is independently estimated by FORM-IS. In the hypothesis of a normal D_{cr} , the most important variable is Θ_{yaw} , while the fluctuations are quite stable in the hypothesis of a lognormal D_{cr} .

When perturbing the standard deviation of the resistance variable D_{cr} , the same phenomenon is witnessed in Fig. 6.8. The perturbations have nearly no consequences, assuming that $D_{cr} \sim \text{LogNormal}$, but a lot of influence when $D_{cr} \sim \text{Normal}$. As a perspective, this study could be completed by a joint perturbation of both the standard deviation and mean of the resistance variable.

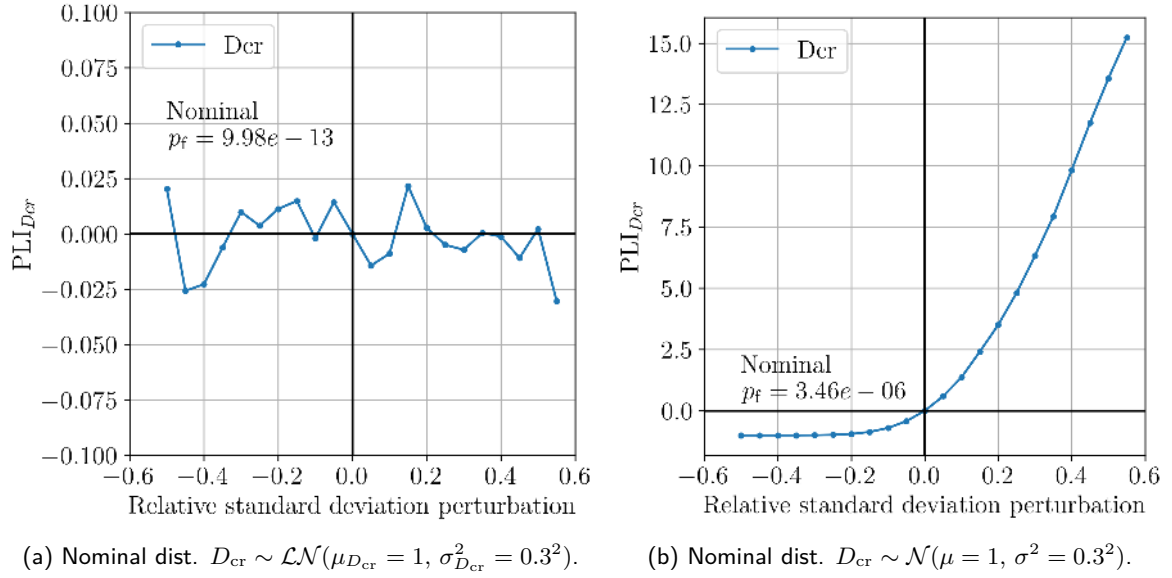


Figure 6.8: Perturbed-law based indices for relative perturbations of the standard deviation of D_{cr} . The failure probabilities studied are each estimated by FORM-IS method on the surrogate model (with sample size $N = 5 \times 10^4$).

6.4 Synthesis

In this section, a fully probabilistic reliability analysis was performed on the monopile foundation of an OWT located in Teesside, UK. This first reliability study required a significant computational effort (i.e., over 10^5 Turbsim-DIEGO simulations), made possible with the development of a tailored wrapper deployed on a HPC facility. These simulations served the construction of a surrogate model emulating the lifetime damage function.

The second part of this section addressed the estimation of a failure probability using this surrogate model. A general conclusion is that this probability is highly dependent on the probabilistic model describing the critical damage. To challenge the modeling assumptions of the critical damage and the system variables, a robustness analysis was realized. It consists in studying the impact of perturbing the input distributions on an output quantity. The robustness of the failure probability was evaluated with the formalism of the perturbed-law based indices, introduced by Lemaître et al. (2015). This additional study mostly confirmed the importance of the critical damage definition.

From an industrial point of view, the failure probabilities obtained are much lower than the risk levels targeted in current standards (Wang et al., 2022), which is in favor of lifetime extension. However, the results of this demonstrative study may be questionable with respect to several effects that were not considered:

- The periods in parked position which could significantly increase the damage (see Velarde et al., 2020). Several studies showed that the aerodynamic damping created by the rotation of the turbine reduces tower vibrations, and therefore fatigue (Liu et al., 2017);
- The rapid transition phases occasioned by emergency stopping, which might increase fatigue;
- The early damage produced during the installation of the monopile foundations by hydraulic impact piling (i.e., hammering);
- The stress-concentration resulting from soldering the structure.
- The consideration of wake effects at a farm scale, affecting the environmental solicitations on the turbines.

The next section reproduces the previous results, for wind conditions pertubated by the wake effect inside the Teesside farm. This complementary study aims at quantifying the impact of the wake on the cumulated fatigue damage.

As a methodological perspective, a stochastic surrogate model could be built by considering all the damages without averaging them over the environmental repetitions. The risk assessment would then be conducted on a

stochastic function. In this particular context, quantile estimation was studied by [Browne et al. \(2016\)](#), however, rare event estimation on stochastic functions remains an open question recently discussed by [Pires et al. \(2023\)](#).

7 Reliability and robustness analyses including the wake

7.1 Introduction

The wake arises from extracting kinetic energy from the wind, leading to a decrease in wind speed and an increase in turbulence downstream of the turbines. Inside a wind farm, it mostly depends on the turbines' layout, the ambient wind speed, and the ambient turbulence intensity. The resulting heterogeneous wind field in a wind farm can be simulated by numerical models with different fidelities (see e.g., [Doubrawa et al., 2020](#)).

This work exploits the results from another HIPERWIND deliverable, which propagated Teesside's ambient metocean conditions through a simplified wake model ([Ardillon et al., 2023](#)). Such wake models (sometimes called "dynamic wake meandering", or "engineering" models, see the seminal paper of [Jensen, 1983](#)) are used to simulate the wind speed deficit and the added turbulence at each turbine. In [Ardillon et al. \(2023\)](#), the ambient metocean distribution from Teesside has been propagated through a dynamic wake meandering model using a Monte Carlo sample (with size $N = 8000$), as a result, one can retrieve a wake-perturbed wind distribution at each turbine. Having different wind distributions naturally impacts the loading and should be considered during fatigue assessment at a farm scale. Note that this heterogeneity is mostly considered in international standards via empirical coefficients (also called "effective turbulence"). [Doubrawa et al. \(2023\)](#) compares the use of the effective turbulence approach with dynamic wake meandering models and studies their impact on loading.

As illustrated in Section 6, assessing the fatigue reliability on a turbine represents a significant computational effort. However, as a consequence of the wake modeling, each turbine presents a different wind distribution, which implies repeating a fatigue assessment for every turbine. To make this computation tractable at the farm scale, [Lovera et al. \(2023\)](#) proposed a way to build clusters of turbines similarly affected by the wake. The MMD is then used as a statistical dissimilarity measure to compare the perturbed distributions induced by the wake and to perform the clustering. Finally, fatigue assessment could be computed on a few turbines, each representing a cluster of turbines facing similar wake-modified wind loading.

After assessing the fatigue reliability of a WT tower w.r.t. the ambient metocean conditions, this section additionally examines the effect of the wake on fatigue reliability. Because of the massive number of TurbSim-DIEGO simulations required by the reliability assessment in Section 6 (over 10^5 simulations, each requiring around 45 minutes of CPU time), the fatigue reliability analysis is only repeated for one wake-affected metocean distribution. Naturally, the turbine most affected by the wake is the one where this additional study is carried out. According to the map of dissimilarity measures between wake-affected and ambient conditions, illustrated in Fig. 7.1, WT14 is the most affected by the wake. This hypothesis seems reasonable, considering its central location within the farm.

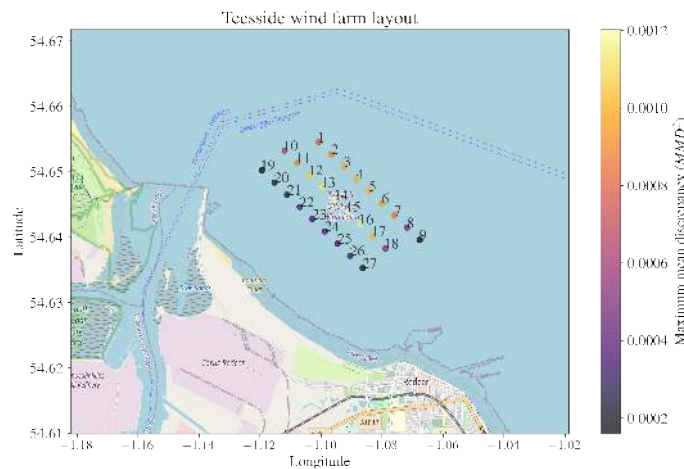


Figure 7.1: Dissimilarity measure between Teesside's ambient metocean distribution and the wake-perturbed distributions at each wind turbine (see [Lovera et al., 2023](#) for further details)

First, let us denote by \mathbf{X}_{w14} the wake-perturbed metocean random vector at the wind turbine #14. Then, similarly to the formalism introduced in Eq. (6.2b), one can defined the wake-perturbed cumulative damage over 10 minute as:

$$D_{w14}(\mathbf{z}) = N_T \mathbb{E} [d_c^{10\min}(\mathbf{X}_{w14} | \mathbf{Z} = \mathbf{z})], \quad (7.1)$$

this quantity can be estimated by KH, identically to Eq. (6.4), as:

$$D_{w14}(\mathbf{z}) \approx D_{w14}^{KH}(\mathbf{z}) = N_T \frac{1}{n_{\mathbf{X}} n_{\text{rep}}} \sum_{i=1}^{n_{\mathbf{X}}} \sum_{r=1}^{n_{\text{rep}}} d_c^{10\min}(\mathbf{x}_{w14}^{(i)}, \omega^{(r)} | \mathbf{Z} = \mathbf{z}), \quad (7.2)$$

where a KH design with size $n_{\mathbf{X}} \in \mathbb{N}$ is denoted by $\{\mathbf{x}_{w14}^{(i)}\}_{i=1}^{n_{\mathbf{X}}}$. Once again, the KH size is fixed at $n_{\mathbf{X}} = 200$ and the repetitions at $n_{\text{rep}} = 11$, which implies a total of 2200 Turbsim-DIEGO simulations per evaluation of the function $D^{KH}(\cdot)$. Note that the kernel herding sample is directly built using the Monte Carlo evaluations of the wake model (with size $N = 8000$) as candidate set (see Lovera et al., 2023 for further details).

7.2 Surrogate model for reliability analysis

In the same fashion as the expression in Eq. (6.5), the uncertainty related to the S-N curve ε is simply introduced by a coefficient applied at the post-processing phase. Therefore, considering a nominal value $\varepsilon = \varepsilon_0 = 1$, one can express the wake-perturbed cumulative damage as:

$$D_{w14}(k_{\text{soil}}, \theta_{\text{yaw}}, \varepsilon) = \frac{1}{\varepsilon} D_{w14}(k_{\text{soil}}, \theta_{\text{yaw}}, \varepsilon_0) \approx \frac{1}{\varepsilon} D_{w14}^{KH}(k_{\text{soil}}, \theta_{\text{yaw}}, \varepsilon_0). \quad (7.3)$$

Once again, a GP regression $\tilde{D}_{\varepsilon_0, w14} : \mathbb{R}^2 \rightarrow \mathbb{R}$ based on the same input design of experiments $\mathbf{Z}_{n_{\mathbf{Z}}}$ (with a size of $n_{\mathbf{Z}} = 50$), and same tuning as in subsection 6.2 is built to emulate the computationally costly function $D_{w14}^{KH}(\cdot, \cdot, \varepsilon_0)$. After analyzing the results presented in Fig. 7.2, and Fig. 7.3, one can notice that the evaluation of the function D_{w14}^{KH} over the design $\mathbf{Z}_{n_{\mathbf{Z}}}$ presents similar trends as the results in the ambient case (illustrated in Fig. 6.3). However, the cumulative damage values appear to be smaller in the wake-perturbed case of turbine 14 (see Fig. 7.2) than for the ambient case (see Fig. 6.2).

The new surrogate model is again validated by LOO, allowing us to estimate its predictivity coefficient $\hat{Q}_{\text{LOO}}^2 = 0.92$, and to represent its quantile-quantile plot in Fig. 7.5.a. The improvement in terms of predictivity coefficient can be explained by analyzing jointly the map of LOO squared residuals in Fig. 7.5 and the cross-section visualizations of the learning set in Fig. 7.4. In the previous LOO squared residuals map, most of the underperformance was due to the points on the upper border of the domain in Fig. 6.4. While the new LOO squared residuals map in Fig. 7.5.b presents the same pattern, the values of the residuals in this area are significantly lower. This might be explained by the stronger regularity of the wake-perturbed function which appears clearly in the cross-section represented in Fig. 7.4.b.




Figure 7.2: Mean damage evaluated on the mixed design illustrated in Fig. 6.1.b.




Figure 7.3: Three-dimensional plot of the surrogate model $\tilde{D}_{\varepsilon_0}$ (in blue) and learning set (in black).

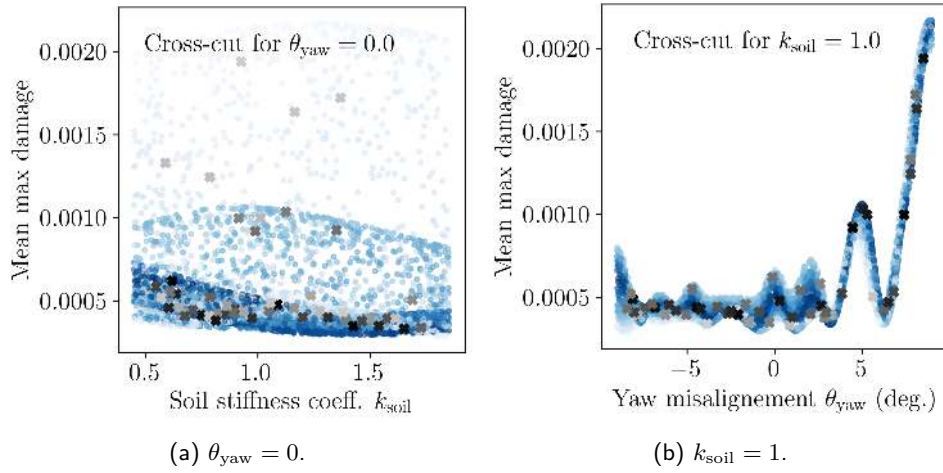


Figure 7.4: Cross-section of the surrogate model $\tilde{D}_{\varepsilon_0}$ (in shades of blue) for given values of K_{soil} and Θ_{yaw} . The darker the shade, the closer to the cross-section.

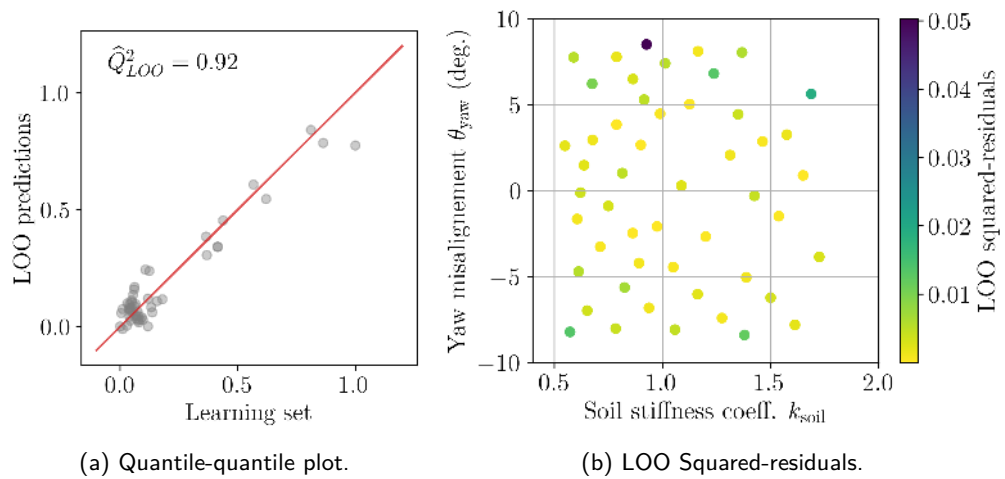


Figure 7.5: Leave-one-out validation results of the surrogate model $\tilde{D}_{\varepsilon_0}$.

7.3 Nominal reliability and robustness analysis

In this section, the reliability and robustness analysis conducted in Section 6 are reproduced using the surrogate model $\hat{D}_{\varepsilon_0, w14}$, built w.r.t. a wake-perturbed metocean distribution at the turbine #14. For the reliability analysis, Table 7.1 presents the updated values of failure probability for different methods. One can notice that the values of failure probability estimated in the wake-perturbed case are very close to the ones in the ambient case (presented in Table 6.1). This results may be counter-intuitive as the wake-perturbed cumulative damages in Fig. 7.2 were lower than the ambient ones in Fig. 6.2. However, this result can be explained by the prime importance of the resistance variable in this reliability analysis. For the robustness analysis, the PLI values obtained in the wake-perturbed case in Fig. 7.6 and Fig. 7.7 are once again similar to the ones from Section 6.

Table 7.1: Nominal reliability analysis (size $N = 5 \times 10^4$ for IS and SS. $p_0 = 0.1$ for SS).

Reliability method	$D_{cr} \sim \mathcal{LN}(\mu_{D_{cr}} = 1, \sigma_{D_{cr}}^2 = 0.3^2)$		$D_{cr} \sim \mathcal{N}(\mu = 1, \sigma^2 = 0.3^2)$	
	\hat{p}_f'	\widehat{cov}	\hat{p}_f'	\widehat{cov}
FORM	9.97×10^{-13}	–	2.84×10^{-6}	–
FORM-IS	9.88×10^{-13}	1%	2.87×10^{-6}	1%
SS	1.08×10^{-12}	7%	2.84×10^{-6}	4%

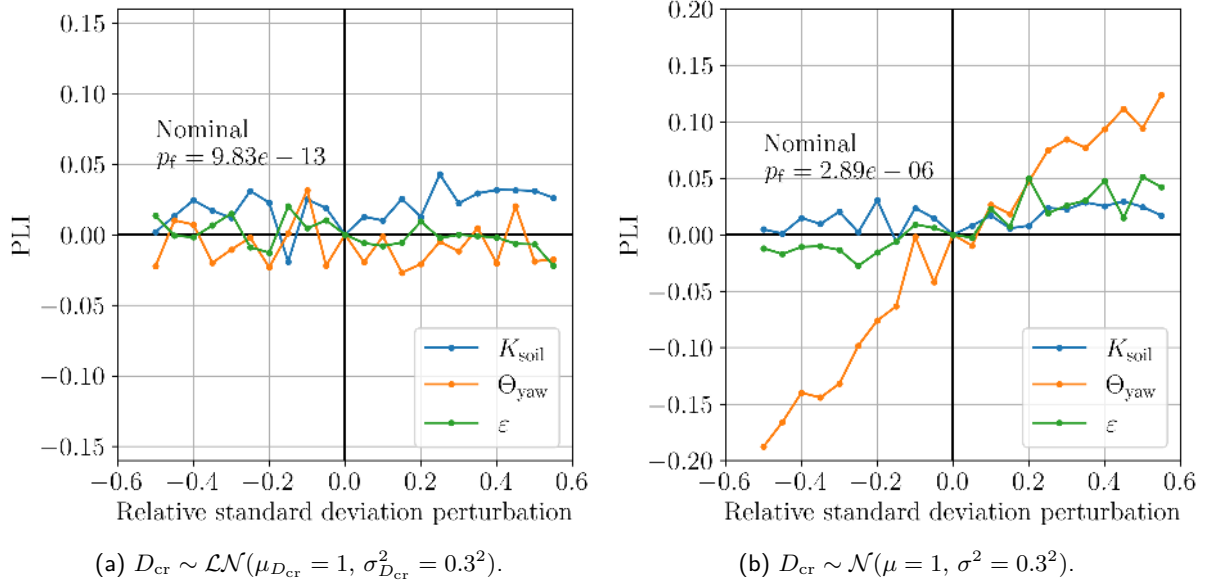


Figure 7.6: Perturbed-law based indices for relative perturbations of the standard deviations of $(K_{soil}, \Theta_{yaw}, \varepsilon)$. The failure probabilities studied are each estimated by FORM-IS method on the surrogate model (with sample size $N = 5 \times 10^4$).

7.4 Synthesis

In this section the reliability analysis and robustness analysis performed for the ambient metocean conditions in Section 6 were repeated by taking the most wake-perturbed metocean distribution inside Teesside's wind farm. After building a surrogate model of the mean cumulative damage at the mudline level of the foundation, the failure probability w.r.t. the critical damage has been estimated using FORM, FORM-IS and subset simulation algorithms. The values of failure probability when considering the wake-perturbed conditions appeared to be very close to the ones estimated previously. Then, the robustness of this quantity has been studied with the PLI approach for two distribution assumptions for D_{cr} , which showed again similar conclusions as the previous section. The critical damage D_{cr} is the uncertain variable which has by far the most impact on the reliability in this case.

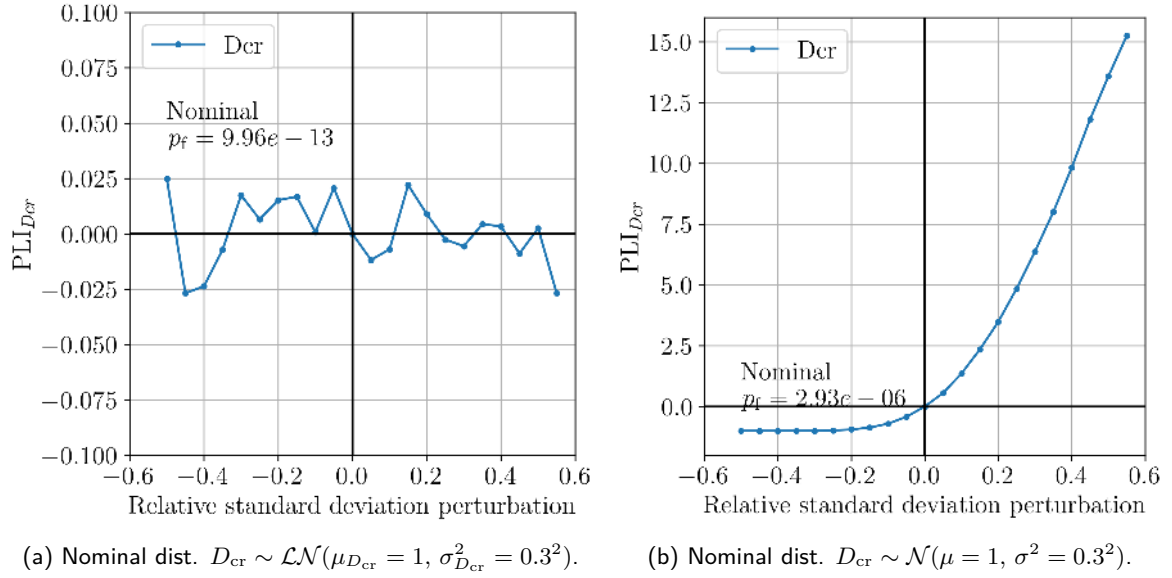


Figure 7.7: Perturbed-law based indices for relative perturbations of the standard deviation of D_{cr} . The failure probabilities studied are each estimated by FORM-IS method on the surrogate model (with sample size $N = 5 \times 10^4$).

8 Conclusion

This work is related to uncertainty propagation in the multi-physics numerical model of an OWT, called DIEGO. A first contribution concerns the estimation of lifetime cumulative fatigue damage as a mean against the environmental conditions (see e.g., Müller and Cheng, 2018). In Section 4, Bayesian quadrature methods such as kernel herding (Chen et al., 2010; Huszár and Duvenaud, 2012) proved to be an efficient and flexible solution for given-data uncertainty propagation (i.e., directly subsampling a large dataset without inference). Our numerical experiments showed conclusive results in comparison to Monte Carlo and quasi-Monte Carlo sampling. This development was gathered in a Python package named `otkerneldesign`.

As a perspective, the damage surrogate modeling could be further studied (e.g., extending the work of Slot et al., 2020). Recent techniques for stochastic emulators could be tested on this case (Baker et al., 2022; Zhu and Sudret, 2023; Lüthen et al., 2023). However, the method used should also be robust to the highly skewed distribution of the output variable. The mechanics could be approximated more accurately by considering other cumulative models for fatigue rule rather than Miner's rule. For example, nonlinear models (Rocher et al., 2020) take the order of occurrence of the cyclic solicitations into account.

On top of averaging the cumulative damage over the environmental conditions, other variables related to the system can be considered random in Section 6. Therefore, a second contribution deals with the reliability analysis based on a few system variables: soil stiffness, S-N curve parameters, yaw misalignment, and damage resistance. This study was made possible by the use of HPC facilities and a surrogate model built on a space-filling design over the system variables. Beyond the evaluation of a monopile foundation's failure probability, this work studied its robustness to the perturbation of input distributions. The perturbed-law based sensitivity indices (Lemaître et al., 2015) were applied to assess the relative influence of perturbing the inputs' moments. This demonstrative study on the Teesside wind farm showed that the failure probabilities estimated are much lower than the recommended value.

An industrial perspective regarding the reliability analysis is its application to a floating wind turbine case. This setup could introduce additional nonlinearities arising from the stronger impact of the waves on a floating structure. Along with fatigue reliability, the connection with other failure modes such as extreme loading could be studied (see, for example, the recent study on ultimate limit states by Wang et al., 2023).

References

- K. Aas, C. Czado, A. Frigessi, and H. Bakken. Pair-copula constructions of multiple dependence. *Insurance: Mathematics and Economics*, 44(2):182–198, 2009.
- I. Abdallah, C. Lataniotis, and B. Sudret. Parametric hierarchical kriging for multi-fidelity aero-servo-elastic simulators – Application to extreme loads on wind turbines. *Probabilistic Engineering Mechanics*, 55:67–77, 2019.
- A. Ajenjo. *Info-gap robustness assessment of reliability evaluations for the safety of critical industrial systems*. PhD thesis, Université Bourgogne Franche-Comté, 2023.
- E. Ardillon, M. Paskyabi, A. Cousin, N. Dimitrov, M. Dupoirson, S. Eldevik, E. Fekhari, C. Ferreira, M. Guiton, B. Jezequel, et al. Turbine loading and wake model uncertainty. Technical report, European Union, 2023.
- H. Bai, L. Shi, Y. Aoues, C. Huang, and D. Lemosse. Estimation of probability distribution of long-term fatigue damage on wind turbine tower using residual neural network. *Mechanical Systems and Signal Processing*, 190: 110101, 2023.
- E. Baker, P. Barbillon, A. Fadikar, R. Gramacy, R. Herbei, D. Higdon, J. Huang, L. Johnson, A. Mondal, B. Pires, et al. Analyzing Stochastic Computer Models: A Review with Opportunities. *Statistical Science*, 37(1):64 – 89, 2022.
- M. Baudin, A. Dutfoy, B. Iooss, and A. Popelin. OpenTURNS: An industrial software for uncertainty quantification in simulation. In R. Ghanem, D. Higdon, and H. Owhadi, editors, *Handbook of Uncertainty Quantification*, pages 2001–2038. Springer, 2017.
- E. Beauregard, E. Bérille, N. Berrabah, M. Berthelot, J. Burrows, M. Capaldo, S. Cornet, V. Costan, M. Duchet, E. Dufossé, E. Dupont, M. Franchet, E. Gouze, A. Grau, A. Joly, N. Kell, V. de Laleu, F. Latraube, A. Lovera, A. de Bazelaire, E. Monnot, G. Nogaro, J. Pagot, R. Pérony, C. Peyrard, C. Piguet, A. Régner, E. Santibanez, C. Senn, C. Smith, F. Soriano, P. Stephan, N. Terte, P. Veyan, D. Vizireanu, and L. Yeow. *L'éolien en mer : un défi pour la transition énergétique*. Lavoisier, 2022.
- M. Binois, J. Huang, R. Gramacy, and M. Ludkovski. Replication or exploration? Sequential design for stochastic simulation experiments. *Technometrics*, 61(1):7–23, 2019.
- E. Bitner-Gregersen. Joint met-ocean description for design and operations of marine structures. *Applied Ocean Research*, 51:279–292, 2015.
- J.-B. Blanchard, R. Chocat, G. Damblin, M. Baudin, N. Bousquet, V. Chabridon, B. Iooss, M. Keller, J. Pelamatti, and R. Sueur. Fiches pédagogiques sur le traitement des incertitudes dans les codes de calcul. Technical report, EDF, 2023.
- G. Blatman and B. Sudret. Adaptive sparse polynomial chaos expansion based on least angle regression. *Journal of Computational Physics*, 230(6):2345–2367, 2011.
- F. Briol, C. Oates, M. Girolami, M. Osborne, and D. Sejdinovic. Probabilistic Integration: A Role in Statistical Computation? *Statistical Science*, 34(1):1 – 22, 2019.
- F.-X. Briol. *Statistical computation with kernels*. PhD thesis, University of Warwick, 2019.
- F.-X. Briol, C. Oates, M. Girolami, and M. Osborne. Frank-Wolfe Bayesian Quadrature: Probabilistic Integration with Theoretical Guarantees. In *Advances in Neural Information Processing Systems*, 2015.
- T. Browne, B. Iooss, L. Le Gratiet, J. Lonchamp, and E. Remy. Stochastic simulators based optimization by Gaussian process metamodels—application to maintenance investments planning issues. *Quality and Reliability Engineering International*, 32(6):2067–2080, 2016.
- T. Burton, N. Jenkins, E. Bossanyi, D. Sharpe, and M. Graham. *Wind Energy Handbook*. John Wiley & Sons, 2021.
- M. Capaldo, M. Guiton, G. Huwart, E. Julan, N. Dimitrov, T. Kim, A. Lovera, and C. Peyrard. Design brief of HIPERWIND offshore wind turbine cases: bottom fixed 10MW and floating 15MW. Technical report, EDF ; DTU ; IFPEN, June 2021. URL <https://ifp.hal.science/hal-04033059>.

- M. Carmassi. *Uncertainty quantification and calibration of a photovoltaic plant model: warranty of performance and robust estimation of the long-term production*. PhD thesis, Université Paris-Saclay, 2018.
- T. Chen, X. Wang, G. Yuan, and J. Liu. Fatigue bending test on grouted connections for monopile offshore wind turbines. *Marine Structures*, 60:52–71, 2018.
- Y. Chen, M. Welling, and A. Smola. Super-samples from kernel herding. In *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence*, pages 109–116. AUAI Press, 2010.
- S. Da Veiga. Global sensitivity analysis with dependence measures. *Journal of Statistical Computation and Simulation*, 85:1283 – 1305, 2015.
- S. Da Veiga, F. Gamboa, B. Iooss, and C. Prieur. *Basics and Trends in Sensitivity Analysis: Theory and Practice in R*. Society for Industrial and Applied Mathematics, 01 2021.
- G. Damblin. *Contributions statistiques au calage et à la validation des codes de calcul*. PhD thesis, AgroParisTech, 2015.
- E. De Rocquigny, N. Devictor, and S. Tarantola. *Uncertainty in industrial practice: a guide to quantitative uncertainty management*. John Wiley & Sons, 2008.
- N. Dimitrov. *Structural reliability of wind turbine blades: Design methods and evaluation*. PhD thesis, Technical University of Denmark, 2013.
- N. Dimitrov, M. Kelly, A. Vignaroli, and J. Berg. From wind to loads: wind turbine site-specific load estimation with surrogate models trained on high-fidelity load databases. *Wind Energy Science*, 3:767 – 790, 2018.
- T. Dirlik and D. Benasciutti. Dirlik and Tovo-Benasciutti Spectral Methods in Vibration Fatigue: A Review with a Historical Perspective. *Metals*, 11(9):1333, 2021.
- DNV-RP-C203. DNV-RP-C203: Fatigue design of offshore steel structures. Technical report, Det Norske Veritas, 2016.
- P. Doubrawa, E. Quon, L. Martinez-Tossas, K. Shaler, M. Debnath, N. Hamilton, T. Herges, D. Maniaci, C. Kelley, A. Hsieh, et al. Multimodel validation of single wakes in neutral and stratified atmospheric conditions. *Wind Energy*, 23(11):2027–2055, 2020.
- P. Doubrawa, K. Shaler, and J. Jonkman. Difference in load predictions obtained with effective turbulence vs. a dynamic wake meandering modeling approach. *Wind Energy Science*, 2023:1–28, 2023.
- N. E. Dowling. Fatigue Failure Predictions for Complicated Stress-Strain Histories. *Journal of Materials, JMLSA*, 7:71–87, 1972.
- S. Drexler and M. Muskulus. Reliability of an offshore wind turbine with an uncertain S-N curve. *Journal of Physics: Conference Series*, 2018:012014, 2021.
- V. Dubourg. *Adaptive surrogate models for reliability analysis and reliability-based design optimization*. PhD thesis, Université Blaise Pascal, 2011.
- M. Eldred. Recent advances in non-intrusive polynomial chaos and stochastic collocation methods for uncertainty analysis and design. In *50th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, 4–7 May 2009, Palm Springs, California, USA*, page 2274, 2009.
- K. Fang, R. Li, and S. A. *Design and Modeling for Computer Experiments*. Computer Science and Data Analysis Series. Chapman & Hall / CRC, 2006.
- K. Fang, M.-Q. Liu, H. Qin, and Y.-D. Zhou. *Theory and Application of Uniform Experimental Designs*, volume 221. Springer, 2018.
- A. Fatemi and L. Yang. Cumulative fatigue damage and life prediction theories: a survey of the state of the art for homogeneous materials. *International Journal of Fatigue*, 20(1):9–34, 1998.
- E. Fekhari. *Uncertainty quantification in multi-physics models for wind turbine asset management*. PhD thesis, Université Nice Côte d’Azur, 2024.

- E. Fekhari, B. Iooss, V. Chabridon, and J. Muré. Efficient techniques for fast uncertainty propagation in an offshore wind turbine multi-physics simulation tool. In *Proceedings of the 5th International Conference on Renewable Energies Offshore*, pages 837–846, 2022.
- E. Fekhari, B. Iooss, J. Muré, L. Pronzato, and J. Rendas. Model predictivity assessment: incremental test-set selection and accuracy evaluation. In N. Salvati, C. Perna, S. Marchetti, and R. Chambers, editors, *Studies in Theoretical and Applied Statistics*, pages 315–347. Springer, 2023.
- C. Freyssinet, V. Rey, F. Schoefs, and T. Moro. Bayesian calibration of a non linear damage model of steel structures with random material property: Sensitivity analysis and reliability assessment. *Engineering Structures*, 295:116853, 2023.
- J. Fuhg., A. Fau, and U. Nackenhorst. State-of-the-Art and Comparative Review of Adaptive Sampling Methods for Kriging. *Archives of Computational Methods in Engineering*, 28:2689–2747, 2021.
- C. Gauchy, J. Stenger, R. Sueur, and B. Iooss. An information geometry approach to robustness analysis for the uncertainty quantification of computer codes. *Technometrics*, 64(1):80–91, 2022.
- P. Graf, G. Stewart, M. Lackner, K. Dykes, and P. Veers. High-throughput computation and the applicability of Monte Carlo integration in fatigue load estimation of floating offshore wind turbines. *Wind Energy*, 19(5): 861–872, 2016.
- A. Gretton, K. M. Borgwardt, M. Rasch, B. Schölkopf, and A. Smola. A Kernel Method for the Two-Sample-Problem. In *Proceedings of the 19th International Conference on Neural Information Processing Systems*, pages 513–520, 2006.
- Z. Guédé, B. Sudret, and M. Lemaire. Life-time reliability based assessment of structures submitted to thermal fatigue. *International Journal of Fatigue*, 29(7):1359–1373, 2007.
- Q. Han, Z. Hao, T. Hu, and F. Chu. Non-parametric models for joint probabilistic distributions of wind speed and direction data. *Renewable Energy*, 126:1032–1042, 2018.
- M. Hansen and L. Henriksen. *Basic DTU Wind Energy controller*. Number 0028 in DTU Wind Energy E. DTU Wind Energy, 2013.
- K. Hasselmann, T. Barnett, E. Bouws, H. Carlson, D. Cartwright, K. Enke, J. Ewing, A. Gienapp, D. Hasselmann, P. Kruseman, et al. Measurements of wind-wave growth and swell decay during the Joint North Sea Wave Project (JONSWAP). *Ergänzungsheft zur Deutschen Hydrographischen Zeitschrift, Reihe A*, 1973.
- F. Hickernell. A generalized discrepancy and quadrature error bound. *Mathematics of Computation*, 67(221): 299–322, 1998.
- Q. Huchet. *Kriging based methods for the structural damage assessment of offshore wind turbines*. PhD thesis, Université Blaise Pascal, 2018.
- Q. Huchet, C. Mattrand, P. Beaurepaire, N. Relun, and N. Gayton. AK-DA: An efficient method for the fatigue assessment of wind turbine structures. *Wind Energy*, 22(5):638–652, 2019.
- F. Huszár and D. Duvenaud. Optimally-Weighted Herding is Bayesian Quadrature. In *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence*, pages 377 – 386, 2012.
- IEC-61400-1. IEC 61400-1: Wind energy generation systems - Part 1: Design requirements. Technical report, International Electrotechnical Commission (IEC), 2019.
- B. Iooss, V. Vergès, and V. Larget. BEPU robustness analysis via perturbed law-based sensitivity indices. *Proceedings of the Institution of Mechanical Engineers, Part O: Journal of Risk and Reliability*, 236(5):855–865, 2022.
- JCSS. JCSS PROBABILISTIC MODEL CODE PART 3: RESISTANCE MODELS. Technical report, Joint Committee on Structural Safety, 2011.
- N. Jensen. A note on wind generator interaction. Technical report, DTU, 1983.

- B. Jonkman. Turbsim User's Guide: Version 1.50. Technical report, NREL, 2009.
- J. Kaimal, J. Wyngaard, Y. Izumi, and O. Coté. Spectral characteristics of surface-layer turbulence. *Quarterly Journal of the Royal Meteorological Society*, 98(417):563–589, 1972.
- M. Kanagawa and P. Hennig. Convergence Guarantees for Adaptive Bayesian Quadrature Methods. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- M. Kanagawa, P. Hennig, D. Sejdinovic, and B. Sriperumbudur. Gaussian Processes and Kernel Methods: A Review on Connections and Equivalences. arXiv:1807.02582, 2018.
- S. Kanner, A. Aubault, A. Peiffer, and B. Yu. Maximum dissimilarity-based algorithm for discretization of metocean data into clusters of arbitrary size and dimension. In *International Conference on Offshore Mechanics and Arctic Engineering*, volume 51319, 2018.
- Z. Kaplan, Y. Li, M. Nakayama, and B. Tuffin. Randomized quasi-Monte Carlo for quantile estimation. In *2019 Winter Simulation Conference (WSC)*, pages 428–439, 2019.
- G. Katsikogiannis, S. Sørnum, E. Bachynski, and J. Amdahl. Environmental lumping for efficient fatigue assessment of large-diameter monopile wind turbines. *Marine Structures*, 77:102939, 2021.
- T. Kim, A. Natarajan, A. Lovera, E. Julan, E. Peyrard, M. Capaldo, G. Huwart, P. Bozonnet, and M. Guiton. A comprehensive code-to-code comparison study with the modified IEA15MW-UMaine Floating Wind Turbine for H2020 HIPERWIND project. *Journal of Physics: Conference Series*, 2265(4):042006, 2022.
- I. Klebanov, I. Schuster, and T. Sullivan. A Rigorous Theory of Conditional Mean Embeddings. *SIAM Journal on Mathematics of Data Science*, 2(3):583–606, 2020.
- S. Kucherenko, B. Feil, N. Shah, and W. Mauntz. The identification of model effective dimensions using global sensitivity analysis. *Reliability Engineering and System Safety*, 96:440–449, 2011.
- S. Lacoste-Julien, F. Lindsten, and F. Bach. Sequential Kernel Herding: Frank-Wolfe Optimization for Particle Filtering. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, volume 38, pages 544–552, 2015.
- M. Lasserre. *Apprentissages dans les réseaux bayésiens à base de copules non-paramétriques*. PhD thesis, Sorbonne Université, 2022.
- C. Lataniotis. *Data-driven uncertainty quantification for high-dimensional engineering problems*. PhD thesis, ETH Zürich, 2019.
- O. Le Maître and O. Knio. *Spectral methods for uncertainty quantification: with applications to computational fluid dynamics*. Springer Science & Business Media, 2010.
- R. Lebrun and A. Dutfoy. A generalization of the Nataf transformation to distributions with elliptical copula. *Probabilistic Engineering Mechanics*, 24(2):172–178, 2009.
- M. Lemaire, A. Chateauneuf, and J.-C. Mitteau. *Structural Reliability*. John Wiley & Sons, 2009.
- P. Lemaître, E. Sergienko, A. Arnaud, N. Bousquet, F. Gamboa, and B. Iooss. Density modification-based reliability sensitivity analysis. *Journal of Statistical Computation and Simulation*, 85(6):1200–1223, 2015.
- G. Leobacher and F. Pillichshammer. *Introduction to Quasi-Monte Carlo Integration and Applications*. Springer, 2014.
- X. Li and W. Zhang. Long-term fatigue damage assessment for a floating offshore wind turbine under realistic environmental conditions. *Renewable Energy*, 159:570–584, 2020.
- H.-P. Lieudarde. *La Pratique des essais de fatigue – méthodes expérimentales et analyse des résultats*. Pyc édition, 1982.
- Y. Lin and S. Dong. Wave energy assessment based on trivariate distribution of significant wave height, mean period and direction. *Applied Ocean Research*, 87:47–63, 2019.

- X. Liu, C. Lu, G. Li, A. Godbole, and Y. Chen. Effects of aerodynamic damping on the tower load of offshore horizontal axis wind turbines. *Applied Energy*, 204:1101–1114, 2017.
- A. Lovera, E. Fekhari, B. Jézéquel, M. Dupoiron, M. Guiton, and E. Ardillon. Quantifying and clustering the wake-induced perturbations within a wind farm for load analysis. In *Journal of Physics: Conference Series*, volume 2505, page 012011. IOP Publishing, 2023.
- N. Lüthen, S. Marelli, and B. Sudret. Sparse polynomial chaos expansions: Literature survey and benchmark. *SIAM/ASA Journal on Uncertainty Quantification*, 9(2):593–649, 2021.
- N. Lüthen, S. Marelli, and B. Sudret. Automatic selection of basis-adaptive sparse polynomial chaos expansions for engineering applications. *International Journal for Uncertainty Quantification*, 12(3), 2022.
- N. Lüthen, S. Marelli, and B. Sudret. A spectral surrogate model for stochastic simulators computed from trajectory samples. *Computer Methods in Applied Mechanics and Engineering*, 406:115875, 2023.
- S. Mak and V. Joseph. Support points. *The Annals of Statistics*, 46:2562 – 2592, 2018.
- S. Marelli and B. Sudret. UQLab: A framework for uncertainty quantification in Matlab. In *Vulnerability, uncertainty, and risk: quantification, mitigation, and management*, pages 2554–2563. 2014.
- S. Marelli, N. Lüthen, and B. Sudret. UQLab user manual – Polynomial chaos expansions. Technical report, Chair of Risk, Safety and Uncertainty Quantification, ETH Zurich, Switzerland, 2024. Report UQLab-V2.1-104.
- A. Marrel and V. Chabridon. Statistical developments for target and conditional sensitivity analysis: Application on safety studies for nuclear reactor. *Reliability Engineering and System Safety*, 214:107711, 2021.
- M. McKay, R. Beckman, and W. Conover. A Comparison of Three Methods for Selecting Vales of Input Variables in the Analysis of Output From a Computer Code. *Technometrics*, 21:239 – 245, 05 1979.
- D. Milano. *Numerical prototype for floating offshore wind turbines*. PhD thesis, The University of Edinburgh, 2021.
- D. Milano, C. Peyrard, M. Capaldo, D. Ingram, Q. Xiao, and L. Johanning. Impact of High Order Wave Loads on a 10 MW Tension-Leg Platform Floating Wind Turbine at Different Tendon Inclination Angles. In *Proceedings of the ASME 2019 38th International Conference on Ocean, Offshore and Arctic Engineering*, 2019.
- R. Montes-Iturrizaga and E. Heredia-Zavoni. Multivariate environmental contours using C-vine copulas. *Ocean engineering*, 118:68–82, 2016.
- J. Morio and M. Balesdent. *Estimation of Rare Event Probabilities in Complex Aerospace and Other Systems: A Practical Approach*. Woodhead Publishing, Elsevier, 2015.
- W. J. Morokoff and R. E. Caflisch. Quasi-Monte Carlo Integration. *Journal of Computational Physics*, 122(2): 218–230, 1995.
- J. Murcia, P. Réthoré, N. Dimitrov, A. Natarajan, J. Sørensen, P. Graf, and T. Kim. Uncertainty propagation through an aeroelastic wind turbine model using polynomial surrogates. *Renewable Energy*, 119:910–922, 2018.
- K. Müller and P. Cheng. Application of a Monte Carlo procedure for probabilistic fatigue design of floating offshore wind turbines. *Wind Energy Science*, 3:149 – 162, 2018.
- T. Nagler, C. Schellhase, and C. Czado. Nonparametric estimation of simplified vine copula models: comparison of methods. *Dependence Modeling*, 5:99–120, 2017.
- J. Nielsen and J. Sørensen. Risk-based derivation of target reliability levels for life extension of wind turbine structural components. *Wind Energy*, 24(9):939–956, 2021.
- C. J. Oates. Minimum Discrepancy Methods in Uncertainty Quantification, 2021. Lecture Notes at École Thématique sur les Incertitudes en Calcul Scientifique (ETICS21), <https://www.gdr-mascotnum.fr/etics.html>.
- W. Oberkampf and C. Roy. *Verification and validation in scientific computing*. Cambridge university press, 2010.
- A. O’Hagan. Bayes-Hermite quadrature. *Journal of Statistical Planning and Inference*, 29:245–260, 1991.

- A. Owen. The dimension distribution and quadrature test functions. *Statistica Sinica*, 13:1–17, 2003.
- E. Petrovska. *Fatigue life reassessment of monopile-supported offshore wind turbine structures*. PhD thesis, University of Edinburgh, 2022.
- A. Pires, M. Moustapha, S. Marelli, and B. Sudret. Surrogate-based reliability analysis for noisy models. In *14th International Conference on Applications of Statistics and Probability in Civil Engineering (ICASP14)*. Trinity College Dublin, 2023.
- W. Popko, A. Robertson, J. Jonkman, F. Wendt, P. Thomas, K. Müller, M. Kretschmer, T. Hagen, C. Galinos, J.-B. Le Dreff, et al. Validation of numerical models of the offshore wind turbine from the alpha ventus wind farm against full-scale measurements within OC5 Phase III. *Journal of Offshore Mechanics and Arctic Engineering*, 143(1):012002, 2021.
- L. Pronzato and M.-J. Rendas. Validation of Machine Learning Prediction Models. *The New England Journal of Statistics in Data Science*, 1:394–414, 2023.
- L. Pronzato and A. Zhigljavsky. Bayesian quadrature and energy minimization for space-filling design. *SIAM/ASA Journal on Uncertainty Quantification*, 8:959 – 1011, 2020.
- P. Ragan and L. Manuel. Comparing estimates of wind turbine fatigue loads using time-domain and spectral methods. *Wind engineering*, 31(2):83–99, 2007.
- C. Rasmussen and C. Williams. *Gaussian processes for machine learning*, volume 1. Springer, 2006.
- A. Robertson, S. Gueydon, E. Bachynski, L. Wang, J. Jonkman, D. Alarcon, E. Amet, A. Beardsell, P. Bonnet, B. Boudet, et al. OC6 Phase I: Investigating the underprediction of low-frequency hydrodynamic loads and responses of a floating wind turbine. In *Journal of Physics: Conference Series*, volume 1618, page 032033, 2020.
- B. Rocher, S. Schoefs, M. François, A. Salou, and A.-L. Caouissin. A two-scale probabilistic time-dependent fatigue model for offshore steel wind turbines. *International Journal of Fatigue*, 136:105620, 2020.
- M. Rosenblatt. Remarks on a multivariate transformation. *The annals of mathematical statistics*, 23(3):470–472, 1952.
- A. Saltelli, M. Ratto, T. Andres, F. Campolongo, J. Cariboni, D. Gatelli, M. Saisana, and S. Tarantola. *Global sensitivity analysis: the primer*. John Wiley & Sons, 2008.
- A. Sancetta and S. Satchell. The Bernstein copula and its applications to modeling and approximations of multivariate distributions. *Econometric Theory*, 20(3):535–562, 2004.
- T. J. Santner, B. Williams, and W. Notz. *The Design and Analysis of Computer Experiments*. Springer Series in Statistics. Springer-Verlag, 2003.
- D. Sejdinovic, B. Sriperumbudur, A. Gretton, and K. Fukumizu. Equivalence of distance-based and RKHS-based statistics in hypothesis testing. *The Annals of Statistics*, 41:2263–2291, 2013.
- R. M. Slot, J. D. Sørensen, B. Sudret, L. Sørensen, and M. L. Thøgersen. Surrogate model uncertainty in wind turbine reliability assessment. *Renewable Energy*, 151:1150 – 1162, 2020.
- R. C. Smith. *Uncertainty Quantification: Theory, Implementation, and Applications*. Computational Science & Engineering. SIAM, 2014.
- B. Sriperumbudur, A. Gretton, K. Fukumizu, B. Schölkopf, and G. Lanckriet. Hilbert Space Embeddings and Metrics on Probability Measures. *Journal of Machine Learning Research*, 11:1517–1561, 2010.
- B. Sudret. Probabilistic Design of Structures Submitted to Fatigue. In C. Bathias and A. Pineau, editors, *Fatigue of Materials and Structures*, pages 223–263. John Wiley & Sons, Ltd, 2013.
- R. Sueur, B. Iooss, and T. Delage. Sensitivity analysis using perturbed-law based indices for quantiles and application to an industrial case. In *10th International Conference on Mathematical Methods in Reliability (MMR 2017)*, Grenoble, France, 2017.

- T. Sullivan. *Introduction to Uncertainty Quantification*, volume 63. Springer, 2015.
- S. Suresh. *Fatigue of Materials*. Cambridge University Press, 1998.
- G. J. Székely and M. L. Rizzo. Energy statistics: A class of statistics based on distances. *Journal of Statistical Planning and Inference*, 143:1249 – 1272, 2013.
- R. Teixeira, M. Nogal, A. O'Connor, J. Nichols, and A. Dumas. Stress-cycle fatigue design with Kriging applied to offshore wind turbines. *International Journal of Fatigue*, 125:454–467, 2019a.
- R. Teixeira, N. O'Connor, and M. Nogal. Probabilistic sensitivity analysis of offshore wind turbines using a transformed Kullback-Leibler divergence. *Structural Safety*, 81:101860, 2019b.
- E. Torre, S. Marelli, B. Embrechts, and B. Sudret. Data-driven polynomial chaos expansion for machine learning regression. *Journal of Computational Physics*, 388:601–623, 2019a.
- E. Torre, S. Marelli, P. Embrechts, and B. Sudret. A general framework for data-driven uncertainty quantification under complex input dependencies using vine copulas. *Probabilistic Engineering Mechanics*, 55:1–16, 2019b.
- E. Torre, S. Marelli, and B. Sudret. UQLab user manual – Statistical inference. Technical report, Chair of Risk, Safety and Uncertainty Quantification, ETH Zurich, Switzerland, 2024. Report UQLab-V2.1-114.
- L. Van den Bos. *Quadrature Methods for Wind Turbine Load Calculations*. PhD thesis, Delft University of Technology, 2020.
- E. Vanem. Joint statistical models for significant wave height and wave period in a changing climate. *Marine Structures*, 49:180–205, 2016.
- E. Vanem, E. Fekhari, N. Dimitrov, M. Kelly, A. Cousin, and M. Guiton. A Joint Probability Distribution for Multivariate Wind-Wave Conditions and Discussions on Uncertainties. *Journal of Offshore Mechanics and Arctic Engineering*, 146(6):061701, 2024.
- J. Velarde, C. Kramhøft, and J. Sørensen. Global sensitivity analysis of offshore wind turbine foundation fatigue loads. *Renewable Energy*, 140:177 – 189, 2019.
- J. Velarde, C. Kramhøft, J. Sørensen, and G. Zorzi. Fatigue reliability of large monopiles for offshore wind turbines. *International Journal of Fatigue*, 134:105487, 2020.
- H. Wang, O. Gramstad, S. Schär, S. Marelli, and E. Vanem. Comparison of Probabilistic Structural Reliability Methods for Ultimate Limit State Assessment of Wind Turbines. *ArXiv preprint arXiv:2312.04972*, 2023.
- L. Wang, A. Kolios, X. Liu, D. Venetsanos, and R. Cai. Reliability of offshore wind turbine support structures: A state-of-the-art review. *Renewable and Sustainable Energy Reviews*, 161:112250, 2022.
- D. Wilkie and C. Galasso. Gaussian process regression for fatigue reliability analysis of offshore wind turbines. *Structural Safety*, 88:102020, 2021.
- D. Xiu and G. E. Karniadakis. The wiener–askey polynomial chaos for stochastic differential equations. *SIAM Journal on Scientific Computing*, 24(2):619–644, 2002.
- X. Zhu. *Surrogate Modeling for Stochastic Simulators Using Statistical Approaches*. PhD thesis, ETH Zurich, 2022.
- X. Zhu and B. Sudret. Stochastic polynomial chaos expansions to emulate stochastic simulators. *International Journal for Uncertainty Quantification*, 13(2):31–52, 2023.
- D. Zwick and M. Muskulus. The simulation error caused by input loading variability in offshore wind turbine structural analysis. *Wind Energy*, 18:1421 – 1432, 2015.