



# Surrogate models for the blade element momentum aerodynamic model using non-intrusive polynomial chaos expansions

**Rad Haghi and Curran Crawford**

Department of Mechanical Engineering, Institute for Integrated Energy Systems,  
University of Victoria, British Columbia, Canada

**Correspondence:** Rad Haghi (rhaghi@uvic.ca)

Received: 22 September 2021 – Discussion started: 1 November 2021  
Revised: 29 April 2022 – Accepted: 9 June 2022 – Published: 27 June 2022

**Abstract.** In typical industrial practice based on IEC standards, wind turbine simulations are computed in the time domain for each mean wind speed bin using a few unsteady wind seeds. Software such as FAST, BLADED, or HAWC2 can be used to capture the unsteadiness and uncertainties of the wind in the simulations. The statistics of these aeroelastic simulation outputs are extracted and used to calculate fatigue and extreme loads on the wind turbine components. The minimum requirement of having six seeds does not guarantee an accurate estimation of the overall statistics. One solution might be running more seeds; however, this will increase the computation cost. Moreover, to move beyond blade element momentum (BEM)-based tools toward vortex/potential flow formulations, a reduction in the computational cost associated with the unsteady flow and uncertainty handling is required. This study illustrates the unsteady wind aerodynamic statistics' stationary character based on the standard turbulence models. This character is shown based on the output of National Renewable Energy Lab (NREL) 5MW reference machine BEM simulations. Afterwards, we propose a non-intrusive polynomial chaos expansion (PCE) to build a surrogate model of the loads' statistics, the rotor thrust, and torque, at each time step, to estimate the extreme statistics more accurately and efficiently.

## 1 Introduction

The process of calculating loads on wind turbine components is one of the core parts of wind turbine aerodynamic and structural design and optimization. In the last few decades, international organizations have developed different aeroelastic codes such as Fatigue, Aerodynamics, Structures, and Turbulence (FAST) (Jonkman and Buhl, 2005), BLADED (DNV GL, 2018), and HAWC2 (Larsen and Hansen, 2007) to accurately calculate load time series based on the standardized or site-specific environmental conditions. Engineers and researchers use wind turbine aeroelastic simulation output statistics to calculate extreme and fatigue loads on wind turbine structures and estimate the unsteady power. To take into account the randomness in the unsteady wind, according to IEC standards (IEC 61400-1, 2019), the simulation process must use a semi-Monte Carlo (MC) method. Therefore,

a full simulation set should include a limited number of seeds for generating multiple wind speed time series of 600 s.

In normal practice, for each mean wind speed, at least six different seeded unsteady wind time series are required as the minimum to take into account the uncertainties. This limited number of unsteady simulations does not yield an entirely accurate estimation of the statistics. Gradient-based optimization algorithms may not deal with these inaccuracies well. One option to solve this problem is running more seeds, which will increase the computational cost. The increase in computational cost will play a more critical role in our decision making if we want to move towards vortex (van Garrel, 2003) and potential flow codes for load calculations as they require more computation resources inherently. An alternative approach to direct simulation is to use a surrogate model that can provide us with an accurate statistical estimation set based on a limited number of simulations.

The origin of the *surrogate model* lies in uncertainty quantification (UQ) analysis (Sudret, 2007). There are many uncertainty quantification implementations in wind energy. More specifically, surrogate models show much potential in wind farm load estimation, wind turbine optimization, and reliability analysis. Many researchers have investigated these potentials. For example, Dimitrov et al. (2018), Schröder et al. (2018), van den Bos et al. (2018), and Dimitrov (2019) used surrogate models to estimate the loads on a wind turbine based on the stochastic variables' gross parameters such as turbulence intensity, mean wind, or wind direction. Ashuri et al. (2016), Murcia et al. (2018), and Schröder et al. (2020a) used surrogate models for uncertainty propagation through the wind turbine models. More recently, the surrogate models have been used for the wind turbines' reliability assessments (Slot et al., 2020; Schröder et al., 2020b). Also, Wang et al. (2020) and Barlas et al. (2021) showed the application of the surrogate model in wind turbine optimization. However, very few have looked at building a surrogate model of the aerodynamic model of a wind turbine based on the random phases as the input. Fluck and Crawford (2018) showed an initial attempt to build a surrogate model based on intrusive polynomial chaos expansion (PCE) on simple lifting line and blade element momentum (BEM) models (Fluck and Crawford, 2016a, 2018). As they were quickly faced with the *curse of dimensionality*, they showed it is possible to reduce the number of random variables in the unsteady wind model of Veers (1988) significantly. Afterwards, they used this reduced-dimension wind model to propagate stochasticity through a simple lifting line (Fluck and Crawford, 2016b) or BEM (Fluck and Crawford, 2018) model. However, with intrusive PCE it is necessary to change the model implementation fundamentally to incorporate the random variables (Sudret, 2007). This might work for a simple model, but when we want to utilize commercially available aeroelastic codes, this will be challenging or even impossible.

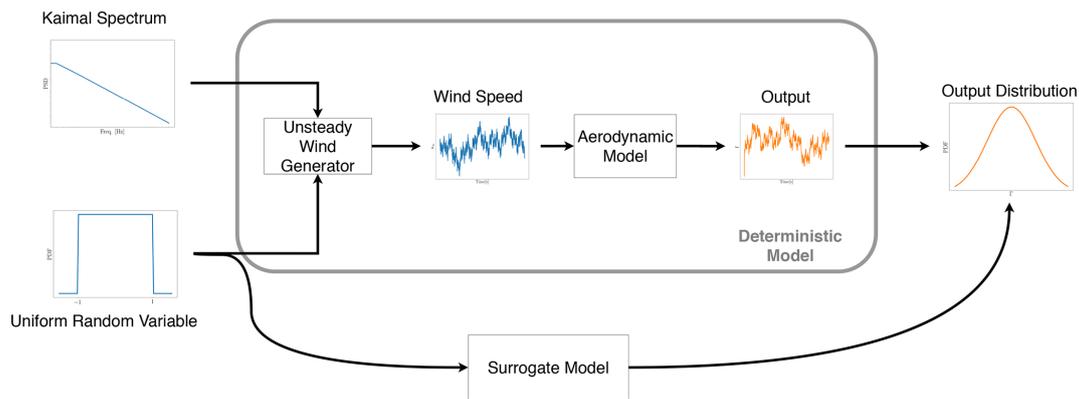
This paper's goal is to build a non-intrusive PCE surrogate model of a deterministic aerodynamic model driven by stochastic unsteady wind. This study's implemented aerodynamic model takes wind time series as input and calculates thrust and torque on a National Renewable Energy Lab (NREL) 5MW turbine rotor using BEM. The motivation is to build a surrogate model based on a limited amount of simulation data to estimate the statistics of the aerodynamic model output at each time step of the time series quickly. Having this surrogate model at hand helps us explore and experience the opportunities it can provide. This output guides future research in the surrogate model realms for us in the long run. The surrogate model investigation presented is an exploration of the potential benefits and limitations of PCE-based time-domain surrogate models to help researchers and practitioners develop future surrogate modeling approaches.

As the surrogate models are inherently cheap to run, we take this surrogate model through a Monte Carlo simulation (MCS) a large number of times. The input of these MCSs

are the samples drawn from the uniform random variables. The unsteady wind generator uses the same random variable distribution to make sure the generated time series will match a Gaussian process (Veers, 1988). This process is presented in Fig. 1 schematically. By this method, we can reduce the computational cost and time for the aerodynamic simulation without compromising the validity of the results. One can interpret this model as a tool to map the input distribution (in this case, a uniform distribution of random seed phases) to the output distribution (in this case, distribution of thrust and torque on the rotor).

Fitting a surrogate model at each time step of 600 s of the aerodynamic output times series, using the random phases (Fig. 1), is computationally expensive and redundant to current practice. Therefore, we start by showing that the aerodynamic simulation results based on Veer's reduced model (Fluck and Crawford, 2017) statistically converge. We also show that the unsteady wind aerodynamic process's statistical properties are constant in time (stationary process). Therefore, by keeping the computational effort the same, it is possible to run more simulations while shortening the length of the simulations. Furthermore, more simulations with the same computational effort give us the chance to fit higher-degree PCEs, which provides a more accurate estimation of the statistics. We built four different PCE surrogate models with four different polynomial degrees (degrees two to five) to pick the best in terms of accuracy and computational cost trade-off. These surrogate models have been used for MCSs for a large number of runs (cheaply). The results of the MC runs of the surrogate models are compared with 48 000 unsteady wind aerodynamic simulation results. In this case, the simulation results are the thrust and torque forces induced on the NREL 5MW in an unsteady wind. We compare the results by looking at the thrust and torque distribution from both the deterministic and surrogate models. Finally, we show how the extreme loads extracted from the MCS can converge to the extreme loads from the results from 48 000 unsteady wind aerodynamic simulation results.

This paper is organized as follows. Section 2 describes the unsteady wind generation and aerodynamic model. Sections 3 and 4 explain the statistical elements and PCE method used in the study. Next, in Sect. 5, we set out the approach to tackle the challenge. Section 6 provides the BEM simulation results, convergence of the sectional statistics the PCE fit on the sectional statistics, and the emulation output. At the end of the "Results" section, we discussed the accuracy and efficiency of the surrogate models developed in this study. This paper concludes in Sect. 7, reiterating the key findings of the study and offering suggestions for fruitful future work in the area of wind turbine surrogates.



**Figure 1.** The common deterministic process of aerodynamic modeling vs. the suggested surrogate model method schematic flow chart.

## 2 Models

This section provides an overview of the unsteady wind generation basics and aerodynamic model used in this study.

### 2.1 Reduced Veers unsteady wind model

One famous unsteady wind model in the wind turbine community is the Veers model (Veers, 1988). The history of the model goes back to the late 1980s, and it has a long record in wind turbine load calculation practice. Very briefly, the Veers unsteady wind model is inherently an inverse Fourier transformation. The one-dimensional unsteady wind time series at location  $P$  is generated via

$$u_{\infty}(t_n) = \sum_m \sqrt{S_m} e^{i(\omega_m t_n + 2\pi \xi_m)}. \quad (1)$$

For this inverse Fourier transformation in Eq. (1), the frequencies  $\omega_m$  are taken from the Kaimal spectrum (Fig. 1). The random phases are based on the independent random variable  $\xi_m$  drawn from a uniform distribution over  $[0, 1]$ . Finally, the amplitude  $S_m$  is specified based on the power spectrum at the frequencies  $\omega_m$  (Bossanyi et al., 2011).

In load calculation practice, the Veers' model for the unsteady wind is the method to generate turbulent boxes, commonly implemented in TurbSim (Jonkman, 2009). The method is briefly explained in Fluck and Crawford (2017) and extensively in Veers (1988). To make the unsteady wind in TurbSim, this method uses a large number of random variables on the order of thousands as required by Jonkman (2009). This large number of random variables pushes the surrogate model problem into the curse of dimensionality very quickly. Therefore, building a PCE surrogate model is almost impossible. To tackle this issue, Fluck and Crawford (2016b) showed that using only 10 uniformly distributed independent random variables with 10 frequencies logarithmically sampled from the Kaimal spectrum (Veers, 1988) is enough for building unsteady wind time series. This *reduced Veers' model* generated unsteady wind that can capture the

same level of randomness and probability distribution as the full model. This study used this reduced Veers model to generate unsteady wind time series. This method does not lead to a model that fully replaces high-fidelity TurbSim outputs but rather a surrogate model necessary to study trade-offs of various accuracy and assumption aspects.

The randomness in the generated unsteady wind comes from the 10 random variables,  $\phi_j$ , in the  $\xi$  vector describing the frequency components' phases  $2\pi \xi_m$  in the reduced Veers model in Eq. (1) (Fluck and Crawford, 2017). Based on the Veers method (Veers, 1988) and in TurbSim (Jonkman, 2009), the employed sampling method is a pseudo-random number generator (pRNG) which is the basis of MCSs. However, the problem with this way of sampling for MCSs is the lack of control over the random variables' domain as it may fill some voids in the domain and may leave some of it empty (Niederreiter, 1992). Therefore, the random domain may not be filled evenly for the same reason. For this study, a low-discrepancy quasi-Monte-Carlo (QMC) sampling method, namely the Sobol sequence (Sobol, 1967) is used to draw samples from the random variables in this work to calculate the PCE coefficients via the point collocation method. The main reasons to select the Sobol sequence over other sampling methods are the samples' consistency and computational efficiency (Kucherenko et al., 2015). A custom random wind generator based on the reduced Veers model used these samples to generate unsteady wind fields.

### 2.2 Aerodynamic model

The aerodynamic model for this study is a BEM model (Bossanyi et al., 2011) with frozen wake based on the work of Lupton (2019). This non-linear BEM model is used to run simulations on a NREL 5MW (Jonkman et al., 2009) rotor to calculate thrust and torque on the rotor. For the simulations, the rotor speed was kept constant based on the mean wind speed of the simulations. Also, the pitch angle was set according to the data provided in Jonkman et al. (2009). There is no controller involved in the simulations. The un-

steady wind defined in the previous section is set at 100 m hub height and remains the same on the rotor. The Python package for BEM is *bemused* (Lupton, 2019). The NREL 5MW model characteristics and properties are extracted from Jonkman et al. (2009). The model employed in this study for the simulations and surrogate model is essentially equivalent to the NREL 5MW model in any wind turbine aerodynamic code (e.g. FAST, BLADED, and HAWC) but is nicely formulated in Python. The model and analysis code used in this work has been previously verified against the NREL 5MW full model using BLADED by Lupton (2015).

### 3 Statistical convergence metric

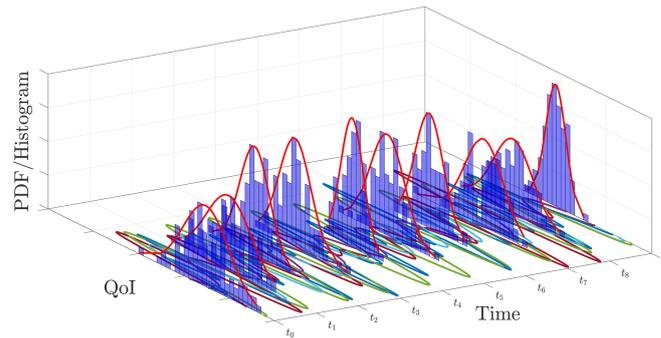
For this study, we want to investigate the null hypothesis that cross-sectional statistics (statistics at each time step) of a combination of a large number of aerodynamic simulation outputs are similar. In other words, we want to investigate if the statistical properties of the output at each time step converge as a function of the number of simulations (stationary process) for the non-linear stochastically autocorrelated system. Figure 2 presents a generic example of distributions (histogram fits) at each time step for a set of realizations of one random process. (The figure shows a schematic plot; therefore, histograms and fitted distributions *do not* represent the plotted trajectories.)

To show that the sectional statistics of a large number of simulations are convergent, we need a metric to quantify the difference between the distributions at each time step. There are different metrics for this purpose (Basu et al., 2011); for this study, we use *Hellinger distance* (Hellinger, 1909) as a metric due to its ease of application and interpretation. The Hellinger distance is the metric to quantify how similar two probability distributions are to each other. The distance is zero if they are the same and is one when the two distributions are disjoint. The Hellinger distance for two *discrete* probability distributions  $P$  and  $Q$ , which have an equal number of bins, can be formulated as follows:

$$H(P, Q) = \frac{1}{\sqrt{2}} \sqrt{\sum_{i=1}^k (\sqrt{p_i} - \sqrt{q_i})^2}. \quad (2)$$

In Eq. (2),  $p_i$  and  $q_i$  are the probabilities for  $P$  and  $Q$  at every bin. In our case, to make the comparison fair, not only the number of bins are the same, but also the bin width is the same for both  $P$  and  $Q$ . This assures us that there is no artificial distance reduction in the results.

In this study, we use Hellinger distance to show that the cross-sectional statistics changes for a large number of simulations are minimal. Therefore, we can shorten the simulations without losing the accuracy of the sectional statistics. We also use the Hellinger distance as an error metric to compare the accuracy of the MCS with the reference case. The reason behind choosing this measure, instead of simply look-



**Figure 2.** A schematic drawing presenting possible distributions at each time step based on a set of time trajectories for a quantity of interest (QoI).

ing at the mean and standard deviation difference, is the distribution of the aerodynamic model output. As the next section we will show, the distribution of the aerodynamic output is a Weibull distribution. Therefore, comparing mean and standard deviation would not provide us the full statistical picture.

### 4 Polynomial chaos expansion fundamentals

Uncertainty propagation of mathematical models has been the subject of many studies in the last 30 years. One method to propagate uncertainty is using models of models, called surrogate models. A surrogate model is a cheap-to-run approximation of the actual model (Kim and Boukouvala, 2020). Among surrogate models, the polynomial chaos expansion (PCE) has gained attention especially after the work of Ghanem and Spanos (2003) and Xiu and Karniadakis (2003). PCE is a method that uses a variable described by its statistical distribution (random variable  $\xi$ ) and projects the model onto a basis of polynomials. In this way, the uncertainty can be propagated through the model with a limited number of simulations (Tyson et al., 2015). In other words, PCE is a technique to estimate the response of a mathematical or numerical model based on a series of orthogonal polynomials, which are functions of a random variable  $\xi$ . The solution is expanded and described in stochastic space spanned by  $\xi$  and the associated polynomial basis.

The main reasons to use PCE instead of other surrogate model methods are as follows: (a) with minimum computational effort, one can extract statistical moments directly from PCEs; (b) PCEs are easy to integrate into deterministic linear and non-linear mathematical models; and (c) one can build a PCE surrogate model by treating the model as a black box (Kaintura et al., 2018; Sudret, 2015) using a non-intrusive formulation.

In order to illustrate the application of PCE surrogates, assume  $Y(t_n) = M(t_n, \xi)$ , where  $t_n$  is time step  $n$ ,  $\xi$  is the random variable vector,  $M(t_n, \xi)$  is our deterministic time

marching mathematical model, and  $Y(t_n)$  is the output of the model at time step  $n$ . Therefore, the stochastic output of the model  $Y(t_n, \xi)$  can be expanded as follows:

$$Y(t_n, \xi) = \sum_{i=1}^{\infty} y_i(t_n) \Psi_i(\xi), \quad (3)$$

where  $y_i(t_n)$  are PCE coefficients at *each time step*, and  $\Psi_i(\xi)$  is a member of an orthogonal polynomial class. These polynomials are orthogonal with respect to the probability space of random variable  $\xi$ . The selection of the polynomial type is a function of the probability distribution on the random variable  $\xi$ . For example, if a random variable  $\xi$  has a normal distribution, then a Hermite polynomial is selected (Xiu and Karniadakis, 2002).

The polynomials do not necessarily need to be selected from the specific polynomial family as long as they are orthogonal polynomials. For instance, Fluck and Crawford (2018) showed exponential components worked best for their purposes. As the randomness in this study comes in the form of a uniform distribution for wind frequency component phases  $\phi_j$ , the surrogate model is based on the Legendre polynomials (Xiu and Karniadakis, 2003). In practice, the PCE summation in Eq. (3) is truncated at a reasonably high order  $p$ . The task of fitting the expansion in Eq. (3) is finding the coefficients  $y_i(t)$ . There are two main approaches to solve this problem.

- For the *intrusive* approach, the model is projected on the orthogonal polynomials using a Galerkin projection (Ghanem et al., 2017). This approach requires building a detailed stochastic model from the deterministic model governing equations. The intrusive approach was used by Fluck and Crawford (2016b, 2018) to build a surrogate model on lifting line and BEM models.
- The *non-intrusive* approach allows for calculating the PCE coefficients from a series of deterministic model evaluations. This approach considers the model as a black box and does not require any model modification (Sudret, 2007; Eldred et al., 2008). There are two sub-categories to calculate the coefficients, namely *simulation methods* and *quadrature methods* (Sudret, 2007).

The presented work uses the *non-intrusive* approach to calculate the PCE coefficients. In the non-intrusive approach category we primarily used a *simulation method* to calculate the PCE coefficients. In mathematical form, the output of the aerodynamic model  $M_{\text{aero}}(t_n, \xi)$  at time step  $n$  is thrust  $\text{Trt}(t_n)$  and torque  $\text{Trq}(t_n)$ . Therefore, one can re-write Eq. (3) as follows:

$$\text{Trt}(t_n, \xi) \approx \sum_{i=1}^m T_i(t_n) \Psi_i(\xi), \quad (4)$$

$$\text{Trq}(t_n, \xi) \approx \sum_{i=1}^m \tau_i(t_n) \Psi_i(\xi), \quad (5)$$

where the goal is to calculate the polynomial coefficients  $T_i(t_n)$  and  $\tau_i(t_n)$  at each time step. This surrogate model's input is the random variable  $\xi$  vector used in the reduced Veers model to generate the unsteady wind. The surrogate model's output is the thrust and torque at a specific time step for which the surrogate model is built. The main difference between Eq. (3) and Eqs. (4) and (5) is the approximation with finite polynomial series expansion as it is not feasible to take into account an infinite number of polynomials. This work's surrogate model is built employing the Python toolbox `chaospy` (Feinberg and Langtangen, 2015). `chaospy` is a numerical tool for uncertainty quantification using different methods, including PCE. For this study, we used the *point collocation* method to calculate the coefficients due to the ease of implementation in the `chaospy` toolbox. This method has been explained well in the literature (see Xiu et al., 2002; Ghanem and Spanos, 2003; Sudret, 2007).

## 5 Surrogate modeling methodology

In this piece of work, whenever we talk about simulations, we mean aerodynamic simulations in time using the BEM aerodynamic model introduced above. The input of the aerodynamic simulations is what we call wind time series or unsteady wind and is auto-correlated by construct. This wind time series is generated based on the reduced Veers model explained above.

This study starts with running an extensive set of simulations based on the reduced Veers model at  $12 \text{ m s}^{-1}$  mean wind speed, aerodynamic simulation model, and Sobol sampling, explained previously. This wind speed is the rated power wind speed for NREL 5MW. Afterwards, as we have a large number of simulation outputs in our database, we can show that the thrust and torque statistics with time are not changing significantly. Therefore the statistical process properties at each time step (mean, standard deviation, etc.) would be significantly similar to other time steps. Knowing the process statistics stays the same in time, we conclude that only building a single surrogate model, i.e. a single time step or a few, would suffice for our purpose. The accuracy of the PCEs depends on the polynomial degree. However, an increase in the polynomial degree pushes the problem further toward the curse of dimensionality. The number of required coefficients to build the surrogate model and the required number of simulations are presented in Table 1. Equation (6) presents the formula to calculate the number of PCE coefficients. In Eq. (6),  $M$  is the required number of coefficients,  $N$  is the number of random variables, and  $P$  is the polynomial order.

$$M + 1 = \frac{(N + P)!}{P!N!} \quad (6)$$

The table shows we need a large number of simulations to build an accurate PCE. According to the recommendation of Hosder et al. (2012), twice the number of simulations ( $M +$

**Table 1.** The number of coefficients and required data points to calculate the coefficients for 10 random variables and using the point collocation method. This number of coefficients should be calculated for every time step. The last column demonstrates the simulation (sims) length for the fitted PCEs as explained in Sect. 6.

Polynomial degree	No. of coefficients + 1	No. of required simulations	Sims. length to fit PCEs [s]
1	11	22	163.6
2	66	132	27.3
3	268	536	6.7
4	1001	2002	1.8
5	3003	6006	0.6

1) can provide acceptable accuracy for the point collocation method (Hosder et al., 2012). That recommendation is the basis for the number of simulations in this study.

We want to limit the computational cost for a single average wind speed to 6 times 600 s simulations (in total 3600 s) to be competitive with the standard practice in wind turbine aerodynamic simulation. Combining a large number of simulations and 3600 s cumulative simulation length leads to a large number of short simulations instead of a few long ones. We kept the simulation's cumulative length at 3600 s to make this trade-off fair. This means that as the simulations' length decreases, the number of simulations increases. Sobol sampling is the base of the unsteady wind generation and input to the aerodynamic simulation setup. For every set of the required number of simulations in Table 1, the random phases are drawn independently from the rest of the sets. For example, for the second row of Table 1, when we need 132 simulations, 132 unique samples of  $\xi$  are drawn from the random domain. These  $\xi$  have not been used for other simulation sets. By having a large number of data points at each time step, we built a few surrogate models in time and compared the results with the simulations' reference case. For the sake of accuracy, in this study, we do not build any surrogate model based on first-degree polynomials.

Another approach to calculate the coefficients is Gaussian quadrature (GQ). This method has been extensively explained in the literature (e.g. Le Maître and Knio, 2010). There are also extensions to GQ referred to as sparse Gaussian quadrature (SGQ) methods that seek to reduce the number of simulations required to fit the surrogate (e.g. Smolyak). Our tests show that for a standard GQ method with 10 random variables and polynomial orders 2, 3, and 4, we need 59 049, 1 048 576, and 9 765 625 evaluation points, respectively. On the other hand, the Smolyak sparse approach for GQ (Le Maître and Knio, 2010; Smolyak, 1963) will reduce the number of evaluation points drastically. We tested the SGQ method for polynomial orders of 2, 3, and 4 with 10 uniform random variables and the Smolyak sparse approach for SGQ. The number of the required evaluations

for each polynomial order is 221, 1581, and 8761, respectively. We ran the evaluations for the SGQ method, calculated the weights, and built the polynomials. However, the results were not as promising as expected. The results from the SGQ method are shown and briefly discussed in Sect. 6.4

For a stationary input (reduced Veers model), the sample statistics of output converges at the rate of  $1/\sqrt{n}$ , where  $n$  is the number of data points (in this case, 48 000 data points at a one-time instance). Consequently, it is possible to *estimate* the statistical parameters of the output distribution by different methods. One possible approach is using the maximum likelihood estimator (Rao, 2008). A question that then arises is why we go through the complication of building a surrogate model. The research goal we present here is to build a surrogate model of an aerodynamic model, whether the aerodynamic model is simple or complex, with the model capable of resolving the form of the performance statistics, as an alternative to maximum likelihood estimation methods. We chose an aerodynamic model that is easy to simulate while complex enough to capture the inherent non-linear behaviour. Hence, the specific aerodynamic example model does not compromise the validity of the method we introduce here to later more complex aeroelastic simulations with, for example, FAST or BLADED.

## 6 Results

This section presents the results of our numerical experiments. We start by looking at the Hellinger distance of a large number of aerodynamic simulation output, thrust, and torque values and show that the distance does not change significantly. Therefore, the sectional statistics are almost the same across time steps. Afterwards, based on that conclusion, we built a number of surrogate models for polynomial orders of 2 to 5 from a limited number of simulations and show the statistics match the reference case. Finally, we show how extracted extreme thrust and torque are comparable with the reference case.

### 6.1 Sectional statistical convergence

As mentioned before, in Sect. 5, we started by running a broad set of reference simulations. For this case, we ran 48 000 simulations for a  $12 \text{ m s}^{-1}$  wind speed and turbulence intensity of 0.16. The wind generator code took 48 000 samples from a 10-dimensional uniform distribution domain based on the quasi-random sampling method. Each sample is a 10 by 1 vector of  $\xi_j$ , and we have 48 000 of them. A total of 48 000 wind speed time histories were generated, and simulations on the aerodynamic models run with a time step of 0.1 s for 630 s (in total 6300 time steps per simulation), with a 30 s transient period. We discarded the transient period for all the processes in this study.

This simulation setup builds a database for the investigation and shows that the process distributions at each time step

change are insignificant. Initially, we started calculating the histogram at each time step with identical binning for all of them. Afterwards, using the Hellinger distance formulation, each histogram's distance to the other histograms (5999 other histograms) was calculated and stored in a matrix. Each row of this matrix shows the histogram difference at one time step compared to the other ones. Therefore, this is a symmetric matrix with zeros on the diagonal. What is important is the maximum of all of the data in the matrix; in Fig. 3a and b, we show the max of the Hellinger distance at each time step for the aerodynamic model simulation outputs. The Hellinger distance is a normalized metric, and the distances are shown in the percentage. The plot shows a comparison of all the 18 million possible combinations to calculate the Hellinger distance for each model output. For the simulation outputs (thrust and torque), the distributions' difference does not exceed 2.21%. This shows a sound coherence in the statistics at each time step. Therefore, we can conclude that building a surrogate model on a limited number of time steps, or even one time step, is enough, and it is not necessary to create a surrogate model on every time step as predicted by the aerodynamic model form.

## 6.2 PCE surrogate model construction

We use the same simulation setup as explained above for the reference case to run the specified number of simulations in Table 1. These simulations are input for building the surrogate models. The number of samples drawn from the 10-dimensional uniform random space is equal to the number of simulations in Table 1. The employed sampling method is Sobol as tests show it provides a better convergence for the PCE.

Referring to Fig. 3 and the discussion in Sect. 6.1, the changes in statistical properties at each time step are minimal. Therefore, one surrogate model that can accurately emulate the sectional statistics of the aerodynamic simulations' output would suffice. Knowing this means building surrogate models is more feasible from a computational cost perspective.

As explained in Sect. 5, we fit surrogate models at every time step of a large set of short simulations instead of a few long ones for increasing polynomials of order  $P$ . The number of simulations is based on the polynomial order as shown in Table 1. The length of the simulations that the surrogate model is built on is the last column in Table 1 to keep the cumulative length of the simulations at 3600 s. Although it was unnecessary, for polynomial degrees 3 to 5, we built surrogate models at every time step of the whole 10 s worth of simulations to have an acceptable sample size for direct comparison. The `chaospy` (Feinberg and Langtangen, 2015) toolbox was used to perform the task of building these surrogate models.

Figure 4 compares the descriptive statistics (first quartile  $Q_1$ , second quartile  $Q_2$ , and third quartile  $Q_3$ ) for both thrust

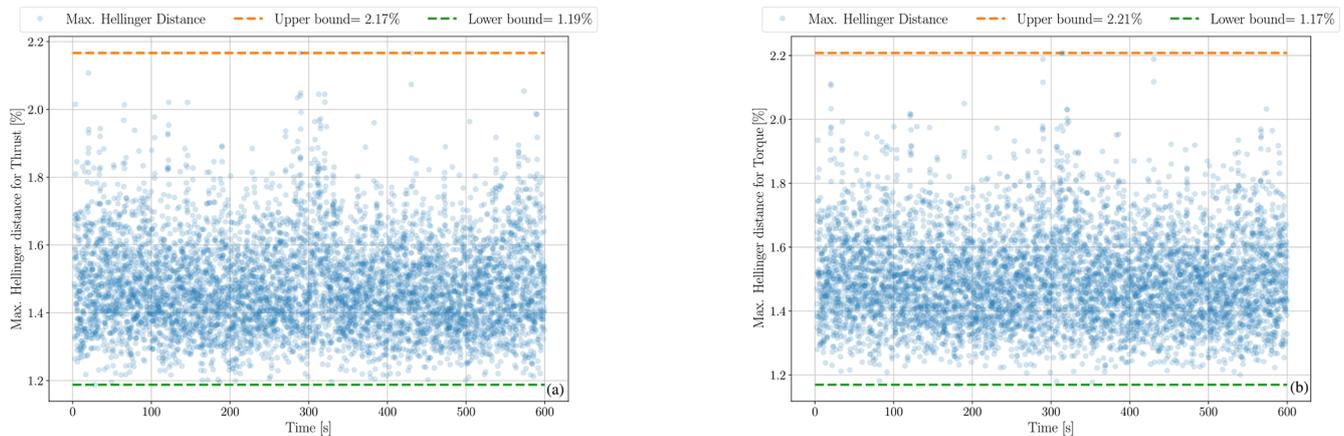
and torque from the reference case with 48 000 simulation outputs and 48 000 MCSs of the surrogate model built at each time step. The results in Fig. 4 show the PCE fits for four polynomial degrees;  $P$  on each plot indicates the polynomial degree. As the polynomial degree increases from  $P = 2$  to  $P = 4$ , the fit to the reference case improves, as is expected. However, it seems there is more error in the mean value and quantiles when we move from  $P = 4$  to  $P = 5$  for both thrust and torque. This increase in the error is explained further in Sect. 6.3 and Fig. 6. We calculate the average difference for the MCS and reference case over time for  $Q_1$ ,  $Q_2$ , and  $Q_3$  from Fig. 4 for polynomials  $P = 4$  and  $P = 5$ . This is presented in Table 2.

## 6.3 Surrogate model MCS

PCE surrogate models were constructed in the previous section. Those surrogates can now be exercised via MCS to quickly provide output-sampled statistics without actually running further simulations. We initially ranked the surrogate models constructed for each time step based on their mean values and standard deviations for each polynomial order separately to choose a single PCE surrogate model for the MCS carried out next. As the surrogate models are constructed as PCEs, mean and standard deviation extraction is a simple step from the PCE coefficients (Owen et al., 2017). After ranking the surrogate models, we selected the middle mean surrogate model from the ranked succession for each polynomial order. That provided us four surrogate models for thrust and four surrogate models for torque, one per polynomial order.

Next, we took the selected surrogate models through MCSs of the surrogate models  $10^6$  times. Essentially, we took random samples from our 10-dimensional random domains  $10^6$  times and inserted those in the PCEs (Eqs. 4 and 5). The MCS outputs are then used to check the surrogate model's accuracy. One can argue that this method is cherry-picking the surrogate models. This argument is valid for the low-order ( $P = 2, 3$ ) polynomial surrogate models. However, from Fig. 4 we know these polynomials are not accurate regardless of which one we choose. This inaccuracy is more visible in Fig. 6. For polynomials of orders 4 and 5, referring to Fig. 4, the polynomial selection procedure induces an insignificant effect on the statistics.

To verify the surrogate models' accuracy, we use the Hellinger distance explained in Sect. 5. This time, the Hellinger distance shows the difference between the surrogate model's  $10^6$  MCS outputs per polynomial order with the reference case at every time step. This procedure provides a vector of Hellinger distances for each polynomial degree, in which the vector's length is the same as the number of time steps in the simulations. As Hellinger distance is sensitive to binning, the bins are identical for each polynomial order surrogate model. The same bins were used to calculate the reference simulations' histograms at each time step. Figure 5



**Figure 3.** Maximum Hellinger distance for thrust and torque at each time step. The upper and lower bounds for the *extreme* of the Hellinger distance are indicated.

shows the average Hellinger distance changes within a narrow band for each polynomial degree. For the order 4 and 5 polynomials, we calculate the average of the Hellinger distances over 600 s. The averaged Hellinger distance presented in Table 2 serves as an error metric for the surrogate models.

Afterwards, we compare the histogram of those with one arbitrary time step of the reference case of 48 000 simulations. For each polynomial degree, regardless of the reference case time-step location in the time series, the difference between the reference case and the MCS only depends on the polynomial order. In other words, the difference between the MCS result histogram and the reference case histogram was only dependent on the polynomial degree and not the position of the time step in the time series as expected for a stationary process. Figure 6 compares the histogram of  $10^6$  MCS for the middle mean ranked surrogate model to the reference case at one arbitrary time step for four polynomial degrees.

Figure 6 shows how the surrogate models match the simulations' output histograms. It is visible that the polynomial orders  $P = 3$  and  $P = 4$  can cover the non-linearity in the results, while the second-order polynomial cannot. Polynomial order  $P = 5$  does not work well for both thrust and torque. Although we met the rule of thumb for the number of simulations as mentioned in Table 1, this shows an *under fit* for  $P = 5$ . This means more simulations are required to make the fit feasible. Both Figs. 5 and 6 show that  $P = 4$  provides an acceptable accuracy for the surrogate models. Therefore, further tests on the surrogate models with  $P = 5$  appear unnecessary. From Figs. 5 and 6 we can conclude that the PCE surrogate model with polynomial order  $P = 4$  is accurate enough to emulate the aerodynamic model with an acceptable accuracy while covering its non-linearity.

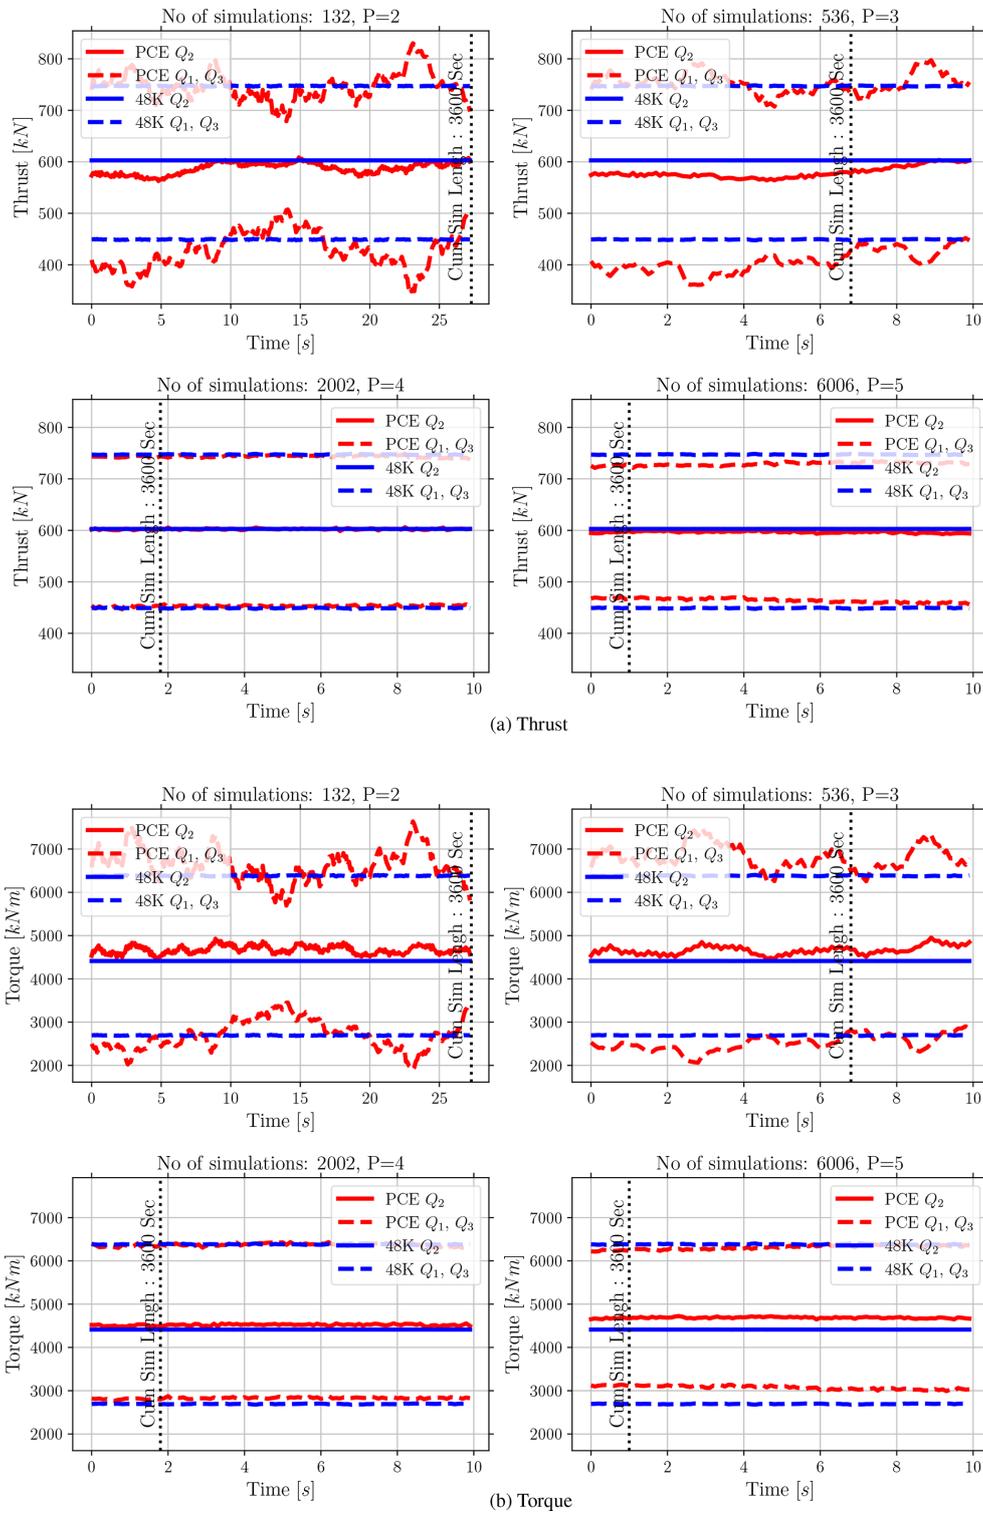
Another metric to show the accuracy of the surrogate model is normalized root mean square error (NRMSE), also known as  $L_2$  norm error. The NRMSE is calculated for every time step for the first 10 s by running MCS for the surrogate

model  $P = 4$  for the same input as the reference case simulations. This means we use the same samples we took from the  $10 - D$  random variables to generate the unsteady wind and then calculate the 48 000 reference case to run the surrogate model MCS. Figure 7 shows the error against time. As expected and visible from Fig. 4 the NRMSE is higher for torque and lower for thrust. In both cases, the maximum NRMSE is less than 10 %. This error is deemed acceptable as the surrogate model aims to provide overall accurate statistics and not point-to-point accuracy in the estimation and is necessarily a trade-off between speed and accuracy in the intended applications. Recall from Eqs. (4) and (5) that the PCE method is formulated as an expansion over the space formed by polynomials which are functions of random variables. The simulation method of fitting the PCE coefficients is essentially performing a statistical fit across the summative set of simulation results rather than optimizing the surrogate fit to a specific simulation. The NRMSE comparison here is therefore perhaps unfair to the intent of the PCE model, the earlier comparisons of MCS histograms and Hellinger distances being more appropriate metrics for the proposed PCE surrogate model approach.

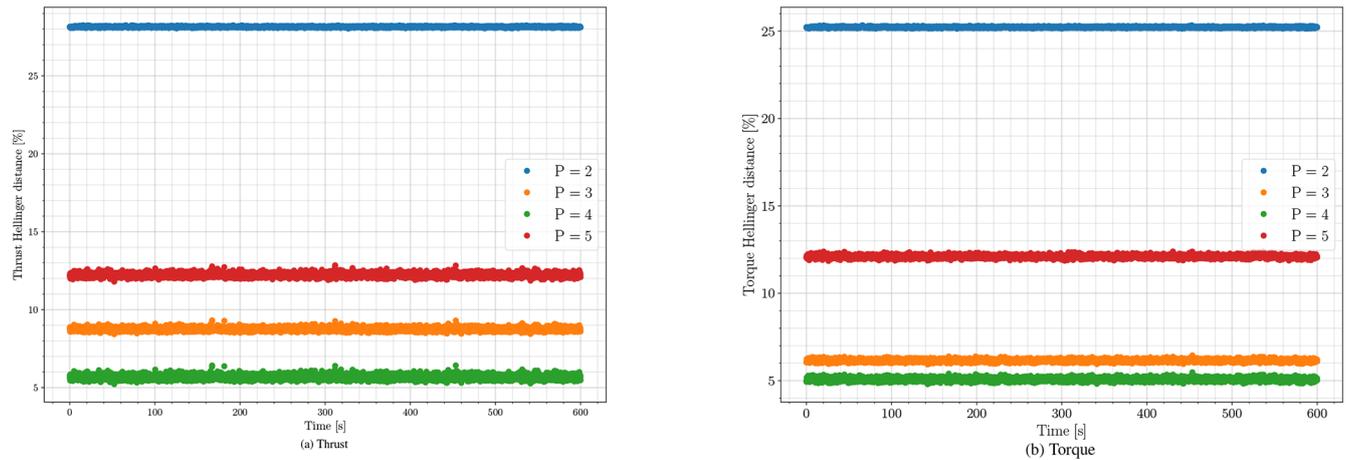
Figures 5, 6, and 7 show the PCE surrogate model has succeeded with the samples from the  $10 - D$  uniform distribution and converts them to an approximate Weibull distribution for both thrust and torque. This result highlights the ability of the PCE surrogate in this study to deal with the inherent non-linearity of the combination of unsteady wind generation and aerodynamic model together.

#### 6.4 SGQ PCE results

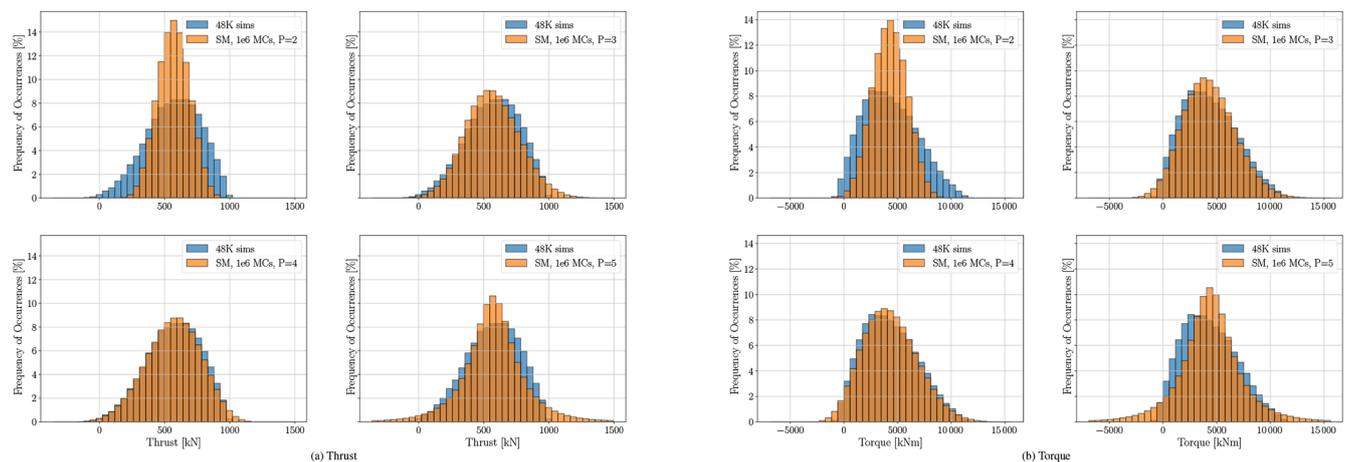
As mentioned in Sect. 5, we also calculated the coefficients for the PCE using the SGQ method for the polynomial orders 2, 3, and 4. We used the same procedure described in the previous sections to run the simulations, build the surrogate models (using `chaospy`), and select the surrogate



**Figure 4.** The  $Q_1$ ,  $Q_2$ , and  $Q_3$  comparison from the reference case (48 000 simulations) and extracted values from PCEs for both thrust and torque. The number of simulations used to build the PCEs and polynomial degree  $P$  are mentioned on the plots. The cumulative data length of 3600 s sufficient to build surrogates is shown with the vertical line.



**Figure 5.** Hellinger distance between the different polynomial order surrogate models for 1 million MCSs of the selected thrust and torque surrogates and the 48 000 reference case simulations at every time step.



**Figure 6.** Thrust and torque surrogate models with 1 million MCSs vs. the reference case histogram.

models. This results in six surrogate models (three for thrust and three for torque). We checked the accuracy of the surrogate models in the same manner as explained in the previous section. We ran  $10^6$  MCS on the six selected surrogate models and compared the histograms of the results with the histogram from one arbitrary time step from the reference case of 48 000 simulations. The results of the investigation are presented in Fig. 8.

The results in Fig. 8 show the SGQ method is less accurate than the point collocation method used in the previous section. Although the SGQ method requires more evaluation points than the point collocation method, the under-performance of the SGQ method is consistent for all the polynomial orders. As the initial accuracy test for the SGQ method did not provide comparable results with the point collocation method, we did not pursue further investigation of the SGQ method in this study. This finding is in line with the literature that shows point collocation typically out-

performs the SGQ in accuracy and efficiency (Eldred and Burkardt, 2009).

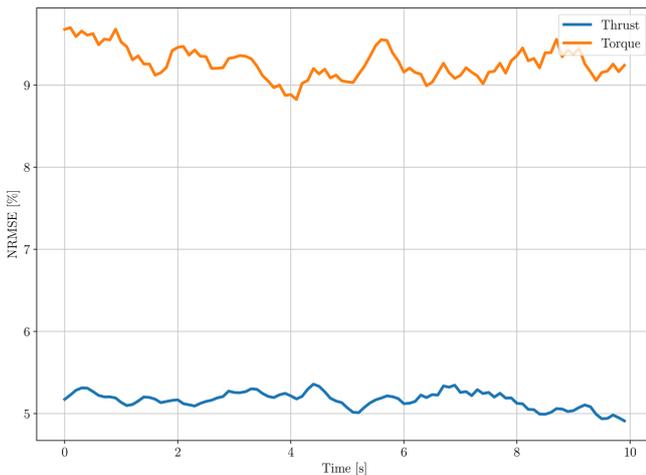
The Hellinger distance in Fig. 9a and b shows the difference between the SGQ surrogate model's  $10^6$  MCS outputs per polynomial order and the reference case at every time step. The Hellinger distance is much larger than what we showed in Fig. 5 using the point collocation method to build the surrogate models.

### 6.5 Surrogate model efficiency

Building the surrogate models aims to emulate the output of the actual model in an accurate and computationally efficient fashion. To inspect success in this respect, we start with computation time required to run 2002, 6006, and reference case 48 000 aerodynamic simulations of 600 s. The previous section shows that we do not need the 600 s length simulations to build the surrogate models. Based on what we showed in Sect. 6.3, the required time to run 1 s simulations (10 time

**Table 2.** The average Hellinger distance for polynomials  $P = 4$  and  $P = 5$  and the quantile value difference with respect to the reference case.

Case	No. of aerodynamic sims	Avg. Hellinger distance (thrust, torque) [%]	Thrust 1 million MCSs vs. 48 000 sims $Q_1$ , $Q_2$ , and $Q_3$ [%]	Torque 1 million MCSs vs. 48 000 sims $Q_1$ , $Q_2$ , and $Q_3$ [%]
Sims for PCE $P = 4$	2002	5.6, 5.0	0.90, 0.05, 0.50	4.94, 2.44, 0.21
Sims for PCE $P = 5$	6006	8.7, 6.1	3.36, 1.03, 2.42	14.12, 6.22, 1.10

**Figure 7.** Surrogate model  $P = 4$  NRMSEs for both thrust and torque with respect to the reference case.

steps) and build one surrogate model is provided. The computational time in Table 3 includes the unsteady wind generation. Additionally, we record the time required to build one surrogate model. The time to build one surrogate model for both thrust and torque is similar, and the average is provided here. As IEC standards (IEC 61400-1, 2019) ask for at least six aerodynamic simulations per average wind speed, we register the time for that set of aerodynamic simulations also. We perform aerodynamic simulations and surrogate model building on Compute Canada WestGrid clusters. The CPU time for the aerodynamic simulations and building the surrogate models is presented in Table 3.

The computational time to build one surrogate model is long for  $P = 5$ . This is due to employing the point collocation method to calculate the PCE coefficients. The point collocation method is inherently a regression method, using least squares to minimize the error (Feinberg and Langtangen, 2015). For a more complex aeroelastic model, the simulation times would be increased, potentially substantially, shifting the balance of computational time from PCE construction toward aeroelastic simulations. Of course, the aeroelastic simulations may be parallelized on available computing infrastructure.

After building the surrogate models, we ran large sets of MCSs for the PCE surrogate models with the polynomial or-

der 4 (as it is the most accurate one) and tracked the required time to run the MCSs. All the MCSs were performed on Compute Canada WestGrid clusters. The computation time for the MCSs is shown in Table 4. As the computational time difference between thrust and torque is insignificant, the one which took longer is mentioned here.

The number of time steps in the reference case is 288 million (6000 time steps multiplied by 48 000 aerodynamic simulations). Therefore, to have a fair comparison we can compare computational time for the reference case in Table 3 with 288 million MCSs in Table 4. Adding up the computational time required for the surrogate model input simulations and building the surrogate model, still the MCS is more efficient by a big margin.

The ease of running MCSs provides the ability to have more samples from the random domain. More samples from the random domain cover more of the statistical domain and capture the extreme loads more efficiently than running time marching aerodynamic simulations and extrapolation. Figures 10 and 11 present the comparison between different MCS set sizes and the reference case aerodynamic simulations, maximum, 99th percentile  $P_{99}$ , 95th percentile  $P_{95}$ , and 90th percentile  $P_{90}$ . The maximum load and the percentiles extracted from the  $P = 4$  surrogate models run for both thrust and torque are shown in relation to the number of MCSs.

According to IEC standards (IEC 61400-1, 2019), the maximum load needs to be calculated based on the mean of the max (mean-max) of at least six seeds of unsteady wind aerodynamic simulations per average wind speed. Therefore, we randomly grouped the reference case simulations (48 000 simulations) into 8000 unique groups to have a fair comparison with the common practice. Afterwards, the mean of the max per group, as well as the 90th percentile  $P_{99}$ , 95th percentile  $P_{95}$ , and 90th percentile  $P_{90}$  of each group, is calculated. These data are presented in Figs. 10 and 11 as clouds of dots for both thrust and torque.

To extract the maximum value that matches the full reference case aerodynamic simulations, we need to run a large number of MCSs. Figures 10 and 11 show for the PCE order 4 the maximum thrust and torque from the MCS matches the reference case around almost 1000 MCSs. Looking at  $P_{99}$ ,  $P_{95}$ , and  $P_{90}$  for both surrogate models, the convergence happen at around the same number of MCSs. Figures 10 and

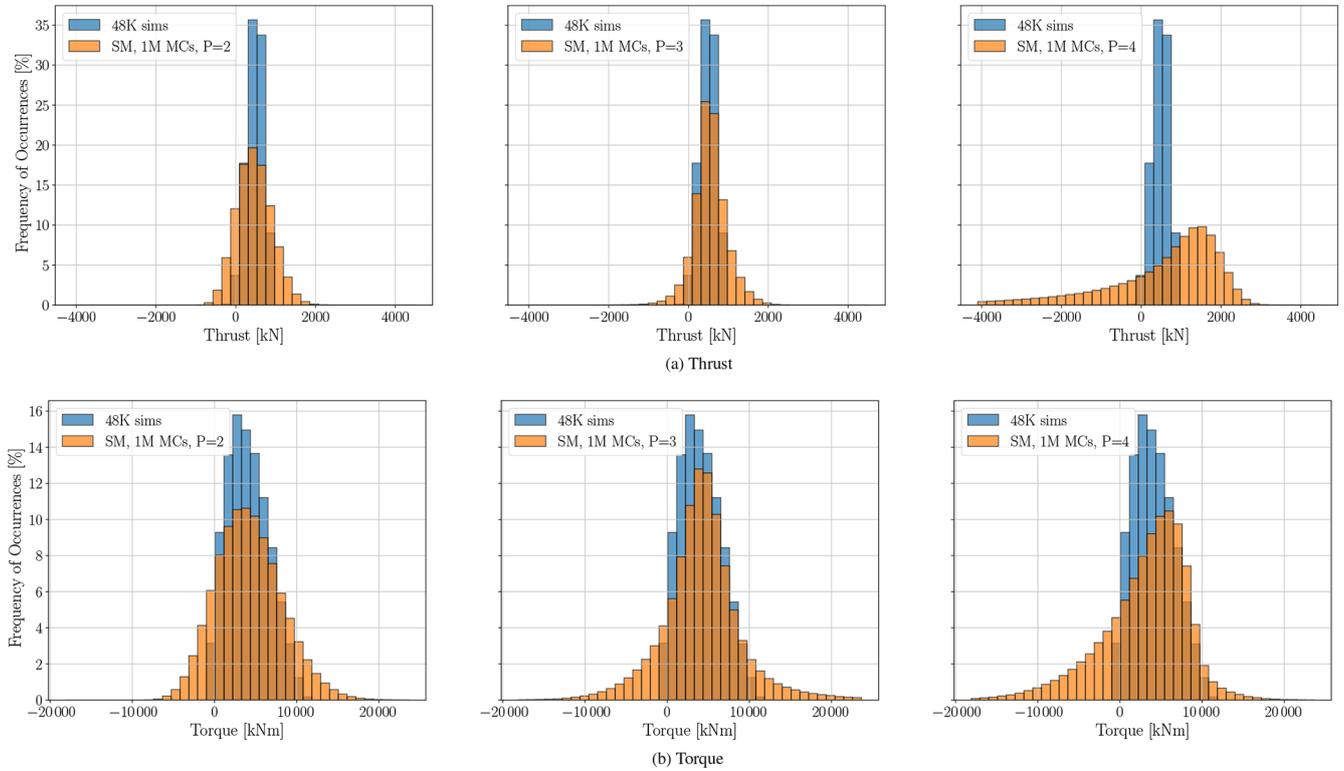


Figure 8. Thrust and torque surrogate models based on SGQ with 1 million MCSs vs. the reference case histogram.

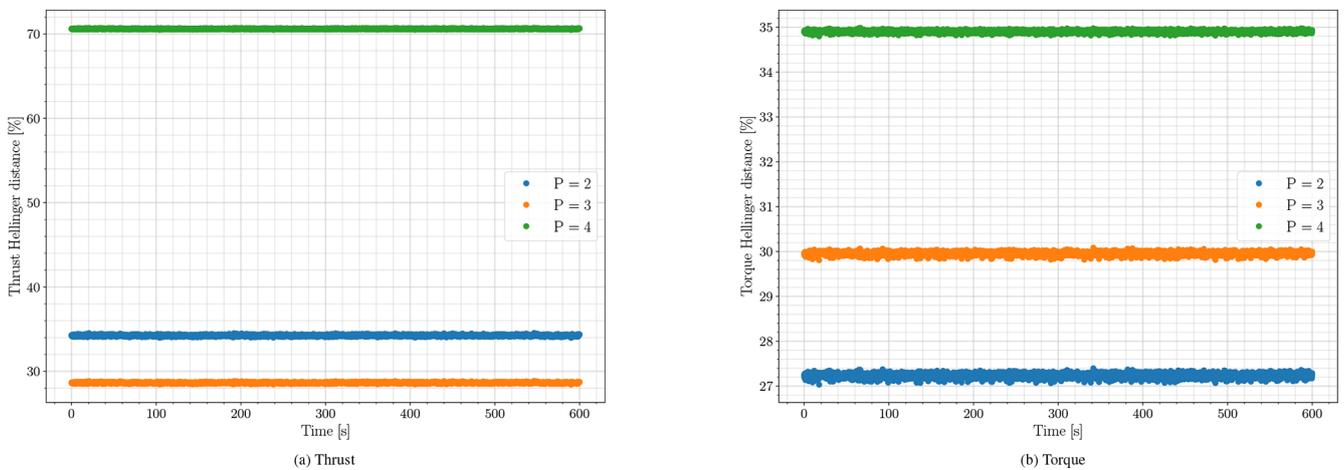


Figure 9. Hellinger distance between the different polynomial order surrogate models for 1 million MCSs of the selected thrust and torque SGQ surrogates and the 48 000 reference case simulations at every time step.

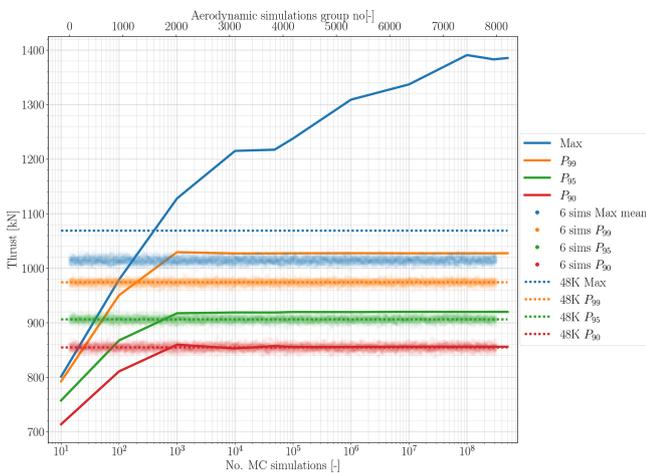
**Table 3.** The computational time required to run aerodynamic simulations, building the surrogate models, and the average Hellinger distance.

Case	No. of aerodynamic sims	Simulation length	Computation time	Time required to build one surrogate model
Sims for PCE $P = 4$	2002	1 s	10 s	1 min 2 s
Sims for PCE $P = 5$	6006	1 s	31 s	15 min 54 s
Common practice	6	600 s	18.62 s	n/a
Reference case	48 000	600 s	41 h 55 min 48 s	n/a

n/a – not applicable.

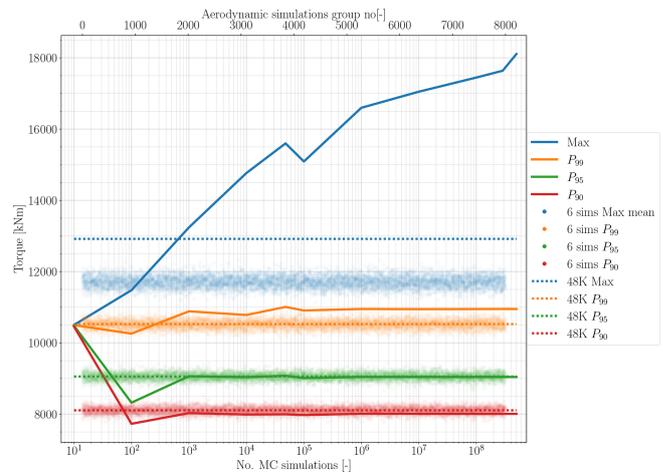
**Table 4.** Computational time to run MCSs on the surrogate models with polynomial order 4.

No. of MCSs	Surrogate model $P = 4$
10, 100, 1000	< 0.1 s
10 000	< 0.25 s
48 000	0.85 s
100 000	1.70 s
1 million	30.84 s
10 million	5 min 51 s
100 million	55 min 33 s
288 million	2 h 40 min 1 s
500 million	4 h 46 min 36 s



**Figure 10.** Maximum load and percentile comparison between the aerodynamic simulation reference case, the MCSs, and group aerodynamic simulations for thrust.

11 show that after 1 or 100 000 MCSs the percentiles are close to the reference case. Interestingly, the mean-max output from the grouped aerodynamic simulations has a wide distribution. This shows the inaccuracy of looking at a small number of simulations to calculate the extreme loads. This distribution is smaller for the percentile data; however, it is not comparable with the convergence of MCS outputs. Also, looking at the grouped simulation output, and com-



**Figure 11.** Maximum load and percentile comparison between the aerodynamic simulation reference case, the MCSs, and group aerodynamic simulations for torque.

pared to the standard practice (mean-max) and  $P_{99}$ , illustrates the conservative design approach of the IEC standards (IEC 61400-1, 2019).

Referring to the computational time required to build the surrogate model and then run the MCS, these plots show promising results to extract accurate extreme loads from the surrogate models in a computationally efficient manner. Here again, it is emphasized that the utility of the proposed PCE-based surrogate, with MCSs of the constructed surrogate and examination of the statistical load distributions, is the key contribution of the work. The point-to-point accuracy of the model for a single run of the surrogate, as discussed earlier, is not the focus of the surrogate, but rather our focus is the overall computational costs and accuracy in spread of loading conditions.

## 7 Conclusion

This paper’s aim is to build a non-intrusive surrogate model of time marching aerodynamic simulations. The form of the surrogate model used in this paper is a PCE. In Sect. 5, we explained the aerodynamic model used for this study. Then,

we briefly described the method that we are using to build the PCEs. One major challenge with the building of the surrogate models is the curse of dimensionality, which we tried to tackle by using a reduced Veers model.

We showed how by increasing the number of simulations, the results' statistics converge and do not change in time. As a result, building a few accurate surrogate models, or even one, for a small length of time would suffice for our purpose. In other words, time does not have any meaning in the sectional statistics. Therefore, to build an accurate surrogate model, we can significantly reduce the simulation length while increasing the number of simulations. In the results section, we showed the surrogate model using a fourth-order polynomial built on 2002 simulations with a length of 2 s gives us sufficiently accurate results in large MC runs to obtain output statistics. Afterwards, we demonstrate the surrogate model's efficiency by comparing the computational time required to run the aerodynamic simulations and build the surrogate model with the required time for running MCSs to have accurate statistics. Also, we show that the high percentile values extracted from the MCSs match the reference case with a relatively low number of MCSs, and thus the model is computationally efficient.

The BEM-based aerodynamic model approach is well known in the literature and research. We choose 30 s transient time for the simulations to ensure they do not include any transition results. As the model is a less complex BEM which is quick to run, this is not a challenge. However, for future work with actual aeroelastic codes (e.g. FAST), a smart way to deal with initialization time is essential; otherwise, increasing the number of simulations and the model complexity would be very expensive. For example, if the required initialization time is 60 s (default in FAST), and we want to increase the number of simulations from six 600 s simulations (minimum requirement according to IEC 61400-1, 2019) to 6006 2 s simulations, we are not doing any good in terms of computational cost. Aeroelastic and longer wakes will be studied for this challenge, and a blended or common initialization period will be trialed.

Another challenge is the practical application of this surrogate model. The surrogate model that we built in this study is one or a few time steps, each inherently the same due to the statistical similarity. If we want to build one time series from this surrogate model, we have to sample the 10-dimensional random domain for the number of time steps to have a time series to post-process. For example, suppose we want to have a 600 s time series of thrust or torque with the time step of 0.1 s for the aerodynamic model that we developed in this study. In that case, we need to take 6000 samples from the 10-dimensional uniform distribution random domain and run MCSs for each. However, this would provide us with 6000 thrust and torque values, which will miss the auto-correlation, which is inherent in the generated unsteady wind, in the results. This drawback is crucial if we, for instance, want to calculate fatigue loads from the surrogate

model. This challenge will require a surrogate form capable of resolving the correlation between time steps. Fluck and Crawford (2018) did this previously for intrusive PCEs of an aerodynamic model; however, as mentioned before, that can get very complicated for more advanced models.

Using non-conventional polynomials, such as what Fluck and Crawford (2018) did, might result in a more efficient polynomial that requires fewer simulations to build the surrogate model. Finally, we want to implement the method on commercial wind turbine simulation packages such as FAST to test the approach in future work. It is important to again note, however, that the physics model used in the current work is equivalent to FAST, just conveniently formulated in Python for our surrogate model development efforts. The notion of the "reduced Veers model" worked for the aerodynamic simulations we used in this study. However, this reduced model would not be efficient and sufficient to move towards commercial wind turbine packages. Therefore, a new approach to reduce the data in a "turbulence box" (Jonkman, 2009) spatially and temporally would be necessary. Similar work has been done in Guo and Ganapathysubramanian (2017) and will be explored together with expansions of the methods investigated by Fluck and Crawford (2017) using velocity increments across the wind field.

**Code and data availability.** The code and data that support the findings of this study are available from the corresponding author, Rad Haghi, upon request.

**Author contributions.** RH developed the necessary computer code and wrote the paper in consultation with and under the supervision of CC.

**Competing interests.** The contact author has declared that neither of the authors has any competing interests.

**Disclaimer.** Publisher's note: Copernicus Publications remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Acknowledgements.** We greatly acknowledge the funding for this study by the Natural Sciences and Engineering Research Council of Canada (NSERC). This research was enabled in part by support provided by Digital Research Alliance of Canada (<https://alliancecan.ca/>, last access: 22 June 2022).

**Financial support.** This research has been supported by the Natural Sciences and Engineering Research Council of Canada.

**Review statement.** This paper was edited by Athanasios Kolios and reviewed by two anonymous referees.

## References

- Ashuri, T., Zhang, T., Qian, D., and Rotea, M.: Uncertainty quantification of the levelized cost of energy for a 20 MW research wind turbine model, in: 34th Wind Energy Symposium, p. 1998, <https://doi.org/10.2514/6.2016-1998>, 2016.
- Barlas, T., Ramos-García, N., Pirrung, G. R., and González Horcas, S.: Surrogate-based aeroelastic design optimization of tip extensions on a modern 10 MW wind turbine, *Wind Energ. Sci.*, 6, 491–504, <https://doi.org/10.5194/wes-6-491-2021>, 2021.
- Basu, A., Shioya, H., and Park, C.: *Statistical Inference: The Minimum Distance Approach*, Chapman & Hall/CRC Monographs on Statistics & Applied Probability, CRC Press, <https://doi.org/10.1201/b10956>, 2011.
- Bossanyi, E., Burton, T., Sharpe, D., and Jenkins, N.: *Wind energy handbook*, Wiley, New York, <https://doi.org/10.1002/9781119992714>, 2011.
- Dimitrov, N.: Surrogate models for parameterized representation of wake-induced loads in wind farms, *Wind Energy*, 22, 1371–1389, 2019.
- Dimitrov, N., Kelly, M. C., Vignaroli, A., and Berg, J.: From wind to loads: wind turbine site-specific load estimation with surrogate models trained on high-fidelity load databases, *Wind Energ. Sci.*, 3, 767–790, <https://doi.org/10.5194/wes-3-767-2018>, 2018.
- DNV GL: *Bladed User Manual Version 4.9*, Garrad Hassan & Partners Ltd, Bristol, UK, <https://www.dnv.com/services/wind-turbine-design-software-bladed-3775> (last access: 22 June 2022), 2018.
- Eldred, M. and Burkardt, J.: Comparison of non-intrusive polynomial chaos and stochastic collocation methods for uncertainty quantification, in: 47th AIAA Aerospace Sciences Meeting including The New Horizons Forum and Aerospace Exposition, p. 976, <https://doi.org/10.2514/6.2009-976>, 2009.
- Eldred, M., Webster, C., and Constantine, P.: Evaluation of non-intrusive approaches for Wiener-Askey generalized polynomial chaos, in: 10th AIAA Non-Deterministic Approaches Conference, p. 1892, <https://doi.org/10.2514/6.2008-1892>, 2008.
- Feinberg, J. and Langtangen, H. P.: Chaospy: An open source tool for designing methods of uncertainty quantification, *J. Comput. Sci.*, 11, 46–57, 2015.
- Fluck, M. and Crawford, C.: Fast analysis of unsteady wing aerodynamics via stochastic models, *AIAA J.*, 55, 719–728, 2016a.
- Fluck, M. and Crawford, C.: A stochastic aerodynamic model for stationary blades in unsteady 3D wind fields, *J. Phys. Conf. Ser.*, 753, 082009, <https://doi.org/10.1088/1742-6596/753/8/082009>, 2016b.
- Fluck, M. and Crawford, C.: An engineering model for 3-D turbulent wind inflow based on a limited set of random variables, *Wind Energ. Sci.*, 2, 507–520, <https://doi.org/10.5194/wes-2-507-2017>, 2017.
- Fluck, M. and Crawford, C.: A fast stochastic solution method for the Blade Element Momentum equations for long-term load assessment, *Wind Energy*, 21, 115–128, 2018.
- Ghanem, R., Higdon, D., and Owhadi, H.: *Handbook of uncertainty quantification*, Springer, <https://doi.org/10.1007/978-3-319-12385-1>, 2017.
- Ghanem, R. G. and Spanos, P. D.: *Stochastic finite elements: a spectral approach*, Courier Corporation, <https://doi.org/10.1007/978-1-4612-3094-6>, 2003.
- Guo, Q. and Ganapathysubramanian, B.: Incorporating a stochastic data-driven inflow model for uncertainty quantification of wind turbine performance, *Wind Energy*, 20, 1551–1567, 2017.
- Hellinger, E.: Neue begründung der theorie quadratischer formen von unendlichvielen veränderlichen, *J. Reine Angew. Math.*, 1909, 210–271, 1909.
- Hosder, S., Walters, R., and Balch, M.: Efficient sampling for non-intrusive polynomial chaos applications with multiple uncertain input variables, in: 48th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, p. 1939, <https://doi.org/10.2514/6.2008-1892>, 2012.
- IEC 61400-1: *Wind energy generation systems – Part 1: Design requirements*, Standard, International Electrotechnical Commission, 2019.
- Jonkman, B. J.: *TurbSim user’s guide: Version 1.50*, 2009.
- Jonkman, J., Butterfield, S., Musial, W., and Scott, G.: *Definition of a 5-MW reference wind turbine for offshore system development*, Tech. rep., National Renewable Energy Lab.(NREL), Golden, CO, USA, <https://doi.org/10.2172/947422>, 2009.
- Jonkman, J. M. and Buhl Jr., M. L.: *FAST User’s Guide – Updated August 2005*, U.S. Department of Energy Office of Scientific and Technical Information, <https://doi.org/10.2172/15020796>, 2005.
- Kaintura, A., Dhaene, T., and Spina, D.: Review of Polynomial Chaos-Based Methods for Uncertainty Quantification in Modern Integrated Circuits, *Electronics*, 7, 30, <https://doi.org/10.3390/electronics7030030>, 2018.
- Kim, S. H. and Boukouvala, F.: Machine learning-based surrogate modeling for data-driven optimization: a comparison of subset selection for regression techniques, *Optim. Lett.*, 14, 989–1010, <https://doi.org/10.1007/s11590-019-01428-7>, 2020.
- Kucherenko, S., Albrecht, D., and Saltelli, A.: Exploring multi-dimensional spaces: a Comparison of Latin Hypercube and Quasi Monte Carlo Sampling Techniques, *arXiv*, <https://doi.org/10.48550/arXiv.1505.02350>, 2015.
- Larsen, T. J. and Hansen, A. M.: *How 2 HAWC2, the user’s manual*, DTU, ISBN 978-87-550-3583-6, 2007.
- Le Maître, O. and Knio, O. M.: *Spectral methods for uncertainty quantification: with applications to computational fluid dynamics*, Springer Science & Business Media, <https://doi.org/10.1007/978-90-481-3520-2>, 2010.
- Lupton, R.: *Frequency-domain modelling of floating wind turbines*, PhD thesis, University of Cambridge, <https://doi.org/10.17863/CAM.14119>, 2015.
- Lupton, R.: *bemused*, GitHub [code], <https://github.com/ricklupton/bemused> (last access: 22 June 2022), 2019.
- Murcia, J. P., Réthoré, P.-E., Dimitrov, N., Natarajan, A., Sørensen, J. D., Graf, P., and Kim, T.: Uncertainty propagation through an aeroelastic wind turbine model using polynomial surrogates, *Renew. Energ.*, 119, 910–922, 2018.
- Niederreiter, H.: *Random Number Generation and Quasi-Monte Carlo Methods*, Society for Industrial and Applied Mathematics, SIAM, <https://doi.org/10.1137/1.9781611970081>, 1992.
- Owen, N. E., Challenor, P., Menon, P. P., and Bennani, S.: Comparison of surrogate-based uncertainty quantification methods for computationally expensive simulators, *SIAM/ASA Journal on Uncertainty Quantification*, 5, 403–435, 2017.

- Rao, S. S.: A course in time series analysis, [https://web.stat.tamu.edu/~suhasini/teaching673/time\\_series.pdf](https://web.stat.tamu.edu/~suhasini/teaching673/time_series.pdf) (last access: 22 June 2022), 2008.
- Schröder, L., Dimitrov, N. K., Verelst, D. R., and Sørensen, J. A.: Wind turbine site-specific load estimation using artificial neural networks calibrated by means of high-fidelity load simulations, *J. Phys. Conf. Ser.*, 1037, 062027, <https://doi.org/10.1088/1742-6596/1037/6/062027>, 2018.
- Schröder, L., Dimitrov, N. K., and Sørensen, J. A.: Uncertainty propagation and sensitivity analysis of an artificial neural network used as wind turbine load surrogate model, *J. Phys. Conf. Ser.*, 1618, 042040, <https://doi.org/10.1088/1742-6596/1618/4/042040>, 2020a.
- Schröder, L., Dimitrov, N. K., and Verelst, D. R.: A surrogate model approach for associating wind farm load variations with turbine failures, *Wind Energ. Sci.*, 5, 1007–1022, <https://doi.org/10.5194/wes-5-1007-2020>, 2020b.
- Slot, R. M., Sørensen, J. D., Sudret, B., Svenningsen, L., and Thøgersen, M. L.: Surrogate model uncertainty in wind turbine reliability assessment, *Renew. Energ.*, 151, 1150–1162, 2020.
- Smolyak, S.: Quadrature and interpolation formulas for tensor products of certain classes of functions, in: *Doklady Akademii Nauk, Russian Academy of Sciences*, 148, 1042–1045, 1963.
- Sobol, I. M.: On the distribution of points in a cube and the approximate evaluation of integrals, *Zhurnal Vychislitel'noi Matematiki i Matematicheskoi Fiziki*, 7, 784–802, [https://doi.org/10.1016/0041-5553\(67\)90144-9](https://doi.org/10.1016/0041-5553(67)90144-9), 1967.
- Sudret, B.: Uncertainty propagation and sensitivity analysis in mechanical models – Contributions to structural reliability and stochastic spectral methods, PhD thesis, Université Blaise Pascal – Clermont II, <https://ethz.ch/content/dam/ethz/special-interest/baug/ibk/risk-safety-and-uncertainty-dam/publications/reports/HDRSudret.pdf> (last access: 22 June 2022), 2007.
- Sudret, B.: Polynomial chaos expansions and stochastic finite element methods, *Risk and reliability in geotechnical engineering*, edited by: Phoon, K.-K. and Ching, J., 265–300, <https://doi.org/10.1201/b17970>, 2015.
- Tyson, S., Donovan, D., Thompson, B., Lynch, S., and Tas, M.: Uncertainty Modelling with Polynomial Chaos Expansion: Stage 1 – Final Report, The University of Queensland, ISBN 978-1-74272-173-6, 2015.
- van den Bos, L., Sanderse, B., Blonk, L., Bierbooms, W., and van Bussel, G.: Efficient ultimate load estimation for off-shore wind turbines using interpolating surrogate models, *J. Phys. Conf. Ser.*, 1037, 062017, <https://doi.org/10.1088/1742-6596/1037/6/062017>, 2018.
- van Garrel, A.: Development of a Wind Turbine Aerodynamics Simulation Module, Tech. rep., Energy research Centre of the Netherlands, <https://doi.org/10.13140/RG.2.1.2773.8000>, 2003.
- Veers, P. S.: Three-dimensional wind simulation, Tech. rep., Sandia National Labs., Albuquerque, NM, USA, <https://www.osti.gov/biblio/6633902> (last access: 22 June 2022) 1988.
- Wang, H., Jiang, X., Chao, Y., Li, Q., Li, M., Chen, T., and Ouyang, W.: Numerical optimization of horizontal-axis wind turbine blades with surrogate model, *P. I. Mech. Eng. A-J. Pow.*, 235, 1173–1186, <https://doi.org/10.1177/0957650920976743>, 2020.
- Xiu, D. and Karniadakis, G. E.: The Wiener–Askey polynomial chaos for stochastic differential equations, *SIAM J. Sci. Comput.*, 24, 619–644, 2002.
- Xiu, D. and Karniadakis, G. E.: Modeling uncertainty in flow simulations via generalized polynomial chaos, *J. Comput. Phys.*, 187, 137–167, 2003.
- Xiu, D., Lucor, D., Su, C.-H., and Karniadakis, G. E.: Stochastic modeling of flow-structure interactions using generalized polynomial chaos, *J. Fluids Eng.*, 124, 51–59, 2002.