https://doi.org/10.5194/wes-2025-112
Preprint. Discussion started: 1 July 2025

eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

Load Estimation in Onshore Wind Farms Using Surrogate Modeling and Generic Turbine Models

Alexander Mönnig[1,2], Ansgar Hahn[1], Astrid Lampert[3], and Ulrich Römer[2]

[1]Alterric GmbH, Holzweg 87, 26605 Aurich, Germany
[2]TU Braunschweig, Institute for Acoustics and Dynamics, Langer Kamp 19, 38106 Braunschweig, Germany
[3]TU Braunschweig, Institute of Flight Guidance, Hermann-Blenk-Str. 27, 38108 Braunschweig, Germany

Correspondence: Alexander Mönnig (alexander.moennig@alterric.com)

Abstract. This article investigates the development and application of surrogate models, based on slightly adapted generic turbine models, for predicting loads on real-world wind turbines. A small set of aeroelastic simulations provided training data for both Polynomial Chaos expansion and Gaussian Process regression models, which were trained to predict blade loads, tower accelerations, and their respective seed-to-seed variability. To evaluate the practical suitability of these models a case study was performed. Here, the surrogate models were applied to predict blade loads and tower accelerations respectively, using five years of SCADA data from an onshore wind farm. While the models approximated the real-world turbine behavior with a reasonable accuracy, the prediction quality varied across the different turbines in the park and was further influenced by factors such as the turbine's operational years and diurnal patterns suggesting a correlation with the turbulence intensity. Despite some limitations, the findings support the practicality of developing surrogate models for enabling efficient load estimations.

1 Introduction

Reducing downtimes and extending the lifetime of wind turbines are two important means for improving the profitability and hence, for increasing the future share of wind energy. These goals could be reached by assessing historic and managing future turbine loads based on aeroelastic simulations, which in turn could enable the optimization of operation and maintenance (O&M) strategies.

Two factors can be indentified, which limit the large scale application of aeroelastic simuators to these specific use cases:

1. **Computational burden:** In most use cases, a large number of evaluations are required to cover the range of operating conditions and obtain results that are statistically meaningful, where the latter problem originates in the inherent stochastic nature of turbulence modeling and an associated seed-to-seed variability in the simulations.

2. **Limited access to accurate turbine models:** High-resolution turbine models are usually proprietary to turbine manufacturers. This particularly limits the application of aeroelastic simulators in operational settings where load assessments could benefit turbine operators to improve their O&M strategies.

To address the computational burden, surrogate models have emerged as a promising solution to reduce the computational cost, while still providing accurate representations of complex simulation frameworks.

https://doi.org/10.5194/wes-2025-112
Preprint. Discussion started: 1 July 2025

WIND ENERGY SCIENCE DISCUSSIONS

Various types of surrogate models, including the Polynomial Chaos Expansion (PCE), Murcia et al. (2018), Artificial Neural Networks (ANN), Gasparis et al. (2020), Response Surface Methods (RSM), Toft et al., and Gaussian Process Regression (GPR), Wilkie and Galasso (2021), have shown significant potential in accelerating the estimation of aeroelastic loads from wind turbine simulators.

Previous research has thereby explored surrogate models with different purposes. For example, Murcia et al. (2018) employed PCE to predict fatigue loads and power generation for isolated wind turbines under site-specific inflow conditions. In contrast, Dimitrov (2019) developed surrogate models capable of predicting loads within a farm layout, incorporating parameters such as row spacing and the number of upstream turbines.

While surrogate models offer significant computational advantages in their execution, the stochastic nature of aeroelastic simulation due to turbulence-related seed-to-seed variability can make their development still cost intensive. In the literature, two main approaches have been used to handle this variability. The first class of methods requires realizations with different seeds using the same vector of operational parameters. While only six realizations are recommended for estimating the mean of the loads by the IEC (International Electrotechnical Commission) 61400-1:2019, capturing the variance across these realizations requires a significantly larger sample size. For instance, Murcia et al. (2018) utilized 100 realizations in their study. Because of the high cost associated with to the seed-to-seed variability, authors often focus solely on the mean prediction with a limited set of realizations. For instance, Dimitrov et al. (2018) used eight replications per sample while focusing primarily on the mean loads. On the other hand, Murcia et al. (2018) employed two PCE models to predict both the mean and standard deviation of loads, capturing the heteroscedastic nature of load variance, though at a high computational cost. In a more general context, a surrogate model based on the generalized Lambda distribution and regression data with replication has been proposed in Zhu and Sudret (2020).

The other category of approaches, instead, does not rely on replications. GPR is to be mentioned here, as a classical example. More recently, replication-free versions of the generalized Lambda surrogate model, Zhu and Sudret (2021), and random field based surrogates, Lüthen et al. (2023), have been proposed.

To address the second factor - limited access to accurate turbine models - Reference Wind Turbine (RWT) models could provide a potential workaround. These pre-implemented generic turbine models, cf. Gambier (2022), are openly available and have been used extensively in research. While RWT models are not exact representations of specific commercial turbines, they capture the fundamental physics and design principles of modern wind turbines. This could make them potentially suitable for preliminary load assessments and as a starting point for model refinements to better match specific turbine characteristics.

In this article we investigate whether RWT-models, in combination with a surrogate-based simulation of turbine loads, can overcome the challenges of computational burden and limited access to accurate turbine models identified earlier. The study builds on a large database of historic SCADA data that can be used to assess the accuracy of the combined surrogate - RWT model. Specifically, we will analyze the viability of the generated surrogates in a case study by applying them for the prediction of selected blade bending moments and tower vibrations recorded in an onshore wind farm in Germany. Given that the turbulence-induced variance sets a natural limit on predictive accuracy, our surrogate models - based on PCE and GPR - will estimate both the mean and the standard deviation of the loads.

https://doi.org/10.5194/wes-2025-112
Preprint. Discussion started: 1 July 2025

eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

The remainder of this paper is structured as follows: Section 2 presents the methodology. Section 3 outlines the case study setup which serves to evaluate the proposed approach. Section 4 presents the results of the model training and application. Finally, some conclusions are drawn and directions for future research are suggested.

2 Surrogate Modeling

Complex numerical models cannot be used directly for carrying out parametric studies, uncertainty quantification and other multi-query tasks. It is therefore common to employ surrogate models to reduce the computational complexity while simultaneously controlling the approximation error. This section will present theory and methods employed for constructing the surrogate models in this study. The process of developing the surrogate models is here divided into four key steps:

1. Identification of features for characterizing load variations.
2. Characterization of site-specific operating conditions through appropriate distributions and sample generation.
3. Creation of a load database linking operational conditions to specific turbine loads via aeroelastic simulations.
4. Construction of the surrogate models.

These steps are detailed in the following sections, where we present our modeling choices. An implementation is then presented in Section 3, where surrogates are built for reconstructing historic loads of an onshore wind farm over a time period of 5 years.

2.1 Features for Characterizing Load Variations

The surrogate models are developed using four features identified from similar load studies, see e.g. Dimitrov et al. (2018). Note that, unlike some recent work by Dimitrov (2019), we exclude wind farm layout features and instead use the turbulence intensity of each (wake-affected) inflow to account for farm effects. As this study focuses on reconstructing loads from 10-minute SCADA data, we selected features that are available (directly or indirectly) from industry-standard SCADA systems. The features considered here are as follows and always refer to the 10-minute average values:

Horizontal wind speed V_{mean} [m/s]: It is directly available in typical SCADA systems.

Air density ρ_{mean} [kg/m^3]: While the air density is seldom directly measured in a turbine, an approximation is possible using the ideal gas law as

$$\rho_{\text{mean}} = \frac{p}{RT_{\text{mean}}} = 3.4837\frac{p}{T_{\text{mean}}} \tag{1}$$

with the pressure p, the average air temperature T_{mean} measured at the hub height and the general gas constant $R = 8.314\,\mathrm{J\,mol^{-1}\,K^{-1}}$. The pressure p at hub height can be approximated using

$$p = 101.29 - (0.011837)z + (4.793\times10^{-7})z^2, \tag{2}$$

https://doi.org/10.5194/wes-2025-112
Preprint. Discussion started: 1 July 2025

WIND ENERGY SCIENCE DISCUSSIONS

where z is the hub height in meters, see Manwell et al. (2009, Chap. 2). This is a simplified approach that does not capture any weather-related variations in the air pressure, for instance due to high or low pressure systems, temperature inversion or different moisture content.

Turbulence intensity I $[-]$: The turbulence intensity I can be directly estimated from the available SCADA measurements of wind, utilizing 10-minute SCADA statistics as

$$I = \frac{V_{\text{std}}}{V_{\text{mean}}}, \tag{3}$$

where V_{std} represents the standard deviation and V_{mean} the average of the wind speed in any given 10-minute period. However, the usual anemometer placement behind the rotor distorts these measurements. While the 10-minute average wind speed is commonly internally adjusted to account for the deceleration across the rotor plane, such corrections are not applied to the standard deviation of wind speed. We therefore adopt the method implemented by Barthelmie et al. and Jørgensen et al. (2003) to approximate V_{std} for below-rated conditions as

$$V_{\text{std}} = \frac{P_{\text{std}}}{\left(\frac{dP}{dV}\right)_v \cdot B}, \tag{4}$$

where P_{std} is power's standard deviation, $\left(\frac{dP}{dV}\right)_v$ is the power derivative with respect to wind speed, and B is an empirical correction factor. For above-rated conditions, we extend the method to use the pitch angle β_{std} instead of the power as

$$V_{\text{std}} = \frac{\beta_{\text{std}}}{\left(\frac{d\beta}{dV}\right)_v \cdot B}. \tag{5}$$

To obtain the necessary power and pitch curves, data from Original Equipment Manufacturers (OEM) or regression models applied to historical operational data may be used.

Yaw misalignment angle ψ_{mean} $[^\circ]$: It is directly available in typical SCADA systems.

2.2 Sampling from Site-Specific Distributions

Depending on the scope of the surrogate (i.e., single wind farm or wider region), distributions of the turbines' operating conditions have to be derived, so that training and validation samples may be drawn from the resulting joint distribution. Table 1 presents an example of how these distributions could be defined. For the wind speed, a Beta distribution could be employed, as suggested by Dimitrov et al. (2018), with the aim of generating more samples at low wind speeds where the variance of the remaining parameters is usually larger. The shape parameters α and β may be found empirically through iterative adjustments until the desired shape is achieved and the minimum ($\min_V$) and maximum wind speed ($\max_V$) would align with the cut-in and cut-out wind speed of the considered turbines.

Here, the turbulence intensity, air density, and the yaw misalignment angles are modeled using distributions that solely depend on the wind speed. The bounds for these distributions can be either established using physical laws, as shown in Dimitrov et al. (2018), or derived directly from available historic measurements, for example, through an interpolation approach. In the

https://doi.org/10.5194/wes-2025-112
Preprint. Discussion started: 1 July 2025

Table 1. Probability distributions for the chosen feature variables.

Parameter	Distribution
Wind Speed (V_{mean})	$\text{Beta}(\alpha, \beta, \min_{V_{\text{mean}}}, \max_{V_{\text{mean}}})$
Air Density (ρ_{mean})	$\text{Uniform}(\min_\rho(V_{\text{mean}}), \max_\rho(V_{\text{mean}}))$
Turbulence Intensity (I)	$\text{Uniform}(\min_I(V_{\text{mean}}), \max_I(V_{\text{mean}}))$
Yaw Misalignment (ψ_{mean})	$\text{Uniform}(\min_\psi(V_{\text{mean}}), \max_\psi(V_{\text{mean}}))$

example of Table 1, uniform distributions were selected, for simplicity.

Once the distributions are defined, training samples can be drawn from the input space. A common approach is to use the Sobol sequence to generate samples within a 4D-hypercube, where each dimension represents a uniformly distributed random variable. Subsequently, the samples must be transformed to the correlated distributions defined in the previous section. As we deal with a set of dependent variables, the inverse Rosenblatt transformation is a suitable function for this task, Dimitrov et al. (2018). Assuming just two independent variables u_1 and u_2 for simplicity, it is defined as

$$\begin{cases} \theta_1 = F_{\theta_1}^{-1}(u_1), \\ \theta_2 = F_{\theta_2|\theta_1}^{-1}(u_2|u_1), \end{cases} \tag{6}$$

see Mara and Becker (2021). In this example, $F_{\theta_1}^{-1}$ represents the inverse Cumulative Distribution Function (CDF) of θ_1, and $F_{\theta_2|\theta_1}^{-1}$ is the inverse CDF of θ_2 conditioned on θ_1.

In addition to the training samples, a set of validation samples is required for measuring model performance. These samples can be drawn randomly.

2.3 Generating the Load Database

The next step involves establishing a load database that links each sampled operating condition to the corresponding turbine loads through an aeroelastic simulator f as

$$y = f(\boldsymbol{\theta}, \omega), \tag{7}$$

where $\boldsymbol{\theta}$ represents the vector of input features (wind speed, air density, turbulence intensity, yaw misalignment) and y any given target variable (e.g., blade root bending moments). Moreover, ω represents an elementary random outcome to account for the seed-to-seed variability. Please note that temporal averages and standard deviations are denoted with mean, std, whereas mean and standard deviation with respect to randomness (ω) are denoted with μ, σ, respectively. These simulations require high-resolution turbine models, often unavailable to operators. This study explores the option of using open-source RWTs to approximate loads on real-world turbines. RWTs cover various turbine types across different power classes, rotor diameters, and hub heights, see Gambier (2022). While they cannot perfectly match commercial turbines, they may serve as a good

https://doi.org/10.5194/wes-2025-112
Preprint. Discussion started: 1 July 2025

starting point for further adaptations. A workflow for such an approach could look, as suggested by Barter in the NREL forum[1], as follows:

1. Select an appropriate RWT model which best approximates the real-world wind turbine.
2. Adjust key parameters (e.g. tower height, rotor diameter, rated power).
3. Reoptimize and recalculate the modified structures using WISDEM (Wind-plant Integrated System Design and Engineering Model), see Dykes et al. (2021).
4. Retune the controller using ROSCO (Reference OpenSource Controller), see NREL (2021).
5. Generate new OpenFAST (Fatigue, Aerodynamics, Structures, and Turbulence) input files via WEIS (Wind Energy with Integrated Servo-control), see Abbas et al. (2022).
6. Simulate the loads (e.g. blade bending moments) for each sampled operating condition using OpenFAST, see Buhl et al. (2023).

Therefore, emulating an OEMs design process in a simplified manner should allow turbine operators to obtain turbine models sufficient for initial load assessments on real wind turbines.

2.4 Surrogate Methods

With the generated training and validation data, regression models can be applied to capture the input-output relationships. The case study presented in Section 3 evaluates PCE and GPR. These techniques were selected for their proven ability in capturing the load response of wind turbines (Dimitrov et al. (2018); Murcia et al. (2018); Slot et al.; Teixeira et al.; Gasparis et al. (2020)) also with a moderate size of the training data. While additional approaches were investigated - including stochastic polynomial chaos expansion (SPCE) in a follow-up study - these methods did not yield significant improvements over the approaches presented here. The following sections provide a brief theoretical background for both PCE and GPR.

2.4.1 Polynomial Chaos Expansion

The PCE is a methodology aims at approximating models $y = f(\boldsymbol{\theta}, \omega)$ featuring a finite number of random input variables M summarized in the vector $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_M)$. For ω fixed, the model response y is considered to be a random variable with finite variance. Hence, a set of basis polynomials over $\boldsymbol{\theta}$ can be found for its effective approximation. The polynomials must thereby be orthogonal with respect to the distribution of each input variable, see Xiu and Karniadakis (2002). Once such a set

[1]Garret Barter in Reference Turbines for Scaling - Wind & Water, https://forums.nrel.gov/t/reference-turbines-for-scaling/3601/4, 2022

of polynomials is identified, the moments of y may be expanded as

$$\mu(\boldsymbol{\theta}) = \sum_{j=0}^{\infty} c_j \phi_j(\boldsymbol{\theta}), \tag{8}$$

$$\sigma(\boldsymbol{\theta}) = \sum_{j=0}^{\infty} d_j \phi_j(\boldsymbol{\theta}), \tag{9}$$

where $\mu(\boldsymbol{\theta}), \sigma(\boldsymbol{\theta})$ denote the input-dependent mean value and standard deviation of y, averaged over ω. We focus on the second order moments because of the limited sample size available. Moreover, $\phi_j(\boldsymbol{\theta})$ is a multidimensional polynomial basis function

$$\phi_j(\boldsymbol{\theta}) = \phi_{j_1}(\theta_1) \cdots \phi_{j_M}(\theta_M). \tag{10}$$

The c_j, d_j are the polynomial coefficients. Moreover, j is the enumeration of a multidimensional index element, i.e., $\mathbf{j} = (j_1, \ldots, j_M)$ with $\mathbf{j} \leftrightarrow j \in \mathbb{N}$, that denotes the degree of the polynomial in each dimension, see Murcia et al. (2018). In practice, this expansion is truncated to a finite number of N_c terms, for instance the mean value approximation reads

$$\mu(\boldsymbol{\theta}) \approx \sum_{j=0}^{N_c-1} c_j \phi_j(\boldsymbol{\theta}). \tag{11}$$

One way of finding the N_c values of the coefficients is through the use of regression methods, e.g. with the least-squares or the Least Absolute Shrinkage and Selection Operator (LASSO) approach, as proposed by Tibshirani (1996).
The method as described so far usually assumes that the M input variables are uncorrelated. In cases where this assumption does not hold, the input space should be first transformed into an uncorrelated one using a suitable transformation function (e.g. the Rosenblatt function). Alternatively, a dedicated PCE for dependent variables can be used, see Jakeman et al. (2019).

2.4.2 Gaussian Process Regression

GPR is a non-parametric approach where, in our context, the functions $\mu(\boldsymbol{\theta}), \sigma(\boldsymbol{\theta})$ are assumed to be drawn from a Gaussian process. Each process is defined by a mean and covariance function. While the mean is often assumed to be zero, a common covariance function for two points $\boldsymbol{\theta}$ and $\boldsymbol{\theta}_*$ is the Radial Basis Function (RBF) kernel

$$\mathcal{K}(\boldsymbol{\theta}, \boldsymbol{\theta}_*) = \sigma_\mathrm{f}^2 \exp\left(-\frac{\|\boldsymbol{\theta}_i - \boldsymbol{\theta}_*\|^2}{2\ell^2}\right). \tag{12}$$

Here ℓ corresponds to the length scale of the kernel, i.e. the distance for which random variables are correlated. The standard deviation of the GP is controlled with σ_f. Assuming a vanishing mean for simplicity, the training and test samples of are jointly distributed as

$$\begin{pmatrix} \boldsymbol{\mu}_{\boldsymbol{\Theta}} \\ \boldsymbol{\mu}_* \end{pmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} \mathbf{K}_{\boldsymbol{\Theta},\boldsymbol{\Theta}} + \sigma_n^2 \mathbf{I} & \mathbf{K}_{\boldsymbol{\Theta},*} \\ \mathbf{K}_{\boldsymbol{\Theta},*}^\top & \mathbf{K}_{*,*} \end{bmatrix}\right), \tag{13}$$

https://doi.org/10.5194/wes-2025-112
Preprint. Discussion started: 1 July 2025

where we focus on the GP approximation of μ for simplicity. Through pairwise evaluation of the covariance function $\mathcal{K}$, the covariance matrices between the training points ($\mathbf{K}_{\boldsymbol{\Theta},\boldsymbol{\Theta}}$), between the test points ($\mathbf{K}_{*,*}$) and between the training and test points ($\mathbf{K}_{\boldsymbol{\Theta},*}$ and $\mathbf{K}_{*,\boldsymbol{\Theta}}$) are computed. The term $\sigma_n^2\mathbf{I}$ models the noise in the observed samples $\mathcal{N}(0,\sigma_n^2)$, where $\mathbf{I}$ is the identity matrix.

The posterior distribution is obtained by conditioning the joint Gaussian on the training data $\mathcal{D} = \{\boldsymbol{\Theta}, \mathbf{f}_{\boldsymbol{\Theta}}\}$, yielding the predictive mean and covariance

$$\begin{aligned} \mathbb{E}[\boldsymbol{\mu}_*|\mathcal{D}] &= \mathbf{K}_{*,\boldsymbol{\Theta}}[\mathbf{K}_{\boldsymbol{\Theta},\boldsymbol{\Theta}} + \sigma_n^2\mathbf{I}]^{-1}\mathbf{f}_{\boldsymbol{\Theta}}, \\ \mathrm{Cov}[\boldsymbol{\mu}_*|\mathcal{D}] &= \mathbf{K}_{*,*} - \mathbf{K}_{*,\boldsymbol{\Theta}}[\mathbf{K}_{\boldsymbol{\Theta},\boldsymbol{\Theta}} + \sigma_n^2\mathbf{I}]^{-1}\mathbf{K}_{\boldsymbol{\Theta},*}. \end{aligned} \tag{14}$$

3 Case Study

To evaluate the effectiveness of generic turbine models in predicting real turbine loads, the previously discussed methodology is applied in a case study. This study focuses on a wind farm in northern Germany operated by Alterric. The farm consists of eight turbines in the 3 MW range. Further specifications cannot be disclosed in this study. Historical SCADA data at 10-minute intervals are available for each turbine from 2017 to 2022. Figure 1 illustrates the farm's layout and the mean wind conditions derived from measurements at the nacelle across all turbines.

Surrogate models are developed, as described in Section 2, to efficiently predict 10-minute statistics of the flapwise blade bending moments and fore-aft tower-top accelerations for the turbines within the farm. These loads are chosen due to the availability of sensor measurements on the actual turbines. Successfully predicting these recorded loads with the surrogate models would demonstrate the applicability of generic turbine models to real-world scenarios.

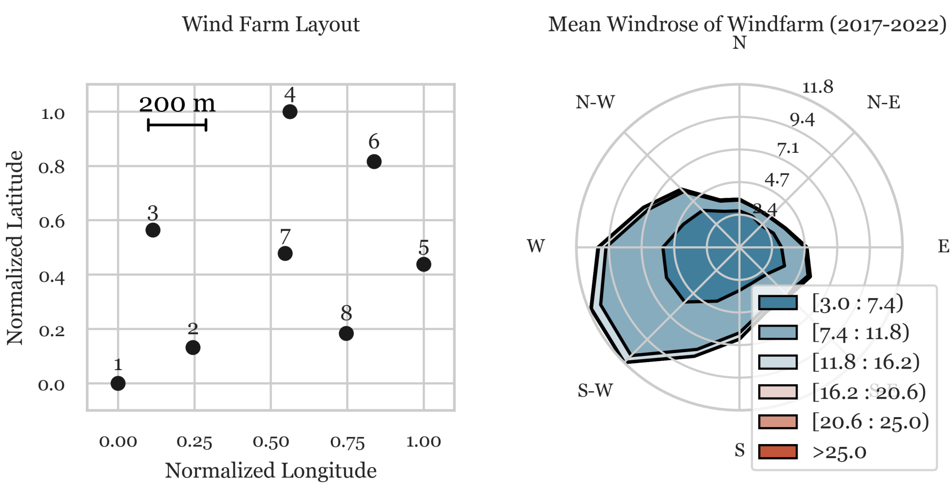

Figure 1. Spatial distribution of wind turbines within the investigated wind farm (left), wind rose of SCADA wind data at the nacelle averaged over the time period 2017-2022.

https://doi.org/10.5194/wes-2025-112
Preprint. Discussion started: 1 July 2025

eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

3.1 Data Availability and Processing

Each turbine in the wind farm is equipped with a SCADA system that records various signals. Specifically, it captures the 10-minute average, standard deviation, minimum, and maximum values for all measurements. The relevant signals for this article include:

- Wind speed measured by behind the rotor,
- Ambient temperature measured at the nacelle height,
- Root bending moment in the flapwise direction for each blade,
- Tower-top acceleration in the fore-aft and side-to-side directions,
- Power output of the turbine,
- Yaw misalignment of the turbine,
- Pitch angle of the blades.

In addition to sensor measurements, each turbine logs its operational status through unique codes that indicate any warnings or alarms. Unlike the 10-minute binned data from the sensors, these status codes are timestamped at the moment they are triggered and when they end.

In this study, only normal operating conditions are considered, focusing specifically on the wind speed range between the cut-in and the cut-out speed of the turbines.
The data cleaning process involves two main steps:

1. Utilizing available status and alarm data from the SCADA system to filter out obvious turbine malfunctions.
2. Employing a set of filter functions from the OpenOA toolkit, see Perr-Sauer et al. (2021), provided by NREL, for a detailed clean-up of the power and pitch curve (see the provided Jupyter notebooks for details).

Having established a clean dataset, the turbulence intensity and air density are approximated using the method described in Section 2.1. In the case of the turbulence intensity, the power and pitch curves are obtained for each turbine and operating year through the application of generalized additive models on the clean datasets. However, due to the absence of mast wind measurements, there is no empirical basis to adjust the parameter B away from 1. Therefore, no correction is applied in this study. The final step in data processing involves averaging the recorded blade loads across the three blades and adjusting the SCADA coordinate system to align with the simulations.

3.2 Obtaining the Turbine Model

As stated in Section 2.3, a generic high resolution turbine model needs to be identified which can serve as a basis for simulating the real world wind turbine dynamics. Here, the IEA-3.4-130 RWT model is chosen as the generic high-resolution turbine

https://doi.org/10.5194/wes-2025-112
Preprint. Discussion started: 1 July 2025

model, closely approximating the real turbines with a 130 m rotor diameter, 110 m hub height, and 3.37 MW rated power. However, to match the real turbine's specifications, the hub height and rotor diameter are adjusted using WEIS, Abbas et al. (2022), which employs WISDEM, Dykes et al. (2021), for blade and tower parametrization and ROSCO, NREL (2021), for controller gain adjustment. Within the limited scope of this study, the model was not reoptimized, i.e. the spanwise geometries were maintained after adjusting the rotor diameter and hub height, therefore deviating from the method outlined in Section 2.3. However, simulations were conducted to verify that no critical resonances occur in the blades or tower during normal operation.

3.3 Definition of the Target Variables

The loads analyzed in this study include the blade's root bending moment in the flapwise direction, as well as the tower top acceleration in the fore-aft direction. The side-to-side tower acceleration was excluded as it could not be reliably modeled using the adapted generic turbine model, as determined by comparing SCADA measurements against aeroelastic simulations in a preliminary analysis. Each target is fully specified in Table 2. We treat each target as a stochastic variable, modeling both the mean response and its standard deviation to account for seed-to-seed variability.

Table 2. Lookup table for the symbols used in our case study.

Target Variable	Symbol
10-min mean root bending moment in flapwise direction	$M_{\text{blade,mean}}$
10-min maximum root bending moment in flapwise direction	$M_{\text{blade,max}}$
10-min standard deviation of tower top acceleration in fore-aft direction	$\ddot{x}_{\text{tower,std}}$
10-min maximum acceleration of tower top in fore-aft direction	$\ddot{x}_{\text{tower,max}}$

3.4 Sampling of Feature Vectors for the Training and Validation

In the case study, the sampling of feature vectors for training and validation of the surrogate models adheres to the procedure outlined in the Section 2.2. A Beta distribution is used for the wind speed between 3 m s^{-1} (cut-in) and 25 m s^{-1} (cut-out), with empirically determined shape parameters. For the remaining features, we used uniform distributions as presented in Section 2.2, with the boundary functions obtained using polynomial fits to the cleaned SCADA data. This limits the surrogates scope to the site-conditions of the investigated wind farm. For the model training a total of 300 samples are generated for the joint distribution using the Sobol sequence. Another 30 samples are generated randomly for the validation dataset. Each sample is simulated with 30 turbulent seeds, enabling the estimation of both the mean and standard deviation of the turbine responses. This seed number aligns with TurbSim's documentation Jonkman as a compromise between the small and large sample sizes found in the literature.

https://doi.org/10.5194/wes-2025-112
Preprint. Discussion started: 1 July 2025

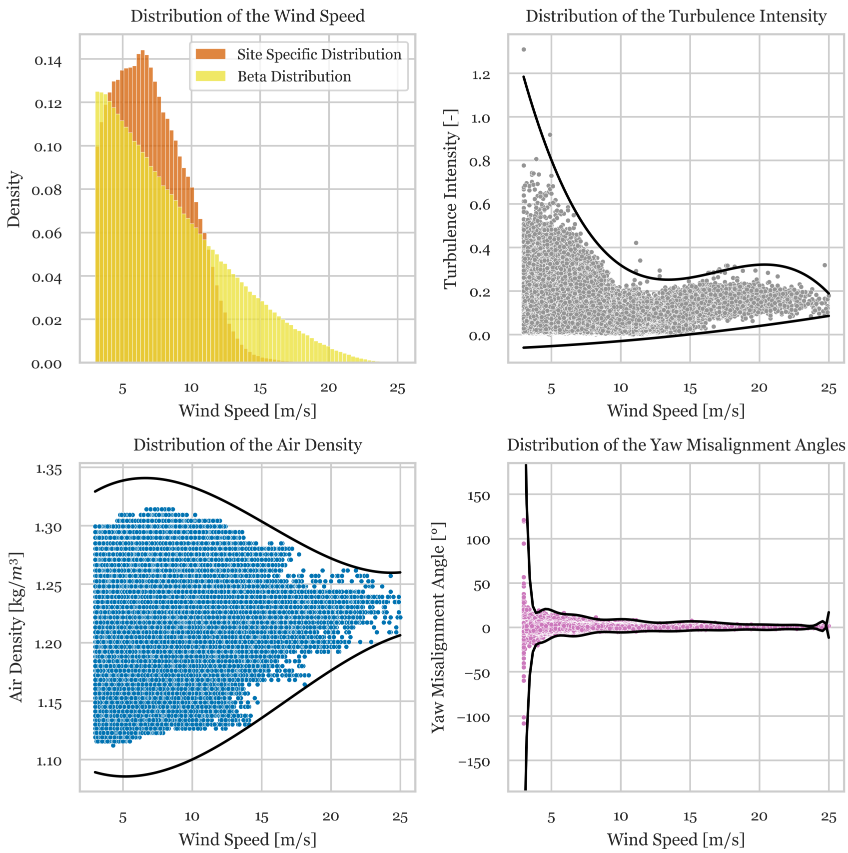

Figure 2. Historical distributions of the four input features from the wind farm, with black lines showing the boundary functions for sampling. The site-specific Weibull distribution of the wind speed is overlaid with the Beta distribution used in the sampling process. These plots visualize the sampling ranges for each variable. They do not aim to suggest physical dependencies between wind speed and the other parameters. The wind speed based boundaries were applied solely to enable a simple and efficient sampling of feature combinations.

3.5 Generation of Load Data

All aeroelastic simulations in this study are conducted within OpenFAST, see Buhl et al. (2023). Each simulated operating condition is characterized only by the four dimensional feature vector. All other parameters are kept constant. This includes one key simplification, in that the wind shear is modeled using the power law while applying a constant shear exponent of 1/7, due to the lack of precise measurements of the vertical distribution of the wind speed for the wind farm, see Letcher (2023).

Each simulation is performed for a duration of 630 s to align with the 10-minute period of each SCADA record. An additional 30 seconds at the beginning are included to account for any initial transient effects, which are subsequently excluded from the analysis. In TurbSim, the wind field is simulated on a grid 10% larger than the rotor diameter of the turbine model with 25 grid points in each direction. Exemplary configuration files are available in the provided data of this article.

3.6 Surrogate Model Training

With the training and validation datasets in place, the surrogate models are implemented based on PCE and GPR. Following the approach of Murcia et al. (2018), two separate models are constructed for each target variable: one to predict its mean and another for its standard deviation. Each model is trained using the respective statistics derived from the 30 turbulent realizations per sample.

The PCE models are implemented using the Chaospy package, Feinberg and Langtangen (2015), where the training samples are first transformed into independent uniform variables using the Rosenblatt transformation, thereby enabling the application of Legendre polynomials. Polynomial orders ranging from one to five are tested, employing both the least-squares and LASSO regression methods for fitting the coefficients. For the LASSO approach, regularization coefficients are determined through 10-fold cross-validation.

The GPR models are implemented using the Scikit-learn library, Pedregosa et al. (2011). As with the PCE models, the training samples are first normalized using the Rosenblatt transformation. The models are then trained using RBF kernels modified with an active noise term. Scikit-learn employs the log marginal likelihood method for tuning the hyperparameters.

To evaluate model performance, the validation set is employed and the normalized Mean Absolute Error (nMAE) and normalized Root Mean Squared Error (nRMSE) are computed as

$$\mathrm{nMAE} = \frac{1}{N \cdot \bar{y}} \sum_{i=1}^{N} |y_i - \hat{y}_i|, \tag{15}$$

$$\mathrm{nRMSE} = \frac{\sqrt{\frac{1}{N}\sum_{i=1}^{N} (y_i - \hat{y}_i)^2}}{\bar{y}} \tag{16}$$

where y_i represents the true value, $\hat{y}_i$ is the predicted value, $\bar{y}$ is the mean of the true values, and N is the number of samples in the validation set.

4 Results

4.1 Surrogate Model Training

This section evaluates the training performance of the PCE and GPR surrogate models described above. The objective is to provide a high-level overview that informs the decision-making process for model application in subsequent analyses.

4.1.1 Polynomial Chaos Expansion

As anticipated, increasing the sample size and polynomial order generally improved accuracy, provided that sufficient training data were available. LASSO regression consistently outperformed or matched least squares regression, by benefiting from model simplifications, thereby enhancing training efficiency. The results presented in Table 3 reflect the validation scores obtained for PCE models trained on all available data using the LASSO regressor and a maximum polynomial order of 5, which

https://doi.org/10.5194/wes-2025-112
Preprint. Discussion started: 1 July 2025

was the highest tested in our study. In this table and throughout the following discussion, μ refers to the mean response of each target, while σ describes its standard deviation resulting from the seed-to-seed variability in the simulations. Examining the various targets and statistical moments listed, PCE demonstrated particular effectiveness in modeling the blade loads. However, its performance in predicting tower acceleration was comparatively lower.

A notable observation is the discrepancy between the predictions of mean responses and their standard deviations. While mean responses were generally well predicted, standard deviations showed significantly larger errors. This is likely due to the fact that the response surfaces of the standard deviation are much less smooth compared to the mean responses. Using a larger dataset, with more than 30 turbulent realizations, could potentially help in smoothing the surface of the training data. Similarly, we found that the models for the maximum blade loads and tower accelerations perform worse compared to the respective other model of blade and tower. In this case, increasing the number of replications may not necessarily improve the results, since maximum values can change significantly even with small parameter variations, leading to inherently discontinuous response surfaces.

Table 3. Performance summary of the two surrogate modeling approaches. The mean of each target is denoted by μ, the standard deviation by σ.

		PCE (5th order and LASSO)		GPR (RBF kernel)	
Target	Moment	nMAE	nRMSE	nMAE	nRMSE
$M_{\text{blade,mean}}$	μ	0.02	0.03	0.02	0.03
$M_{\text{blade,mean}}$	σ	0.13	0.17	0.12	0.16
$M_{\text{blade,max}}$	μ	0.07	0.08	0.22	0.25
$M_{\text{blade,max}}$	σ	0.87	1.05	0.83	0.94
$\ddot{x}_{\text{tower,std}}$	μ	0.05	0.09	0.03	0.06
$\ddot{x}_{\text{tower,std}}$	σ	0.36	0.54	0.27	0.44
$\ddot{x}_{\text{tower,max}}$	μ	0.44	0.64	0.23	0.31
$\ddot{x}_{\text{tower,max}}$	σ	1.07	1.20	3.94	7.09

4.1.2 Gaussian Process Regression

The convergence behavior of the GPR models largely followed expected trends, with accuracy generally increasing as more training data were provided. This improvement was particularly evident in models that predicted the mean responses of the target variables.

However, the analysis of standard deviation models proved more complex. The error curves for these models exhibited diverse behaviors, including increasing, decreasing, or fluctuating accuracy. This variability likely stems from the challenges in modeling seed-to-seed uncertainty, already observed in the PCE approach.

When comparing the final performance of the GPR models to that of the respective PCE models (see Table 3), they generally demonstrate an enhanced performance in modeling the tower acceleration, with an exception being the standard deviation model of the maximum acceleration.

Based on these results, the following decisions are made for the case study:

- PCE models, trained using the LASSO method with a maximum polynomial order of 5, are applied for predicting the wind turbine's blade bending moments.
- GPR models are applied for predicting the turbine's tower accelerations.

4.2 Application of the Surrogates

In this section, the previously selected surrogate models are applied to predict the historic blade loads and tower oscillations on the real-world wind turbines. First, the results are visually inspected. Next, the prediction errors are quantified and investigated in detail.

4.2.1 Prediction of the Blade Loads and Tower Acceleration

Figure 3 illustrates a selection of randomly sampled predictions for each target variable, overlaid with measured SCADA records and nearby training samples. Error bars represent the predicted standard deviation, visualizing the seed-to-seed uncertainty. The surrogates for the mean blade loads generally capture trends well, but show limitations in resolving the full spread of measurements, particularly at high wind speeds. Maximum blade load predictions align well with measurements overall, but underestimate at low turbulence intensities. The standard deviations of tower acceleration show decreasing accuracy at turbulence intensity extremities and some outliers at high intensities. For the maximum tower acceleration, no satisfactory fit could be achieved during model training, so it will not be analyzed further here. Finally, this visualization highlights the importance of considering the seed-to-seed uncertainty when assessing the accuracy of historic load predictions.

In the next sections, we conduct a detailed examination of prediction errors to quantify them, understand their causes, and identify any systematic patterns.

4.2.2 Quantification of Prediction Errors

Figure 4 provides a high-level overview of the surrogate models' performance, showing the mean and median of the relative prediction errors across all available operating years (2017-2022) for each of the eight turbines.

The surrogate models for the blade loads display considerably lower mean and median errors when compared to those for the tower acceleration. Specifically, for the maximum blade loads the median of the errors is consistently around 10% for all turbines, suggesting a uniform behavior across the entire wind farm. The mean of the errors is only slightly elevated compared to the median statistics (around 15% for most turbines), indicating that this level of performance is achieved in most conditions without many significant outliers. The mean blade load predictions vary in accuracy among the turbines. The turbines 2, 6,

https://doi.org/10.5194/wes-2025-112
Preprint. Discussion started: 1 July 2025

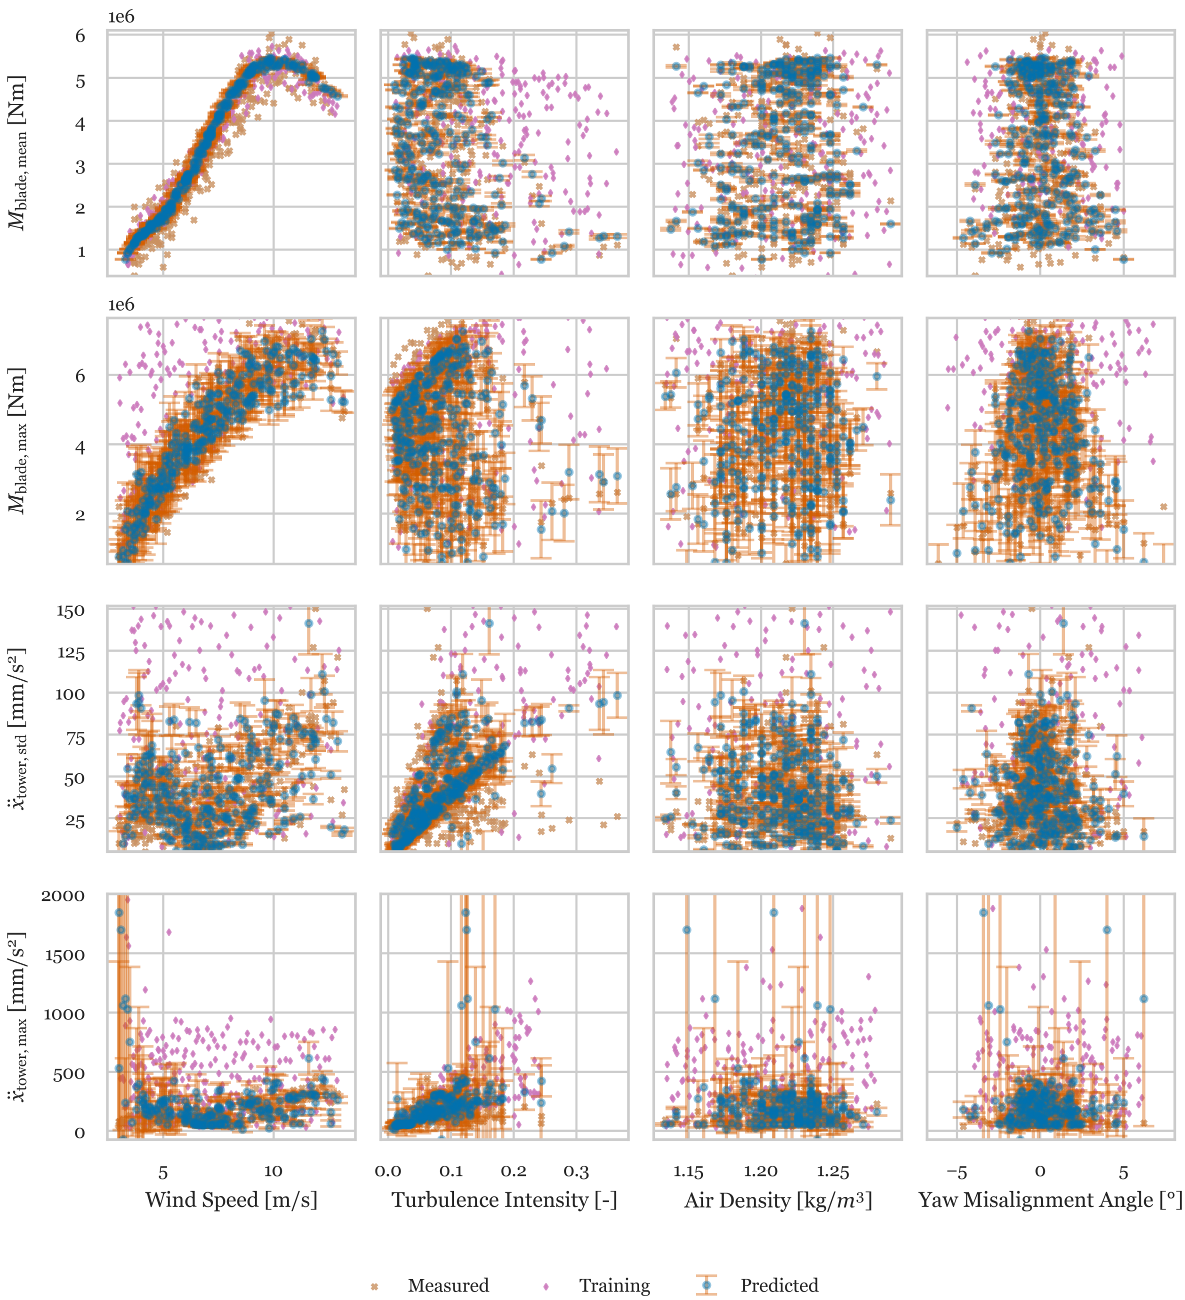

Figure 3. Scatter plots of predicted and measured SCADA records with selected training samples across various input features.

and 8 show median errors below 10%, while the others range between 15-19%. This variance could be due to factors like the turbine location, ageing effects or sensor inaccuracies. Lastly, while the error magnitudes for the mean and maximum blade

https://doi.org/10.5194/wes-2025-112
Preprint. Discussion started: 1 July 2025

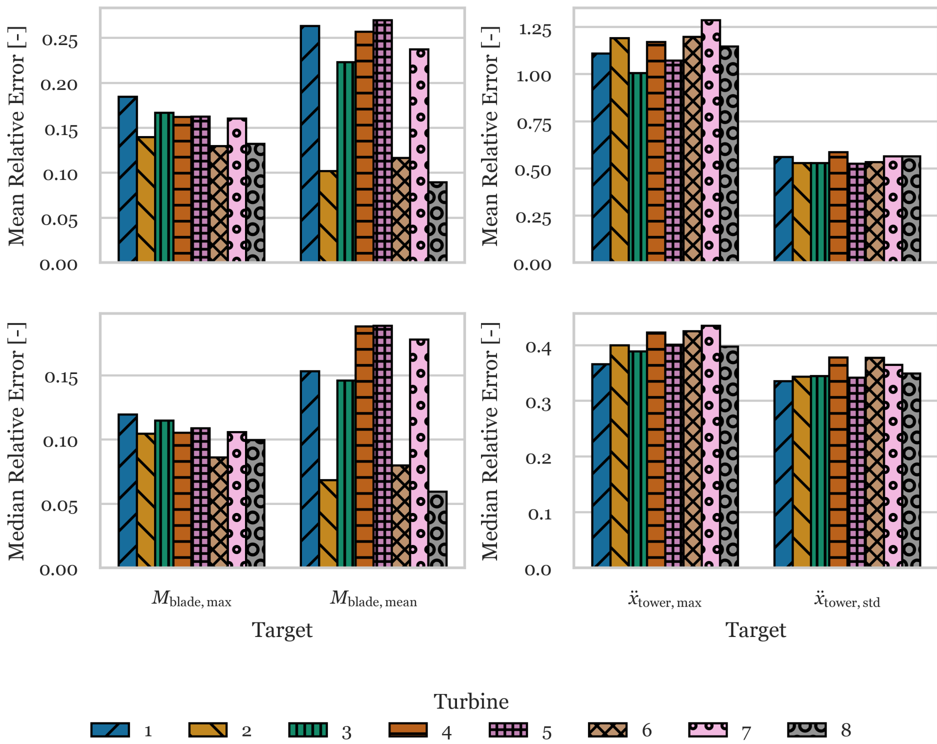

Figure 4. Mean and median of the relative prediction errors for the four targets across all turbines between 2017 and 2022.

loads are comparable in scale, the higher seed-to-seed uncertainty for the maximum loads, as observed in Figure 3, suggests that it significantly contributes to the observed errors, representing an irreducible contribution.

The models for the tower acceleration targets display significantly higher prediction errors. The median error for the acceleration's standard deviation ranges from 34% to 39%, and for the maximum acceleration from 36% to 42%. Regarding the mean of the errors, they are only slightly elevated for the standard deviation model, similar to the blade load cases. Yet, for the maximum acceleration they are substantially higher, ranging from approximately 100% to 128%.

The next section will focus on the temporal variance of the prediction errors. While the site dependence was investigated as well, no conclusive patterns could be identified.

4.2.3 Time Dependence

Resolving the observed mean and median errors over the discrete years of available data provides an opportunity to investigate potential trends, such as ageing effects. For the blade loads, a distinct trend of increasing errors is noticeable, as seen in Figure 5, especially for turbines 1, 3, 4, 5, and 7. Notably, their errors start to significantly diverge from those of the other turbines after

https://doi.org/10.5194/wes-2025-112
Preprint. Discussion started: 1 July 2025

2019, which results in the overall elevated error seen in Figure 4. To further investigate this behavior, Figure 6 illustrates the

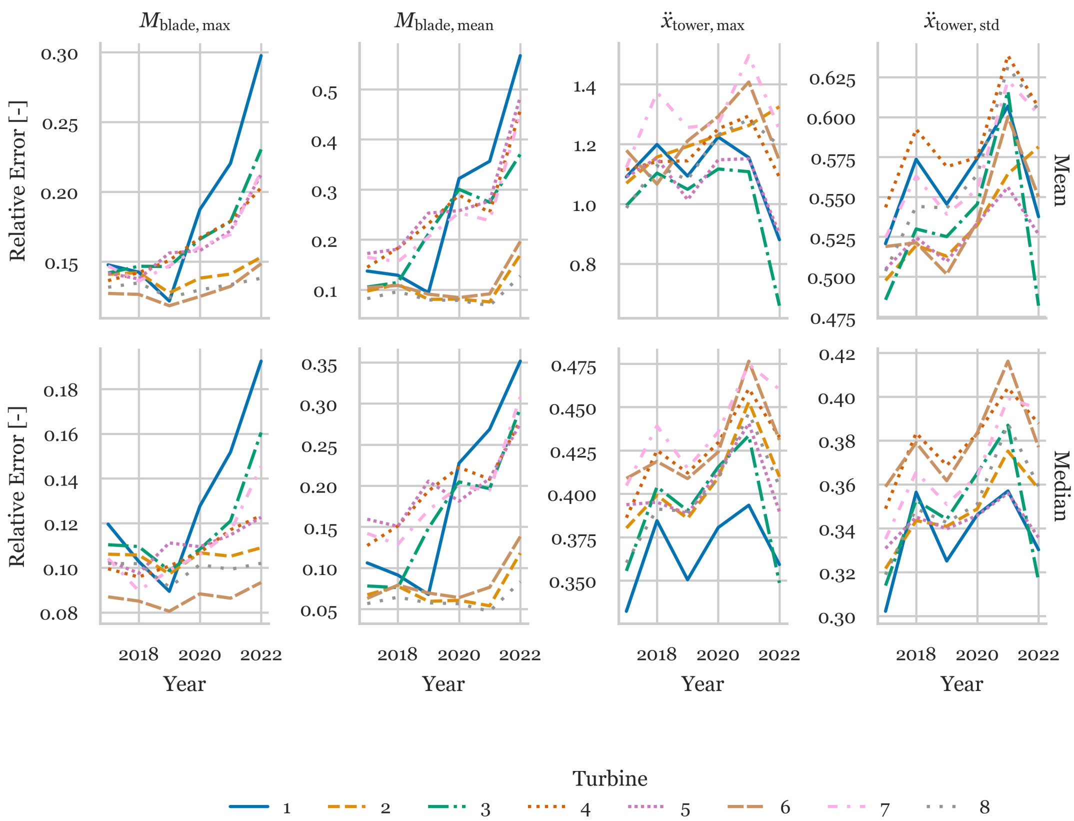

Figure 5. Annual mean and median of the relative prediction errors for each turbine.

mean and median of the measured blade loads across the available operating years of the wind turbines. Interestingly, there is a consistent decrease in both the average and peak blade loads over time. This trend might be attributed to ongoing adjustments in the pitch controller's parameterization, aimed at reducing the loads to extend the turbines' lifespan. Given that this trend starts in 2017 and affects all turbines uniformly, it seems unlikely to be the only cause of the significant error divergence highlighted above.

In contrast, the errors associated with the tower acceleration do not show a clear annual trend, which is also the case for the measured tower accelerations of each turbine.

Besides the annual trends, we also analyze the diurnal patterns of the errors. Figure 7 displays the mean and median of the prediction errors, grouped by the hour of the day and each turbine. In most cases, the errors seem to be low during the night and in the early morning, while peaking in the afternoon. This aligns with the diurnal evolution of the atmospheric boundary layer in flat terrains. With the heating during daytime a convective bodunary layer grows, increasing vertical mixing

https://doi.org/10.5194/wes-2025-112
Preprint. Discussion started: 1 July 2025

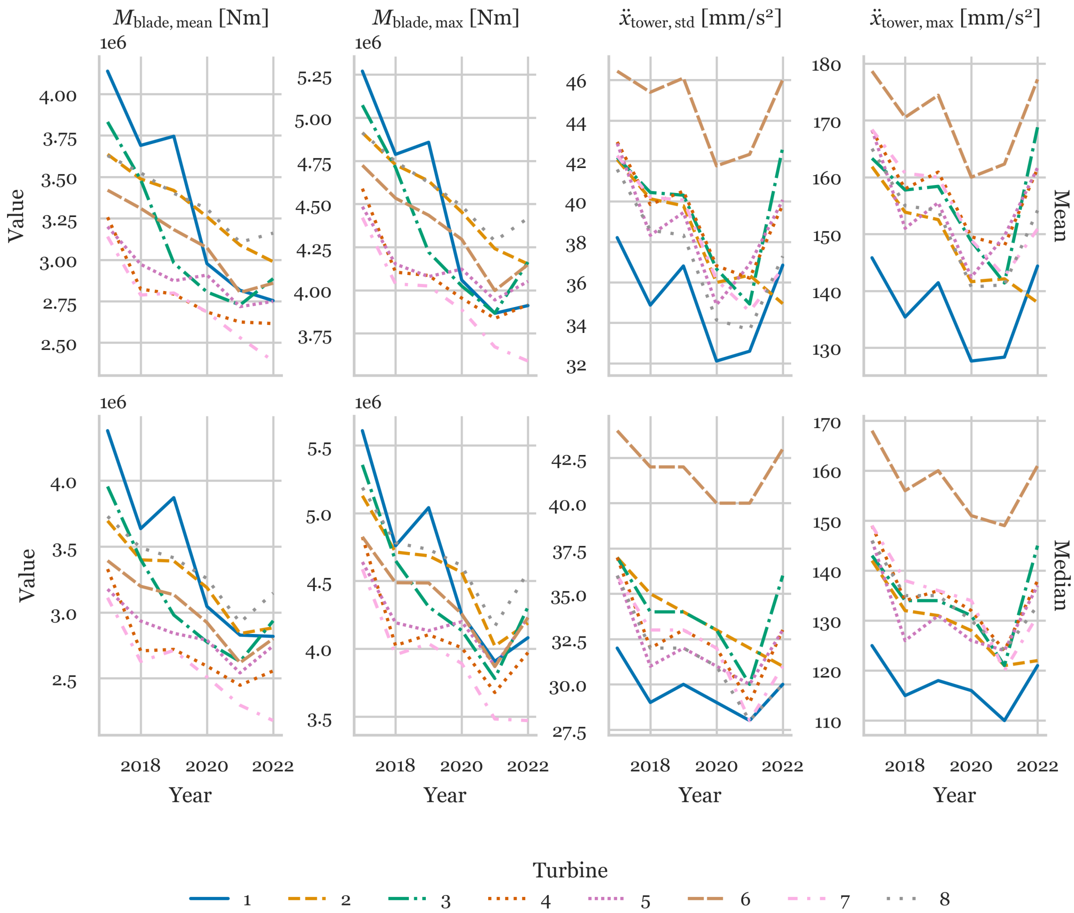

Figure 6. Annual mean and median of the recorded blade load and tower acceleration signals.

and turbulence. At night-time, the cooling effects generally lead to a more stable boundary layer near the ground, characterized by low turbulence, Emeis (2018). Therefore, one possible interpretation of the errors is a decreasing model performance at higher turbulence intensitities.

5 Discussion

375 Concerning the developed surrogate models, it can be noted that both the PCE and GPR proved to be effective and efficient at capturing turbine responses of the training data. We explained the challenges in the fit of the standard deviation models with the noisy response surfaces which could be solved by increasing the number of turbulent realizations. Given the lower than anticipated costs for generating the current set of training data, an increase in turbulent realizations to 100 is not unrealistic. Therefore, the development of surrogates as a predictive tool is definitely feasible for industry practitioners, and the current

https://doi.org/10.5194/wes-2025-112
Preprint. Discussion started: 1 July 2025

WIND ENERGY SCIENCE DISCUSSIONS

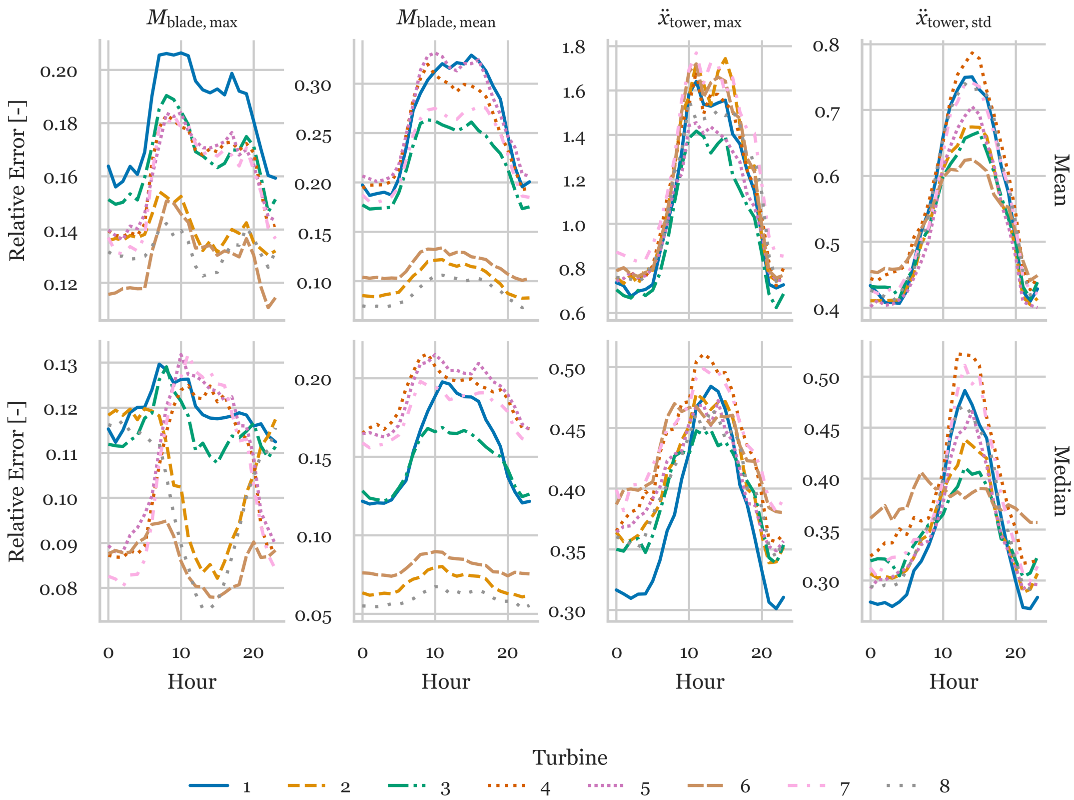

Figure 7. Analysis of diurnal patterns in the prediction errors.

models provide a good baseline for more complex models that may be tested in the future.

However, while these results only confirm the findings from existing literature, the main purpose of this study was to assess the usability of generic turbine models as a basis for the surrogates. In applying the developed surrogates for predicting the historic loads of 8 turbines within one onshore wind farm, the initial results were promising. For the blade loads, the median of the prediction errors were consistently within a range of only 10% and 15%. For the tower accelerations, the errors were higher, though still within acceptable bounds. These results support that the method could be used for generating initial load assessments, e.g., for comparative qualitative analyses, but not for use-cases demanding high levels of accuracy.

A limitation of the proposed method became obvious after observing the deterioration of the prediction accuracy for the mean blade loads over the analyzed operating years. This highlights the inability to fully capture all events, such as controller modifications or other operational changes, either due to modeling constraints or a lack of detailed knowledge. Another weakness of the surrogate models is their lacking capability in achieving a constant accuracy throughout the day, due to the varying levels of turbulence in the wind.

It should be noted that this study focused on a single wind farm with one turbine type, limiting the generalizability of our findings. Moreover, our simplified process for adapting the generic turbine model, while sufficient for this initial assessment, could be significantly improved, as pointed out in the case study. We expect this to be the simplification with the most impact on the prediction accuracy achieved in the presented case study. Other simplifications, including the limitation to four input parameters, their calculation (e.g. in the case of the air density) and the relatively coarse grid for the simulation of the wind field likely had less significant effects on the results.

6 Conclusions

This study investigated the feasibility of using generic turbine models coupled with surrogate modeling techniques, to efficiently predict loads on real-world wind turbines. Using approximately 10.0000 aeroelastic simulations of a slightly adapted generic turbine model, we trained surrogate models to predict historic flapwise blade bending moments and fore-aft tower accelerations for an onshore wind farm in northern Germany over five years. The models accounted for the seed-to-seed uncertainty by predicting both the expected mean and standard deviation of the targets.

Aligning with existing literature, both applied methods, PCE and GPR, were effective in capturing the dynamics of wind turbines. Here, PCE was found to be more suitable for modeling the blade loads, while GPR was more suitable for predicting the tower acceleration. However, it should be noted that the small number of turbulent realizations limited the quality of fit for the standard deviation models.

Applying these best-performing models for the prediction of historic turbine loads revealed promising, though varying levels of accuracy. For the blade loads, the surrogates achieved median relative prediction errors of around 10% for maximum loads and 10-20% for the mean loads. In contrast, the tower acceleration models performed notably worse, with median prediction errors typically between 30-40%, and even higher errors for maximum accelerations. In analyzing the temporal variation of the prediction errors, we found a diurnal correlation, likely related to the changes in atmospheric stability over the day. Furthermore, the prediction quality of the blade loads deteriorating with time, likely due to changes in the turbines operating regimes.

The current study is limited by its focus on only one single commercial wind turbine type and generic turbine model, as well as the employed method for adapting the generic turbine model. More work is required to extend this methodology to diverse wind farm settings and turbine types, improve the adaptation of the generic turbine and investigate methods to enhance the models' adaptability to operational changes and environmental variations. Additionally, increasing the number of turbulent realizations is required for improving the prediction quality of the seed-to-seed uncertainty. It is also yet to be explored to what extent simpler turbine models could achieve comparable results, at significantly reduced computational cost.

Despite these challenges, this study provides a promising foundation for a cost-effective approach to estimating loads in wind farms when detailed models are unavailable. This methodology could support turbine operators in their ongoing efforts to improve current farm planning, operation, and maintenance procedures.

https://doi.org/10.5194/wes-2025-112
Preprint. Discussion started: 1 July 2025

425 # Appendix A

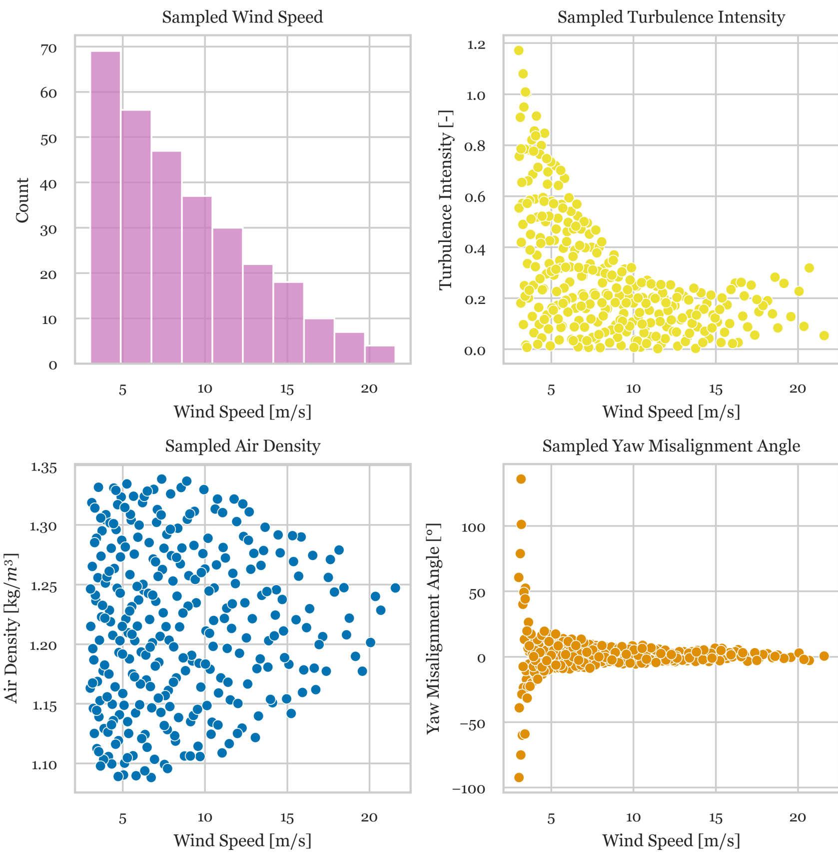

Figure A1. Three hundred pseudo-random input samples generated using the Sobol sequence.

https://doi.org/10.5194/wes-2025-112
Preprint. Discussion started: 1 July 2025

Code and data availability. The analysis code and simulation scripts are archived at https://doi.org/10.5281/zenodo.15446360 (see Mönnig and Römer (2025a)). The datasets required for reproducing our results are available at https://doi.org/10.5281/zenodo.15380253 (see Mönnig and Römer (2025b)).

Author contributions. A.M. performed the aeroelastic simulations and surrogate modeling in the framework of his master thesis supervised by A.H., A.L., and U.R. A.M. wrote the main text. A.H. contributed both industry-specific contextual knowledge and expertise on wind turbine operations. A.L. contributed to the interpretation of meteorological phenomena. U.R. contributed to the development of the methodology. All authors edited the text and contributed to the publication.

Competing interests. This research was conducted as part of A.M's Master's thesis in cooperation with Alterric Deutschland GmbH, where A.M. was employed as a working student. Co-author A.H., an employee of Alterric Deutschland GmbH, supported the thesis as an industry supervisor. The company contributed by providing data access and domain expertise. The authors declare no other competing interests.

References

Windenergieanlagen Teil 1: Auslegungsanforderungen (IEC 61400-1:2019); Deutsche Fassung EN IEC 61400-1:2019.

Abbas, N. J., dzalkind, Barter, G., Jasa, J., Bortolotti, P., Mudafort, R. M., Nunemaker, J., paulf81, Quon, E., Ramos, D. S., Sundarrajan, A. K., Lee, Y. H., Chetan, M., Gaertner, E., Branlard, E., Mulders, S., Key, A., davidheff, Rinker, J., Meunier, P.-E., Du, X., and amoratoc: WISDEM/WEIS: WEIS v1.1, https://doi.org/10.5281/zenodo.7150136, 2022.

Barthelmie, R. J., Frandsen, S. T., Nielsen, M. N., Pryor, S. C., Rethore, P.-E., and Jørgensen, H. E.: Modelling and Measurements of Power Losses and Turbulence Intensity in Wind Turbine Wakes at Middelgrunden Offshore Wind Farm, 10, 517–528, https://doi.org/10.1002/we.238.

Buhl, M., Hayman, G., Jonkman, J., Jonkman, B., Mudafort, R., Platt, A., and Sprague, M.: OpenFAST, https://github.com/OpenFAST/openfast, 2023.

Dimitrov, N.: Surrogate Models for Parameterized Representation of Wake-Induced Loads in Wind Farms, 22, 1371–1389, https://doi.org/10.1002/we.2362, 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, 3, 767–790, https://doi.org/10.5194/wes-3-767-2018, 2018.

Dykes, K., Ning, S. A., Scott, G., Graf, P., of Energy Efficiency, U. O., and Energy, R.: WISDEM® (Wind-Plant Integrated System Design and Engineering Model) [SWR-14-05], https://doi.org/10.11578/dc.20211208.2, 2021.

Emeis, S.: Wind Energy Meteorology, Green Energy and Technology, Springer International Publishing, ISBN 978-3-319-72858-2 978-3-319-72859-9, https://doi.org/10.1007/978-3-319-72859-9, 2018.

Feinberg, J. and Langtangen, H. P.: Chaospy: An Open Source Tool for Designing Methods of Uncertainty Quantification, 11, 46–57, https://doi.org/10.1016/j.jocs.2015.08.008, 2015.

Gambier, A.: Control of Large Wind Energy Systems: Theory and Methods for the User, Advances in Industrial Control, Springer International Publishing, ISBN 978-3-030-84895-8, https://doi.org/10.1007/978-3-030-84895-8, 2022.

Gasparis, G., Lio, W. H., Wai Hou Lio, Fanzhong Meng, and Meng, F.: Surrogate Models for Wind Turbine Electrical Power and Fatigue Loads in Wind Farm, 13, https://doi.org/10.3390/en13236360, 2020.

Jakeman, J. D., Franzelin, F., Narayan, A., Eldred, M., and Plfüger, D.: Polynomial chaos expansions for dependent random variables, Computer Methods in Applied Mechanics and Engineering, 351, 643–666, 2019.

Jonkman, B. J.: TurbSim User's Guide v2.00.00, https://www.nrel.gov/wind/nwtc/assets/downloads/TurbSim/TurbSim_v2.00.pdf.

Jørgensen, H. E., Frandsen, S., and Vølund, P.: Wake effects on Middelgrund Windfarm, Tech. Rep. Risø-R-1415(EN), Risø National Laboratory, Roskilde, Denmark, 2003.

Letcher, T. M.: Wind Energy Engineering: A Handbook for Onshore and Offshore Wind Turbines, AP, Academic Press, an imprint of Elsevier, 2nd edition edn., ISBN 978-0-323-99353-1, https://doi.org/10.1016/C2021-0-00258-3, 2023.

Lüthen, N., Marelli, S., and Sudret, B.: A spectral surrogate model for stochastic simulators computed from trajectory samples, Computer Methods in Applied Mechanics and Engineering, 406, 115 875, 2023.

Manwell, J. F., McGowan, J. G., and Rogers, A. L.: Wind Energy Explained: Theory, Design and Application, Wiley, 2nd ed edn., ISBN 978-0-470-01500-1, https://doi.org/10.1002/9781119994367, 2009.

Mara, T. A. and Becker, W. E.: Polynomial Chaos Expansion for Sensitivity Analysis of Model Output with Dependent Inputs, 214, 107 795, https://doi.org/10.1016/j.ress.2021.107795, 2021.

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, 119, 910–922, https://doi.org/10.1016/j.renene.2017.07.070, 2018.

Mönnig, A. and Römer, U.: alexandermoennig/wind-farm-load-estimation: Wind Farm Load Estimation v1.0.0, https://doi.org/10.5281/zenodo.15446361, 2025a.

Mönnig, A. and Römer, U.: Load Estimation in Onshore Wind Farms Using Surrogate Modelling and Generic Turbine Models, https://doi.org/10.5281/zenodo.15380254, 2025b.

NREL: ROSCO. Version 2.4.1, https://github.com/NREL/ROSCO, 2021.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., and Cournapeau, D.: Scikit-Learn: Machine Learning in Python, 12, 2825–2830, https://jmlr.csail.mit.edu/papers/v12/pedregosa11a.html, 2011.

Perr-Sauer, J., Optis, M., Fields, J., Bodini, N., Lee, J., Todd, A., Simley, E., Hammond, R., Phillips, C., Lunacek, M., Kemper, T., Williams, L., Craig, A., Agarwal, N., Sheng, S., and Meissner, J.: OpenOA: An Open-Source Codebase For Operational Analysis of Wind Farms, 6, 2171, https://doi.org/10.21105/joss.02171, 2021.

Slot, R. M. M., Sørensen, J. D., Sudret, B., Svenningsen, L., and Thøgersen, M. L.: Surrogate Model Uncertainty in Wind Turbine Reliability Assessment, 151, 1150–1162, https://doi.org/10.1016/j.renene.2019.11.101.

Teixeira, R., Nogal, M., O'Connor, A., Nichols, J., and Dumas, A.: Stress-Cycle Fatigue Design with Kriging Applied to Offshore Wind Turbines, 125, 454–467, https://doi.org/10.1016/j.ijfatigue.2019.04.012.

Tibshirani, R.: Regression shrinkage and selection via the lasso, Journal of the Royal Statistical Society Series B: Statistical Methodology, 58, 267–288, 1996.

Toft, H. S., Svenningsen, L., Moser, W., Sørensen, J. D., and Thøgersen, M. L.: Assessment of Wind Turbine Structural Integrity Using Response Surface Methodology, 106, 471–483, https://doi.org/10.1016/j.engstruct.2015.10.043.

Wilkie, D. and Galasso, C.: Gaussian Process Regression for Fatigue Reliability Analysis of Offshore Wind Turbines, 88, 102 020, https://doi.org/10.1016/j.strusafe.2020.102020, 2021.

Xiu, D. and Karniadakis, G. E.: The Wiener–Askey polynomial chaos for stochastic differential equations, SIAM journal on scientific computing, 24, 619–644, 2002.

Zhu, X. and Sudret, B.: Replication-based emulation of the response distribution of stochastic simulators using generalized lambda distributions, International Journal for Uncertainty Quantification, 10, 2020.

Zhu, X. and Sudret, B.: Emulation of Stochastic Simulators Using Generalized Lambda Models, SIAM/ASA Journal on Uncertainty Quantification, 9, 1345–1380, https://doi.org/10.1137/20M1337302, 2021.