

Virtual sensing via Gaussian Process for bending moment response prediction of an offshore wind turbine using SCADA data

Bridget Moynihan^a, Eleonora M. Tronci^{a,c}, Michael C. Hughes^b, Babak Moaveni^{a,*}, Eric Hines^a

^a Department of Civil and Environmental Engineering, Tufts University, Medford, MA, USA

^b Department of Computer Science, Tufts University, Medford, MA, USA

^c Department of Civil and Environmental Engineering, Northeastern University, Boston, MA, USA

ARTICLE INFO

Keywords:

Virtual sensing
Offshore wind
Gaussian process
Bending moment prediction
Sparse Gaussian process
SCADA

ABSTRACT

Offshore wind turbines (OWTs) can be equipped with two types of monitoring systems: (1) a Supervisory Control and Data Acquisition (SCADA) system that monitors operational data such as wind speed and power generation, and (2) vibration sensors like accelerometers and strain gauges to track structural dynamics. While strain gauges enable fatigue damage calculations, not all OWTs in a wind farm have these sensors installed. This paper proposes a Gaussian process regression (GPR) strategy to predict the bending moment time-histories of an OWT. The model takes into account various SCADA data such as wind speed, power, and nacelle acceleration as input and learns to predict the high and low-frequency dynamic response of the system. The strategy is implemented and tested to predict the bending moment response of a 6 MW offshore wind turbine in the fore-aft and side-side directions of the turbine. The accuracy and reliability of the proposed strategy are evaluated and demonstrated considering different operational conditions and multiple hotspot locations along the height of the turbine. The proposed strategy proves to be an efficient virtual sensing strategy and can be easily transferred to other turbines in the same wind farm without the need for widespread installation of strain gauges.

1. Introduction

Offshore wind energy continues to play a significant role in the build-out of renewable energy infrastructure worldwide. Offshore wind capacity has been able to accelerate with the growth and maturity of Offshore Wind Turbines (OWTs), but the operation and maintenance of such critical structures presents numerous challenges. By improving the reliability of these turbines, it is possible to extend their operating lifetimes, reduce costs, and ultimately assist the transition towards renewable energy. Predictive maintenance methods are useful for anticipating and preventing failure before it occurs. The approach utilizes measurements from structural response sensors to identify deviations in the physical condition of the OWT. Fatigue life estimations are one of the primary methods used in a predictive maintenance strategy, particularly for OWTs which are constantly exposed to cyclic loads from wind and waves, leaving them susceptible to fatigue damage. The measured time-histories of material stress and any other fatigue-based quantity are highly informative for estimating the remaining useful life (RUL) of a structure. Bending moments – measured by strain gauges – can provide the necessary information to estimate a structure's RUL. Therefore, accurate and efficient estimates of the bending moment

time-histories across wind farms are crucial information for extending the lifetime of OWTs [1–3].

One of the primary challenges faced by the offshore wind industry is the high cost associated with the installation and maintenance of strain gauges on OWTs [4], so it is rare to find them installed on every turbine in a given wind farm. On the other hand, a Supervisory Control and Data Acquisition (SCADA) system is typically installed in all OWTs in a farm. The SCADA system is designed to capture and record operational and environmental measurements such as wind speed and power output, which enables operators to keep track of key metrics in real time and keep turbines functioning at optimal levels. Given the challenging task of directly measuring strains at many critical spots on the structure, the indirect extrapolation of the response measurements without the need for prior sensor installation becomes an appealing alternative solution. This concept, known as virtual sensing [5,6], involves using a limited set of response measurements such as accelerations, strains, or SCADA data to reconstruct the full-field response of the structure.

Several methods for estimating stresses and strains using response estimation techniques have been presented in various literature, in some cases also quantifying the uncertainty associated with the virtual

* Corresponding author.

E-mail address: babak.moaveni@tufts.edu (B. Moaveni).

<https://doi.org/10.1016/j.renene.2024.120466>

Received 12 December 2023; Received in revised form 5 April 2024; Accepted 7 April 2024

Available online 9 April 2024

0960-1481/© 2024 Elsevier Ltd. All rights reserved.

sensing technique adopted [7]. Virtual sensing techniques can be divided into five categories: model-based robust observers, Kalman filter-based techniques, joint input-state filtering techniques, modal decomposition and expansion techniques, and finally, data-driven methodologies. Gillijns and Moor first proposed the joint-state input estimation [8] and further developed by Lourens et al. [9]. Later, this algorithm was extended to be applied when accelerations are measured. Similarly, the Kalman filter-based techniques estimate states based on a limited number of response measurements. This approach was first applied to structural dynamics by Papadimitriou et al. [10]. This strategy was also implemented to develop an optimal sensor placement approach for offshore wind turbines where the value of information is maximized for virtual sensing purposes [11]. The Kalman filters have proven effective in simultaneously integrating various types of measurements such as accelerations and strains [12]. The modal expansion technique assumes that structural response is a linear combination of the modal response [13,14]. A comparison of the applicability in offshore monopile wind turbines of the Kalman filter, joint input-state filter and modal expansion is carried out in [15], which concludes on an acceptable performance in the prediction of acceleration responses in offshore wind turbine foundations. Virtual sensing formulations based on modal expansion and a Kalman filter for a monopile-supported OWT were compared in different literature works [16,17] which found both methods to provide acceptable virtual sensing performance. In the context of purely data-driven virtual sensing strategies, Dimitrov et al. [18] adopted a Long Short-Term Memory neural network architecture to numerically generate a wind turbine response time series using data types compatible with data available from high-frequency SCADA data combined with a blade and tower load measurement system. All of the methods described depend on a model of the dynamical system, commonly built thanks to available dynamic measurements collected on the structure itself.

This work focuses on developing a purely data-driven, model-free regression strategy to achieve virtual sensing of the strain-related quantities for an offshore wind turbine. The model chosen to carry out the regression strategy is Gaussian Processes (GPs). GP regression stands out as a powerful and flexible approach that is useful in scenarios where a model-free, data-driven strategy is needed. In the context of virtual sensing across a wind farm, GPs are a strategic choice that allows for modeling complex physical quantities without relying on computationally expensive structural models. GPs are characterized by their ability to provide predictions as well as estimates of prediction uncertainty, which provide additional information to wind farm operators that is crucial for decision-making. The proposed approach illustrates the practical application of machine learning methodologies in addressing engineering challenges alongside advanced research in the field of artificial intelligence.

GP modeling has demonstrated impressive performance regarding time-series regression tasks [19,20], and are highly flexible, allowing us to incorporate our knowledge of the physical structure into the model by selecting the kernel function. While other studies have developed physics-informed kernel functions for fatigue assessment, what sets our approach apart is that it utilizes real data from an OWT and relies only on SCADA data to make high-frequency predictions of the structural response. As noted in [21,22], a few studies have shown the potential of GP-based strategies for offshore wind applications. In the context of virtual sensing, Papatheou et al. [23] explored the potential of using SCADA data from the Lillgrund offshore wind farm for constructing the power curves for the turbines in the farm using GP and neural networks. Avendano et al. [24] use Gaussian Processes to estimate the fatigue damage equivalent loads of a wake-affected wind turbine based on the wind speed and direction measurements from an upwind wind turbine. The paper also compares the performance of the model with different kernel functions and input features and validates the model with experimental data from a wind farm in Denmark. Bilbao et al. [25] implemented a Gaussian process latent force model for an onshore

turbine to estimate both dynamic loading and strain response using the acceleration data, comparing the predicted strains with the measured quantities and assessing fatigue by comparing the damage equivalent loads calculated with the predicted as opposed to the measured strains. Zou et al. [26] proposed a Gaussian process latent force model to estimate the strain response of the soil-foundation system of offshore wind turbines, which is difficult to measure directly with physical sensors. Pimenta et al. [27] proposed a bending moment estimation strategy using the Monte Carlo Markov Chain algorithm to constrain the parameters that best describe the covariance structure shared between accelerations and displacements of a linear structure under stochastic loading. Then, they incorporated that structure into a GP strategy to convert accelerations into displacements and these into bending moments.

This paper aims to provide a data-driven virtual sensing strategy for reconstructing bending moments using only the SCADA information, providing an alternative solution to direct measurement with strain gauges. Specifically, this paper uses Gaussian process regression (GPR) to predict bending moment time-histories of an instrumented 6 MW OWT using measured SCADA data parameters as input features. The probabilistic nature of GPR provides uncertainty quantification on predicted bending moments. Future work can incorporate these levels of uncertainty into fatigue estimations of the RUL. Successful models could be used on other turbines in the same wind farm to understand the RUL of each OWT without the need for widespread installation of strain gauges.

The workflow of this paper is described in Fig. 1. Section 2 describes the dataset used for model training and validation in this work, which comes from an instrumented OWT with structural response measurements from strain gauges and a SCADA system measuring operational and environmental conditions. In Section 2, the basic operational regions of an OWT are described, and measurements from the OWT are classified into operational regions according to SCADA parameters. Then, response measurements are investigated with respect to available SCADA, defining the input and target data features in the dataset. Section 3 describes the formulation of the GP model, which describes both plain and sparse GP. This section describes the mean and kernel functions created for the modeling task. Section 4 presents modeling results for the target bending moments of interest. Conclusions are shared in Section 5.

2. Dataset

The proposed strategy has been tested and validated using experimental data collected from an operational OWT. The OWT under investigation is a 6 MW OWT instrumented with various structural and operational sensors which are described in this section. The turbine is supported by a monopile driven into the seabed and a transition piece allowing access to the wind turbine. The overall height of the turbine from the mudline to the nacelle is about 140 m.

There are two independent data collection systems in use on the OWT. One installed monitoring system consists of dynamic vibrational response measurements measured by 16 strain gauges and 12 accelerometers distributed along the tower and foundation. This system collects the structural response measurements, which can be used to understand the long-term health and reliability of the structure. Fig. 2 displays the location and distribution of the accelerometers and strain gauges along the structure. The reader is referred to [28] for a more detailed characterization of the structural system, its dynamic response, and modal properties.

The second system, known as the supervisory control and data acquisition (SCADA) system, is located inside the nacelle at the top of the tower. The SCADA system measures operational and environmental information for all OWTs, including power production, rotor speed, pitch and yaw angles, and wind speed. In addition, the SCADA system




Fig. 1. Methodology Flowchart.




Fig. 2. Monitoring instrumentation.

of this OWT collects readings from 2 accelerometers placed inside the nacelle to measure tower-top deformations.

Due to the variable nature of OWT operational and environmental conditions, a single model cannot be used to make all desired predictions. Where necessary, predictive models are tailored to the unique conditions experienced by an OWT over its lifetime. Therefore, this section first describes the basic operational regions of an OWT and explains how the loading conditions and structural dynamics differ in each region. These regions of operation are the basis for independent models that consider the unique OWT properties that vary during operation. After describing the regions of operation, we explore the target and input variables for each region and explain how these variables behave according to operating conditions. Finally, we describe the data preparation for GP model testing and validation.

2.1. Regions of operation

This section describes the regions of operation and how varying environmental conditions and controller settings drive differences in the structural response of the OWT. During operation, wind flows over the blades, which creates lift forces that turn the rotor and, therefore, the electrical generator. Since operation is conditioned on incident wind speed, OWTs experience unique loading conditions and structural dynamics based on the wind speed and resulting power output. This means that the relationship between the target variables of interest in this paper and the available SCADA input variables is different depending on the operational conditions.

In the operation of an OWT, the differences in operational conditions are categorized according to 3 distinct “regions of operation”. The regions of operation are defined by incident wind speed and corresponding power production. Fig. 3 shows the power curve for the OWT

used in this work, which plots the power output as a function of wind speed as 1-minute average values, with the regions of operation noted. Fig. 3 also notes which data points in the power curve correspond to the data sets used to train GP models for each region.

As seen in Fig. 3, an OWT only operates beyond a specific wind speed, known as the “cut-in” wind speed. Region 1 considers conditions when the wind speed is below the cut-in speed. This region is known as idling conditions, which occur when the turbine is ready to begin operation but the wind speed is too low to turn the rotor or produce power sufficiently. During region 1, the blades are pitched at a small angle, which allows lift forces to act on the blades and turn the rotor should the wind speed increase beyond the cut-in wind speed.

Region 2 consists of operation between the cut-in and “rated” wind speed. The rated wind speed refers to the lowest wind speed at which the OWT reaches full power production. In this region, the blades are at a low pitch angle to optimize the lift forces and turn the rotor. Power output in region 2 may be any output below the rated power output of the turbine.

Region 3 considers operation at full power production, which occurs at wind speeds between the rated and the “cut-out” wind speed (the highest wind speed the OWT is allowed to operate at). In this region, the rotor and generator are held at constant rotational speeds in order to keep the power at its rated output, avoiding power levels too high for the generator to produce. The blades’ pitch angles are varied to manage the lift forces and resulting torque. This results in a reduction of forces on the turbine as wind speeds approach the cut-out speed, with the highest forces occurring at the rated wind speed. The cut-out speed is not noted because the data set contains no conditions where the wind speed was high enough to stop the rotor.

The turbine may sometimes pitch the blades into the “parked” condition to stop operation for reasons other than dangerous winds,




Fig. 3. Power curve of instrumented OWT with regions of operation.




Fig. 4. Turbine cross section with locations of instrumentation, and directions of primary bending moments shown.

such as maintenance operations. This work does not consider parked conditions, and modeling focuses on operational or idling conditions.

2.2. Target variables

The goal of prediction is to determine the bending moments at a specific height of the wind turbine in two distinct directions, which are aligned with the local axis of the OWT hub. These two directions are commonly known as fore-aft (FA) and side-side (SS). They correspond to the primary directions of movement for an OWT, as they are parallel and perpendicular, respectively, to the incoming wind flow. It is worth noting that the direction of the local axes changes as the wind direction changes. The hub and blades rotate automatically to face the incoming wind.

The measured bending moments have been computed from installed strain gauges at each elevation noted in Fig. 2. Fig. 4 depicts the orientation of four strain gauges at a generic elevation of installation. Bending moments M_{North} and M_{East} are computed from these sensors directly and are a measure of the global North and East bending moments at the elevation of strain gauges. M_{North} and M_{East} can be

rotated into the local axis of the turbine according to its yaw angle – denoted as γ in Fig. 4 – which can change to any angle between 0 and 360 as the rotor continuously turns to face the incoming wind. The rotated bending moments in the local axis of the rotor are known as fore-aft and side-side moments, M_{FA} and M_{SS} , which are our target variables.

The goal of the modeling task is to build predictive GP models for each of these bending moments in all regions of operation. Example readings of the measured bending moments in each region of operation are presented in Fig. 5. The readings show notable differences in the magnitude and variance of each signal, which can be attributed to the different loading conditions experienced by the OWT in each region. The most significant differences are observed in the FA direction, where the OWT is subjected to more complex loading patterns due to power and wind speed variations. Notably, the FA bending moment displays a varying mean value in regions 2 and 3, which is not observed in region 1 nor the SS direction, as illustrated in Fig. 5.

The SS bending moment is primarily caused by the torque generated at the nacelle as well as environmental loading due to waves, which results in the SS bending moments behaving likewise in each region of operation. While there are slight differences in magnitude, each SS bending moment has a similar dynamic behavior. In particular, the mean value of the SS bending moment does not vary like the FA bending moments in regions 2 and 3.

2.3. Input data features

Input features for each GP model are selected from the available SCADA system measurement channels. The SCADA channels include five operational measurements (wind speed, power output, yaw angle, rotor speed, and pitch angle) and two structural response measurements (FA acceleration and SS acceleration), all measured at 10 Hz.

The relationship between available input features and target variables differs by region of operation and direction (FA or SS). This section explores the input features for the GP models built in this work, highlighting how input/target variable correlations vary in each region of operation. These relationships drive GP model formulation.

2.3.1. Nacelle FA and SS accelerations

The FA and SS accelerations measure the relative deflection of the nacelle in the FA and SS directions. Physically, these readings are closely tied to the target bending moments we wish to predict. The SCADA accelerometers help predict the dynamic behavior in the bending moment readings in both the FA and SS directions.




Fig. 5. Bending moment target variables, example data.




Fig. 6. Correlation of SS Acceleration from SCADA with SS Target Bending Moments, by Region.

In the SS direction, the target bending moments are highly correlated with the SS acceleration in all three regions, as shown in Fig. 6. This is the most crucial predictive input variable needed in the SS direction for SS bending moment prediction.

The relationship between the FA bending moment and FA accelerometers from SCADA is shown in Fig. 7. FA acceleration still shows a high correlation with target bending moments in region 1, when there is little incident wind flow and no power production. In this region, the FA bending moment shows behavior similar to that of the SS direction.

In regions 2 and 3, which are distinguished by higher wind speeds and non-zero power production, the FA bending moment shows a lower correlation with the FA accelerometers, in contrast to the behavior in the SS direction. Notably, the FA accelerometers are not able to capture the varying mean value of the FA bending moment in regions 2 and 3. However, there is still a useful correlation with the FA accelerometers

in these regions, as the reading captures the fundamental frequency of the FA bending moment's dynamic behavior.

2.3.2. Power output

In the FA direction, additional SCADA variables, such as power output, correlate with the target measurements' varying mean. Fig. 8 compares the power output of the OWT to the target bending moments in FA. The low-frequency variation of the power output through time matches the FA bending moment in regions 2 and 3, with a particularly good match to the region 2 FA moment. In region 1 where there is no power production, this reading is not correlated with target bending moments and provides no information on the target variable.

In regions 2 and 3, rotor speed and wind speed have a similar relationship to the FA bending moments as the power output. These are not plotted here but are considered in the GP model formulation.




Fig. 7. Correlation of FA Acceleration from SCADA with FA Target Bending Moments, by Region.




Fig. 8. Correlation of Power Output from SCADA with FA Target Bending Moments, by Region.

2.4. Data preparation and evaluation metrics

This modeling task focuses on the high-frequency reconstruction of bending moments using a 100-minute dataset. We prioritize a high-frequency prediction task because fatigue calculations rely on a high-frequency bending moment reading in order to count the stress cycles experienced in certain ranges of magnitude [29]. For this study, the entire data set is down-sampled to 2.5 Hz, which reduces the number of data instances by a factor of 10 while retaining a suitable sampling frequency to capture the fundamental frequency of the system. The dataset is split into training and testing sets in which the first 70% of data points are allocated as the training set and the last 30% as the testing set. All the target variables and input features are standardized to have 0 mean and unit variance according to the training set, while the same factors scale the testing set.

In the context of evaluating the predictive models, four key metrics are used to assess their performance: the mean absolute error (MAE), the root mean squared error (RMSE), the relative root mean squared error (RRMSE), and the time-response assurance criterion (TRAC) [30].

MAE provides a straightforward measure of the model's predictive accuracy, and it is particularly useful when you want to understand the average magnitude of errors without giving extra weight to large errors, as squared differences do in the case of RMSE.

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^N |\hat{y}_i - y_i| \quad (1)$$

where N is the number of data points, $\hat{y}_i$ is the predicted value at time i and y_i is the true value at time i .

The RMSE is a commonly used metric to measure the accuracy of a predictive model's predictions. It quantifies the average magnitude of

errors between predicted and true values. The RMSE is calculated as the square root of the mean of the squared differences between predicted values ($\hat{y}$) and true values (y):

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N (\hat{y}_i - y_i)^2} \quad (2)$$

where N is the number of data points, $\hat{y}_i$ is the predicted value at time i , and y_i is the true value at time i .

The RRMSE is a variation of RMSE that expresses the error as a relative value. It is useful for comparing models on datasets with different scales or units. RRMSE is calculated by dividing RMSE by the mean of the true values (y) and then multiplying by 100% to express it as a percentage:

$$\text{RRMSE} = \sqrt{\frac{\sum_{i=1}^N (\hat{y}_i - y_i)^2}{\sum_{i=1}^N (y_i)^2}} \quad (3)$$

The TRAC is a metric used to compare time series data and assess the fit of a predicted time series ($\hat{y}$) to the true time series (y). A TRAC value of 1 indicates a perfect fit. The TRAC metric is computed using the following equation:

$$\text{TRAC} = \frac{(y^T \hat{y})^2}{(y^T y)(\hat{y}^T \hat{y})} \quad (4)$$

where y and $\hat{y}$ are vectors representing the true time series and predicted time series, respectively. T denotes the transpose of a vector. The equation compares the inner product of the true time series and predicted time series to the product of the inner products of each with itself.

In summary, RMSE measures the absolute error between predicted and true values, RRMSE normalizes the error relative to the mean of the true values, and TRAC assesses the fit of predicted time series to true time series data, with a value of 1 indicating a perfect match. These metrics quantify the performance of predictive models, especially in the context of time series data.

3. Virtual sensing using Gaussian process

3.1. Gaussian process formulation

Gaussian processes have emerged as a powerful data-driven strategy for prediction tasks based on prior knowledge [31]. Primarily used for regression problems [32], GPs involve fitting a function to data. However, Gaussian processes have also demonstrated utility for classification and clustering tasks.

In a regression problem, given a set of data training observations, $X = [x_1, \dots, x_n]$ and a target output y , there are countless functions, $f(X)$, that can fit the data. In traditional regression tasks such as linear or polynomial regression, a single function is assumed to describe the underlying data-generating process, and the function parameters can be estimated using different strategies (least squares method, etc.). In these types of regression, the goal is to estimate the parameters of a parametric function.

By contrast, Gaussian processes acknowledge the existence of countless solutions and offer a sophisticated strategy to infer a distribution over possible functions $f(X)$ that fit a set of points, modeling them as a multivariate normal distribution and assigning them a probability $p(f|X)$ [31].

Before observing any data, a prior distribution represents the expected outputs of f over inputs X . GPs aim to learn this underlying distribution and update the prior assumption based on observation data, keeping only the functions that fit the observed data points.

Therefore, given any set of observations X and by defining the possible function outputs $f(X)$ as a random variable, a GP is fully characterized by a mean function, $m(X)$ and covariance function $k(X, X')$ following Eq. (5). Consequently, the function values $f(X)$, corresponding to data points, X , have a distribution according to Eq. (6).

$$f(X) \sim \mathcal{GP}(m(X), k(X, X')) \quad (5)$$

$$p(f|X) = \mathcal{N}(f|m(X), K(X, X)) \quad (6)$$

The mean function, $m(X)$, is an $n \times 1$ vector, representing the most likely characterization of the data (where n is the size of the training set). The covariance matrix, $K(X, X)$, is an $n \times n$ matrix, which influences the shape of the prediction distribution. The covariance matrix is generated according to the kernel function, which is often also called covariance function, pairwise on all the points as described in Eq. (7)

$$K_{ij} = k(x_i, x_j) \quad (7)$$

where the kernel function, $k(x_i, x_j)$, receives two points, x_i and x_j , as an input and returns a similarity measure between those points in the form of a scalar value. As long as a valid kernel function is chosen, K will be a valid covariance matrix for a multivariate Gaussian process [33]. If points x_i and x_j are considered similar by the kernel, the function values at these points, $f(x_i)$ and $f(x_j)$, can be expected to be similar too.

As in many applications of GPs, here we assume a zero-mean prior for simplicity [31,33]. This should be a suitable assumption since (1) all our data measurements are standardized before modeling, and (2) the power of GPs largely lies in the kernel functions.

To fully perform regression via GP, the problem is treated with Bayesian inference, where the current hypothesis—or prior—is to be updated as new information becomes available. In the case of Gaussian processes, this information is the training data. Then, a GP prior is converted into a posterior by conditioning it on a set of training input X and output y observations corresponding to noise-free or noisy realizations of function values f .

In real-world scenarios, the assumption of noise-free or perfect training observation points y is unrealistic since most of our data is afflicted with measurement errors or uncertainty. Gaussian processes offer a simple solution to this problem by modeling the error of the measurements. That means that an error term is independently added to each training point $y = f(X) + \epsilon$, where noise $\epsilon \sim \mathcal{N}(0, \tau^2)$.

Given the observed data and a mean function f estimated by these observed data points, it is possible to make predictions at new points X_* as $f(X_*)$. The joint distribution of f and f_* can be formalized according to Eq. (8), which represents the prior over f_n and f_* , the function values at training and testing points, X_n and X_* .

$$p\left(\begin{bmatrix} f_n \\ f_* \end{bmatrix}\right) = \mathcal{N}\left(\begin{bmatrix} 0_n \\ 0_* \end{bmatrix}, \begin{bmatrix} K_{nn} & K_{n*} \\ K_{*n} & K_{**} \end{bmatrix}\right) \quad (8)$$

It is possible to define the marginal likelihood of y by recognizing that the distribution is a sum of two independent Gaussians: the distribution of function values plus the distribution of the noise term.

$$p\left(\begin{bmatrix} y_n \\ y_* \end{bmatrix}\right) = \mathcal{N}\left(\begin{bmatrix} 0_n \\ 0_* \end{bmatrix}, \begin{bmatrix} K_{nn} + \tau^2 I_n & K_{n*} \\ K_{*n} & K_{**} + \tau^2 I_* \end{bmatrix}\right) \quad (9)$$

Then, we can transform the above joint probability into a conditional probability to obtain our posterior predictive distribution for y_* . This is accomplished by applying the joint-to-conditional transformation of Gaussian variables [31].

$$p(y_*|y_n) = \mathcal{N}(\mu_*, \Sigma_*) \quad (10)$$

where

$$\mu_* = K_{*n}(\tau^2 I_n + K_{nn})^{-1} y_n \quad (11)$$

$$\Sigma_* = \tau^2 I_* + K_{**} - K_{*n}(\tau^2 I_n + K_{nn})^{-1} K_{n*} \quad (12)$$

Additional details and extensive background on the above derivations can be found at [31,33,34]. Computation of μ_* and Σ_* requires the inversion of $(\tau^2 I_n + K_{nn})$, an $n \times n$ matrix, which becomes computationally intractable for larger n .

To outline a solution to this problem, we note that the posterior can also be defined as

$$p(f_*|y) = \int p(f_*|f)p(f|y)df \quad (13)$$

where the conditioning on inputs X and X_* has been made implicit. The second term inside the integral is the posterior over the training latent variables f conditioned on observations y ; the first term is the posterior over predictions f_* conditioned on latent training variables f . For the reasons explained above, both terms are intractable when computing larger training datasets.

3.2. Sparse Gaussian process

Exact Gaussian processes cannot be applied to larger training datasets because their time complexity scales with $O(n^3)$ where n is the size of the training set. Approximate or Sparse Gaussian Processes (SGPs) are based on a small set of m inducing variables that reduce the time complexity to $O(nm^2)$.

The assumption behind SPGs is that there is a small set of m inducing variables f_m evaluated at inputs X_m that describe the function to be modeled sufficiently well so that it is possible to use them as an approximation to f and X and define an approximate posterior

$$q(f_*) = \int p(f_*|f)\phi(f_m)df_m \quad (14)$$

where $\phi(f_m)$ is an approximation to the intractable $p(f_m|y)$:

$$\phi(f_m) = \mathcal{N}(f_m|\mu_m, A_m) \quad (15)$$

The goal is to find optimal values for mean μ_m and covariance matrix A_m . The quality of $\phi(f_m)$ also depends on the location of the inducing inputs X_m . Hence, the goal is to find their optimal values as well. The mean and covariance matrix of the Gaussian approximate posterior $q(f_*)$ are defined in terms of μ_m , A_m and X_m

$$q(y_*|y_n) = \mathcal{N}(\mu_*^q, \Sigma_*^q) \quad (16)$$

where

$$\mu_*^q = K_{*m}K_{mm}^{-1} \quad (17)$$

$$\Sigma_*^q = K_{**} - K_{*m}K_{mm}^{-1}A_mK_{mm}^{-1}K_{m*} \quad (18)$$

For a single test input, this approximate posterior can be computed in $O(nm^2)$ after finding optimal values for μ_m , A_m , and X_m . A popular way for optimizing these parameters is a variational approach focused on minimizing the Kullback–Leibler divergence between the approximate posterior $q(f)$ and the exact posterior $p(f|y)$ over training latent variables f .

The two critical parameters for defining a SGP strategy are the number of Inducing Points m and the location z of those points. The choice of m is a trade-off. A larger m leads to a more accurate approximation but increases computational complexity. A smaller m speeds up computations but might lead to an inadequate approximation. The strategy is often empirical, guided by experimentation and the available computational resources. The location of inducing points should ideally be chosen in regions representative of the data. A subset of data points is commonly placed based on domain knowledge.

3.3. Kernel formulation

Selecting the appropriate model is essential when using Gaussian Processes for prediction. It involves making informed decisions on various aspects of the GP prior model based on the available measurements. These decisions can have a significant impact on the accuracy and performance of the regression model. Model selection is essentially a set of decisions based on a priori knowledge to configure the regression model, such as selecting the covariance form and its corresponding hyperparameters.

The kernel or covariance function is crucial in Gaussian Processes, mainly when the mean function is kept at zero. It specifies the underlying data structure, including its smoothness, periodicity, and noise characteristics. The first step in model selection is choosing the appropriate covariance form that accurately represents the inherent structure of the data. The selection process typically involves domain expertise, data exploration, and visual analysis of the data's features.

After observing the variable relationship and correlations between the target and the input variables in Section 2.3, it was evident that more than one GP model may be needed. In particular, two distinct GP model formulations have been developed to predict our target variables. A baseline model was trained first and evaluated on all target variables. We found that an upgraded model with a more complex kernel function is required to capture the behavior of the region 2 and 3 FA bending moments.

3.3.1. Baseline reference model

This section describes the initial model built for each prediction task. The baseline model has a mean function of zero, so the covariance structure is relied on to represent the relationships between input and target variables. The baseline kernel function is built using just two input features from SCADA: time and acceleration. This kernel function is chosen to simultaneously capture the dynamic periodicity of the measurements from acceleration readings while modeling the linear dependencies with time by including the time vector.

In detail, let k_p be the periodic kernel operating on the time dimension (denoted as x_t) modified via a Matérn kernel and let k_{lin} be the linear kernel operating on the acceleration dimension (denoted as x_{acc}). The combination of the periodic kernel wrapping a Matérn kernel is achieved by modulating the distance r in the Matérn kernel to reflect the periodic structure.

Given the first dimension inputs x_t and x'_t , the composite Periodic–Matérn kernel can be expressed as:

$$k_{periodic-Matern}(x_t, x'_t) = \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\sqrt{2\nu\tilde{r}} \right)^\nu K_\nu \left(\sqrt{2\nu\tilde{r}} \right) \quad (19)$$

where $\nu > 0$ is the smoothness parameter of the Matérn kernel, controlling the smoothness of the modeled function, K_ν is the modified Bessel function of the second kind of order ν , Γ is the gamma function, p is the period of the periodic modulation, indicating how the periodicity is introduced into the kernel, ℓ is the length scale parameter, influencing how quickly similarity (or correlation) between points decreases with distance and $\tilde{r}$ is the modified Euclidean distance between points, incorporating the periodic structure to capture periodic behavior within the Matérn kernel framework

$$\tilde{r} = \sqrt{\frac{2 \sin^2(\pi|x_t - x'_t|/p)}{\ell^2}} \quad (20)$$

On the other hand the linear kernel can be expressed as

$$k_{lin}(x_{acc}, x'_{acc}) = x_{acc}^\top x'_{acc} \quad (21)$$

This kernel captures the linear relationship between the data points, making it suitable for problems where the target variable is expected to change linearly with the input variables.

The composite kernel tailored for the baseline model K is defined as the sum of these kernels:

$$K(x, x') = k_{periodic-Matern}(x_t, x'_t) + k_{lin}(x_{acc}, x'_{acc}) \quad (22)$$

The Matérn represents a popular kernel choice for modeling spatial and temporal data, which captures the smooth trend underlying the data. It models long-term behavior, which may include gradual changes or non-periodic variations. On the other hand, the periodic kernel captures the periodic patterns or oscillations that are superimposed on top of the smooth trend. It models short-term, repetitive behavior with a fixed period. By combining these two kernels, the composite kernel can capture complex patterns in the data that may have both smooth

and periodic components. This can be particularly useful when data exhibits both long-term trends and periodic fluctuations.

Given the more or less pronounced linear correlation exhibited between the target quantities (bending moments) and the acceleration measurements along the same locations (FA and SS), a linear kernel is used to mimic that direct linear relationship.

This kernel function can be successfully used to represent the moment in the SS direction for all the operating regions and the FA moment in the first region of operation, given the absence of any quasi-static influence in those scenarios. In the case of the FA and SS moments in the first operating region, the smoothness coefficient is picked equal to 5/2 for the Matérn kernel. This parameter determines how smooth or rough the modeled functions can be. By looking at the moments in the first region (Figs. 7, 8 and 6), it is clear how, in that case, a smoother function would better fit the regression variables. Meanwhile, a Matérn kernel with a smoothness coefficient equal to 1/2 is used to represent the SS moment in the second and third operational regions, allowing it to capture functions with minimal smoothness. It is particularly suitable for modeling data with non-smooth or irregular behavior, such as data with sharp transitions or discontinuities, which tend to emerge with higher wind speed conditions. This kernel allows for more flexibility in capturing abrupt changes compared to kernels with higher smoothness parameters.

3.3.2. Upgraded model

An upgraded model needs to be built to predict three of the six target variables: FA and SS bending moments in regions 2, and FA in region 3. These variables show increased complexity and correlations with additional operational conditions of the turbine. In addition to time and acceleration, SCADA parameters of wind speed, power output, and rotor speed (rpm) are used as inputs in the upgraded models. Including parameters like wind speed or power output is expected to improve the prediction performance for M_{FA} in particular.

As in the baseline models, the kernel function is created by adding sub-kernels corresponding to each selected input feature. Input parameters are selected according to physics-based knowledge of the system. Specifically, wind speed, power, and rotor speed are each expected to have a relatively high correlation with M_{FA} . The proposed upgrade for this model includes the addition of each of these input parameters as additive linear kernels for the FA moment in the second operational region, according to Eq. (23). This model also increased accuracy in the SS direction for region 2. On the other hand, the model built for the FA moment in the third operational region includes the additional SCADA inputs with a Gaussian kernel function or Radial Basis Function (RBF), as described in Eq. (24). This choice is driven by the nonlinear correlation trends exhibited by the FA moment with the power output and the remaining input parameters in this region. The RBF kernel is used to capture these nonlinear relationships in data and is characterized by its smooth, radial symmetry.

$$K(x, x') = k_{\text{periodic-Matern}}(x_t, x'_t) + k_{\text{lin}}(x_{\text{acc}}, x'_{\text{acc}}) + k_{\text{lin}}(x_{\text{power}}, x'_{\text{power}}) + k_{\text{lin}}(x_{\text{rpm}}, x'_{\text{rpm}}) + k_{\text{lin}}(x_{\text{wind}}, x'_{\text{wind}}) \quad (23)$$

$$K(x, x') = k_{\text{periodic-Matern}}(x_t, x'_t) + k_{\text{lin}}(x_{\text{acc}}, x'_{\text{acc}}) + k_{\text{RBF}}(x_{\text{power}}, x'_{\text{power}}) + k_{\text{RBF}}(x_{\text{rpm}}, x'_{\text{rpm}}) + k_{\text{RBF}}(x_{\text{wind}}, x'_{\text{wind}}) \quad (24)$$

4. Results

This section presents the results of the proposed virtual sensing strategy based on Gaussian Process regression. As mentioned in the previous section, independent models were built and trained for six target variables—the FA and SS bending moment for each of the three

Table 1

Gaussian process vs. Sparse Gaussian process performance on a 10-minute example dataset.

	Gaussian process	Sparse Gaussian process
CPU Time (s)	12.25	3.30
Execution Time (s)	10.69	5.15
	Training/Testing	Training/Testing
RMSE (MNm)	0.879/0.759	0.879/0.759
RRMSE (%)	14.04/12.23	14.04/12.23
TRAC	0.980/0.985	0.980/0.985
MAE (MNm)	0.493/0.476	0.493/0.476

regions of operation, as measured by strain gauges at a particular elevation on an OWT. The elevation chosen for the presentation of the following results is the bending moment at the elevation of 0.34H (see Fig. 2). This elevation is of special interest because it is located in a spot sensitive to fatigue damage. For each target variable to be predicted, the selected GP model is trained on a selection of the data for the target direction and region of operation. The fully trained GP model is then evaluated on unseen measurements from the same region.

Due to the probabilistic nature of GP, the reported prediction of target variables is the mean prediction of 2000 model generations. The confidence interval is computed using the standard deviation of the same 2000 model predictions. Evaluation metrics for prediction accuracy are computed using the mean model prediction. These metrics, as described in Section 2.4, include the MAE, the TRAC score, the RMSE, and the RRMSE on training and testing data sets.

The first investigation of GP modeling for the bending moment regression task concerns usingSGP to enable efficient, high-frequency, long-term predictions. Two GP models – one SGP and one plain GP – were trained on a 10-minute dataset using our defined baseline kernel and mean formulations (the latter assumed zero in this study). Results show that the accuracy of plain and sparse GP model predictions was equal, while the sparse GP model displays superior computational efficiency. Table 1 reports the training computational time and model performance of a regular GP model and a sparse GP model trained on the same 10-minute dataset. Sparse GP models are used moving forward to build the regression models for the six target variables (the FA and SS bending moments for the three regions of operations). For all regression tasks, a confidence interval of 95% was chosen.

For each target variable, either the baseline model or the upgraded model – as defined in Section 3.3 – was trained on the data. The upgraded model is necessary depending on the signal's dynamics for each task of interest. For some tasks, the baseline model – which requires just one SCADA input variable – performs very well. Meanwhile, the upgraded model was necessary for the remaining tasks. The baseline model formulation is applicable for the region 1 SS, region 1 FA, and region 3 SS bending moments. For these target variables, the accelerometer measurements from SCADA are highly informative; there is either no power production (region 1) or steady power production (region 3), which means there is less low-frequency variation in these signals over time. In contrast, the region 2 SS and region 2 and 3 FA bending moments require the upgraded model, likely due to how the variable power production can affect the bending moment in both the FA and SS directions in region 2. In particular, the region 2 FA moment shows a varying mean magnitude, which matches the trends of the power output, as noted in Fig. 8. In region 3, the varying dynamic behavior of the power output has a more relevant effect on the FA bending moment (as compared to the SS), requiring the upgraded model for this target, too. Model selection for each target variable is summarized in Table 2. Model results presented in the remainder of this paper correspond to the model formulation noted in this table.

The first results presented are the regression of the SS bending moment in region 1. Fig. 9 shows the model training results for a 100-minute data set, split into training and testing regions. Fig. 10 provides




Fig. 9. Sparse Gaussian Process predictions for the SS moment in region 1, training and testing set separated by vertical line.




Fig. 10. Sparse Gaussian Process predictions for the SS moment in region 1, testing dataset.

Table 2
Model formulation selected for each target variable.

	Side-Side	Fore-Aft
Region 1	Baseline Model	Baseline Model
Region 2	Upgraded Model	Upgraded Model
Region 3	Baseline Model	Upgraded Model

Table 3
Sparse GP results for the prediction of the side-to-side moment.

Metric	Region 1		Region 2		Region 3	
	Training	Testing	Training	Testing	Training	Testing
RMSE (MNm)	0.843	0.754	0.835	1.186	1.805	1.654
RRMSE (%)	14.05	12.39	16.37	24.34	30.15	27.26
TRAC	0.980	0.985	0.998	0.973	0.943	0.926
MAE (MNm)	0.475	0.489	0.629	0.833	1.396	1.307

a 5-minute excerpt of model predictions on the testing data set, with the computed confidence interval also shown. Evaluation metrics are shown in Table 3, which show a very good match between GP model predictions and true SS bending moment for this task.

Given the homogeneous and consistent behavior in the SS bending moments across regions (see Fig. 5), the model result for region 1 SS bending moment is presented as a representative example of the accuracy and model behavior for all regions of operation. The error metrics for models trained on all regions can be found in Table 3.

In the FA direction, the baseline model is only used on region 1, while both regions 2 and 3 require the upgraded model, as noted in Table 2. During the idling conditions found in region 1, the FA bending

Table 4
Sparse GP results for the prediction of the fore-aft moment.

Metric	Region 1		Region 2		Region 3	
	Training	Testing	Training	Testing	Training	Testing
RMSE (MNm)	0.727	0.868	1.546	1.753	3.807	4.811
RRMSE (%)	4.201	4.955	6.495	6.073	7.912	9.683
TRAC	0.998	0.997	0.996	0.996	0.994	0.991
MAE (MNm)	0.432	0.5420	1.223	1.366	2.999	3.797

moment behaves very similarly to the SS direction due to the low winds and zero power production (see Fig. 5). As a result, the baseline model performs well on this target because the FA accelerometer reading is informative enough to capture the bending moment. Due to the SS direction similarities, region 1 model results are not plotted here, but error metrics are provided in Table 4.

Fig. 11 shows the model training and testing datasets when the baseline model was used on region 2 FA bending moments. The model predictions on the testing set show that the lack of additional SCADA information in the model fails to capture the slow low-frequency behavior of the signal. This behavior is closely related to power output and displays the need for an upgraded model to fit this data.

In regions 2 and 3, the upgraded model formulation can better capture the FA bending moments. The upgraded model involves using other SCADA input features, which increases model complexity but is required to capture the target variable in these regions. The FA bending moment prediction model results for regions 2 and 3 are shown in Figs. 12–15. Table 4 summarizes the results for predicting the FA bending moments using the described GP models in the three operational regions.




Fig. 11. Sparse Gaussian Process predictions for the FA moment in region 2, training and testing set separated by vertical line considering the baseline kernel function.




Fig. 12. Sparse Gaussian Process predictions for the FA moment in region 2, training and testing set separated by vertical line.




Fig. 13. Sparse Gaussian Process predictions for the FA moment in region 2, testing dataset.

The final analysis conducted concerns the model behavior for the prediction of bending moments at different elevations on the OWT. As noted in Section 2.2, bending moments are measured at four elevations of strain gauge installation on the OWT. The above results focus on predicting bending moments at one chosen elevation (elevation three at the transition piece). However, it is interesting to understand the models' predictive capabilities for different elevations.

Following the model design described in the previous sections, differentiated by operational region (Table 2), individual sparse GP models were trained and tested for the prediction of each target bending moment at the four elevations. The results presented in terms of

mean absolute error show that the error on model predictions generally increases with lower elevation bending moments Fig. 16.

Comparing the mean absolute error results with the mean and standard deviation computed over the observed period used for prediction reveals how small the error magnitude is. The error metric and the bending moment statistics are presented in Tables 5 and 6. In particular, for the SS bending moment, the values remain below 1.5 MNm for all operational regions except for region 3 at elevation 4, which is still below 2.5 MNm. Compared with the mean values, which are very small, commonly centered, and revolving around zero for the SS moment, it could be perceived as a high error. However, the




Fig. 14. Sparse Gaussian Process predictions for the FA moment in region 3, training and testing set separated by vertical line.




Fig. 15. Sparse Gaussian Process predictions for the FA moment in region 3, testing dataset.




Fig. 16. Testing mean absolute error obtained in the prediction of the for aft-aft and side-to-side moments for the four different measured elevations.

dispersion of the bending moment provided by the standard deviation indicates a sufficiently wide variation around the mean, contextualizing the MAE metric and showing the optimal predictive capabilities of the sparse GP model. The predictive capabilities are more efficient with respect to the FA bending moment. In operational region 1, the mean

absolute error remains below 1 MNm, where the mean of the bending moment is negative and higher than 15 MNm in absolute magnitude. Finally, in the two most challenging operational regions, regions 2 and 3, the MAE oscillates between 0.4 MNm and 2.5 MNm and 1 MNm and 5MNm, respectively, for the four elevations. Compared with the

Table 5

FA bending moment (MNm): MEA metric for the sparse GP results and the mean and standard deviation values of the bending moments computed over the observed period of time used for prediction.

Elevation	Region 1		Region 2		Region 3	
	MAE	Mean (STD)	MAE	Mean (STD)	MAE	Mean (STD)
1	0.078	−15.846 (0.827)	0.444	−6.831 (3.205)	1.059	−5.016 (1.516)
2	0.354	−17.171 (3.460)	1.240	20.706 (13.337)	2.666	33.361 (4.920)
3	0.542	−16.667 (4.678)	1.366	35.144 (16.945)	3.371	48.743 (6.482)
4	0.997	−18.487 (9.061)	2.193	47.620 (22.319)	4.836	63.212 (9.021)

Table 6

SS bending moment (MNm): MEA metric for the sparse GP results and the mean and standard deviation values of the bending moments computed over the observed period of time used for prediction.

Elevation	Region 1		Region 2		Region 3	
	MAE	Mean (STD)	MAE	Mean (STD)	MAE	Mean (STD)
1	0.105	−0.931 (1.056)	0.187	−4.661 (1.223)	0.370	−5.890 (0.765)
2	0.364	−0.081 (3.952)	0.753	−3.041 (3.368)	1.007	−6.312 (3.153)
3	0.489	1.238 (5.412)	0.833	−0.848 (4.699)	1.307	−4.512 (4.265)
4	1.224	−1.603 (7.775)	1.299	−4.846 (6.444)	2.133	−2.568 (5.990)

mean values that vary between −7 MNm and 63 MNm, it proves the outstanding performance of the GP model, considering the error in the prediction of the bending moments is extremely small.

5. Conclusions

This work developed a virtual sensing strategy for predicting the bending moment time histories of an offshore wind turbine using Gaussian Process regression. The work aims to overcome the need to install strain sensors along the tower and foundation by building and training a virtual model. The model can reproduce the bending moment quantities along the local primary directions of the tower using only readily available SCADA data as model input features.

A baseline model formulation is built, which only considers the time index and the nacelle accelerometer reading as input variables. It is shown that this baseline model can be trained on data for the SS bending moment in the first and third regions and the FA bending moment in region 1 and provides accurate predictions on testing datasets with high confidence in the prediction. The modeling task considers the varying structural dynamics of the OWT along the principal directions as the operational conditions (wind speed, power, etc.) change over time. Then, it is shown how the baseline model fails to predict the bending moment in the FA direction as the wind speed and power output ramp up in the operational region 2. Therefore, an upgraded model was built for the region 2 and 3 FA bending moments which considered additional SCADA data input variables – power output, wind speed, and rotor speed – for the modeling task. The upgraded model appropriately captures the tower's overall high and low-frequency dynamics with a promising level of accuracy and magnitude error.

The model can successfully predict the bending moment at different elevations along the height of the tower. The prediction error increases as the tower moves towards the bottom for hotspot elevations located far away from the SCADA data acquisition location, where the detailed dynamics tend to be different from the operational conditions measured at the top. However, the overall behavior is accurately captured. This work currently only predicts the bending moment response at locations that are instrumented with strain gauges, which may be used in model training. In order to extrapolate these predictions to non-instrumented elevations, future work will consider a physics/data-driven hybrid approach in which bending moment readings along the structure could be used in tandem with modal expansion techniques to simplify the individual GP models for each elevation into a single model.

Model results show a good fit for all target variables. However, the bending moments for which the baseline model applies perform slightly better than the targets that required the upgraded model. The accuracy of all models is deemed suitable for fatigue calculations. Future work

will focus on the sensitivity of fatigue calculations to model predictions and developing a flexible model that can be trained and applied to all operational regions.

CRedit authorship contribution statement

Bridget Moynihan: Writing – original draft, Validation, Methodology, Formal analysis, Conceptualization. **Eleonora M. Tronci:** Writing – review & editing, Validation, Methodology, Investigation, Formal analysis. **Michael C. Hughes:** Supervision, Methodology, Conceptualization. **Babak Moaveni:** Writing – review & editing, Supervision, Funding acquisition. **Eric Hines:** Writing – review & editing, Supervision.

Declaration of competing interest

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

Acknowledgments

Partial support of this study by National Science Foundation, United States grant 2230630. A special acknowledgment goes to Finn Rüdinger and Ross McAdam at Ørsted for the technical support. The opinions, findings, and conclusions expressed in this document are those of the authors and do not necessarily represent the views of the sponsors and organizations involved in this project.

References

- [1] F. de N Santos, P. D'Antuono, K. Robbelein, N. Noppe, W. Weijtjens, C. Devriendt, Long-term fatigue estimation on offshore wind turbines interface loads through loss function physics-guided learning of neural networks, *Renew. Energy* 205 (2023) 461–474, <http://dx.doi.org/10.1016/j.renene.2023.01.093>, URL <https://www.sciencedirect.com/science/article/pii/S0960148123001027>.
- [2] A. Iliopoulos, W. Weijtjens, D. Van Hemelrijck, C. Devriendt, Remaining useful life assessment of offshore wind turbines: Validation of virtual sensing on long term measurements, in: *NDT.Net*, Vol. 21, eighth ed., NDT.net, 2016, pp. 1–8, URL <http://www.ndt.net/app/EWSHM2016>. 8th European Workshop On Structural Health Monitoring, EWSHM 2016 ; Conference date: 05-07-2016 Through 08-07-2016.
- [3] J. Schijve (Ed.), Introduction to fatigue of structures and materials, in: *Fatigue of Structures and Materials*, Springer Netherlands, Dordrecht, 2009, pp. 1–9, http://dx.doi.org/10.1007/978-1-4020-6808-9_1.
- [4] M. Bezziccheri, P. Castellini, P. Evangelisti, C. Santolini, N. Paone, Measurement of mechanical loads in large wind turbines: Problems on calibration of strain gage bridges and analysis of uncertainty, *Wind Energy* 20 (12) (2017) 1997–2010, <http://dx.doi.org/10.1002/we.2136>, arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1002/we.2136>.

- [5] L. Liu, S.M. Kuo, M. Zhou, Virtual sensing techniques and their applications, in: 2009 International Conference on Networking, Sensing and Control, IEEE, 2009, pp. 31–36.
- [6] D. Martin, N. Kühl, G. Satzger, Virtual sensors, *Bus. Inform. Syst. Eng.* 63 (2021) 315–323.
- [7] J.F. Toftekar, J.T. Vestermark, M.S. Jepsen, Uncertainty of virtually sensed stress ranges in offshore wind support structures, in: International Conference on Offshore Mechanics and Arctic Engineering, vol. 86830, American Society of Mechanical Engineers, 2023, V001T01A011.
- [8] S. Gillijns, B. De Moor, Unbiased minimum-variance input and state estimation for linear discrete-time systems, *Automatica* 43 (1) (2007) 111–116.
- [9] E. Lourens, C. Papadimitriou, S. Gillijns, E. Reynders, G. De Roeck, G. Lombaert, Joint input-response estimation for structural systems based on reduced-order models and vibration data from a limited number of sensors, *Mech. Syst. Signal Process.* 29 (2012) 310–327.
- [10] C. Papadimitriou, C.-P. Fritzen, P. Kraemer, E. Ntotsios, Fatigue predictions in entire body of metallic structures from a limited number of vibration sensors using Kalman filtering, *Struct. Control Health Monit.* 18 (5) (2011) 554–573.
- [11] A. Mehrjoo, M. Song, B. Moaveni, C. Papadimitriou, E. Hines, Optimal sensor placement for parameter estimation and virtual sensing of strains on an offshore wind turbine considering sensor installation cost, *Mech. Syst. Signal Process.* 169 (2022) 108787.
- [12] M. Song, B. Moaveni, H. Ebrahimian, E. Hines, A. Bajric, Joint parameter-input estimation for digital twinning of the block island wind turbine using output-only measurements, *Mech. Syst. Signal Process.* 198 (2023) 110425.
- [13] A. Iliopoulos, R. Shirzadeh, W. Weijtjens, P. Guillaume, D. Van Hemelrijck, C. Devriendt, A modal decomposition and expansion approach for prediction of dynamic responses on a monopile offshore wind turbine using a limited number of vibration sensors, *Mech. Syst. Signal Process.* 68 (2016) 84–104.
- [14] D. Augustyn, R.R. Pedersen, U.T. Tygesen, M.D. Ulriksen, J.D. Sørensen, Feasibility of modal expansion for virtual sensing in offshore wind jacket substructures, *Mar. Struct.* 79 (2021) 103019.
- [15] K. Maes, A. Iliopoulos, W. Weijtjens, C. Devriendt, G. Lombaert, Dynamic strain estimation for fatigue assessment of an offshore monopile wind turbine using filtering and modal expansion algorithms, *Mech. Syst. Signal Process.* 76 (2016) 592–611.
- [16] M.-S. Nabiyan, F. Khoshnoudian, B. Moaveni, H. Ebrahimian, Mechanics-based model updating for identification and virtual sensing of an offshore wind turbine using sparse measurements, *Struct. Control Health Monit.* (2021) <http://dx.doi.org/10.1002/stc.2647>.
- [17] N. Noppe, K. Tatsis, E. Chatzi, C. Devriendt, W. Weijtjens, Fatigue stress estimation of offshore wind turbine using a Kalman filter in combination with accelerometers, in: Proceedings of International Conference on Noise and Vibration Engineering (ISMA 2018), International Conference on Uncertainty in Structural Dynamics, USD 2018, KU Leuven, Department of Mechanical Engineering, 2018, pp. 4693–6701.
- [18] N. Dimitrov, T. Göçmen, Virtual sensors for wind turbines with machine learning-based time series models, *Wind Energy* 25 (9) (2022) 1626–1645.
- [19] J. Han, X.-P. Zhang, F. Wang, Gaussian process regression stochastic volatility model for financial time series, *IEEE J. Sel. Top. Sign. Proces.* 10 (6) (2016) 1015–1028, <http://dx.doi.org/10.1109/JSTSP.2016.2570738>.
- [20] E.J. Cross, T.J. Rogers, Physics-derived covariance functions for machine learning in structural dynamics** the authors would like to acknowledge the support of the EPSRC, particularly through grant reference number EP/S001565/1, IFAC-PapersOnLine 54 (7) (2021) 168–173, <http://dx.doi.org/10.1016/j.ifacol.2021.08.353>, URL <https://www.sciencedirect.com/science/article/pii/S2405896321011277>. 19th IFAC Symposium on System Identification SYSID 2021.
- [21] E.J. Cross, S.J. Gibson, M.R. Jones, D.J. Pitchforth, S. Zhang, T.J. Rogers, Physics-informed machine learning for structural health monitoring, in: A. Cury, D. Ribeiro, F. Ubertini, M.D. Todd (Eds.), *Structural Health Monitoring Based on Data Science Techniques*, Springer International Publishing, Cham, 2022, pp. 347–367, http://dx.doi.org/10.1007/978-3-030-81716-9_17.
- [22] S.J. Gibson, T.J. Rogers, E.J. Cross, Integrating physical knowledge into Gaussian process regression models for probabilistic fatigue assessment, in: P. Rizzo, A. Milazzo (Eds.), *European Workshop on Structural Health Monitoring*, Springer International Publishing, Cham, 2023, pp. 472–481.
- [23] E. Papatheou, N. Dervilis, A.E. Maguire, I. Antoniadou, K. Worden, A performance monitoring approach for the novel Lillgrund offshore wind farm, *IEEE Trans. Ind. Electron.* 62 (10) (2015) 6636–6644.
- [24] L.D. Avendaño-Valencia, I. Abdallah, E. Chatzi, Virtual fatigue diagnostics of wake-affected wind turbine via Gaussian process regression, *Renew. Energy* 170 (2021) 539–561.
- [25] J. Bilbao, E.-M. Lourens, A. Schulze, L. Ziegler, Virtual sensing in an onshore wind turbine tower using a Gaussian process latent force model, *Data-Centric Eng.* 3 (2022) e35.
- [26] J. Zou, E.-M. Lourens, A. Cicirello, Virtual sensing of subsoil strain response in monopile-based offshore wind turbines via Gaussian process latent force models, *Mech. Syst. Signal Process.* 200 (2023) 110488.
- [27] F. Pimenta, J. Pacheco, S. Pereira, F. Magalhães, Reconstructing the bending moments time history of wind turbine tower from acceleration measurements using Gaussian processes, 2265, (3) IOP Publishing, 2022, 032080.
- [28] B. Moynihan, A. Mehrjoo, B. Moaveni, R. McAdam, F. Rüding, E. Hines, System identification and finite element model updating of a 6 MW offshore wind turbine using vibrational response measurements, *Renew. Energy* 219 (2023) 119430, <http://dx.doi.org/10.1016/j.renene.2023.119430>, URL <https://www.sciencedirect.com/science/article/pii/S0960148123013459>.
- [29] I. Rychlik, Fatigue cycle counting, in: Q.J. Wang, Y.-W. Chung (Eds.), *Encyclopedia of Tribology*, Springer US, Boston, MA, 2013, pp. 1032–1041, http://dx.doi.org/10.1007/978-0-387-92897-5_1076.
- [30] B. Nabuco, T. Friis, M. Tarpø, S. Amador, E.I. Katsanos, R. Brincker, Nonlinear Strain Estimation Based on Linear Parameters, in: International Conference on Offshore Mechanics and Arctic Engineering, Volume 3: Structures, Safety, and Reliability, 2018, <http://dx.doi.org/10.1115/OMAE2018-77785>, arXiv:<https://asmedigitalcollection.asme.org/OMAE/proceedings-pdf/OMAE2018/51227/V003T02A039/2535535/v003t02a039-omae2018-77785.pdf>. V003T02A039.
- [31] C.E. Rasmussen, C.K. Williams, et al., *Gaussian Processes for Machine Learning*, vol. 1, Springer, 2006.
- [32] M. Seeger, Gaussian processes for machine learning, *Int. J. Neural Syst.* 14 (02) (2004) 69–106.
- [33] C.B. Do, H. Lee, *Gaussian processes*, 2008.
- [34] J. Wang, An intuitive tutorial to Gaussian processes regression, 2022, arXiv:2009.10862.