

Virtual sensing in an onshore wind turbine tower using a Gaussian process latent force model

Bilbao Nieva, J.A.; Lourens, E.; Schulze- Bonhage, Andreas; Ziegler, Lisa

DOI

[10.1017/dce.2022.38](https://doi.org/10.1017/dce.2022.38)

Publication date

2022

Document Version

Final published version

Published in

Data-Centric Engineering

Citation (APA)

Bilbao Nieva, J. A., Lourens, E., Schulze- Bonhage, A., & Ziegler, L. (2022). Virtual sensing in an onshore wind turbine tower using a Gaussian process latent force model. *Data-Centric Engineering*, 3(3), Article e35. <https://doi.org/10.1017/dce.2022.38>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

RESEARCH ARTICLE

Virtual sensing in an onshore wind turbine tower using a Gaussian process latent force model

Joaquin Bilbao^{1,2} , Eliz-Mari Lourens^{1,*} , Andreas Schulze² and Lisa Ziegler²

¹Faculty of Civil Engineering and Geosciences, Delft University of Technology, 2628 CN Delft, The Netherlands

²EnBW Energie Baden-Württemberg AG, 20095 Hamburg, Germany

*Corresponding author: E-mail: E.Lourens@tudelft.nl

Received: 04 November 2021; **Revised:** 18 August 2022; **Accepted:** 03 November 2022

Keywords: Fatigue load monitoring; Gaussian process; input estimation; latent force models; state estimation

Abstract

Wind turbine towers are subjected to highly varying internal loads, characterized by large uncertainty. The uncertainty stems from many factors, including what the actual wind fields experienced over time will be, modeling uncertainties given the various operational states of the turbine with and without controller interaction, the influence of aerodynamic damping, and so forth. To monitor the true experienced loading and assess the fatigue, strain sensors can be installed at fatigue-critical locations on the turbine structure. A more cost-effective and practical solution is to predict the strain response of the structure based only on a number of acceleration measurements. In this contribution, an approach is followed where the dynamic strains in an existing onshore wind turbine tower are predicted using a Gaussian process latent force model. By employing this model, both the applied dynamic loading and strain response are estimated based on the acceleration data. The predicted dynamic strains are validated using strain gauges installed near the bottom of the tower. Fatigue is subsequently assessed by comparing the damage equivalent loads calculated with the predicted as opposed to the measured strains. The results confirm the usefulness of the method for continuous tracking of fatigue life consumption in onshore wind turbine towers.

Impact Statement

As the first large-scale wind farms reach their end-of-lifetime, attention is drawn to the importance of having accurate knowledge on structural reserves that can enable wind farm operators to make reliable decisions about lifetime extension in the future. In this contribution, a methodology for fatigue load monitoring on the basis of a sensor network is proposed and validated. The loads exciting the tower are estimated in conjunction with its dynamic response using a Gaussian process latent force model. Estimates of the full-field strain response can be derived, and therewith fatigue consumption at critical locations in the structure. The results confirm the potential of the proposed methodology for continuous tracking of fatigue life consumption in the structures supporting wind turbine rotors.

1. Introduction

Towers of onshore wind turbines are highly susceptible to fatigue damage given the cyclic wind loading, where the cyclic nature of the experienced wind loading is a consequence of rotational sampling of the wind field by the rotor. In recent years, a lot of effort has been spent on indirect measurement of the true

fatigue life consumption by means of so-called virtual sensing techniques. Virtual sensing refers to the prediction of response quantities at unmeasured, but critical locations on a structure, based on limited data from sensors installed at easily accessible locations. In the case of wind turbines, also data from the Supervisory Control and Data Acquisition (SCADA) system, present in all modern wind turbines, can be utilized for the purpose of virtual sensing.

Several different techniques for virtual sensing have been proposed. A distinction can be made between data-driven models, physics-based models, and hybrid techniques. Data-driven techniques use artificial intelligence to train black-box models from measured or simulated data. For example, Cosack (2010), Smolka and Cheng (2013), Vera-Tudela and Kühn (2017), and Schröder et al. (2018) used artificial neural networks with SCADA as input to estimate damage equivalent loads at different locations of the structure. Ziegler et al. (2017, 2019) applied a k-nearest neighbor regression algorithm to extrapolate strain measurements along a monopile substructure. Recently, d N Santos et al. (2021) estimate thrust loads from 1s SCADA data through artificial neural networks, which are then combined with high-resolution acceleration measurements to estimate tower loads. Gulgec et al. (2020) apply neural networks to estimate time-history of strain responses directly from acceleration measurements. Other types of machine learning have also been applied in the context of virtual sensing, such as a Gaussian process regression by Avendaño-Valencia et al. (2021).

In the physics-based approach, a complete numerical model of the asset is used for virtual sensing. This requires an update of the model with a limited set of measurement data. Augustyn et al. (2020) recently used *in situ* acceleration measurements from a jacket substructure of an offshore wind turbine to update its numerical counterpart. The updated model can then be used for aero-hydro-elastic simulations to obtain loads on unmeasured locations. Currently, there is a trend toward framing these concepts as “digital twins.” A digital twin refers to a digital replicate of the physical asset continuously updated with real-time measurements (Wagg et al., 2020). Recent works promoting the concept of digital twins for wind turbines have been published by Branlard et al. (2020), Schröder (2020), and Augustyn (2021).

In hybrid schemes, measurements are combined with simplified numerical models of the asset for virtual sensing. This includes model-based extrapolation of responses and techniques where the forces driving the response are also estimated. In the former category, one finds the modal decomposition and expansion approach (MD&E), first applied in the field of (offshore) wind turbine foundation monitoring by Iliopoulos et al. (2014, 2016) and Noppe et al. (2016). Using MD&E, a classical inverse problem is solved where modal coordinates of the measured response are estimated through a pseudo-inverse of the model-based eigenvector matrix. To improve accuracy, a system identification experiment followed by a model update can be performed to obtain the eigenvectors. Also, the inverse problems can be solved separately for different frequency bands using multi-band MD&E (Iliopoulos et al., 2017), where the set of eigenvectors used in the inversion can be more carefully tailored to the expected response from each mode within the frequency band.

The second category of virtual sensing techniques encompasses the Kalman filter based methods (Kalman, 1960), in numerous variants. In this case, the response extrapolation is performed by simultaneously estimating the driving forces and the states of the system, where the latter is often defined to be the modal coordinates. This is done in a stochastic setting, accounting for both modeling and measurement errors, and applying principles of minimum-variance unbiased estimation. In this category falls, in no particular order, the augmented Kalman filter (Lourens, Reynders, et al., 2012), the dual Kalman filter (Eftekhar Azam et al., 2015; Tatsis and Lourens, 2016), and the joint input-response estimator (Gillijns and De Moor, 2007; Lourens, Papadimitriou, et al., 2012; Maes et al., 2016). Although strongly related in terms of their basic principles, these variants differ in their modeling of the input force. In the dual and augmented filter, the force is modeled as a random walk, and its variance is tuned based on the measurement data. In the joint input-response estimator, however, no assumptions on the evolution of the force in time are made, and its time history is found solely through minimum-variance unbiased estimation in the context of a set of uncertain model equations. Although excellent

results have been obtained with the Kalman filter based methods in various applications, both in the laboratory and full-scale, a major drawback remains their instability. Especially in the case of acceleration-only measurements, the more classical filters require modifications to improve their stability (Eftekhar Azam et al., 2015; Naets et al., 2015). Often these modifications come at a cost, for example, the need to include additional parameters that require a cumbersome and fragile tuning process.

In this contribution, the virtual sensing problem is solved using a Gaussian process latent force model (GPLFM). The GPLFM relieves instability problems and/or fragile tuning processes that afflict the Kalman filter based methods, while at the same time preserving the benefits of a simultaneous estimation of both the driving forces and the system states. The GPLFM was first applied to the joint input-state estimation problem by Nayek et al. (2019) and extended for the solution of joint input-state-parameter identification by Rogers et al. (2020), both drawing on work by Hartikainen and Särkkä (2010), Särkkä (2013), and Särkkä and Solin (2019). Unlike the Kalman filter based methods, it models the driving forces as Gaussian processes, driven by specified covariance functions. The parameters defining the covariance functions can be optimized for using the measured data, resulting in a flexible, yet robust framework for incorporating prior information on the time-varying behavior of the force.

In the following sections, an application of the GPLFM for virtual sensing on a full-scale wind turbine tower is presented. Section 2 provides the theoretical framework, followed by the full-scale case study in Section 3, including the validation through strain gauges installed near the bottom of the tower. For finding the parameters that define the covariance functions of the unknown forcing, a new procedure exploiting the physics underlying the GPLFM is suggested in Section 2.6 and applied in Section 3.3. It is noted here that the estimation of these unknown parameters is commonly realized by finding the parameters such that the posterior distribution is maximized. More theoretical background can be found in Särkkä (2013), where several strategies for parameter estimation are discussed, such as maximum *a posteriori* (applied in Nayek et al., 2019) and Markov Chain Monte Carlo (applied in Rogers et al., 2020). The downside of these approaches is the need for a robust and expensive optimization process, and furthermore, the loss of physical relevance of this process. It is shown that the suggested simpler, more intuitive procedure can be effective for finding the hyperparameters needed to characterize the covariance of the unknown forces.

2. Theoretical Framework

The GPLFM is built by combining a mechanical model of the wind turbine structure and a stochastic model of the unknown load. Both models are expressed in state-space form and are subsequently combined into an augmented state-space model. This augmented model is a stochastic representation of the system, including the response and the load. The following subsections describe each model and show how the resulting augmented model is handled. In particular, the computation of the prior and the posterior as well as the estimation of the hyperparameters of the model are presented.

2.1 Mechanical model

Consider the governing equation of motion of a linear system discretized in space,

$$\mathbf{M}\ddot{\mathbf{u}}(t) + \mathbf{C}\dot{\mathbf{u}}(t) + \mathbf{K}\mathbf{u}(t) = \mathbf{S}_p \mathbf{p}(t),$$

where the mass, damping, and stiffness matrices are represented by $\mathbf{M}, \mathbf{C}, \mathbf{K} \in \mathbb{R}^{n_u \times n_u}$, respectively, and the load locations are defined through $\mathbf{S}_p \in \mathbb{R}^{n_u \times n_p}$. The response of the system is defined as $\mathbf{u}(t) \in \mathbb{R}^{n_u}$, and the loads as $\mathbf{p}(t) \in \mathbb{R}^{n_p}$. The mechanical model employed in this work considers only a single load: $n_p = 1$.

The classical modal representation of this equation of motion reads

$$\mathbf{I}\ddot{\mathbf{u}}_m(t) + \mathbf{\Gamma}\dot{\mathbf{u}}_m(t) + \mathbf{\Omega}^2 \mathbf{u}_m(t) = \mathbf{\Phi}^T \mathbf{S}_p \mathbf{p}(t),$$

with $\mathbf{u}_m \in \mathbb{R}^{n_u}$ the modal coordinates, $\Phi \in \mathbb{R}^{n_u \times n_u}$ the mass-normalized eigenvector matrix (i.e., $\Phi^T \mathbf{M} \Phi = \mathbf{I}$), and $\Gamma \in \mathbb{R}^{n_u \times n_u}$ and $\Omega \in \mathbb{R}^{n_u \times n_u}$ diagonal matrices formed by the natural frequencies ω_n and damping ratios ζ_n so that $\Omega = \text{diag}(\omega_n)$ and $\Gamma = \text{diag}(2\zeta_n \omega_n)$, respectively.

A subset of modes $n_m \leq n_u$ is assumed to be sufficient to accurately approximate the response. In other words, modal reduction is employed as

$$\mathbf{u}(t) \approx \Phi_r \mathbf{u}_{mr}(t),$$

with the reduced mode shape matrix $\Phi_r \in \mathbb{R}^{n_u \times n_m}$ and modal coordinates $\mathbf{u}_{mr} \in \mathbb{R}^{n_m}$. The reduced system matrices are written as $\Omega_r \in \mathbb{R}^{n_m \times n_m}$ and $\Gamma_r \in \mathbb{R}^{n_m \times n_m}$.

Finally, the reduced-order modal model is written in state-space form

$$\underbrace{\begin{Bmatrix} \dot{\mathbf{u}}_{mr} \\ \ddot{\mathbf{u}}_{mr} \end{Bmatrix}}_{\dot{\mathbf{x}}_m(t)} = \underbrace{\begin{bmatrix} \mathbf{0} & \mathbf{I} \\ -\Omega_r^2 & -\Gamma_r \end{bmatrix}}_{\mathbf{A}_{cm}} \underbrace{\begin{Bmatrix} \mathbf{u}_{mr} \\ \dot{\mathbf{u}}_{mr} \end{Bmatrix}}_{\mathbf{x}_m(t)} + \underbrace{\begin{bmatrix} \mathbf{0} \\ \Phi_r^T \mathbf{S}_p \end{bmatrix}}_{\mathbf{B}_{cm}} p(t), \tag{1}$$

with $\mathbf{x}_m \in \mathbb{R}^{2n_m}$, $\mathbf{A}_{cm} \in \mathbb{R}^{2n_m \times 2n_m}$, $\mathbf{B}_{cm} \in \mathbb{R}^{2n_m \times n_p}$.

2.2. Stochastic load model

The stochastic load model is defined as a stationary Gaussian process with zero mean and a Matérn covariance function with $\nu = 3/2$ (see Rasmussen and Williams, 2006), hereafter referred to as the Matérn32 covariance function:

$$p(t) \sim \mathcal{GP}(0, \kappa_{M32}(\tau; \sigma, l_{sc})),$$

$$\kappa_{M32}(\tau; \sigma, l_{sc}) = \sigma^2 \left\{ 1 + \frac{\sqrt{3}|\tau|}{l_{sc}} \right\} e\left(-\frac{\sqrt{3}|\tau|}{l_{sc}}\right),$$

where τ corresponds to the time lag $\tau = |t - t'|$. The covariance function requires the definition of two hyperparameters: σ and l_{sc} , where σ controls the amplitude of the covariance function, and l_{sc} the extent of the correlation in time. It is highlighted that σ has force units (N), as it is related to the amplitude of the load, and l_{sc} has time units (s), since it is directly related to the expected frequency content of the load: a lower l_{sc} result in the presence of higher frequencies and vice versa. For further details, see Rasmussen and Williams (2006) and Hartikainen and Särkkä (2010).

A Gaussian process with a Matérn32 covariance function can be expressed as the output of the following linear time invariant stochastic differential equation (Hartikainen and Särkkä, 2010):

$$\ddot{p}(t) + \left(2\frac{\sqrt{3}}{l_{sc}}\right)\dot{p}(t) + \left(\frac{3}{l_{sc}^2}\right)p(t) = w(t). \tag{2}$$

Note that it is a stochastic differential equation as it is driven by a white Gaussian noise $w(t)$ with a known spectral density

$$w(t) \sim \mathcal{N}(0, \sigma_w^2); \sigma_w^2 = \frac{12\sqrt{3}\sigma^2}{l_{sc}^3}.$$

Equation (2) can again be written in continuous-time state-space form as

$$\underbrace{\begin{Bmatrix} \dot{p}(t) \\ \ddot{p}(t) \end{Bmatrix}}_{\dot{\mathbf{z}}(t)} = \underbrace{\begin{bmatrix} 0 & 1 \\ -\frac{3}{l_{sc}^2} & -2\frac{\sqrt{3}}{l_{sc}} \end{bmatrix}}_{\mathbf{F}_c} \underbrace{\begin{Bmatrix} p(t) \\ \dot{p}(t) \end{Bmatrix}}_{\mathbf{z}(t)} + \underbrace{\begin{bmatrix} 0 \\ 1 \end{bmatrix}}_{\mathbf{L}_c} w(t), \tag{3}$$

with $\mathbf{z}(t) \in \mathbb{R}^2$, $\mathbf{F}_c \in \mathbb{R}^{2 \times 2}$ and $\mathbf{L}_c \in \mathbb{R}^2$. The load $p(t)$ is directly related to the state vector $\mathbf{z}(t)$ by

$$p(t) = \mathbf{H}_c z(t),$$

with $\mathbf{H}_c = [1 \ 0]$.

Note that the performance of the GPLFM in virtual sensing depends on an appropriate choice of the GP kernel and its hyperparameters. A large number of convenient kernels are available, where a selection is made based on prior knowledge of the time-varying behavior of the modeled process, for example, periodicity, smoothness, or differentiability. The Matérn kernel chosen here is used to characterize processes that are finitely differentiable (Hartikainen and Särkkä, 2010), and is often employed in engineering applications because of its flexibility in accommodating a broad range of process traits.

2.3. Augmented model

Both state space models (equations 1 and 3) can be combined into an augmented system of equations as

$$\underbrace{\begin{bmatrix} \dot{x}_m(t) \\ \dot{z}(t) \end{bmatrix}}_{\dot{z}^a(t)} = \underbrace{\begin{bmatrix} \mathbf{A}_{cm} & \mathbf{B}_{cm}\mathbf{H}_c \\ 0 & \mathbf{F}_c \end{bmatrix}}_{\mathbf{F}_c^a} \underbrace{\begin{bmatrix} x_m(t) \\ z(t) \end{bmatrix}}_{z^a(t)} + \underbrace{\begin{bmatrix} \mathbf{0} \\ \mathbf{L}_c w(t) \end{bmatrix}}_{w^a(t)} \tag{4}$$

with

$$w^a(t) \sim \mathcal{N}(0, \mathbf{Q}_c^a); \mathbf{Q}_c^a = \begin{bmatrix} 0 & 0 \\ 0 & L_c \sigma_w^2 \mathbf{I}_c^T \end{bmatrix},$$

Where $\mathbf{Q}_c^a \in \mathbb{R}^{2 \times 2}$ is the augmented noise covariance matrix, $z^a(t) \in \mathbb{R}^{2n_m+2}$, and $\mathbf{F}_c^a \in \mathbb{R}^{(2n_m+2) \times (2n_m+2)}$.

Note that the augmented state-space model is a stochastic differential equation. The augmented state has mean $\hat{z}^a(t) = \mathbb{E}[z^a(t)]$ and covariance $\mathbf{P}^a(t) = \mathbb{E}[z^a(t)z^a(t)^T]$. Expressions that define the derivatives of these statistics can be found through equation (4). The details can be found from Särkkä and Solin (2019):

$$\dot{\hat{z}}^a(t) = \mathbf{F}_c^a \hat{z}^a(t), \tag{5}$$

$$\dot{\mathbf{P}}^a(t) = \mathbf{F}_c^a \mathbf{P}^a(t) + \mathbf{P}^a(t) \mathbf{F}_c^{aT} + \mathbf{Q}_c^a, \tag{6}$$

with $\hat{z}^a(t) \in \mathbb{R}^{2n_m+2}$, $\mathbf{P}^a(t) \in \mathbb{R}^{(2n_m+2) \times (2n_m+2)}$.

It is highlighted that when combining the (deterministic) mechanical model (equation 1) with the stochastic model for the load (equation 3) into the augmented model (equation 4), a stochastic definition for the state of the system is obtained (augmented state z^a). In particular, z^a is now a Gaussian process: the response of the structure is recognised to be stochastic since it is driven by a stochastic load.

2.4. The prior distribution

The prior distribution of the augmented state $p(z^a(t))$ can be obtained from the stationary solutions of the expressions for the statistics (equations 5 and 6, first derivatives equal to zero):

$$\mathbf{F}_c^a \hat{z}_\infty^a = \mathbf{0}, \tag{7}$$

$$\mathbf{F}_c^a \mathbf{P}_\infty^a + \mathbf{P}_\infty^a \mathbf{F}_c^{aT} + \mathbf{Q}_c^a = \mathbf{0}, \tag{8}$$

Thus, as a result of equation (7), the stationary expected value is zero $\hat{z}_\infty^a = \mathbf{0}$. The stationary covariance \mathbf{P}_∞^a can be found numerically by solving the continuous-time Lyapunov equation, equation (8). The prior distribution is then defined as

$$p(z^a(t)) = \mathcal{N}(\mathbf{0}, \mathbf{P}_\infty^a).$$

This is a prior for the augmented state which contains information regarding the response of the system as well as the load. Therefore, when defining a prior for the load, a prior for the response is also defined.

2.5. The posterior distribution

The posterior distribution of the augmented state refers to the conditional distribution upon some observations $\mathbf{y}_{[1:N]}$ of the response

$$p(\mathbf{z}_{[k]}^a | \mathbf{y}_{[1:N]}) = \mathcal{N}(\hat{\mathbf{z}}_{[k]}^a, \mathbf{P}_{[k]}^a),$$

where discrete time notation is used. Kalman filtering and smoothing can be applied to efficiently compute the posterior distribution, as suggested by Hartikainen and Särkkä (2010). This requires the definition of the discrete version of the augmented model and the observation equation.

2.5.1. Discrete augmented model

The expressions found for the statistics (equations 5 and 6) can be solved, as these are ordinary differential equations. The solution can be expressed using discrete notation, evaluating between time steps, as

$$\hat{\mathbf{z}}_{[k+1]}^a = \underbrace{\mathbf{e}^{\mathbf{F}_c^a \Delta t}}_{\mathbf{F}^a} \hat{\mathbf{z}}_{[k]}^a,$$

$$\mathbf{P}_{[k+1]}^a = \mathbf{F}^a \mathbf{P}_{[k]}^a \mathbf{F}^{aT} + \underbrace{\int_{k\Delta t}^{(k+1)\Delta t} \mathbf{e}^{\mathbf{F}_c^a((k+1)\Delta t - \tau)} \mathbf{Q}_c^a \mathbf{e}^{\mathbf{F}_c^a((k+1)\Delta t - \tau)T} d\tau}_{\mathbf{Q}_d^a}.$$

The expressions obtained for the statistics define an equivalent discrete version of the system

$$\mathbf{z}_{[k+1]}^a = \mathbf{F}^a \mathbf{z}_{[k]}^a + \mathbf{w}_{[k]}^a,$$

with $\mathbf{w}_{[k]}^a \sim \mathcal{N}(0, \mathbf{Q}_d^a)$. Defining the discrete version of the noise covariance \mathbf{Q}_d^a through the integral is difficult and impractical. An alternative method of computing it is by realizing that the stationary solution is also valid for the discrete formulation.

$$\mathbf{P}_\infty^a = \mathbf{F}^a \mathbf{P}_\infty^a \mathbf{F}^{aT} + \mathbf{Q}_d^a$$

Since the stationary covariance can be computed by solving the continuous-time Lyapunov equation (equation 8), the noise covariance can then be obtained as

$$\mathbf{Q}_d^a = \mathbf{F}^a \mathbf{P}_\infty^a \mathbf{F}^{aT} - \mathbf{P}_\infty^a.$$

2.5.2. Observation equation

Let $\mathbf{y}(t) \in \mathbb{R}^{n_y}$ be a subset of the system response,

$$\mathbf{y}(t) = \mathbf{S}_a \ddot{\mathbf{u}}(t) + \mathbf{S}_v \dot{\mathbf{u}}(t) + \mathbf{S}_d \mathbf{u}(t).$$

The matrices $\mathbf{S}_a, \mathbf{S}_v, \mathbf{S}_d \in \mathbb{R}^{n_y \times n_u}$ are the selection matrices for accelerations, velocities, and displacements, respectively. These relate the observed (measured) response $\mathbf{y}(t)$ of the structure to the mechanical model.

Considering the modal formulation, and the equations of motion, the observations $\mathbf{y}(t)$ can be defined in terms of the state and load applied to the system, see for example (Lourens, Papadimitriou, et al., 2012),

$$\mathbf{y}(t) = \underbrace{\begin{bmatrix} [-\mathbf{S}_a \Phi_r \Omega_r^2 + \mathbf{S}_d \Phi_r] & [-\mathbf{S}_a \Phi_r \Gamma_r + \mathbf{S}_v \Phi_r] \end{bmatrix}}_{\mathbf{G}_{cm}} \mathbf{x}_m(t) + \underbrace{\mathbf{S}_a \Phi_r \Phi_r^T \mathbf{S}_p}_{\mathbf{J}_{cm}} p(t),$$

with dimensions: $\mathbf{G}_{cm} \in \mathbb{R}^{n_y \times 2n_m}$; $\mathbf{J}_{cm} \in \mathbb{R}^{n_y \times n_p}$. Using the state space representation of the load, this can be written in terms of the augmented state

$$\mathbf{y}(t) = \underbrace{\begin{bmatrix} \mathbf{G}_{cm} & \mathbf{J}_{cm} \mathbf{H}_c \end{bmatrix}}_{\mathbf{H}_c^a} \mathbf{z}^a(t).$$

Table 1. Kalman filter and smoother equations for the augmented model.

Time update	Measurement update	Smoothing
$\mathcal{P}\left(\mathbf{z}_{[k]}^{a-} \mathbf{y}_{[1:k-1]}\right) = \mathcal{N}\left(\hat{\mathbf{z}}_{[k]}^{a-}, \mathbf{P}_{[k]}^{a-}\right)$	$\mathcal{P}\left(\mathbf{z}_{[k]}^{a+} \mathbf{y}_{[1:k]}\right) = \mathcal{N}\left(\hat{\mathbf{z}}_{[k]}^{a+}, \mathbf{P}_{[k]}^{a+}\right)$ $\mathcal{P}\left(\mathbf{y}_{[1:k]} \mathbf{y}_{[1:k-1]}\right) = \mathcal{N}\left(\mathbf{H}_c^a \hat{\mathbf{z}}_{[k]}^{a-}, \mathbf{S}_{[k]}\right)$	$\mathcal{P}\left(\mathbf{z}_{[k]}^{a+} \mathbf{y}_{[1:N]}\right) = \mathcal{N}\left(\hat{\mathbf{z}}_{[k]}^{a+}, \mathbf{P}_{[k]}^{a+}\right)$
$\hat{\mathbf{z}}_{[k]}^{a-} = \mathbf{F}^a \hat{\mathbf{z}}_{[k-1]}^{a+}$	$\mathbf{r}_{[k]} = \mathbf{y}_{[k]} - \mathbf{H}_c^a \hat{\mathbf{z}}_{[k]}^{a-}$	$\mathbf{C}_{[k]} = \mathbf{P}_{[k]}^{a+} \mathbf{F}^{aT} \left(\mathbf{P}_{[k+1]}^{a-}\right)^{-1}$
$\mathbf{P}_{[k]}^{a-} = \mathbf{F}^a \mathbf{P}_{[k-1]}^{a+} \mathbf{F}^{aT} + \mathbf{Q}_d^a$	$\mathbf{S}_{[k]} = \mathbf{H}_c^a \mathbf{P}_{[k]}^{a-} \mathbf{H}_c^{aT} + \mathbf{R}$	$\hat{\mathbf{z}}_{[k]}^{a+} = \hat{\mathbf{z}}_{[k]}^{a+} + \mathbf{C}_{[k]} \left(\hat{\mathbf{z}}_{[k+1]}^{a+} - \hat{\mathbf{z}}_{[k+1]}^{a-}\right)$
	$\mathbf{K}_{[k]} = \mathbf{P}_{[k]}^{a-} \mathbf{H}_c^a \left(\mathbf{S}_{[k]}\right)^{-1}$	$\mathbf{P}_{[k]}^{a+} = \mathbf{P}_{[k]}^{a+} + \mathbf{C}_{[k]} \left(\mathbf{P}_{[k+1]}^{a+} - \mathbf{P}_{[k+1]}^{a-}\right) \mathbf{C}_{[k]}^T$
	$\hat{\mathbf{z}}_{[k]}^{a+} = \hat{\mathbf{z}}_{[k]}^{a-} + \mathbf{K}_{[k]} \mathbf{r}_{[k]}$	
	$\mathbf{P}_{[k]}^{a+} = \mathbf{P}_{[k]}^{a-} - \mathbf{K}_{[k]} \mathbf{S}_{[k]} \mathbf{K}_{[k]}^T$	

Rewriting to discrete-time, and accounting for measurement noise, the following form of the observation equation is employed:

$$\mathbf{y}_{[k]} = \mathbf{H}_c^a \mathbf{z}_{[k]}^a + \mathbf{r}_{[k]}, \tag{9}$$

with $\mathbf{r}_{[k]}$ a stochastic variable representing the measurement error $\mathbf{r}_{[k]} \sim \mathcal{N}(0, \mathbf{R})$, where the noise matrix $\mathbf{R} \in \mathbb{R}^{n_y \times n_y}$ is assumed a diagonal matrix populated with the noise levels for each sensor $\sigma_{R,s}^2$ where “s” refers to sensor “s”.

2.5.3. Kalman filtering and smoothing

With the discrete version of the augmented model and the observation equation well defined, Kalman filtering and smoothing can be applied to recursively compute the posterior. The relevant expressions are summarized in Table 1.

2.6. Hyperparameter estimation

So far, it has been assumed that the hyperparameters σ, l_{sc} of the Matérn32 covariance function that represents the load are well defined. However, given that the load is unknown, so is the covariance function that defines it stochastically. Furthermore, the noise levels $\sigma_{R,s}^2$ in \mathbf{R} are also unknown.

As presented in Särkkä (2013), the Bayesian manner to obtain these unknown parameters is either through maximum *a posteriori* or maximum likelihood estimation (a method used by Nayek et al., 2019), or by employing Markov Chain Monte Carlo methods (a method used by Rogers et al., 2020). An important disadvantage of these methods lies, however, in their computational cost. In what follows, an alternative, more practical method to compute the optimal values of the hyperparameters is suggested, where empirical qualities of the data are used to directly link the prior model to the measurements.

The suggested procedure for the hyperparameter estimation is divided into two parts: estimating the hyperparameters of the load (σ, l_{sc}), and estimating the noise levels of the sensors $\sigma_{R,s}^2$.

2.6.1. Estimating the load hyperparameters (σ, l_{sc})

In Section 2.4, it was shown that when a prior for the load is defined—through the hyperparameters (σ, l_{sc})—then a prior for the response of the whole system is also defined. This includes the acceleration at the measured locations. The proposed method involves defining (σ, l_{sc}) so as to have a good match between the prior of the accelerations and the measurements.

The prior distribution of the accelerations can be easily computed from the prior distribution of the augmented state through the observation equation, disregarding the noise, so that

$$\mathbf{y}_{(\text{noise-free})}(t) = \mathbf{H}_c^a \mathbf{z}^a(t),$$

and

$$\mathcal{P}(\mathbf{y}_{(\text{noise-free})}(t)) = \mathcal{N}(\hat{\mathbf{y}}_\infty, \mathbf{P}_{y,\infty}),$$

$$\hat{\mathbf{y}}_\infty = \mathbf{H}_c^a \mathbf{z}_\infty^a = \mathbf{0},$$

$$\mathbf{P}_{y,\infty} = \mathbb{E}[\mathbf{y}_{(\text{noise-free})}(t)\mathbf{y}_{(\text{noise-free})}(t)^T] = \mathbf{H}_c^a \mathbf{P}_\infty^a \mathbf{H}_c^{aT},$$

where $\hat{\mathbf{y}}_\infty = \mathbf{0}$ is the expected value and $\mathbf{P}_{y,\infty}$ the associated covariance of the accelerations. In particular, the variance for the acceleration at each measured location “ s ”, $\sigma_{y,s}^2$, will be defined by the diagonal elements of this covariance matrix $\mathbf{P}_{y,\infty}$. The correlation between the observed points is defined through the off-diagonal elements.

It is highlighted that, when calculating $\mathbf{y}_{(\text{noise-free})}(t) \sim \mathcal{N}(\hat{\mathbf{y}}_\infty, \mathbf{P}_{y,\infty})$ through the augmented model, no information about the measurements has been used other than their location within the mechanical model. This is done through \mathbf{H}_c^a , which is constructed using the selection matrices for the measurement locations as well as the dynamic characteristics of the system (cf., Section 2.5.2.).

Since the load is modeled as a Gaussian process, the response obtained from the model is also a Gaussian process as shown above. Therefore, it is possible to compare the variance obtained from the model $\sigma_{y,s}^2$ with the variance estimated from the signal $\sigma_{y,s}^{*2}$. The best hyperparameters for the load (σ, l_{sc}) are such that $\sigma_y^2 \approx \sigma_y^{*2}$. An exact match may, however, not be possible. This may for instance be due to (a) the model not being a perfect representation of the real structure, (b) an inaccurate estimation of the variance from the records, or (c) due to a high amount of noise in the signals.

To match the variances, an individual objective function is defined for each sensor “ s ”: $q_s = q(\sigma_{y,s}^2)$. The individual objective function is defined as a lognormal distribution function. The peak will correspond to the computed variance of the measured records $\sigma_{y,s}^{*2}$, where $q_s^* = q(\sigma_{y,s}^{*2})$, and the spread of the distribution is defined so as to have the 95% interval of confidence between half and twice the measured variance. The full definition and illustration can be found in Figure 1.

When the hyperparameters (σ, l_{sc}) are defined, the variance of the accelerations at each sensor location $\sigma_{y,s}^2$ are defined as a consequence, from which each individual objective function can be evaluated: $q_s = q(\sigma_{y,s}^2)$. The optimum hyperparameters (σ, l_{sc}) will be such that the individual objective function for each sensor is maximized: $q_s = q_s^*$. Instead of evaluating each individual objective function, a combined objective function is defined by their multiplication: $q_{s1, s2, \dots, sn} = q_{s1} \cdot q_{s2} \cdot \dots \cdot q_{sn}$. Note that this combined objective function is the same as defining it as the joint (independent) lognormal distribution.

The best hyperparameters will therefore maximize this objective function $q_{s1, s2, \dots, sn}$. The benefit of using this function is that it has some properties that are convenient for the analysis: it will take high values for variances relatively close to the estimated variance from the signal, and low values for variances that are far from either half or twice the estimated variance. Furthermore, each estimation can be assessed individually through each individual objective function. Alternative ways of defining the objective function or error measure between $\sigma_{y,s}^2$ and $\sigma_{y,s}^{*2}$ may be adopted and or explored in future research.

2.6.2. Estimating the noise levels $\sigma_{R,s}^2$

The noise levels for each signal, $\sigma_{R,s}^2$, contained in \mathbf{R} , also needs to be defined. This will be realized by studying the posterior distributions of the measured accelerations. The acceleration measurements are related to the model through the observation equation which includes noise (equation 9). An estimation for the noise process can therefore be obtained by using the expected value for the augmented state, obtained from the posterior




Figure 1. Definition of the objective function as a log-normal distribution.

$$\hat{\mathbf{r}}_{[k]} = \mathbf{y}_{[k]} - \mathbf{H}_c^a \hat{\mathbf{z}}_{[k]}^a.$$

The measurements $\mathbf{y}_{[k]}$ are deterministic, as the values are directly obtained from the accelerometers. An estimation of the variance of the noise can now be obtained as

$$\mathbf{R}^* = \text{var}(\hat{\mathbf{r}}_{[k]}).$$

Assuming that the mechanical model (equation 1) and load definition (equation 3) are representative of the true structure and loading, and furthermore assuming that the sensor disturbances can be modeled as white Gaussian noise (equation 9), then the noise levels $\sigma_{R,s}^2$, defined for the system in \mathbf{R} should match with the noise levels observed from the posterior $\sigma_{R,s}^{*2}$ in \mathbf{R}^* . Note that in order to estimate the noise levels \mathbf{R}^* a noise level must first be defined in \mathbf{R} .

The methodology employed to match the noise levels is iterative, as illustrated in Figure 2, and described as follows:

1. An arbitrary amount of noise is defined for the signals. A starting point is taken equal to the variance of the signal $\mathbf{R} = \text{var}(\mathbf{y}_{[k]})$.
2. The posterior distribution is computed, from which an estimation for the noise is obtained: \mathbf{R}^* .
3. The noise is defined as the estimated noise $\mathbf{R} = \mathbf{R}^*$.
4. Steps 2 and 3 are repeated until the maximum relative difference between the variances is less than a user-defined tolerance: $\max(\text{diag}\{\frac{|\mathbf{R}^* - \mathbf{R}|}{\mathbf{R}}\}) < \text{tol}$. If the relative difference between the variances starts to increase, the process is stopped and the achieved tolerance is reported.

3. Case Study: Onshore Wind Turbine

In what follows, the GPLFM is used for virtual sensing in an onshore wind turbine tower, with hyperparameter optimization performed as described in the previous section. First, specifics of the measurement campaign are introduced, followed by a presentation of the mechanical model employed, and its validation through system identification. Results are presented in Section 3.3.4 and onward.

3.1. Measurement campaign

Measurement data from one of EnBW’s modern onshore wind turbines were analyzed in this case study. It is a variable-speed, pitch-controlled, and three-bladed horizontal axis wind turbine located in Germany. The wind turbine is equipped with a dedicated structural health monitoring system. The sensor locations and local axes are illustrated in Figure 3. All sensors are mounted on the inner surface of the tower. The




Figure 2. Iterative process employed to match the noise levels $\sigma_{R,s}^2$ defined for the system in \mathbf{R} with the noise levels $\sigma_{R,s}^2$ observed in the system through \mathbf{R}^* .




Figure 3. Sensor locations and local axes.

acceleration sensors are bidirectional, as shown in Figure 3. Four unidirectional strain gauges are mounted at the bottom of the tower; these are combined to obtain bending strains. All sensors measure at a sampling frequency of 20 Hz. The records considered are from December 9, 2020 to January 30, 2021. Both acceleration and strain measurements are divided into 10 min records for analysis. All records are filtered using a fourth-order Butterworth high-pass filter with a cut-off frequency of 0.1 Hz as the quasi-static information cannot be extracted from the acceleration measurements (quasi-static amplitudes are lower than the noise threshold).

Ten-minute statistics from the SCADA system are not used in the methodology itself, but for the interpretation of results obtained in different operational conditions in the following subsections.

Table 2. Selected records.

ID	Design situation	Record	ID	Design situation	Record
OP1	Power production ($v = 4$ m/s)	2020-12-12 07:10:00	SU2	Start-up	2021-01-19 05:30:00
OP2	Power production ($v = 6$ m/s)	2021-01-15 21:40:00	SU3	Start-up	2021-01-14 16:20:00
OP3	Power production ($v = 8$ m/s)	2021-01-20 10:20:00	SD1	Normal shut-down	2021-01-16 08:50:00
OP4	Power production ($v = 10$ m/s)	2020-12-22 06:50:00	SD2	Normal shut-down	2021-01-16 19:00:00
OP5	Power production ($v = 12$ m/s)	2021-01-21 10:40:00	SD3	Normal shut-down	2021-01-17 00:10:00
OP6	Power production ($v = 14$ m/s)	2021-01-24 13:00:00	PK1	Idling	2020-12-27 19:10:00
OP7	Power production ($v = 16$ m/s)	2020-12-29 07:30:00	PK2	Idling	2020-12-20 11:30:00
SU1	Start-up	2021-01-19 10:30:00	PK3	Idling	2021-01-01 18:50:00

Records are selected so as to cover the design situations for fatigue as defined in standard design codes (International Electrotechnical Commission, 2019). The design situations depend upon the cut-in and cut-out wind speed. For the wind turbine under consideration, the cut-in wind speed is $V_{in} = 3$ m/s and the cut-out wind speed $V_{out} = 25$ m/s. Table 2 summarizes the specific records considered in the upcoming analyses.

3.2. Mechanical model and system identification

The mechanical model is based on a two-dimensional finite element model of the turbine, consisting of a set of beam elements for the tower and a point mass for the Rotor Nacelle Assembly. The foundation is assumed to clamp the tower at the bottom. The geometry of the tower varies in height. To capture this geometry, $n = 300$ uniform beam elements are used. Each element differs in average diameter and thickness. Figure 4 summarizes the properties of the mechanical model. The exact same model is considered for both the FA and SS directions, implying that model accuracy will deteriorate for higher-order modes that are not dominated by tower deformation. The loading applied to the tower is modeled as a concentrated point load acting at the top of the tower. Note that the point load is defined stochastically as per equation (3), with its hyperparameters yet to be defined. The damping matrix is defined by assuming a damping ratio of 0.75% for all modes. This damping value is chosen on the basis of the system identification results discussed below.

To obtain an estimate of the accuracy of the mechanical model, system identification is employed. Covariance-driven stochastic subspace identification (SSI-Cov, Peeters and De Roeck, 1999) is applied, in combination with the OPTICS clustering algorithm (OPTICS, Ankerst et al., 1999). The approach employed is the one presented in Boroschek and Bilbao (2019). The system identification is applied to distinguish between the tower Fore–Aft (FA) and tower Side–Side (SS) direction of the wind turbine, which are obtained by transforming the local coordinates using the yaw angle provided by the SCADA data. The identified modal properties and their comparison to the mechanical model are shown in Table 3. The three tower modes are well observed by the data. Given the good correspondence with the mechanical model, changes in the mechanical model to better fit the identification results were deemed unnecessary.




Figure 4. Details of the mechanical model.

Table 3. Frequency (f), damping ratio (ζ), and Modal Assurance Criterion (MAC) values between mechanical model and identification results based on record PK1.

Mode	Identification		Finite Element model				
	f_i^{ID} (Hz)	ζ_i^{ID} (%)	f_i^{FE} (Hz)	Error (%)	MAC(%)	ζ_i^{FE}	Error (%)
Tower FA 1	0.17	0.36	0.17	0.00	99.97	0.75	108.33
Tower FA 2	1.23	1.04	1.21	-1.63	99.74	0.75	-27.88
Tower FA 3	3.21	0.88	3.50	9.03	99.59	0.75	-14.77
Tower SS 1	0.17	1.19	0.17	0.00	99.92	0.75	-36.97
Tower SS 2	1.17	0.60	1.21	3.42	99.91	0.75	25.00
Tower SS 3	3.65	0.61	3.50	-4.11	97.82	0.75	22.95

3.3. Strain and fatigue estimation on the onshore turbine tower

The GPLFM is now applied on the mechanical model, informed by each set of acceleration records. The estimated states and forces are used to obtain the strain at the location where the strain gauges are located, to compare the estimations with the measurements.

In what follows, the selected records are first illustrated and described, and subsequently, the hyperparameter estimation procedure is shown for both the prior distribution (to obtain the hyperparameters σ and l_{sc}) and the posterior distribution (to obtain the noise levels $\sigma_{R,s}^2$ in \mathbf{R}).

3.3.1. Description of selected records

The power production records (OP1 to OP7) are characterized by relatively stationary behavior of the response, illustrated in Figure 5a,b. As a general trend, operational records associated with lower wind speeds will have lower amplitudes of vibrations, and vice versa.

Idling records (PK1–PK3) are also stationary but are characterized by very low amplitudes of vibrations given the low wind speeds that the structure is subject to. These records were observed to have quantization issues, illustrated in Figure 5c.

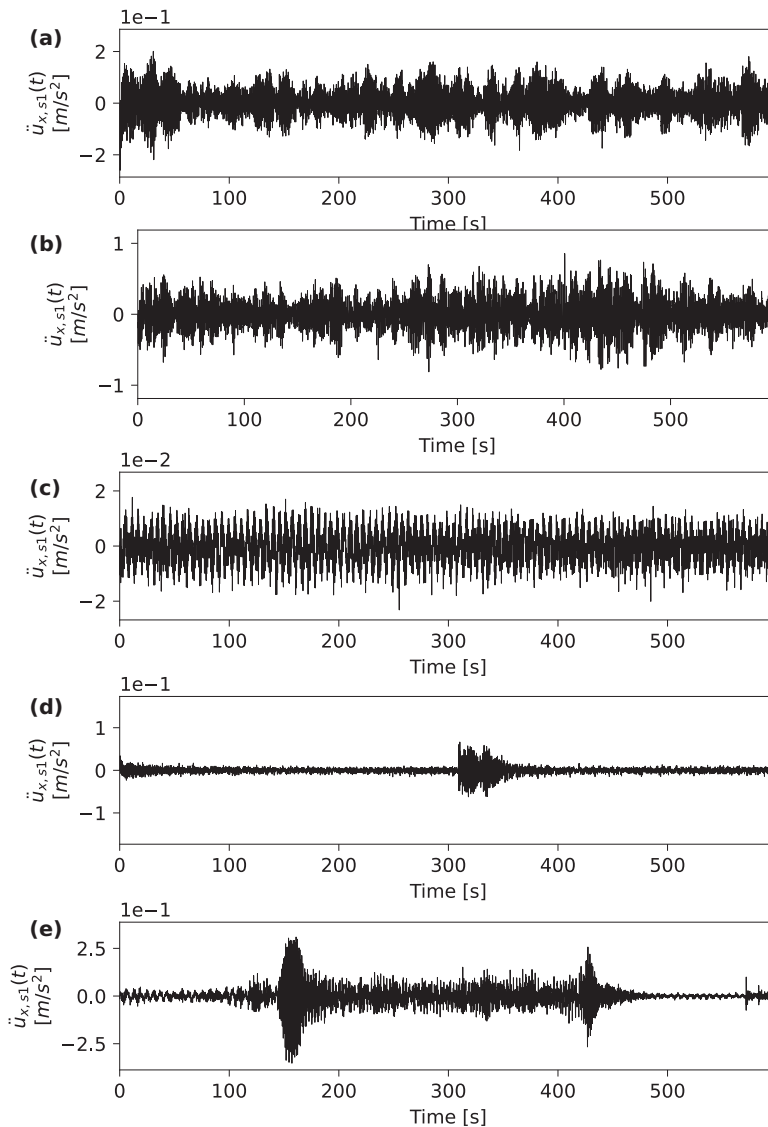


Figure 5. Acceleration records, sensor *s1*. (a) *OP1-x* (mean wind speed 4 m/s). (b) *OP7-x* (mean wind speed 16 m/s). (c) *PK2-x* (quantization issues due to lack of resolution in the time signal). (d) *SU-3* start-up condition. (e) *SD2-x* (shut-down condition).

Start-up and shut-down records (*SU1–SU3* and *SD1–SD3*, respectively) are characterized by transient behavior due to the operational change—compare [Figure 5d](#) and [5e](#).

3.3.2. Estimation of the load hyperparameters (σ , l_{sc})

The hyperparameters σ and l_{sc} are defined for each analyzed record. This is done through the objective function described in [Section 2.6.1](#). The objective function for record *PK1-x* is shown in [Figure 6](#), where a clear maximum is observed, which corresponds to the values used for the hyperparameters. [Figure 7](#) compares the prior defined through the chosen hyperparameters (σ , l_{sc}) with the distribution observed from the acceleration measurements. Because of the good fit observed, the assumption for the load location is validated.




Figure 6. Record PK1-x. σ and l_{sc} estimation. Objective function plot: $q_{s1,s2,s3}$. Optimum found for maximum value: $l_{sc} = 0.318\text{s}$, $\sigma = 2556\text{N}$.




Figure 7. Record PK1-x. Prior distribution visual check using optimum hyperparameters: $l_{sc} = 0.318\text{s}$, $\sigma = 2556\text{N}$. (a) Sensor s1. (b) Sensor s2. (c) Sensor s3.

The comparison of the prior with the distribution observed from the acceleration measurements for record OP6-x is shown in Figure 8. Note that in this case, the fit is not as good as in the idling case. This is justified given the change in the operational condition of the turbine, which implies that the mechanical model—compared only in idling conditions—might not be as representative.

For transient conditions (shut-down and start-up records) the acceleration signal is not stationary. Therefore, the prior distribution will over- and/or underestimate the acceleration response, as illustrated in Figure 9 for record SU3-x.




Figure 8. Record OP6-x. Prior distribution visual check using optimum hyperparameters: $l_{sc} = 0.05s$, $\sigma = 20845N$. (a) Sensor s1. (b) Sensor s2. (c) Sensor s3.

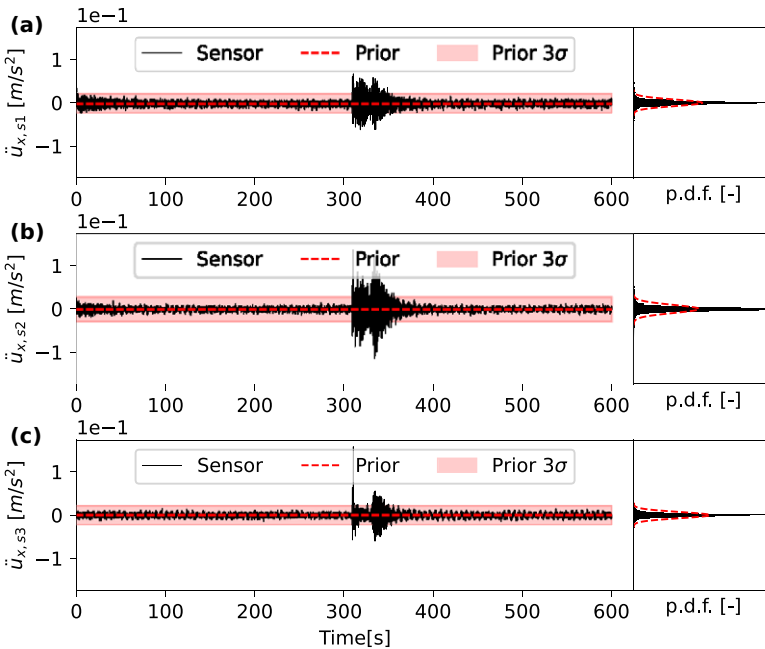


Figure 9. Record SU3-x. Prior distribution visual check using optimum hyperparameters: $l_{sc} = 0.034s$, $\sigma = 866N$. (a) Sensor s1. (b) Sensor s2. (c) Sensor s3.

Table 4. Estimation of load hyperparameters (σ, l_{sc}).

ID	$\sigma(N)$	$l_{sc}(s)$	$\frac{q_{s1}}{q_{s1}^*}(\%)$	$\frac{q_{s2}}{q_{s2}^*}(\%)$	$\frac{q_{s3}}{q_{s3}^*}(\%)$	$\frac{q_{s1,s2,s3}}{q_{s1,s2,s3}^*}(\%)$
OP1-x	3,961	0.046	20.07	94.00	10.04	1.89
OP2-x	2,049	0.045	35.42	94.43	20.53	6.87
OP3-x	10,319	0.048	49.09	99.10	41.41	20.14
OP4-x	13,179	0.048	22.72	97.78	15.42	3.42
OP5-x	14,404	0.045	28.19	96.59	16.27	4.43
OP6-x	20,845	0.050	13.24	85.08	6.18	0.70
OP7-x	15,928	0.050	25.76	93.60	17.20	4.15
SU1-x	5,952	0.052	37.59	99.99	36.87	13.86
SU2-x	7,008	0.048	10.53	95.19	5.38	0.54
SU3-x	866	0.034	98.10	88.18	78.40	67.82
SD1-x	2,092	0.268	98.78	92.77	97.37	89.23
SD2-x	3,849	0.053	51.72	99.99	52.73	27.27
SD3-x	1,742	0.056	95.84	99.96	96.57	92.52
PK1-x	2,556	0.318	99.53	96.03	98.25	93.91
PK2-x	462	0.300	100.00	99.99	100.00	99.99
PK3-x	368	0.050	86.25	99.93	84.48	72.82

Note. Prior distribution match measured through objective function q . Local direction x .

Table 5. Estimation of load hyperparameters (σ, l_{sc}).

ID	$\sigma(N)$	$l_{sc}(s)$	$\frac{q_{s1}}{q_{s1}^*}(\%)$	$\frac{q_{s2}}{q_{s2}^*}(\%)$	$\frac{q_{s3}}{q_{s3}^*}(\%)$	$\frac{q_{s1,s2,s3}}{q_{s1,s2,s3}^*}(\%)$
OP1-y	4,993	0.046	37.99	95.88	24.34	8.87
OP2-y	2,592	0.045	22.75	91.36	10.00	2.08
OP3-y	10,859	0.046	3.90	78.84	0.62	0.02
OP4-y	9,557	0.050	22.11	99.15	18.01	3.95
OP5-y	11,911	0.051	17.08	96.63	12.16	2.01
OP6-y	11,983	0.061	99.84	100.00	99.89	99.73
OP7-y	14,404	0.050	7.01	95.28	3.95	0.26
SU1-y	5,532	0.052	15.58	99.99	15.92	2.48
SU2-y	6,004	0.051	15.35	99.92	14.18	2.18
SU3-y	498	0.194	94.35	83.41	96.64	76.06
SD1-y	2,848	0.021	98.49	98.38	93.88	90.96
SD2-y	3,538	0.056	67.32	99.69	72.00	48.32
SD3-y	1,530	0.069	92.80	99.43	96.15	88.72
PK1-y	1,895	0.285	99.72	98.10	99.26	97.11
PK2-y	354	0.220	97.98	92.05	97.92	88.31
PK3-y	296	0.077	78.63	96.48	91.34	69.29

Note. Prior distribution match measured through objective function q . Local direction y .

Tables 4 and 5 summarizes the objective function values obtained for each record in local x and y directions, respectively. They include, besides the values defined for σ and l_{sc} , the likelihood measured as a fraction of the maximum likelihood defined for each lognormal distribution and for the joint distribution. Therefore, a value close to 100% indicates a very good fit for the prior distribution.

It is observed that for parking conditions (PK1–PK3), the relative likelihood is high, indicating that the mechanical model provides a good representation of the structure in idling conditions. For all power

production conditions (OP1–OP7) the relative likelihood is lower, which indicates that the issue previously discussed for record OP3-x is valid for all other power production records: the mechanical model considered is less representative under operating conditions (possibly due to aerodynamic damping, increased relevance of higher modes, etc.). Finally, for start-up and shut-down records (SU1–SU3 and SD1–SD3, respectively) the relative likelihood varies. For these cases, and as previously mentioned for record SU3-x, the response is not stationary. In the future, one could investigate redefining the load assumption for these particular cases.

3.3.3. Estimating the noise levels $\sigma_{R,s}^2$

The assumed noise levels for each signal are defined iteratively, as described in Section 2.6.2, with a relative tolerance of 1%. This iterative process was observed to diverge in approximately 30% of the records. The process is illustrated using records PK1-x, OP3-x, and SU3-x.

Figure 10a shows the noise levels obtained in each iteration of record PK1-x. Note that the starting value is the variance of the signal, and it quickly decays into a constant value that corresponds to the estimated variance of the noise. This is further illustrated by the relative error obtained in each iteration as shown in Figure 10b. Figure 11 shows a visual check of the resulting noise variance compared with the estimated noise \hat{r} .

Figure 13a shows the noise levels defined for each iteration of record OP3-x. Similar to record PK1-x, the noise levels are observed to decrease to a stable value. Figure 13b shows the relative error found for each iteration. Note that in this case, the relative error for accelerometer s3 increases in the third iteration. The iteration is not stopped, however, as the maximum error at each iteration is still decreasing. It should be noted that in some cases the part of the signal considered as noise may in fact contain a part of the response, complicating the search for an optimal value.

Figure 14a shows the noise levels per iteration of record SU3-x. Note that this record was diverging given the increase in the relative error, as illustrated in Figure 14b. Given the transient nature of the response, the predicted error is not Gaussian, and as discussed in Section 3.3.2, the model employed is not representative for the particular record. The resulting high error variance and lack of convergence attest to inaccurate modeling of the measurements in this case.

A summary is presented in Tables 6 and 7. The noise level obtained for each signal is presented in terms of the noise-to-signal ratio (NSR), defined as follows:

$$NSR(s) = \frac{\sigma_{R,s}^2}{\sigma_{y,s}^2 - \sigma_{R,s}^2},$$




Figure 10. Record PK1-x. Noise variance estimation $\sigma_{R,s}^2$. (a) Noise variance estimation per iteration. (b) Error measure: $\text{diag}\left(\frac{\hat{R}-R}{R}\right)$.




Figure 11. Record PK1-x. Noise variance visual check using optimized noise variances.




Figure 12. Record OP3-x. Noise variance estimation $\sigma_{R,s}^2$. (a) Noise variance estimation per iteration. (b) Error measure: $\text{diag}\left(\frac{\hat{R}-R}{R}\right)$.

with $\sigma_{y,s}^2$ referring to the measured acceleration variance, and the difference $(\sigma_{y,s}^2 - \sigma_{R,s}^2)$ corresponding to the noiseless signal variance. The *NSR* thus indicates the estimated relative amount of noise with respect to the noiseless response. If the *NSR* is larger than 100%, then the variance of the noise is higher than the variance of the noiseless signal, and vice versa. For idling conditions (PK1 to PK3), the process never failed to converge. This can be understood given a representative mechanical model for the structure. The found noise levels are quite high for these records, however, which can be justified by the considerable amount of quantization issues. For power production records (OP1–OP7), it is observed that for most cases the process converges. It is debatable, however, if the defined noise levels are representative of the true measurement

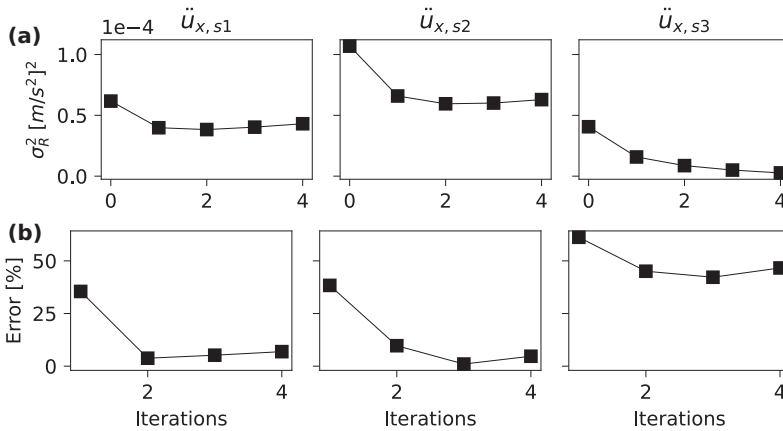


Figure 13. Record SU3-x. Noise variance estimation $\sigma_{R,s}^2$. (a) Noise variance estimation per iteration. (b) Error measure: $\text{diag}\left(\frac{\hat{R}-R}{R}\right)$.

Table 6. Noise variance estimation.

ID	NSR(s1)(%)	NSR(s2)(%)	NSR(s3)(%)	tol(%)
OP1-x	11.51	6.66	74.09	0.67
OP2-x	36.14	22.71	35.01	1.00
OP3-x	12.81	12.48	36.80	0.88
OP4-x	18.35	22.95	22.62	0.86
OP5-x	15.78	20.35	30.94	5.48
OP6-x	39.83	47.56	11.07	19.30
OP7-x	9.39	12.69	57.27	0.90
SU1-x	7.65	9.55	44.51	0.79
SU2-x	8.42	12.47	70.04	0.93
SU3-x	231.26	143.10	7.03	46.64
SD1-x	21.60	21.63	4.67	0.91
SD2-x	9.63	9.32	25.27	0.97
SD3-x	37.17	34.36	3.52	30.01
PK1-x	5.09	8.46	0.94	1.00
PK2-x	92.31	36.46	10.22	0.72
PK3-x	296.96	208.58	57.43	0.92

Note. Noise to signal ratios. Local direction x.

error as the mechanical model is less accurate. The nonconvergence of some of these records can be justified by the inaccuracy of the mechanical model employed, as previously discussed. Finally, for start-up and shut-down records (SU1–SU3 and SD1–SD3, respectively) the response is transient, and therefore the estimation of the error is questionable, as the model employed was shown to be less representative. Nevertheless, as will be shown in the following sections, the procedure provides a computationally inexpensive way to obtain hyperparameter values that allow for accurate estimation results.

3.3.4. Response estimation

Having analyzed and defined the hyperparameters for each record, the estimated response in terms of the dynamic strains is now compared with the measured dynamics strains.

Table 7. Noise variance estimation.

ID	NSR(s1) (%)	NSR(s2) (%)	NSR(s3) (%)	tol (%)
OP1-y	25.31	23.86	21.10	0.97
OP2-y	30.78	41.80	6.18	0.93
OP3-y	5.84	6.12	247.18	0.66
OP4-y	11.74	14.99	19.03	3.36
OP5-y	12.98	17.49	21.58	3.96
OP6-y	7.76	6.10	12.81	0.73
OP7-y	14.86	22.03	11.38	6.57
SU1-y	6.98	9.82	33.41	0.93
SU2-y	8.72	12.89	12.43	0.24
SU3-y	72.68	71.98	31.06	0.80
SD1-y	28.50	16.66	0.23	87.99
SD2-y	7.80	5.98	15.41	0.76
SD3-y	35.65	41.50	1.00	2.21
PK1-y	11.11	11.67	0.77	0.77
PK2-y	171.95	79.95	23.98	0.46
PK3-y	394.12	280.21	96.73	0.88

Note. Noise-to-signal ratios. Local direction y.




Figure 14 Record PK1-x. Dynamic strain estimation.

Figure 16a shows the dynamic strain estimation for record PK1-x, along with the measured and filtered strain for comparison. The mean refers to the expected value obtained for each point in time through the posterior distribution for the strain, and 3σ is the associated 99.7% interval of confidence. Figure 16b shows the associated frequency contents of the measurement and estimation. The unfiltered measurements are also included in the spectra, to show the influence band of the filtering. The dynamic strain estimation is observed to match the measured strain, with the mean prediction closely following the measured strain. The differences in the frequency domain are mainly observed in the low amplitude and high-frequency regions; the estimation for the dominant frequencies is observed to be accurate.




Figure 15. Record OP3-x. Dynamic strain estimation.




Figure 16. Record SU3-x. Dynamic strain estimation.

Figure 11a,b presents the estimation of record OP3-x in the time and frequency domain, respectively. The estimation is again observed to be a close match, though to a lesser degree when compared with the estimation obtained for record PK1-x. This is justified by the analysis exposed in Sections 3.3.2 and 3.3.3: the model employed is less accurate for this operational condition.

Figure 12a,b shows the estimation of record SU3-x in the time and frequency domain, respectively. As expected, larger errors are observed given the transient nature of the response, but the estimations are still reasonably accurate outside these transient parts. It is noted that even though the transient nature is not part of the load assumptions, it is still well estimated in the load, as shown in Figure 17.

To quantify the estimation errors, two error metrics are considered: The Time Response Assurance Criterion (TRAC) and the Mean Absolute Error (MAE). These are defined as follows:

$$TRAC(\boldsymbol{\varepsilon}, \hat{\boldsymbol{\varepsilon}}) = \frac{(\hat{\boldsymbol{\varepsilon}}\boldsymbol{\varepsilon}^T)^2}{(\hat{\boldsymbol{\varepsilon}}\hat{\boldsymbol{\varepsilon}}^T)(\boldsymbol{\varepsilon}\boldsymbol{\varepsilon}^T)},$$




Figure 17. Record SU3-x. Load estimation.

Table 8. Error metrics.

ID	x		y	
	TRAC($\varepsilon, \hat{\varepsilon}$)	MAE($\varepsilon, \hat{\varepsilon}$) $\times 10^{-6}$ (m/m)	TRAC($\varepsilon, \hat{\varepsilon}$)	MAE($\varepsilon, \hat{\varepsilon}$) $\times 10^{-6}$ (m/m)
OP1	91.22	0.45	83.88	1.05
OP2	81.28	0.55	73.60	0.93
OP3	95.69	1.34	93.34	1.33
OP4	94.38	2.14	94.55	1.60
OP5	95.80	2.40	94.58	2.09
OP6	86.08	4.72	99.69	2.87
OP7	96.12	2.14	90.93	2.52
SU1	99.14	0.87	98.72	0.79
SU2	96.01	0.78	96.83	0.84
SU3	67.44	0.19	93.67	0.26
SD1	99.67	1.10	99.18	0.61
SD2	94.91	0.66	97.77	0.78
SD3	98.97	0.31	97.45	0.48
PK1	99.80	1.27	99.60	1.11
PK2	99.11	0.27	95.51	0.23
PK3	85.42	0.09	76.02	0.13

$$MAE(\varepsilon, \hat{\varepsilon}) = \frac{1}{N} \sum_k |\hat{\varepsilon}_{[k]} - \varepsilon_{[k]}|,$$

with $\varepsilon \in \mathbb{R}^N$ the measured and filtered strain in vector form representing the values in time, and $\hat{\varepsilon} \in \mathbb{R}^N$ the estimation. The TRAC represents a measure of correlation between the estimation and the measurement. In simple words, a value close to 100% indicates a very strong similarity of the shapes, and a value close to 0% indicates the estimations and measurements are very different. The MAE indicates the average error of the estimation, in units of the measurement. The results for each analyzed record are summarized in Table 8.

3.3.5. Fatigue load estimation

The underlying aim of this work is to be able to predict the damage that the tower has suffered in terms of fatigue. In this section, it is shown how the accuracy obtained for the strain estimates translates to accuracy in fatigue damage predictions. Damage equivalent loads are used to compare design fatigue load levels with fatigue loads the asset experienced on site. The damage equivalent load, in terms of stresses, is defined by, see for example (Cosack, 2010):

$$\Delta\sigma_{eq} = \sqrt[m]{\frac{\sum_i \Delta\sigma_i^m N_i}{N_{ref}}},$$

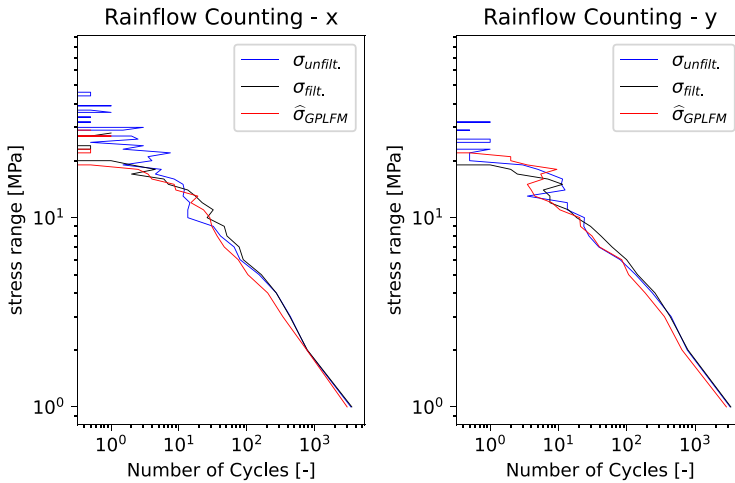


Figure 18. Stress ranges and number of cycles. Rainflow-counting method.

Table 9. Damage equivalent load.

Direction	Measured unfiltered	Measured filtered	GPLFM	
	$\Delta\sigma_{eq}$ (MPa)	$\Delta\sigma_{eq}$ (MPa)	$\Delta\hat{\sigma}_{eq}$ (MPa)	$e_{\Delta\hat{\sigma}_{eq}}$ (%)
<i>x</i>	1.419	0.860	0.820	-4.71
<i>y</i>	1.036	0.733	0.801	+9.36
$\sqrt{x^2 + y^2}$	1.757	1.130	1.146	+1.45

Note. All records summary.

with $\Delta\sigma_i$ stress ranges and N_i the corresponding number of cycles. m is a material constant referring to the idealized straight slope observed on a double-logarithmic scaled SN-curve. $\Delta\sigma_{eq}$ is the damage equivalent load, and N_{ref} the number of cycles of the equivalent load.

The stress ranges and the corresponding number of cycles are determined using the rainflow-counting method. The rainflow-counting method is applied to each record, and subsequently combined to generate the stress ranges and the corresponding number of cycles for all the records analyzed. This was done by counting the number of cycles for each record, and then adding them together. The result is shown in Figure 18. This figure also includes the rainflow-counting results for the unfiltered strain signals showcasing the general lack of the dynamic strains to capture the quasi-static high-amplitude stress ranges. The damage equivalent load and relative errors are presented in Table 9. A more detailed analysis of the damage equivalent load derived for each record is provided in Table 10. Note that stresses are computed by multiplying the strains with the assumed modulus of elasticity $E = 200,000\text{MPa}$, the material constant is set as $m = 4$ and the reference number of cycles is set as $N_{ref} = 10^7$. The Wöhler slope $m = 4$ is commonly used for welded steel as simplification of its SN-curve with two slopes of respectively 3 (applies to high stress / low number of cycles) and 5 (applies to low stress/high number of cycle).

From Table 10 and 11, it is clear that the estimation error in terms of damage equivalent loads varies greatly across the different records, from less than 1% to almost 30%. Averaging over the records, errors of less than 10% are found—compare Table 9.

Table 10. *Damage equivalent load.*

ID	Measured unfiltered	Measured filtered	GPLFM	
	$\Delta\sigma_{eq,x}^{unfilt.}$ (MPa)	$\Delta\sigma_{eq,x}$ (MPa)	$\Delta\hat{\sigma}_{eq,x}$ (MPa)	$e_{\Delta\hat{\sigma}_{eq,x}}$ (%)
OP1-x	0.096	0.090	0.084	-6.35
OP2-x	0.080	0.064	0.061	-4.87
OP3-x	0.400	0.326	0.328	+0.52
OP4-x	0.738	0.478	0.460	-3.74
OP5-x	0.765	0.548	0.551	+0.63
OP6-x	1.129	0.739	0.666	-9.86
OP7-x	0.650	0.530	0.532	+0.50
SU1-x	0.440	0.208	0.230	+10.81
SU2-x	1.104	0.253	0.247	-2.53
SU3-x	0.056	0.031	0.019	-40.22
SD1-x	0.431	0.231	0.266	+14.86
SD2-x	0.365	0.168	0.160	-4.73
SD3-x	0.089	0.086	0.097	+12.22
PK1-x	0.315	0.299	0.343	+14.77
PK2-x	0.047	0.046	0.054	+16.97
PK3-x	0.008	0.008	0.010	+14.81

Note. All records.**Table 11.** *Damage equivalent load.*

ID	Measured unfiltered	Measured filtered	GPLFM	
	$\Delta\sigma_{eq,y}^{unfilt.}$ (MPa)	$\Delta\sigma_{eq,y}$ (MPa)	$\Delta\hat{\sigma}_{eq,y}$ (MPa)	$e_{\Delta\hat{\sigma}_{eq,y}}$ (%)
OP1-y	0.195	0.154	0.109	-29.26
OP2-y	0.116	0.088	0.083	-6.02
OP3-y	0.393	0.323	0.290	-9.92
OP4-y	0.499	0.344	0.349	+1.35
OP5-y	0.625	0.436	0.438	+0.51
OP6-y	0.664	0.644	0.746	+15.90
OP7-y	0.662	0.467	0.425	-9.09
SU1-y	0.475	0.178	0.195	+9.65
SU2-y	0.824	0.218	0.231	+5.98
SU3-y	0.051	0.037	0.042	+16.22
SD1-y	0.235	0.134	0.151	+12.08
SD2-y	0.309	0.189	0.206	+9.00
SD3-y	0.075	0.073	0.086	+17.89
PK1-y	0.208	0.203	0.239	+17.66
PK2-y	0.029	0.028	0.034	+20.45
PK3-y	0.010	0.010	0.012	+18.64

Note. All records. y-direction.

4. Conclusions

The GPLFM was used to estimate the dynamic strains and subsequently predict damage equivalent loads in an onshore wind turbine tower, based only on a set of acceleration measurements. Comparisons were made with measured and filtered strains, and the damage equivalent loads were computed from them. To find the hyperparameters for the GPLFM, a new, intuitive procedure was suggested. It is highlighted that quasi-static loads and strains are not estimated, since quasi-static information cannot be extracted from the acceleration measurements. Since these loads should, however, be considered when assessing the life of a turbine, it is suggested to include quasi-static information through inclinometer measurements in the future.

The results show that the GPLFM in combination with the suggested hyperparameter optimization procedure provides a robust Bayesian framework for response estimation. The dynamic strains were estimated with, on average, high accuracy, with slight differences in accuracy depending on the operating condition. The highest accuracy is achieved for the parked plant due to a good agreement between model assumptions and the operating condition. The shut-down and start-up operating states, on the other hand, are captured with the lowest precision owing to the transient nature of these operating states. The average TRAC value per operational condition was found to be 90.8, 92.0, 98.0, and 92.3 for the power production, start-up, shut-down, and idling conditions, respectively.

The dynamic strains estimated through GPLFM were further utilized to assess the damage suffered by the tower in terms of damage equivalent loads. Contrary to the results of the dynamic strain analysis, the closest match between measured and GPLFM-based damage equivalent loads is observed for the turbine during standard operation.

Factors that were shown to influence the accuracy of the estimation include the accuracy of the mechanical model employed, and the validity of the assumptions made when formulating the stochastic model representing the load. Future research should aim to improve these factors.

List of Abbreviations

GPLFM	Gaussian process latent force model
MAC	Modal Assurance Criterion
MAE	Mean Absolute Error
MD&E	modal decomposition and expansion
NSR	noise-to-signal ratio
OP	operational conditions
PSD	power spectral density
PK	idling conditions
SCADA	Supervisory Control and Data Acquisition
SD	shut-down conditions
SU	start-up conditions
TRAC	Time Response Assurance Criterion

Acknowledgments. We acknowledge the provision of measurement data for this project by EnBW.

Author Contributions. Conceptualization: all authors; Data curation: A.S., L.Z.; Formal analysis: J.B., E.L.; Investigation: all authors; Methodology: J.B., E.L.; Project administration: J.B., A.S., L.Z.; Resources: A.S., L.Z.; Software: J.B.; Supervision: E.L., L.Z.; Visualization: J.B.; Writing—original draft: J.B., E.L.; Writing—review and editing: all authors.

Competing Interests. The authors declare no competing interests exist.

Data Availability Statement. The data that support the findings of this study are available from EnBW. Restrictions apply to the availability of these data, which were used under license for this study. Data are available from the authors with the permission of EnBW.

Ethics Statement. The research meets all ethical guidelines, including adherence to the legal requirements of the study country.

Funding Statement. This work received no specific grant from any funding agency, commercial or not-for-profit sectors.

References

- Ankerst M, Breunig MM, Kriegel H-P and Sander J (1999) OPTICS: Ordering points to identify the clustering structure. *Proceedings of the 1999 ACM SIGMOD International Conference on Management of Data—SIGMOD '99*, 49–60. <https://doi.org/10.1145/304182.304187>
- Augustyn D (2021) *Towards Offshore Wind Digital Twins—Application to Jacket Substructures* [Ph.D. Thesis]. Aalborg University.
- Augustyn D, Smolka U, Tygesen UT, Ulriksen MD and Sørensen JD (2020) Data-driven model updating of an offshore wind jacket substructure. *Applied Ocean Research* 104, 102366. <https://doi.org/10.1016/j.apor.2020.102366>
- Avendaño-Valencia LD, Abdallah I and Chatzi E (2021) Virtual fatigue diagnostics of wake-affected wind turbine via Gaussian process regression. *Renewable Energy* 170, 539–561. <https://doi.org/10.1016/j.renene.2021.02.003>
- Boroschek RL and Bilbao JA (2019) Interpretation of stabilization diagrams using density-based clustering algorithm. *Engineering Structures* 178, 245–257. <https://doi.org/10.1016/j.engstruct.2018.09.091>
- Branlard E, Giardina D and Brown CSD (2020) Augmented Kalman filter with a reduced mechanical model to estimate tower loads on a land-based wind turbine: A step towards digital-twin simulations. *Wind Energy Science* 5(3), 1155–1167. <https://doi.org/10.5194/wes-5-1155-2020>
- Cosack N (2010) *Fatigue Load Monitoring with Standard Wind Turbine Signals* [Ph.D. Thesis]. Universität Stuttgart.
- d N Santos F, Noppe N, Weijtjens W and Devriendt C (2021) Data-driven farm-wide fatigue estimation on jacket foundation OWTs for multiple SHM setups. *Wind Energy Science* 7, 299–321. <https://doi.org/10.5194/wes-7-299-2022>
- Eftekhar Azam S, Chatzi E and Papadimitriou C (2015) A dual Kalman filter approach for state estimation via output-only acceleration measurements. *Mechanical Systems and Signal Processing* 60–61, 866–886. <https://doi.org/10.1016/j.ymssp.2015>.
- Gillijns S and De Moor B (2007) Unbiased minimum-variance input and state estimation for linear discrete-time systems with direct feedthrough. *Automatica* 43(5), 934–937. <https://doi.org/10.1016/j.automatica.2006.11.016>
- Gulgec NS, Takáč M and Pakzad SN (2020) Structural sensing with deep learning: Strain estimation from acceleration data for fatigue assessment. *Computer-Aided Civil and Infrastructure Engineering*, 35(12), 1349–1364. <https://doi.org/10.1111/mice.12565>
- Hartikainen J and Särkkä S (2010) Kalman filtering and smoothing solutions to temporal Gaussian process regression models. *2010 IEEE International Workshop on Machine Learning for Signal Processing*, 379–384. <https://doi.org/10.1109/MLSP.2010.5589113>
- Iliopoulos A, Devriendt C, Guillaume P and Van Hemelrijck D (2014) Continuous fatigue assessment of an offshore wind turbine using a limited number of vibration sensors. *7th European Workshop on Structural Health Monitoring*, 8.
- Iliopoulos A, Shirzadeh R, Weijtjens W, Guillaume P, Hemelrijck DV and Devriendt C (2016) A modal decomposition and expansion approach for prediction of dynamic responses on a monopile offshore wind turbine using a limited number of vibration sensors. *Mechanical Systems and Signal Processing* 68–69, 84–104. <https://doi.org/10.1016/j.ymssp.2015.07.016>
- Iliopoulos A, Weijtjens W, Van Hemelrijck D and Devriendt C (2017) Fatigue assessment of offshore wind turbines on monopile foundations using multi-band modal expansion: Fatigue assessment of monopile OWTs using multi-band modal expansion. *Wind Energy* 20(8), 1463–1479. <https://doi.org/10.1002/we.2104>
- International Electrotechnical Commission (2019) *Wind Energy Generation Systems. Design Requirements Part 1*. International Electrotechnical Commission: Geneva, Switzerland.
- Kalman RE (1960) A new approach to linear filtering and prediction problems. *Journal of Basic Engineering* 12, 35–45.
- Lourens E, Papadimitriou C, Gillijns S, Reynders E, De Roeck G and Lombaert G (2012) Joint input-response estimation for structural systems based on reduced-order models and vibration data from a limited number of sensors. *Mechanical Systems and Signal Processing* 29, 310–327. <https://doi.org/10.1016/j.ymssp.2012.01.011>
- Lourens E, Reynders E, De Roeck G, Degrande G and Lombaert G (2012) An augmented Kalman filter for force identification in structural dynamics. *Mechanical Systems and Signal Processing* 27, 446–460. <https://doi.org/10.1016/j.ymssp.2011.09.025>
- Maes K, Iliopoulos A, Weijtjens W, Devriendt C and Lombaert G (2016) Dynamic strain estimation for fatigue assessment of an offshore monopile wind turbine using filtering and modal expansion algorithms. *Mechanical Systems and Signal Processing* 76–77, 592–611. <https://doi.org/10.1016/j.ymssp.2016.01.004>
- Naets F, Cuadrado J and Desmet W (2015) Stable force identification in structural dynamics using Kalman filtering and dummy-measurements. *Mechanical Systems and Signal Processing* 50–51, 235–248. <https://doi.org/10.1016/j.ymssp.2014.05.042>
- Nayek R, Chakraborty S and Narasimhan S (2019) A Gaussian process latent force model for joint input-state estimation in linear structural systems. *Mechanical Systems and Signal Processing* 128, 497–530. <https://doi.org/10.1016/j.ymssp.2019.03.048>
- Noppe N, Iliopoulos A, Weijtjens W and Devriendt C (2016) Full load estimation of an offshore wind turbine based on SCADA and accelerometer data. *Journal of Physics: Conference Series* 753, 072025. <https://doi.org/10.1088/1742-6596/753/7/072025>
- Peeters B and De Roeck G (1999) Reference-based stochastic subspace identification for output-only modal analysis. *Mechanical Systems and Signal Processing* 13(6), 855–878. <https://doi.org/10.1006/mssp.1999.1249>
- Rasmussen CE and Williams CKI (2006) *Gaussian Processes for Machine Learning*. Cambridge, MA: MIT Press.

- Rogers TJ, Worden K and Cross EJ** (2020) On the application of Gaussian process latent force models for joint input-state-parameter estimation: With a view to Bayesian operational identification. *Mechanical Systems and Signal Processing* 140, 106580. <https://doi.org/10.1016/j.ymssp.2019.106580>
- Särkkä S** (2013) *Bayesian Filtering and Smoothing*. Cambridge, UK: Cambridge University Press.
- Särkkä S and Solin A** (2019) *Applied Stochastic Differential Equations*. Cambridge, UK: Cambridge University Press.
- Schröder L** (2020) *Towards Digital Twin Technology: Wind Farm Operation Analysis and Optimization Using Model-Supported Data Analytics*. <https://doi.org/10.11581/DTU:00000100>
- Schröder L, Krasimirov Dimitrov N, Verelst D R and Sørensen JA** (2018) Wind turbine site-specific load estimation using artificial neural networks calibrated by means of high-fidelity load simulations. *Journal of Physics: Conference Series* 1037, 062027. <https://doi.org/10.1088/1742-6596/1037/6/062027>
- Smolka U and Cheng PW** (2013) On the design of measurement campaigns for fatigue life monitoring of offshore wind turbines. *The Twenty-Third International Offshore and Polar Engineering Conference*, Anchorage, Alaska.
- Tatsis K and Lourens E** (2016) A comparison of two Kalman-type filters for robust extrapolation of offshore wind turbine support structure response. In Bakker J, Frangopol DM and van Breugel K (eds), *Life-Cycle of Engineering Systems*, 1st Edn. Boca Raton, Florida, US: CRC Press, pp. 209–216. <https://doi.org/10.1201/9781315375175-20>
- Vera-Tudela L and Kühn M** (2017) Analysing wind turbine fatigue load prediction: The impact of wind farm flow conditions. *Renewable Energy* 107, 352–360. <https://doi.org/10.1016/j.renene.2017.01.065>
- Wagg DJ, Worden K, Barthorpe RJ and Gardner P** (2020) Digital twins: State-of-the-art and future directions for modeling and simulation in engineering dynamics applications. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems* 6(3), 030901. <https://doi.org/10.1115/1.4046739>
- Ziegler L, Cosack N, Kolios A and Muskulus M** (2019) Structural monitoring for lifetime extension of offshore wind monopiles: Verification of strain-based load extrapolation algorithm. *Marine Structures* 66, 154–163. <https://doi.org/10.1016/j.marstruc.2019.04.003>
- Ziegler L, Smolka U, Cosack N and Muskulus M** (2017) Brief communication: Structural monitoring for lifetime extension of offshore wind monopiles: Can strain measurements at one level tell us everything? *Wind Energy Science* 2(2), 469–476. <https://doi.org/10.5194/wes-2-469-2017>