

Virtual sensing of subsoil strain response in monopile-based offshore wind turbines via Gaussian process latent force models

Zou, Joanna; Lourens, Eliz Mari; Cicirello, Alice

DOI

[10.1016/j.ymsp.2023.110488](https://doi.org/10.1016/j.ymsp.2023.110488)

Publication date

2023

Document Version

Final published version

Published in

Mechanical Systems and Signal Processing

Citation (APA)

Zou, J., Lourens, E. M., & Cicirello, A. (2023). Virtual sensing of subsoil strain response in monopile-based offshore wind turbines via Gaussian process latent force models. *Mechanical Systems and Signal Processing*, 200, Article 110488. <https://doi.org/10.1016/j.ymsp.2023.110488>

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.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.



Virtual sensing of subsoil strain response in monopile-based offshore wind turbines via Gaussian process latent force models

Joanna Zou ^{a,*}, Eliz-Mari Lourens ^b, Alice Cicirello ^b

^a Massachusetts Institute of Technology, Center for Computational Science and Engineering, 77 Massachusetts Ave, Cambridge, MA, USA

^b Delft University of Technology, Faculty of Civil Engineering and Geosciences, Stevinweg 1, 2628 CN, Delft, Netherlands

ARTICLE INFO

Communicated by E. Chatzi

Keywords:

Virtual sensing
Response estimation
Offshore wind turbine
Structural health monitoring
In-situ validation
Latent force model
Gaussian process
Kalman filter

ABSTRACT

Virtual sensing techniques have gained traction in applications to the structural health monitoring of monopile-based offshore wind turbines, as the strain response below the mudline, which is a primary indicator of fatigue damage accumulation, is impractical to measure directly with physical instrumentation. The Gaussian process latent force model (GPLFM) is a generalized Bayesian virtual sensing technique which combines a physics-driven model of the structure with a data-driven model of latent variables of the system to extrapolate unmeasured strain states. In the GPLFM, unknown sources of excitation are modeled as a Gaussian process (GP) and endowed with a structured covariance relationship with response states, using properties of the GP covariance kernel as well as correlation information supplied by the mechanical model. It is shown that posterior inference of the latent inputs and states is performed by Gaussian process regression of measured accelerations, computed efficiently using Kalman filtering and Rauch–Tung–Striebel smoothing in an augmented state-space model. While the GPLFM has been previously demonstrated in numerical studies to improve upon other virtual sensing techniques in terms of accuracy, robustness, and numerical stability, this work provides one of the first cases of in-situ validation of the GPLFM. The predicted strain response by the GPLFM is compared to subsoil strain data collected from an operating offshore wind turbine in the Westermeerwind Park in the Netherlands. A number of test cases are conducted, where the performance of the GPLFM is evaluated for its sensitivity to varying operational and environmental conditions, to the instrumentation scheme of the turbine, and to the fidelity of the mechanical model. In particular, this paper discusses the capacity of the GPLFM to achieve relatively robust strain predictions under high model uncertainty in the soil-foundation system of the offshore wind turbine by attributing sources of model error to the estimated stochastic input.

1. Introduction

With many existing monopile-based offshore wind farms approaching the end of their service life and the frequency of extreme climate events on the rise in recent years, offshore wind turbines (OWTs) face increasingly greater risk of fatigue-induced structural failure. The fatigue life of an offshore wind turbine (OWT) is driven by its dynamic response to the rotor motion of the nacelle in addition to wind and wave loading, which fluctuates over the lifetime of the structure [1]. For newly constructed offshore wind farms, where OWTs are larger and more slender in design in order to maximize power generation capacity, there is a greater risk

* Corresponding author.

E-mail address: jjzou@mit.edu (J. Zou).

<https://doi.org/10.1016/j.ymssp.2023.110488>

Received 23 June 2022; Received in revised form 29 May 2023; Accepted 30 May 2023

Available online 3 July 2023

0888-3270/© 2023 Elsevier Ltd. All rights reserved.

of resonance effects from the rotor frequency coinciding with the fundamental frequency of the structure, leading to accelerated fatigue damage [2].

Structural health monitoring (SHM) techniques offer a means to obtain information necessary for fatigue life evaluation of existing OWTs and design validation of recently installed OWTs. In practice, the strain response at critical locations can be measured over time to predict the accumulation of fatigue damage. However, current approaches to SHM which rely on physical measurements of strains are not well adapted for OWTs, because fatigue damage in monopile-based turbines concentrates in the foundation below the mudline where accurate strain readings are impractical to obtain. The installation and maintenance of sensors at subsoil fatigue hotspots is cost-prohibitive, and for monopiles instrumented prior to installation, the sensors are often damaged during the pile-driving process [3].

To overcome these challenges, *virtual sensing* approaches to SHM have received considerable attention in OWT monitoring applications, since they provide a full-field estimation of an unobserved quantity of interest (e.g. strains) using measurements of a proxy response variable (e.g. accelerations) from physical sensors. These methods can be efficiently implemented using a small number of standard sensors, such as accelerometers, placed at accessible locations of an operating structure. The response at unmeasured locations of the structure is then reconstructed from the measurements through a dynamics model of the system or transfer function. This approach to SHM not only reduces sensor deployment costs significantly, but also improves the quality of data collected, as accelerometers are more reliable, less sensitive to noise, and capable of recording at higher sampling rates compared to strain gauges [4,5]. Virtual sensing approaches may be categorized into two broad classes: (1) deterministic model-based extrapolation, where the estimation of response states beyond measured locations is obtained through a well-calibrated mechanical model, such as a finite element model, and (2) probabilistic state estimation, where strain response states are treated as random variables whose probability distributions are learned from measured data.

Among the deterministic techniques, the predominant method in practice is modal decomposition and expansion (MD&E) by Iliopoulos et al. [6], which decomposes the dynamical system into mode shape components from which the response between measured points may be inferred. The accuracy of MD&E depends on the parameters of the finite element model, which are calibrated using system identification techniques such as operational modal analysis. In applications to operational data from OWTs, MD&E produces a good estimate of the strain response in the time domain but faces difficulty with estimating higher frequency regions of the response spectrum [1]. Moreover, because MD&E derives displacements from the double integration of measured accelerations, the strain estimate is susceptible to low-frequency noise amplification, a numerical issue termed the “drift effect” which is detailed in [5]. To address the difficulty in the frequency domain strain estimate, Iliopoulos et al. [7] introduce a multi-band method where the dynamic frequency range is separated into two sections, such that a variable number of sensors and modes may be chosen for each frequency band. Studies by Noppe et al. [4,8] show that the multi-band method may lead to discontinuities in the estimate at the defined frequency band limits. Henkel et al. [3,9] extends the studies in [1,4,8] to the estimation of subsoil strains, where multi-band MD&E is observed to result in a phase lag in the time domain and an underestimation of the first mode spectral density for periods of high wind speeds. Generally, ill performance of MD&E can be attributed to its fully deterministic nature, such that it remains highly sensitive to discrepancies between the physical system and idealized model. In particular, MD&E does not explicitly account for stochastic measurement noise and model error resulting from uncertainty in soil-foundation interaction properties. Moreover, since calibration of the model must be performed offline, MD&E cannot account for shifts in the modal parameters of the system (e.g. natural frequencies, damping ratios) across operational states and loading conditions, which limits the robustness of its use.

The other class of techniques takes a Bayesian inference approach to virtual sensing and has a number of notable differences from deterministic methods. First, in estimating the trajectory of the unknown state, Bayesian inference methods produce a probability distribution of values of the state over time, reflecting additional information on the level of uncertainty in the predicted state rather than a single estimate of the time series. Secondly, sources of uncertainty from stochastic measurement noise and model error are explicitly incorporated in the construction of the inference problem. Finally, these virtual sensing methods offer a statistical framework to account for the direct dependence of the dynamic response on qualities of the force, or “input”, acting on the system, such as its magnitude, variability, and correlation structure. As no practical means exist for direct measurement of forces, the unknown input is either implicitly modeled as a form of process noise or treated as an additional random variable whose time evolution may be derived simultaneously with that of the strain response. The latter of these approaches, in which the joint distribution of unknown inputs and states is inferred from measurement data, is referred to as “joint input-state estimation” and constitutes the focus of this paper.

Bayesian inference approaches to virtual sensing are commonly implemented using the Kalman filter [10], whose sequential structure to inferring states from measurements is suitable for long time series-type data. Among the Kalman filter-based techniques, one version comprises of minimum variance unbiased filters derived from optimal control theory which do not assume prior knowledge of the time-varying behavior of the input. While early methods addressed the state estimation problem exclusively [11,12], Hsieh et al. [13] recognized the value of performing input estimation to improve the state estimation. The methods of [11,13] were extended in Darouach et al. [14] to include direct feedthrough between the inputs and states. These contributions lead to the development of the joint input-state estimator by Gillijns & de Moor [15,16], abbreviated as GDF or JIS in literature, which formalized a filter for simultaneously estimating inputs and states in linear systems satisfying minimum variance and unbiasedness conditions. Since then, the method has been extended for cases with a time delay using Rauch–Tung–Striebel (RTS) smoothing [17,18], for cases with a reduced order model and equivalent forces [19], and with a method to account for estimation uncertainty and correlation between inputs and states [20]. The GDF is evaluated extensively in literature, using numerical tests in [21–23], experimental tests in [19,24], and operational data in [1,19,25,26]. Notably, the GDF is compared to other state-of-the-art virtual sensing techniques using in-situ data from an offshore wind turbine in [27], where it is concluded that methods

are competitive and interchangeable. However, other work has noted difficulties with the robustness of the GDF in virtual sensing applications. In particular, the GDF is observed to be susceptible to numerical instability [19,22,26] arising from rank deficiency of system matrices, which occurs in overdetermined systems where the number of sensors exceeds the number of modes in the model [19] and where the sensor configuration results in near linear dependence between the inputs and states [26]. A comparison test in [24] demonstrates that the GDF is more susceptible to the drift effect from the accumulation of integration errors than other state-of-the-art joint input-state estimators. While the theoretical performance of the GDF guarantees minimum variance unbiased estimates for systems without model error, Maes et al. [20,25] note that the GDF no longer produces minimum variance estimates in applications where loading, noise, or structural conditions deviate from those assumed in the model. Moreover, the estimated error covariance matrices of the inputs and states no longer correspond to the true error. Optimality of estimates by GDF are with respect to the noise covariance matrices on the inputs and states, which must be known a priori or tuned appropriately for a given application.

In contrast to the GDF, which makes no prior assumption on the input, more recent developments in Kalman filter-based estimators assume that the input takes the form of a random walk model, where the time evolution of the input is driven by zero-mean white Gaussian noise. The augmented Kalman filter, or AKF [28], performs simultaneous estimation of the input and state variables in augmented state-space; in comparison to Maes et al. [1,27], which implements a classical Kalman filter for pure state estimation by assuming the process noise term encompasses contributions from the input, the AKF defines the covariance structure of the unknown input separate from that of process noise. To address issues of unobservability of the augmented state matrix when acceleration-only measurements are used, Naets et al. [29] proposes the AKF-dm, supplementing the AKF with artificial “dummy displacements” as additional measurements in order to stabilize the state estimate. The dual Kalman filter, or DKF [21], has also been proposed to overcome unobservability issues by introducing a successive structure to the filter, where the input estimation is followed by the state estimation. The performance of the Kalman filter-based estimators is compared using numerical data in [22,23], experimental data in [24], and operational data in [8]. A well-known shortcoming of the AKF, AKF-dm, and DKF is that their accuracy largely depends on an appropriate choice of noise covariance matrices, including process noise on dynamic states, measurement noise of observed states, and modeled noise in the unknown input [21,28]. While tuning procedures for covariance matrices have been proposed, including heuristic rules of thumb, a maximum likelihood estimation approach [30], and exhaustive search by the L-curve criterion [21,28], manual adjustment of the covariance matrix is often implemented to estimate an appropriate order of magnitude for the covariances. In several instances, the problem of estimating the covariance matrix is simplified by constraining the matrix to have a diagonal structure with constant values [21,28], placing the strict assumption that all inputs and states are uncorrelated and have equal magnitude variance. All these methods face the additional important limitation that the white noise assumption on the input may be too restrictive for real life applications. In the case of OWTs, which are assumed as linear time-periodic (LTP) systems, cyclic load sources are expected to contain harmonic components which violate the broadband assumptions of white noise input.

Recently, latent force models have been proposed as a generic case of Bayesian estimators which allow for greater flexibility in the treatment of the unknown input. Latent force models mark a connection between the control theory-based perspective and machine learning-based perspective of the state estimation problem; whereas purely data-driven approaches in statistical modeling traditionally suffer in cases of extrapolation and sparse datasets, the hybrid approach of a latent force model supplements the statistical model with a mechanical model of the underlying system dynamics to guide regression in the latent variable space [31]. In virtual sensing applications, the latent force model relates the measured accelerations, or observed variables, to the inputs and states, which comprise the unobserved or “latent” variables of the structural system [23]. Alvarez et al. [31] introduced the concept of the Gaussian process latent force model, or GPLFM, where a linear system is driven by an unknown Gaussian process whose covariance matrix is characterized by a kernel function. The GPLFM exploits properties of Gaussian processes to construct covariance relationships between the inputs and states, using parameters of the assumed kernel as well as correlation information supplied by the mechanical model. Posterior inference of latent variables may then be reduced to a problem of Gaussian process regression of the observed variables with joint inference of the unobserved variables. Whereas [31] provides a convolution-based solution to Gaussian process regression, Hartikainen & Särkkä [32] demonstrate that Gauss-Markov processes may be reformulated in state-space such that regression can be more efficiently performed using Kalman filtering [10] and RTS smoothing [17]. Their study is extended to the joint input-state estimation problem in [33], which provides the state-space framework for the GPLFM.

A numerical comparison of the GPLFM with competing joint input-state estimation algorithms for linear structural systems is carried out in Nayek et al. [23]. This investigation demonstrated that the GPLFM consistently outperforms or is on par with the performance of the AKF, AKF-dm, and DKF for joint input-state estimation across various loading scenarios. While Bayesian joint input-state estimators share many core characteristics, the superior performance of the GPLFM can be attributed to the nature of the statistical model parameters which require tuning [23]. The GDF, AKF, AKF-dm, and DKF all rely on individual definition and inference of full covariance matrices, leading to a complex covariance matrix estimation problem. To simplify the problem, Yuen et al. [30] propose one approach of taking a maximum likelihood estimate of the matrix diagonal using the negative log likelihood expression recovered from the Kalman filter; however, in this approach, the covariance matrix is parameterized using constraining assumptions, such as independent variance in each state dimension, and requires the solution to an optimization problem whose complexity scales with the dimension of the model. In [19,21], a modified L-curve criterion is adopted, which exhaustively searches for a constant variance parameter defining the noise covariance matrix which minimizes estimation error. Both of these procedures for covariance matrix estimation correspond to model fitting based on variance minimization, but either do not explicitly account for frequency content of the input, as in the case of the GDF, or fix an assumption of broadband frequency in the input using a white noise model, as in the case of the AKF, AKF-dm, and DKF. In contrast, the GPLFM characterizes the relationship between inputs and

states in terms of a GP covariance kernel. The kernel function is parameterized by a relatively small set of hyperparameters, the number of which are independent of the state dimension. Consequently, tuning hyperparameters by parameter estimation techniques such as maximum likelihood estimation [23] or Bayesian inference [34] leads to a much more efficient exploration of the tunable parameter space compared to full covariance matrix estimation. Moreover, these hyperparameters often correspond to interpretable properties of the input and response, such as the amplitude, frequency content, and smoothness or differentiability of the signal, in contrast to parameters describing individual covariance terms.

The GPLFM is first applied to structural health monitoring applications in Nayek et al. [23], where it is shown that the GPLFM is a generalization of the AKF which offers greater flexibility in the definition of the input in augmented state-space. Rogers et al. [34] extends the study to address system parameter estimation in addition to the joint input-state estimation problem. While the GPLFM has been evaluated in numerical studies, literature on the experimental or operational validation of the GPLFM in SHM is sparse. Petersen et al. [35] demonstrates the potential of combining analytical wind load models with the latent force model for long-span suspension bridges, but without the availability of validation data. Using data collected from an operating onshore wind turbine, Bilbao et al. [36] evaluates the performance of the GPLFM for strain estimation across the height of the tower above the foundation.

This contribution provides a critical study of the in-situ validation of the GPLFM and associated practical considerations for its use in fatigue monitoring. The fatigue assessment of offshore wind turbines serves as a suitable case study, as the GPLFM is expected to offer particular advantages in state estimation for systems subject to high model uncertainty and variable excitation sources which are not well characterized by the white noise assumption grounding other virtual sensing techniques. Moreover, this work presents the rare opportunity for validation of strain estimates in the OWT foundation below the mudline, as measurements from strain gauges embedded within the monopile are available in the considered case study. The paper is organized as follows: Section 2 presents the formulation of both the mechanical model of the system and the stochastic model of the unknown input in state-space, as well as the theoretical framework of the GPLFM. Section 3 presents the case study for in-situ validation, using subsoil data obtained in a monitoring campaign of an OWT in the Westermeerwind Park in the Netherlands. Section 4 presents the results of strain estimation for fatigue assessment and an evaluation of the sensitivity of the GPLFM to both measurement data fidelity and model fidelity. Section 5 concludes and provides recommendations for future research directions.

2. Mathematical formulation

2.1. Mechanical model of structural system

In structural dynamics, the response states of a physical system under forced vibration are described by a mechanical model which takes the form of a linear second-order differential equation:

$$\mathbf{M}\ddot{\mathbf{u}}(t) + \mathbf{C}\dot{\mathbf{u}}(t) + \mathbf{K}\mathbf{u}(t) = \mathbf{S}_p \mathbf{p}(t) \quad (1)$$

where $\mathbf{u}(t) \in \mathbb{R}^{n_u}$ are the displacement states at n_u degrees of freedom of a system with mass, damping, and stiffness properties defined by $\mathbf{M}, \mathbf{C}, \mathbf{K} \in \mathbb{R}^{n_u \times n_u}$ respectively. The system is externally excited by n_p forces $\mathbf{p}(t) \in \mathbb{R}^{n_p}$, with spatial distribution indicated by the selection matrix $\mathbf{S}_p \in \mathbb{R}^{n_u \times n_p}$ [19]. In this presentation, the system is assumed to be time-invariant, i.e. the model parameters and force locations are constant in time.

For complex models with several degrees of freedom, the solution for the response states may be more efficiently computed using model order reduction as in [19,37]. The eigensolution to $\mathbf{K}\Phi = \mathbf{M}\Phi\Omega^2$ returns the mass-normalized mode shapes ϕ_j , given by the columns of matrix $\Phi \in \mathbb{R}^{n_u \times n_u}$, and the natural frequencies ω_j^2 in rad/s, which are collected in the diagonal matrix $\Omega^2 \in \mathbb{R}^{n_u \times n_u}$. In the modally reduced order formulation, the matrices are truncated into $\bar{\Phi} \in \mathbb{R}^{n_u \times n_m}$ and $\bar{\Omega}^2 \in \mathbb{R}^{n_m \times n_m}$ to retain a subset of modes n_m which dominate the system response. By casting into modal states $\mathbf{r}(t) \in \mathbb{R}^{n_m}$ with the transformation $\mathbf{u}(t) = \bar{\Phi}\mathbf{r}(t)$ and premultiplying by $\bar{\Phi}^T$, Eq. (1) becomes:

$$\bar{\Phi}^T \mathbf{M} \bar{\Phi} \ddot{\mathbf{r}}(t) + \bar{\Phi}^T \mathbf{C} \bar{\Phi} \dot{\mathbf{r}}(t) + \bar{\Phi}^T \mathbf{K} \bar{\Phi} \mathbf{r}(t) = \bar{\Phi}^T \mathbf{S}_p \mathbf{p}(t) \quad (2)$$

With mass-normalized mode shapes, the orthogonality conditions $\bar{\Phi}^T \mathbf{M} \bar{\Phi} = \mathbf{I}$ and $\bar{\Phi}^T \mathbf{K} \bar{\Phi} = \bar{\Omega}^2$ are satisfied. Assuming proportional damping, the modal damping ratios ξ_j are given by $\bar{\Phi}^T \mathbf{C} \bar{\Phi} = \bar{\Gamma}$, where $\bar{\Gamma} = \text{diag}(2\xi_j \omega_j)$ for $j = 1, \dots, n_m$. The forces are transformed into their modal contributions by $\mathbf{f}(t) = \bar{\Phi}^T \mathbf{S}_p \mathbf{p}(t)$, where $\mathbf{f}(t) \in \mathbb{R}^{n_m}$ is the vector of modal force time histories. The modally reduced order model then becomes:

$$\ddot{\mathbf{r}}(t) + \bar{\Gamma} \dot{\mathbf{r}}(t) + \bar{\Omega}^2 \mathbf{r}(t) = \mathbf{f}(t) \quad (3)$$

To formulate the model for filtering and inference, the continuous-time linear differential equation is converted to state-space form. By defining the state vector $\mathbf{x}(t) = [\mathbf{r}(t) \quad \dot{\mathbf{r}}(t)]^T$, Eq. (3) can be expressed as:

$$\begin{bmatrix} \dot{\mathbf{r}}(t) \\ \ddot{\mathbf{r}}(t) \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{0} & \mathbf{I} \\ -\bar{\Omega}^2 & -\bar{\Gamma} \end{bmatrix}}_{\mathbf{A}_c} \begin{bmatrix} \mathbf{r}(t) \\ \dot{\mathbf{r}}(t) \end{bmatrix} + \underbrace{\begin{bmatrix} \mathbf{0} \\ \mathbf{I} \end{bmatrix}}_{\mathbf{B}_c} \mathbf{f}(t) \quad (4)$$

which is succinctly written as:

$$\dot{\mathbf{x}}(t) = \mathbf{A}_c \mathbf{x}(t) + \mathbf{B}_c \mathbf{f}(t) \quad (5)$$

Eq. (5) constitutes the dynamical model of the state-space system, which describes the evolution of response states over time as a function of the system parameters, given by \mathbf{A}_c , and of the input forces, whose influence is given by \mathbf{B}_c .

To facilitate numerical implementation, the dynamical model is converted from continuous-time to discrete-time with time step size Δt using the zero-order hold assumption. The discretized dynamical model is as follows:

$$\begin{aligned} \mathbf{x}_k &= \mathbf{A}\mathbf{x}_{k-1} + \mathbf{B}\mathbf{f}_{k-1} \\ \mathbf{A} &= \exp(\mathbf{A}_c\Delta t), \quad \mathbf{B} = [\mathbf{A} - \mathbf{I}]\mathbf{A}_c^{-1}\mathbf{B}_c \end{aligned} \quad (6)$$

In addition to the dynamical model, a discrete-time measurement model is defined to describe the relation between state variables and observed variables of a distinct sampling frequency. While accelerometers are predominantly used in SHM applications, axial strain readings from strain gauges may be additionally available with certain sensor configurations. Heterogeneous data fusion with displacement measurements is presented in [5,38]; for the case where strain measurements are available at an equal sampling rate as acceleration measurements, the measured quantities $\mathbf{y}_k \in \mathbb{R}^{n_y}$ for $k = 1, \dots, N$ are related to the original displacement, velocity, and acceleration states by:

$$\mathbf{y}_k = \begin{bmatrix} \boldsymbol{\varepsilon}_k \\ \ddot{\mathbf{u}}_k \end{bmatrix} = \begin{bmatrix} \mathbf{S}_s\mathbf{T}_s & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{S}_a \end{bmatrix} \begin{bmatrix} \mathbf{u}_k \\ \dot{\mathbf{u}}_k \\ \ddot{\mathbf{u}}_k \end{bmatrix} \quad (7)$$

where $\mathbf{S}_a \in \mathbb{R}^{n_{acc} \times n_u}$ is the selection matrix of degrees of freedom where n_{acc} accelerations are recorded and $\mathbf{S}_s \in \mathbb{R}^{n_{str} \times n_e}$ is the selection matrix of elements where n_{str} strains are recorded. $\mathbf{T}_s \in \mathbb{R}^{n_e \times n_u}$ is the transformation matrix which computes axial strain deformation $\boldsymbol{\varepsilon} \in \mathbb{R}^{n_e}$ from lateral displacements of a finite element model with n_e elements based on shape functions from beam theory, provided in Appendix A.2. The total dimension of the observation vector is therefore $n_y = n_{acc} + n_{str}$.

The observation equation with respect to the modal states \mathbf{x} is then given by:

$$\begin{bmatrix} \boldsymbol{\varepsilon}_k \\ \ddot{\mathbf{u}}_k \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{S}_s\mathbf{T}_s\bar{\boldsymbol{\Phi}} & \mathbf{0} \\ -\mathbf{S}_a\bar{\boldsymbol{\Phi}}\bar{\boldsymbol{\Omega}}^2 & -\mathbf{S}_a\bar{\boldsymbol{\Phi}}\bar{\boldsymbol{\Gamma}} \end{bmatrix}}_{\mathbf{G}} \begin{bmatrix} \mathbf{r}_k \\ \dot{\mathbf{r}}_k \end{bmatrix} + \underbrace{\begin{bmatrix} \mathbf{0} \\ \mathbf{S}_a\bar{\boldsymbol{\Phi}} \end{bmatrix}}_{\mathbf{J}} \mathbf{f}_k \quad (8)$$

which can be reduced to the following equation, with output influence matrix \mathbf{G} and direct feedthrough matrix \mathbf{J} :

$$\mathbf{y}_k = \mathbf{G}\mathbf{x}_k + \mathbf{J}\mathbf{f}_k \quad (9)$$

Together, Eqs. (6) and (9) give the discrete-time state-space model of the physical system:

$$\mathbf{x}_k = \mathbf{A}\mathbf{x}_{k-1} + \mathbf{B}\mathbf{f}_{k-1} + \boldsymbol{\eta}_{k-1} \quad (10a)$$

$$\mathbf{y}_k = \mathbf{G}\mathbf{x}_k + \mathbf{J}\mathbf{f}_k + \mathbf{v}_k \quad (10b)$$

The process noise term $\boldsymbol{\eta}_k$ with covariance $\mathbb{E}[\boldsymbol{\eta}_k\boldsymbol{\eta}_k^T] = \mathbf{Q}$ and measurement noise term \mathbf{v}_k with covariance $\mathbb{E}[\mathbf{v}_k\mathbf{v}_k^T] = \mathbf{R}$ are added to account for model uncertainty and measurement error, respectively. The noise terms are mutually uncorrelated white Gaussian noise processes, i.e. $\mathbb{E}[\boldsymbol{\eta}_k\mathbf{v}_k^T] = \mathbf{0}$.

In its current form, the linear state-space model in Eq. (10) cannot be used to directly solve for the response states because the forcing process \mathbf{f}_k is unknown. Whereas Maes et al. [27] simplifies the problem by grouping the unknown input with the white noise term \mathbf{w}_k , this approach has limited robustness in capturing the frequency range of dynamic strains for events with non-white input.

2.2. Gaussian process models

This section provides a brief overview of properties of Gaussian processes (GPs) and GP regression, which form the basis of the GPLFM methodology. In the case of a single dimensional input $t \in \mathbb{R}$, a Gaussian process defines a probability density over possible functions $h(t)$ of a random variable t . A GP has the property of being completely characterized by the first two moments of the process, namely its mean and covariance functions, expressed as follows:

$$h(t) \sim GP(\mu(t), \kappa(t, t'; \theta)) \quad (11)$$

where $\mu(t)$ denotes the mean function and $\kappa(t, t'; \theta)$ denotes the covariance function, also referred to as the kernel, representing the covariance as a function of distance between any two points t and t' of the process. Several options for forms of the kernel and its corresponding hyperparameters θ exist to control the smoothness and differentiability properties of $h(t)$, discussed further in Section 2.3.1.

Gaussian processes are commonly employed in regression problems to define a non-parametric mapping between inputs and outputs, e.g. $h : X \rightarrow Y$. Consider a training dataset $D = \{X, Y\}$ with inputs $X = \{t_1, \dots, t_N\}$ and corresponding outputs $Y = \{y_1, \dots, y_N\}$. The observed Y are assumed to be noisy measurements of the underlying process, modeled as $h(X)$:

$$Y = h(X) + \mathbf{v}, \quad \mathbf{v} \sim \mathcal{N}(\mathbf{0}, \sigma_v^2\mathbb{I}) \quad (12)$$

Discrepancy between measurements Y and the model $h(X)$ is represented with additive white Gaussian noise \mathbf{v} with variance σ_v^2 . In regression, h provides an approximation of the unobserved output corresponding to a test set of values X^* which are separate

(or held out) from the observed dataset D . In GP regression, moments of the Gaussian posterior distribution $p(h(X^*)|Y)$ may be computed with closed form expressions, as provided in [32,39]. The main drawback of this solution to GP regression is the high order of computational complexity, scaling with respect to the number of data points N by $\mathcal{O}(N^3)$, due to the inversion of a $N \times N$ matrix. Rasmussen & Williams [39] provide an algorithm for the practical implementation of the closed form solution, by addressing matrix inversion with Cholesky decomposition. However, it can be shown that for stationary Gauss-Markov processes, a recursive solution to GP regression by means of Kalman filtering and RTS smoothing may be implemented with a computational complexity scaling linearly with the dataset size, e.g. $\mathcal{O}(N)$ [32]. For particularly large datasets, as in cases with time series-type data, the recursive solution has substantial computational advantage.

2.2.1. Special case of stationary Gaussian processes

A stationary GP which satisfies Markov properties may be formulated in state-space such that Kalman filtering and RTS smoothing can be implemented to solve for its posterior distribution [32]. For the zero-mean stationary case of the GP, the mean function evaluates to zero, e.g. $\mu(t) \rightarrow 0$, and the covariance function does not change over time shifts $\tau = t - t'$, becoming:

$$h(t) \sim GP(0, \kappa(\tau; \theta)) \tag{13}$$

The temporal GP $h(t)$ of Eq. (13) may be equivalently expressed as the output of a linear time-invariant (LTI) system driven by a white noise process $w(t)$, taking the form of a stochastic differential equation (SDE) of order β :

$$\frac{d^\beta h(t)}{dt^\beta} + a_{\beta-1} \frac{d^{\beta-1} h(t)}{dt^{\beta-1}} + \dots + a_1 \frac{dh(t)}{dt} + a_0 h(t) = w(t), \quad S_w(\omega) = q_w \tag{14}$$

in terms of coefficients $a_0, \dots, a_{\beta-1}$ and the spectral density of the white noise, q_w . Eq. (14) may be converted to state-space through its companion form, in which the β -order linear SDE is represented as a system of β first-order linear SDEs. Defining a vector of $h(t)$ and its derivatives up to order $\beta - 1$ as $\mathbf{z}(t) = \left[h(t) \quad \frac{dh(t)}{dt} \quad \dots \quad \frac{d^{\beta-1}h(t)}{dt^{\beta-1}} \right]^T$, the companion form follows:

$$\begin{bmatrix} \frac{dh(t)}{dt} \\ \frac{d^2h(t)}{dt^2} \\ \vdots \\ \frac{d^\beta h(t)}{dt^\beta} \end{bmatrix} = \underbrace{\begin{bmatrix} 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 1 \\ -a_0 & -a_1 & -a_2 & \dots & -a_{\beta-1} \end{bmatrix}}_{\mathbf{F}_c} \begin{bmatrix} h(t) \\ \frac{dh(t)}{dt} \\ \vdots \\ \frac{d^{\beta-1}h(t)}{dt^{\beta-1}} \end{bmatrix} + \underbrace{\begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix}}_{\mathbf{L}_c} w(t) \tag{15}$$

or simply as:

$$\dot{\mathbf{z}}(t) = \mathbf{F}_c \mathbf{z}(t) + \mathbf{L}_c w(t) \tag{16}$$

The SDE in Eq. (16) constitutes the state-space model of the GP. It must be highlighted that properties of the SDE — namely, the spectral density q_w , the order of the model m , and the value of coefficients $a_0, \dots, a_{\beta-1}$ - are a function of the kernel chosen for the GP. Certain conditions must be met in order to construct a state-space model of the GP with the intended stationary kernel $\kappa(\tau)$: the kernel must satisfy positive semi-definiteness, and the kernel produces a GP $h(t)$ whose spectral density $S_h(\omega)$ is in rational form. For GPs which meet these criteria, the equivalent state-space model (\mathbf{F}_c , \mathbf{L}_c , and q_w) may be computed by spectral factorization, the steps of which are detailed in [23,32].

From Eq. (16), our interest is to estimate the probability density of the state vector, $p(\mathbf{z}(t))$. The Fokker-Planck-Kolmogorov (FPK) equation provides the complete probability density of the solution of a generic SDE, as detailed in [40]. However, due to properties of GPs, computation of the FPK equation can be bypassed and the first two moments of the distribution can be directly solved to provide a complete probabilistic description of the state vector, $p(\mathbf{z}(t)) \sim \mathcal{N}(\hat{\mathbf{z}}(t), \mathbf{P}(t))$. For the case of linear time invariant (LTI) SDEs with the initial condition $\mathbf{z}(t_0) \sim \mathcal{N}(\hat{\mathbf{z}}_0, \mathbf{P}_0)$, the differential equations of the mean and covariance of $\mathbf{z}(t)$ reduce to:

$$\frac{d\hat{\mathbf{z}}(t)}{dt} = \mathbf{F}_c \hat{\mathbf{z}}(t) \tag{17a}$$

$$\frac{d\mathbf{P}(t)}{dt} = \mathbf{F}_c \mathbf{P}(t) + \mathbf{P}(t) \mathbf{F}_c^T + \mathbf{Q}_c \tag{17b}$$

where the spectral density of the stochastic term is transformed as $\mathbf{Q}_c = \mathbf{L}_c q_w \mathbf{L}_c^T$. The solutions are then given by:

$$\hat{\mathbf{z}}(t) = \exp(\mathbf{F}_c(t - t_0)) \hat{\mathbf{z}}_0 \tag{18a}$$

$$\mathbf{P}(t) = \exp(\mathbf{F}_c(t - t_0)) \mathbf{P}_0 \exp(\mathbf{F}_c(t - t_0))^T + \int_{t_0}^t \exp(\mathbf{F}_c(t - \tau)) \mathbf{Q}_c \exp(\mathbf{F}_c(t - \tau))^T d\tau \tag{18b}$$

Rather than computing the integral in Eq. (18b) directly, the covariance $\mathbf{P}(t)$ can be solved for from the matrix Riccati differential equation in Eq. (17b) using the method of matrix fraction decomposition, detailed in [40,41]. However, the solution is further simplified in the case of zero-mean stationary processes, where it may be deduced that the SDE has a *steady-state* solution in which the state vector has a constant distribution at all time points, $\mathbf{z}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{P}_\infty)$. Considering the asymptotic behavior $\frac{d\mathbf{P}(t)}{dt} \rightarrow \mathbf{0}$ and

$\mathbf{P}(t) \rightarrow \mathbf{P}_\infty$ in the steady-state condition, Eq. (17b) becomes the Lyapunov equation, a special case of algebraic Riccati equations [40]:

$$\mathbf{0} = \mathbf{F}_c \mathbf{P}_\infty + \mathbf{P}_\infty \mathbf{F}_c^T + \mathbf{Q}_c \quad (19)$$

The Lyapunov equation is solvable with numerical algorithms [42] to obtain the steady-state covariance matrix \mathbf{P}_∞ . To recover the kernel function $\kappa(\tau)$ of the process, we can define the following relationship to extract $h(t)$ from $\mathbf{z}(t)$:

$$\begin{aligned} h(t) &= \mathbf{H}_c \mathbf{z}(t) \\ \mathbf{H}_c &= [1 \quad 0 \quad \dots \quad 0] \end{aligned} \quad (20)$$

Then, the kernel may be computed from the solution to the LTI SDE of Eq. (16), following the algebra in [40], to obtain:

$$\kappa(\tau) = \begin{cases} \mathbf{H}_c \mathbf{P}_\infty \exp(\mathbf{F}_c \tau)^T \mathbf{H}_c^T, & \tau \geq 0 \\ \mathbf{H}_c \exp(-\mathbf{F}_c \tau)^T \mathbf{P}_\infty \mathbf{H}_c^T, & \tau < 0 \end{cases} \quad (21)$$

2.2.2. Recursive solution to Gaussian process regression

For the special class of Gauss-Markov processes, GP regression may be more efficiently performed using recursive Bayesian filtering. Unlike the batch solution to GP regression, which considers all data at once by constructing the joint distribution across points, the recursive solution processes the data sequentially by exploiting the Markov property, where upon observation of the present state of the process $h(t)$, future instances of the state are independent of past observations [40]. Let $y_{i:j} = \{y_i, \dots, y_j\}$ denote the subset of data Y between time steps t_i and t_j . First, the steady-state solution is placed as the prior, $p(\mathbf{z}_0) = \mathcal{N}(\mathbf{0}, \mathbf{P}_\infty)$. Then, the independence condition between past and future observations allows for the construction of the recursive Bayesian filter as follows:

$$p(\mathbf{z}_k | y_{1:k}) = \frac{p(y_k | \mathbf{z}_k) p(\mathbf{z}_k | y_{1:k-1})}{\int_x p(y_k | \mathbf{z}_k) p(\mathbf{z}_k | y_{1:k-1}) dx} \quad \forall k = 1, \dots, N \quad (22)$$

in which the filtering posterior distribution $p(\mathbf{z}_k | y_{1:k}) \forall k = 1, \dots, N$ is obtained from a sequential ‘‘forward pass’’ through the data via Kalman filtering. What is worthy to note is that the prior distribution at time step k is taken to be the posterior distribution from the previous time step $k - 1$. Hence, Bayes’ rule can be applied recursively as the probability distribution of \mathbf{z}_k is updated incrementally with each data point in $y_{1:k}$.

Kalman filtering and RTS smoothing is a practical algorithm for performing recursive Bayesian estimation for linear Gaussian systems. First, the state-space equations of the GP are discretized to suit numerical implementation:

$$\mathbf{z}_k = \mathbf{F} \mathbf{z}_{k-1} + \mathbf{L} w_{k-1} \quad (23a)$$

$$h(t_k) = \mathbf{H} \mathbf{z}_k \quad (23b)$$

with discrete matrices $\mathbf{F} = \exp(\mathbf{F}_c \Delta t)$, $\mathbf{L} = \mathbf{L}_c$, $\mathbf{H} = \mathbf{H}_c$. In this sense, Eq. (23a) constitutes the dynamics model representing the time evolution of the states, and Eq. (23b) constitutes the measurement model relating samples of $h(t)$ to the states. As the process is a linear SDE with normally distributed states, its posterior distribution can be recovered through the Kalman filter, which provides the prediction $p(\mathbf{z}_k | y_{1:k})$ using past observations in sequence, followed by the RTS smoother, which provides the prediction $p(\mathbf{z}_k | y_{1:N})$ using a backward pass through the full observation dataset. The combination of filtering and smoothing gives the exact solution to GP regression for a set of noisy samples of the process [32]. The filtering and smoothing algorithms are detailed in Appendix A.3.

In order to implement Kalman filtering and RTS smoothing, the covariance matrix \mathbf{Q} of the discrete white noise $\mathbf{L} w_k$ is required. \mathbf{Q} may be calculated from the spectral density of the continuous case of $w(t)$, \mathbf{Q}_c , as:

$$\mathbf{Q} = \int_0^{\Delta t} \exp(\mathbf{F}_c(\Delta t - \tau)) \mathbf{Q}_c \exp(\mathbf{F}_c(\Delta t - \tau))^T d\tau \quad (24)$$

In the steady-state condition, a practical method of computing \mathbf{Q} is by enforcing $\mathbf{P}_k \rightarrow \mathbf{P}_\infty$ in the Kalman filter prediction equation (see Appendix A.3):

$$\mathbf{Q} = \mathbf{P}_\infty - \mathbf{F} \mathbf{P}_\infty \mathbf{F}^T \quad (25)$$

2.3. Gaussian process latent force model

In the GPLFM, modal components of the unknown input driving the dynamic response of a physical system are modeled probabilistically as a GP, and a posterior estimate of their evolution over time may be computed through GP regression. However, in the context of *latent* forces, no direct measurements of the process are available; rather, regression of the unknown input is performed through its relation with observed dynamics of the system captured in a *hybrid model* consisting of a mechanical component and a stochastic component. The hybrid model provides a stochastic representation of the system dynamics, therefore allowing for the joint posterior inference of all latent variables, including unobserved dynamics such as the strain states.

Using the reduced-order formulation of the mechanical model from Section 2.1, unique sources of inputs $\mathbf{p}(t)$ may be identically decomposed into modal components $\mathbf{f}(t) = [f_1(t) \ \dots \ f_{n_m}(t)]^T$. It is assumed that modal forces can be expressed as mutually uncorrelated GPs, leading to the following linear block-diagonal state-space model for multiple modal forces:

$$\begin{bmatrix} \dot{\mathbf{z}}_1(t) \\ \vdots \\ \dot{\mathbf{z}}_{n_m}(t) \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{F}_{c,1} & & \\ & \ddots & \\ & & \mathbf{F}_{c,n_m} \end{bmatrix}}_{\tilde{\mathbf{F}}_c} \begin{bmatrix} \mathbf{z}_1(t) \\ \vdots \\ \mathbf{z}_{n_m}(t) \end{bmatrix} + \underbrace{\begin{bmatrix} \mathbf{L}_{c,1} & & \\ & \ddots & \\ & & \mathbf{L}_{c,n_m} \end{bmatrix}}_{\tilde{\mathbf{L}}_c} \underbrace{\begin{bmatrix} w_1(t) \\ \vdots \\ w_{n_m}(t) \end{bmatrix}}_{\tilde{\mathbf{w}}(t)} \quad (26a)$$

$$\begin{bmatrix} f_1(t) \\ \vdots \\ f_{n_m}(t) \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{H}_{c,1} & & \\ & \ddots & \\ & & \mathbf{H}_{c,n_m} \end{bmatrix}}_{\tilde{\mathbf{H}}_c} \begin{bmatrix} \mathbf{z}_1(t) \\ \vdots \\ \mathbf{z}_{n_m}(t) \end{bmatrix} \quad (26b)$$

where for the j th modal force, $\mathbf{z}_j(t) = [f_j(t) \ \frac{df_j(t)}{dt} \ \dots \ \frac{d^{\beta-1}f_j(t)}{dt^{\beta-1}}]^T$ is the vector of derivatives corresponding to the companion form of the β -order SDE and $\mathbf{F}_{c,j}$, $\mathbf{L}_{c,j}$, $\mathbf{H}_{c,j}$, the white noise forcing $w_j(t)$ and its spectral density q_{w_j} are defined by the GP covariance kernel and associated hyperparameters from Section 2.2. Eqs. (26a), (26b) may be written in shorthand using the block state vector $\mathbf{s}(t) = [\mathbf{z}_1(t) \ \dots \ \mathbf{z}_{n_m}(t)]^T$ as:

$$\dot{\mathbf{s}}(t) = \tilde{\mathbf{F}}_c \mathbf{s}(t) + \tilde{\mathbf{L}}_c \tilde{\mathbf{w}}(t) \quad (27a)$$

$$\mathbf{f}(t) = \tilde{\mathbf{H}}_c \mathbf{s}(t) \quad (27b)$$

The spectral density $\tilde{\mathbf{Q}}_c$ of the term $\mathbf{L}_c \tilde{\mathbf{w}}(t)$ is a block matrix with $\mathbf{L}_{c,j} q_{w_j} \mathbf{L}_{c,j}^T$ along the diagonal. Because the mechanical model and the block GP model of the latent modal forces both take the form of a linear differential equation, it is possible to linearly combine the two into an augmented state-space model. We define the augmented state vector $\mathbf{z}^a(t) = [\mathbf{x}(t) \ \mathbf{s}(t)]^T$, where $\mathbf{x}(t) = [\mathbf{r}(t) \ \dot{\mathbf{r}}(t)]^T \in \mathbb{R}^{2n_m}$ denotes the modal displacement and velocity states. The GPLFM is given as:

$$\begin{bmatrix} \dot{\mathbf{x}}(t) \\ \dot{\mathbf{s}}(t) \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{A}_c & \mathbf{B}_c \\ \mathbf{0} & \tilde{\mathbf{F}}_c \end{bmatrix}}_{\mathbf{F}_c^a} \begin{bmatrix} \mathbf{x}(t) \\ \mathbf{s}(t) \end{bmatrix} + \underbrace{\begin{bmatrix} \mathbf{0} \\ \tilde{\mathbf{L}}_c \tilde{\mathbf{w}}(t) \end{bmatrix}}_{\tilde{\mathbf{w}}^a(t)} \quad (28a)$$

$$\mathbf{y}_k = \underbrace{\begin{bmatrix} \mathbf{G} & \mathbf{J}\tilde{\mathbf{H}}_c \end{bmatrix}}_{\mathbf{H}^a} \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}_k \quad (28b)$$

or in the form:

$$\dot{\mathbf{z}}^a(t) = \mathbf{F}_c^a \mathbf{z}^a(t) + \tilde{\mathbf{w}}^a(t) \quad (29a)$$

$$\mathbf{y}_k = \mathbf{H}^a \mathbf{z}_k^a \quad (29b)$$

where the spectral density \mathbf{Q}_c^a of the augmented white noise vector $\tilde{\mathbf{w}}^a(t)$ is given as:

$$\mathbf{Q}_c^a = \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \tilde{\mathbf{Q}}_c \end{bmatrix} \quad (30)$$

Eq. (29) constitutes the continuous-discrete version of the GPLFM, describing the time evolution of both the latent forces and latent system states. The hybrid model adopts the form of a linear SDE driven by a steady-state Gaussian input, and because the solution to the SDE is a linear operation of the input, the resulting system states are also GPs [43]. Therefore, joint posterior inference of the augmented states $\mathbf{z}^a(t)$ reduces to a GP regression problem, where regression of the measured variables provides evidence to simultaneously extrapolate the latent variable response [33].

Kalman filtering and RTS smoothing may be implemented to solve GP regression. The model is discretized and noise sources are considered for the purpose of filtering and inference, with $\mathbf{F}^a = \exp(\mathbf{F}_c^a \Delta t)$:

$$\mathbf{z}_k^a = \mathbf{F}^a \mathbf{z}_{k-1}^a + \tilde{\mathbf{w}}_{k-1}^a \quad (31a)$$

$$\mathbf{y}_k = \mathbf{H}^a \mathbf{z}_k^a + \mathbf{v}_k \quad (31b)$$

where as in Eq. (10), \mathbf{v}_k is additive white Gaussian measurement noise with covariance \mathbf{R} . Process noise is considered within the term $\tilde{\mathbf{w}}_k^a$, a notion which is discussed further in Section 2.3.2.

We may adopt the same steps as in Section 2.2.2 to derive moments of the LTI SDE solution. In particular, the augmented steady-state covariance \mathbf{P}_∞^a of the augmented states $\mathbf{z}^a(t)$ is computed from the Lyapunov equation:

$$\mathbf{0} = \mathbf{F}^a \mathbf{P}_\infty^a + \mathbf{P}_\infty^a (\mathbf{F}^a)^T + \mathbf{Q}_c^a \quad (32)$$

and the augmented covariance matrix \mathbf{Q}^a of the discrete noise \mathbf{w}_k^a follows as:

$$\mathbf{Q}^a = \mathbf{P}_\infty^a - \mathbf{F}^a \mathbf{P}_\infty^a (\mathbf{F}^a)^T \quad (33)$$

2.3.1. Parameter inference for the GP kernel function

Given a GP kernel, the corresponding hyperparameters θ which govern the length scales of the resulting process may be tuned according to the measurement dataset for a particular system. Nayek et al. [23] proposes to use maximum likelihood estimation for parameter inference, in which the optimal hyperparameters are obtained from the least squares solution which minimizes the residual norm between the observed data and model predictions. For Markov processes, maximization of the likelihood function $p(Y|\theta)$ may be reframed as the minimization of the negative log likelihood and computed iteratively using the Kalman filter equations. However, since such linear least squares problems are frequently ill-posed [44], the negative log likelihood function may not result in real-valued evaluations and its optimization is known to be susceptible to false solutions at local minima [23]. A discussion of the computational issues can be found in [35]. Rogers et al. [34] suggest a random walk Metropolis–Hastings algorithm to enforce a log acceptance ratio and avoid complex values in the objective function. Moreover, they note that initialization of the objective function at the prior $p(\theta)$ produces the maximum a posteriori (MAP) estimate of the hyperparameters.

As the iterative computation of the objective function using the Kalman filter is computationally intensive, particularly for long time-series data, this work implements a parameter fitting procedure for obtaining optimal values of the hyperparameters θ^* prior to filtering. Through properties of linear Gaussian systems, the GPLFM inherently places a Gaussian prior over the both latent and observed variables of the system. As prior assumptions are encoded by the kernel function assumed in the GPLFM, the hyperparameters of the Gaussian prior correspond exactly to those of the augmented model. Empirical qualities of the measurement data, such as the empirical mean and covariance, are then leveraged as prior information which is used to tune the kernel hyperparameters. In this approach, model parameters are fit by minimizing the distance d between the empirical distribution of the measurements, $p(Y) \sim \mathcal{N}(\mathbf{0}, \mathbf{P}_0^y)$, and the modeled Gaussian prior on the observed states $p(Y|\theta) \sim \mathcal{N}(\mathbf{0}, \hat{\mathbf{P}}_\infty^y)$, where by the steady-state condition, the prior covariance is $\hat{\mathbf{P}}_0^y = \hat{\mathbf{P}}_\infty^y$. For shorthand notation, $p(Y|\theta)$ is referred to by $p^\theta(Y)$. The optimal θ^* is then solved from:

$$\theta^* = \arg \min_{\theta} d(p, p^\theta) \tag{34}$$

In this work, the distance metric d between two probability measures is taken to be the Hellinger distance $H(p, p^\theta)$, which is provided in Appendix A.5.

As the distributions are taken to be zero-mean, parameter inference reduces to a problem of matching the covariance predicted by the model to the underlying covariance in the data. The empirical measurement covariance \mathbf{P}^y is computed directly from time-series data of the sensor channels. The modeled measurement covariance, $\hat{\mathbf{P}}_\infty^y$, may be derived from the augmented steady-state covariance encoded by the GP kernel function, \mathbf{P}_∞^a , by passing it through the measurement model \mathbf{H}^a of the GPLFM with the transformation:

$$\hat{\mathbf{P}}_\infty^y = \mathbf{H}^a \mathbf{P}_\infty^a (\mathbf{H}^a)^T \tag{35}$$

Rather than computing the likelihood function iteratively over each measurement, this approach takes advantage of the linear Gaussian nature of the GPLFM, such that the likelihood of θ may be computed efficiently using low-order moments of long time-series data. As the distance between Gaussian distributions in Eq. (34) results in a smooth, unimodal optimization surface across a defined search domain, the minimization problem may be solved efficiently using gradient descent methods.

2.3.2. Treatment of noise in the GPLFM

A feature of the GPLFM which stands in contrast to other Kalman filter-based estimators is that the definition of stochastic process noise on the augmented states, whose covariance \mathbf{Q}^a is solved for by means of the Lyapunov equation in Eqs. (32) and (33), is inherent to the formulation of a stationary GP in state-space. That is, the joint noise covariance of inputs and states are parameterized in terms of a handful of GP hyperparameters. This parameterization provides a more complete characterization of the model compared to the parameterization suggested in [30] for other estimators, where tunable parameters are limited to diagonal elements of the covariance matrix, or in [21,28], where a constant representing the order of magnitude of variance across all states is estimated using the L-curve criterion. Using Eq. (33), an appropriate augmented process noise covariance matrix \mathbf{Q}^a may be solved for by a data-driven technique of parameter inference, where the parameters are the GP kernel hyperparameters, rather than by heuristic tuning or ad-hoc methods. The resultant noise covariance matrix captures correlations between the unknown inputs and states which are generally neglected in the AKF and DKF.

An underlying assumption in the GPLFM of Eq. (28) is that sources of process noise on response states – any dynamical behavior which is unaccounted for in the mechanical model – are treated as additional external input. In this sense, the GPLFM does not distinguish between noise contributions and input excitation, which are characterized together as a GP perturbed by additive white noise. As a consequence, the input estimation is relatively sensitive to large model uncertainty, and the strain estimation may be susceptible to overconfidence. This aspect is generally not a concern in applications where one is less interested in the input estimation for which validation is impractical. However, for applications where an accurate input estimate is necessary, or wider variance bounds on the strain estimate are expected, process noise in the system states \mathbf{x} may be distinguished either by the addition of process noise in the latent system states $[\mathbf{Q}^x \ \mathbf{0}; \ \mathbf{0} \ \mathbf{0}]$, adding $\mathbf{Q}^x \in \mathbb{R}^{2n_m \times 2n_m}$ to the augmented covariance matrix in Eq. (33), as in [23,35], or by the incorporation of the spectral density of the process noise $\eta(t)$ into the vector of hyperparameters θ to solve for, as in [34].

The accuracy of the GPLFM is also sensitive to the choice of measurement noise covariance \mathbf{R} . Rather than choose a prior estimate of the measurement noise covariance, a method for posterior tuning of \mathbf{R} is adopted from [45], where measured data channels are used as an indicator of sensor noise. Upon obtaining a posterior estimate $p(\hat{\mathbf{z}}^a(t)|Y, \theta) = \mathcal{N}(\hat{\mathbf{z}}^a(t), \mathbf{P}^a(t))$ from Kalman filtering/RTS

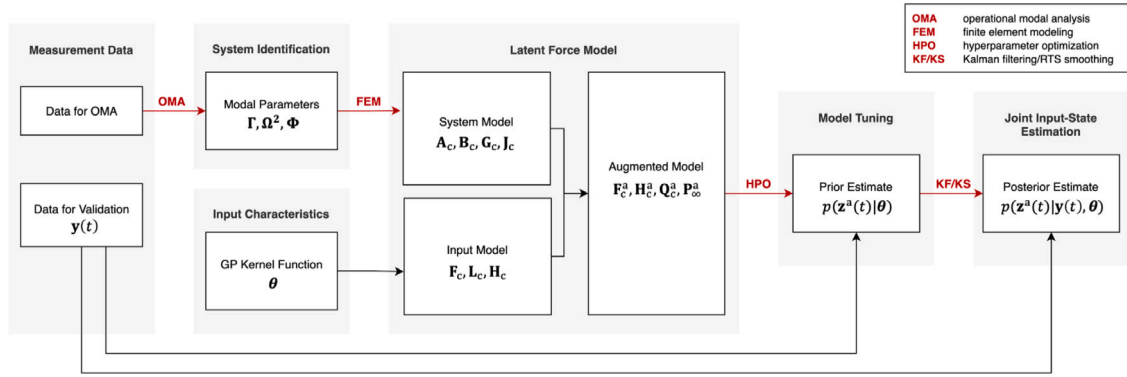


Fig. 1. Diagram of GPLFM methodology.

smoothing, the measurement noise is taken to be the residual r_k between the observed data Y and prediction on the observed variables $H^a \hat{z}_k^a$ as follows:

$$y_k = H^a \hat{z}_k^a + r_k \tag{36a}$$

$$r_k = y_k - H^a \hat{z}_k^a \tag{36b}$$

Measurement noise tuning then becomes a problem of minimizing the error between the assumed measurement noise covariance R in the model and the covariance of the residual time series $\hat{R} = \mathbb{E}[r_k r_k^T]$, executed through multiple iterations of posterior estimation using the GPLFM. One may specify termination criteria for arriving at a final estimate of the measurement noise covariance, such as a tolerance between consecutive iterates of R . It is observed that the measurement noise covariance tuning procedure results in greater accuracy of estimates of latent states, as demonstrated in a numerical example in Appendix A.1.

An overview of the GPLFM methodology is illustrated in Fig. 1. Source code for this work is publicly available at <https://github.com/joannajzou/GPLFM>.

2.3.3. Effect of modeling assumptions

This section discusses the degree of sensitivity of the GPLFM to various sources of model error which may be encountered in virtual sensing applications. In the following, we give emphasis to modeling assumptions chosen for this study and their impact on estimation quality, as well as possible extensions to more general assumptions. To substantiate the discussion on model error, an illustrative example is presented in Appendix A.1 in which the sensitivity of the GPLFM to perturbation in model properties is evaluated on a synthetic model, where the true input and response are known.

Gaussianity of the input and structural response. An underlying assumption in the adoption of the GPLFM for virtual sensing is that the unknown input and response states of the structure are linearly related and Gaussian in distribution. This assumption is deemed valid in structural health monitoring applications which utilize ambient vibrations to determine long-duration changes in structural conditions, a setting where only the linear elastic regime of structural response is engaged. In the case study which follows in Section 3, it is affirmed that the strain and acceleration data collected from an operating offshore wind turbine have ensemble distributions which are very close to Gaussian in nature. Moreover, an illustrative example of the GPLFM implemented on a synthetic model is presented in Appendix A.1, where for a structure excited by a non-Gaussian load, the response is still well captured using the GPLFM. Therefore, the assumption of Gaussianity is not expected to introduce significant error in joint input-state estimation results for the present in-situ validation study. While the present work focuses on an approach to fatigue monitoring over long time scales, Rogers et al. [34] have explored the coupling of input and state estimation with parameter estimation, to track shifts in modal parameters which can indicate nonlinear damage.

Stationarity of the GP. In the present study, the GP adopted in the formulation of the GPLFM is assumed to be zero-mean stationary. This assumption is hypothesized to be sufficient for the virtual sensing of dynamic strains, in which low-frequency components are not considered, for a structure subject to low-intensity ambient vibration, and is further evaluated in subsequent validation experiments. Assumptions of first-order stationarity allow for the GP to be modeled as a steady-state linear time invariant (LTI) SDE; for such systems, the Lyapunov equation may be invoked as an computationally convenient method to solve for the steady-state probability density over the states, as in Eq. (19). The GPLFM may be implemented with assumptions of first-order stationarity removed by computing the time-dependent probability density over states $z(t)$ via Eq. (18) for GPs with time-invariant system matrices, which produces a solution with a non-zero prior mean sequence. A further extension would be to consider time-varying system matrices of the GP, for which the probability density is computed by direct solution to the Fokker–Planck–Kolmogorov equation [40].

Choice of GP covariance kernel. The performance of the GPLFM in joint input-state estimation depends on an appropriate choice of the GP kernel and its hyperparameters. Generally, the kernel is chosen from a standard set of convenient kernels which meet the

criteria for conversion to state-space, as discussed in Section 2.2.1, based on qualitative prior knowledge of the time-varying behavior (e.g. periodicity, smoothness and differentiability) of the stochastic input. One example of a suitable kernel is the Matérn kernel function, which characterizes processes that are finitely differentiable [32]. In this work, we exclusively make use of the Matérn kernel function, as it is flexible enough to accommodate a broad range of process traits and is commonly employed in engineering applications to represent physically realistic processes. In the case of the Matérn kernel with smoothing parameter $\nu = 2$, covariance relationships are governed by two hyperparameters $\theta = [\alpha, l_s]$, where the vertical length scale α controls the amplitude and the horizontal length scale l_s controls the rate of time variability, and therefore the frequency content, of the resulting function. For convenience, the state-space model of the Matérn kernel is provided in Appendix A.4. Other common kernel choices which are valid for state-space implementation include the white noise (or Brownian motion) kernel function, which represents each time step of the GP as independent and identically distributed; the squared exponential kernel function, which describes infinitely differentiable processes; and the periodic kernel function, which captures strong harmonic components in time series data. Both the squared exponential kernel and periodic kernel require an approximation of the spectral density to be converted into state-space form, the details of which are provided in [32,46], respectively. For certain applications, GP regression may be particularly sensitive to the choice of kernel, where the substitution of another kernel function leads to a substantially different posterior prediction. While an evaluation of GP kernel sensitivity is beyond the scope of this work, an optimization-based approach to such a sensitivity study can be found in [47].

Identical GP model for all modal force components. In the virtual sensing case study which follows in Section 3, it is assumed that all modal forces $f_j(t), j = 1, \dots, n_m$ are endowed with the same GP prior. In other words, in Eq. (26), the modal forces are modeled with state-space matrices $\mathbf{F}_{c,j}, \mathbf{L}_{c,j}, \mathbf{H}_{c,j}, w_j(t)$ which are identical for all j and determined by a common set of GP hyperparameters θ . It is important to note that while these hyperparameters determine the prior variance placed on the companion form of the modal forces $\mathbf{z}_j(t)$, the posterior variance on these states is solved for via a Kalman filter update of the prior variance utilizing information from measurement data. Therefore, the posterior distribution will differ between components of the modal force even if the same GP hyperparameters are used to define the prior distribution. Moreover, in the Bayesian framework to posterior inference, with the availability of a sufficiently large set of measurement data, influence over the posterior is dominated by the likelihood and prior assumptions become negligible [44]. While it is possible to assign unique GP hyperparameters θ_j to each modal force, it is observed, but not shown in the paper for the sake of brevity, that this approach does not improve accuracy in the posterior estimate of the latent states in the present case study, and comes at the cost of a more expensive high-dimensional optimization problem which increases the complexity of achieving optimal hyperparameter tuning.

Uncorrelated modal forces. The simplifying assumption is made that all modal forces $f_j(t), j = 1, \dots, n_m$ are uncorrelated with each other, represented by diagonal terms being set to zero in the block-diagonal state-space model in Eq. (26). In reality, structural systems exhibit coupling between dynamic modes, which must be captured for a complete characterization of the covariance relationship between inputs and states. Neglecting information on the correlation between modal forces leads to reduced estimation accuracy and underestimation of uncertainty. In this work, this assumption is taken as necessary in the absence of prior knowledge of the nature of linear correlation between modal components of the unknown forcing. Nevertheless, our results on joint input-state estimation for a numerical system in Appendix A.1 and the case study wind turbine system in Section 4 provide empirical evidence that this assumption still delivers good estimates. Moreover, the state-space form of the GPLFM provides a useful framework for overcoming this assumption; an initial data-driven approach to modeling the correlation between modal forces is to parameterize the covariance term between $(f_i(t), f_j(t))$ for $i \neq j$ in the appropriate matrix entries of Eq. (26) and append these parameters to the set of kernel hyperparameters jointly inferred in an optimization routine.

Fixed modal parameters of the mechanical model. In practical applications of the GPLFM for virtual sensing, the mechanical model is an idealization of the structural system and cannot be known exactly. In general, the mechanical model is constructed to fit measured properties of the in-situ structure — in particular, modal properties such as the modal damping ratios, natural frequencies, and mode shapes which dictate dynamic response. However, for large scale structures, these modal properties must be inferred rather than measured directly. Modal identification of offshore wind turbines by traditional operational modal analysis (OMA) techniques is generally limited to the prediction of modal properties of the OWT in the idling state, when loading conditions most closely match the white noise assumption for ambient loading in OMA [48]. In the run-up and production states, these modal properties shift due to variation in damping sources and the introduction of interaction with harmonic loads from the environment and rotor blade dynamics [49]. Damping identification presents a particular challenge as the level of contribution of aerodynamic damping, in addition to hydrodynamic damping and soil damping in offshore contexts, is expected to fluctuate widely over time [2,49]. Therefore, the results of OMA obtained from data from the idling state do not fully reflect the true modal properties of the OWT in other states and are expected to result in an underestimation of modal damping levels, as well as a mismatch in natural frequencies and mode shapes. It is shown in an illustrative example in Appendix A.1 that errors in modal properties of the system affect input estimation by the GPLFM, but as a measure to result in reduced influence on estimates of response states such as displacements and accelerations. It is expected that error in response estimation resulting from error in modal properties will be present but not necessarily significant, due to compensatory effects of the input estimation.

Beam theory model for the transformation of displacements to axial strains. Since the strain response of a structure is required to assess fatigue damage accumulation, the objective of virtual sensing is to estimate the strain response at fatigue-critical locations where direct measurement of strains is infeasible. When representing the structure with a mechanical model of the form in Eq. (1), dynamic response states of the model correspond to displacements and its higher order derivatives, including velocities and accelerations. Estimates of axial strain response must be computed by an appropriate transformation of the lateral displacement response produced by the model. A simplifying but practically necessary assumption is to perform the transformation using elastic beam theory as

presented in [Appendix A.2](#), which relies on approximations of translational, rotational, and shear deformations described by finite element shape functions. In the present validation study, the shape functions are chosen to be those of Timoshenko beam elements, provided in [50]. Due to this approximation, it is likely that deterministic conversion of displacements to strains will result in mismatch between the predicted strain response and the true strain response of the structure.

2.3.4. Practical notes for numerical implementation

This section briefly discusses numerical sensitivities particular to the presented formulation of the GPLFM and approaches to resolve such issues, which are in part sourced from [41].

High dimensionality. For physical systems with several degrees of freedom, operations on large matrices in the Kalman filter and RTS smoother is susceptible to roundoff error and challenges with the inversion of sparse matrices. With model order reduction of the mechanical model in Section 2.1, the dimension of the system becomes equal to the n_m number of modes considered, thus reducing the cost of computation and controlling for stable matrix inversion. Moreover, the reduced order formulation produces an estimate of the modal contribution of any number of forces rather than the exact forces, such that knowledge of the location of the unknown inputs is not necessary [51,52]. For instance, in the transformation of forces into their modal contributions by $\mathbf{f}(t) = \Phi^T \mathbf{S}_p \mathbf{p}(t)$, \mathbf{S}_p may be taken to be the identity matrix such that the collective effect of all possible forces $\mathbf{p}(t)$ acting at any degree of freedom is considered and no error is introduced from incorrect assumptions on the force locations.

Ill-conditioning of covariance matrices. The Kalman filter algorithm is sensitive to ill-conditioning of the covariance matrices, which more commonly occurs in the companion state-space model where higher derivatives of a process are included in the state vector [41]. The ill-conditioning may be resolved through (a) reformulating the update equation in the Kalman filter using matrix decomposition methods, such that filtering is performed using stable factors of the covariance matrix \mathbf{P}_k , and (b) ensuring the addition of non-zero measurement noise \mathbf{R} to stabilize the update equation in events of near-zero change from \mathbf{P}_{k-1} to \mathbf{P}_k .

Numerical instability from scale difference. Numerical instability may arise if there exists a large scale difference between components of the state vector \mathbf{x} or measurement vector \mathbf{y} , as the filtering of large components may cause cancellation of the significant digits of small components. Scale differences are particularly common when heterogeneous measurement data of different units are used, as is the case with the combination of acceleration and strain measurements. Such numerical errors can be avoided through either a change of units or a rescaling of the vector. In this work, we recommend at a minimum to normalize the measured quantities \mathbf{y} prior to filtering, then applying the reverse operation to the posterior estimate to recover the true signal magnitude.

3. Offshore wind turbine case study

In this contribution, validation of the GPLFM for dynamic strain estimation is performed using measurement data from the SWT 3.0-DD offshore wind turbine on a monopile foundation, referred to as the W27 turbine hereon. The W27 turbine is a member of Westermeerwind Park, an operating near-shore wind farm situated in the IJsselmeer Lake in the Netherlands. The instrumentation scheme for the W27 turbine is rare by industry standards, as it contains strain sensors on the monopile foundation at varying depths below the mudline.

3.1. Measurement setup and data processing

The measurement set used for validation consists of the following data, which was gathered from separate data acquisition systems. All positions of sensors are given with respect to the mean water level (NAP) and illustrated in [Fig. 2](#).

- Rotor-nacelle assembly (RNA) data: data obtained from the SCADA system, where associated sensors are located at the level of the rotor axle at +95.0 m NAP. Accelerations in the fore-aft (FA) and side-side (SS) directions of the turbine are recorded at the RNA, where FA references the direction of the rotor axis and SS is the perpendicular direction. In addition, the RNA data consists of measurement channels for wind direction and wind speed, rotor speed, blade pitch angle, power generation level, and yaw angle (rotational position of the nacelle).
- Support structure (SUS) data: data from accelerometers and strain gauges installed along the height of the turbine tower and monopile foundation. Three accelerometers are located above the waterline: Acc. 1, measuring in the east-west direction at +4.5 m NAP; Acc. 2, measuring in the north-south direction at +4.5 m NAP; and Acc. 3, measuring along the 39.4°–219.4° line from North at +42.5 m NAP. The SUS data includes nine rings of 4 strain gauges labeled A, B, C, D, measuring axial strains at four faces of the monopile; however, several of the strain gauges at positions C and D, as well as all strain gauges in the first ring level, were damaged during the pile-driving process. Therefore, strain gauge readings at positions A and B in the five rings located below the mudline (Rings 2–6) and three rings located above the mudline (Rings 7–9) are retained for validation. The inclinometer reading at +4.5 m NAP is neglected in this study.

The RNA and SUS data are collected continuously and organized into 10-min increments over a 15-month period between 2016 and 2017. The turbine is monitored over various operational states of the OWT, namely the idling state (minimal rotor motion), production state (full rotor speed), and run-up state (the transitional state). Each of the 10-min records is categorized by operational state using state curves, such as the wind speed-rotor speed curve in [53].

Due to the use of different data acquisition systems, signal processing steps are applied to ensure synchronization between the RNA and SUS data. First, the RNA data is downsampled to meet the 50 Hz sampling rate of the SUS data, which is determined to be sufficient for capturing the frequency resolution of all data channels, as the significant frequency content is expected to fall




Fig. 2. Position of 8 strain gauge rings and 3 accelerometers along the support structure of the W27 wind turbine. Source: Figure adopted from [Versteijlen et al.].

below 10 Hz. Resampling is followed by the time series synchronization of the data channels based on cross-correlation sequences. Full details on the time series synchronization procedure are provided in [54]. Because the primary interest in this study is to estimate dynamic strain response which drives fatigue damage accumulation, quasi-static components are removed via a bandpass Butterworth filter of order 6 with stopband frequencies of 0.12 Hz and 10 Hz.

3.2. Modal analysis

To identify in-situ modal parameters of the structural system – such as mode shapes, natural frequencies, and damping ratios – operational modal analysis (OMA) is conducted using the covariance-driven stochastic subspace identification (SSI-Cov) technique [48]. For the W27 turbine, acceleration measurements at three elevations from Acc. 1 and 2, Acc. 3, and the accelerometer in the RNA are transformed into the local RNA coordinate system to allow for separate identification of the FA and SS modes. From the stabilization diagram produced by the SSI-Cov routine, stable columns corresponding to structural modes are identified using the OPTICS clustering algorithm [55,56].

As with the majority of OMA techniques, it is implicitly assumed in SSI-Cov that (1) the system is linear time invariant, meaning the modal parameters remain unchanged over time; (2) the system is driven by unobserved ambient excitation which takes the form of white Gaussian noise; and (3) the system response is not colored by nonstructural damping sources. Therefore, to ensure robust and accurate results, a number of considerations are made in the selection of 10-min records to use for modal analysis. (A) The dataset is filtered for records which correspond to the idling state, as loading conditions in this state most closely match the assumption of white noise excitation in the SSI-Cov method and allow for structural resonance to be discerned from noise [2]. (B) The records are further filtered based on the yaw angle of the turbine in order to allow for distinction between the FA and SS modes. The specific sensor configuration of the W27 turbine presents the limitation that Acc. 3 only provides a uniaxial measurement along the 39.4° – 219.4° line, meaning it cannot measure components of motion perpendicular to its axis. Therefore, the RNA must be oriented parallel to the axis of Acc. 3 for the channel to be utilized in identifying FA modes and perpendicular to its axis for SS modes, such that the measurement corresponds completely to the motion in the respective direction. Records with little to no deviation in yaw angle across the 10-min period are chosen to reduce conflation with out-of-plane modes. (C) Of the remaining records, those which correspond to low wind speeds are prioritized in the selection in order to reduce the effect of aerodynamic damping on the modal estimate. (D) Records in the final selection are truncated to sections of the data with relatively consistent time-frequency content which reflects stationary conditions. Following these criteria, the FA and SS modal parameters are each derived from twelve 10-min records taken across the 15-month duration of the monitoring campaign, with a sample record shown in Fig. 3. It should be clarified that these criteria for record selection only applies to the dataset used for modal identification by SSI-Cov; the validation records used in strain estimation are not subject to such constraints, but rather are chosen broadly across operational states and wind speeds.

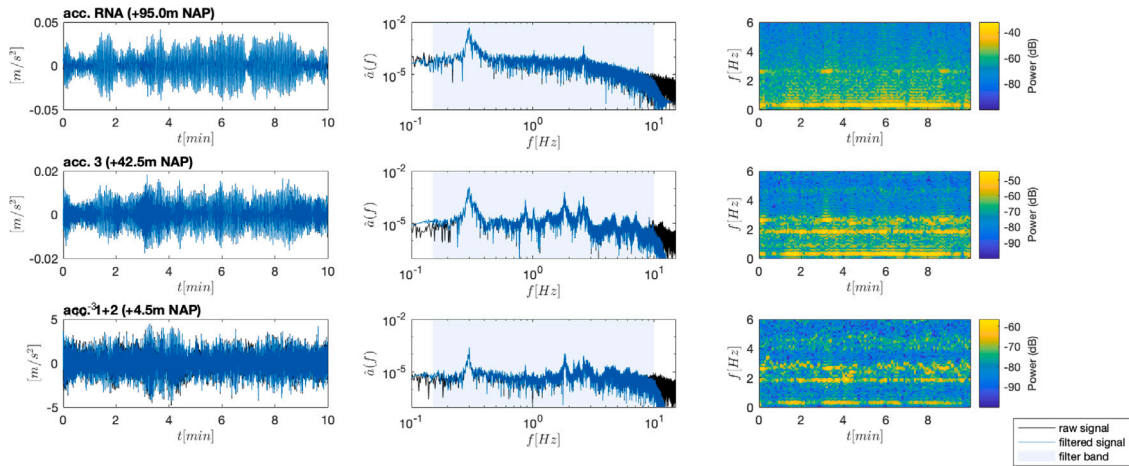


Fig. 3. Sample 10-min record of the three acceleration channels used for modal identification, shown with the filtered time series (left column), discrete Fourier transform (center column), and spectrogram (right column) of the data. The cutoff frequencies of the bandpass filter are indicated in the Fourier spectrum.

3.3. Finite element model

A mechanical state-space model of the W27 turbine is derived from a Finite Element (FE) model which is first constructed based on known geometric and material properties, then tuned to match with the modal parameters identified through OMA. A relatively simple FE model, such as a cantilever beam model, is sufficient for capturing the dominant modal properties of the system. The FE model consists of 124 Timoshenko beam elements, each 1.0 m in length, to span across the 29.0 m monopile depth and 95.0 m tower height. The tower does not have a separate transition piece as the monopile and tower have a bolted connection. The tapered profile of the top section of the turbine is modeled using pipe cross sectional elements which diminish in diameter linearly with height. All elements of the turbine are made of structural steel with Young’s modulus $E = 210$ GPa, Poisson’s ratio $\nu = 0.3$, and mass density $\rho = 7850$ kg/m³. External components such as the RNA, platform, and power unit are represented with concentrated masses at their respective elevations.

A nontrivial source of model uncertainty is in the subsoil properties of the foundation. Mass and stiffness contributions of the soil are modeled by adding a mass density of 1500 kg/m³ and a stiffness profile to the elements located below the mudline. A custom 1D soil stiffness profile for the W27 turbine is taken from Versteijlen et al. [57,58], where detailed studies of soil characterization and modeling were conducted for the site. In their work, geophysical measurements from Cone Penetration Tests (CPT), Seismic Cone Penetration tests (SCPT), and laboratory testing of borehole samples are used to develop an initial 1D effective Winkler profile $k_{eff}(z)$ which fits elastic properties of the soil-structure system as a function of depth z . Then, in-situ shaker tests of the W27 monopile are performed to obtain dynamic strain measurements from Rings 2–7, which are used to solve for an optimal scale factor γ which tunes the magnitude of the profile. It is found that while the optimal γ varies as a function of frequency, a constant factor $\gamma = 0.78$ is most appropriate for the low-frequency regime which governs the response in OWTs. The resultant $\gamma k_{eff}(z)$ is illustrated as a distributed springs model in Fig. 4. Note that in the absence of high-fidelity soil characterization data, the p-y curve [57] is an alternative soil stiffness profile which is commonly adopted in practice.

Table 1 compares five modes of the final FE model, denoted with subscript fem , with those identified by OMA, denoted with subscript id . From OMA, natural frequencies are identified with relatively high confidence, reflected by low sample standard deviation, whereas there is greater sample standard deviation in estimates of the modal damping ratios. Between the FE model and OMA results, there is a very good match in the natural frequencies for the first four modes, with the relative error well below 1%. The high modal assurance criterion (MAC) above 0.9 indicates a strong fit between the mode shapes, as is visible in Fig. 5. In general, the FA and SS modes are close to symmetric; any deviations from total symmetry are likely due to differential local soil damping and stiffness. There is greater difficulty in fitting the third SS mode of the FE model to the OMA results, which is due to the fact that the dynamic response of the OWT is controlled primarily by the first two modes. Therefore, higher modes are identified with significantly greater sample variance and reduced certainty in the OMA procedure; for instance, it was not possible to identify the third FA mode from the dataset. It is expected that the first two modes in each of the FA and SS directions dominate modal contributions to the dynamic response of the OWT across operating states.

4. Application to virtual sensing

This section discusses validation of the GPLFM for virtual sensing using operational data from the W27 turbine. First, the fitting of kernel hyperparameters using the approach of Eq. (34) is presented. The results of dynamic strain and modal force estimation using acceleration-only measurement data are shown for sample records from each of the turbine’s operational states. Next, the sensitivity




Fig. 4. Diagram of Winkler foundation model and soil stiffness profile. Source: Figure adopted from [Versteijlen et al.].

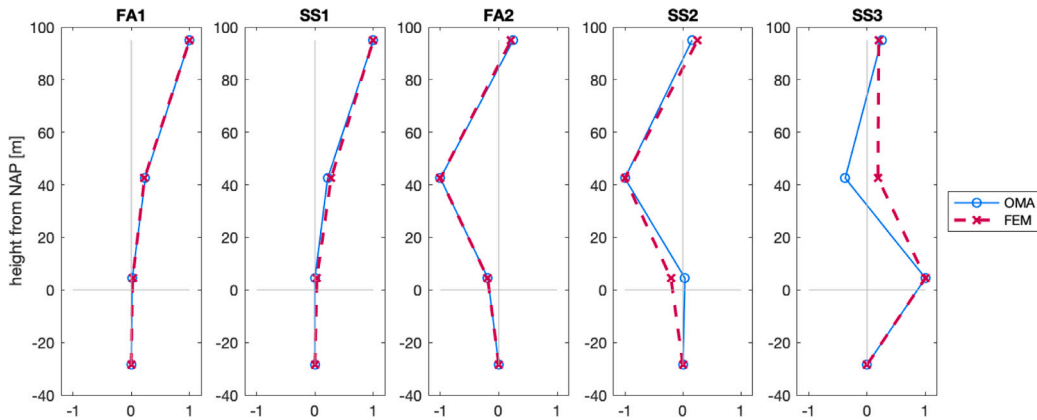


Fig. 5. Mode shapes identified by OMA (solid blue) and estimated by the finite element model (dashed red).

Table 1

Mean and standard deviation of modal damping ratios and natural frequencies identified by OMA. Error in natural frequencies (e_f) and mode shapes (MAC) of the FE model are computed with respect to the mean modal parameters from OMA.

Mode	ζ_{id}		f_{id} [Hz]		f_{fem} [Hz]	e_f [%]	MAC
	mean	std.	mean	std.			
FA1	1.932	2.051	0.296	0.015	0.296	0.01%	0.999
SS1	3.011	1.703	0.290	0.009	0.289	-0.12%	0.996
FA2	2.379	2.126	1.869	0.058	1.874	0.23%	0.988
SS2	2.887	1.931	1.885	0.029	1.883	-0.09%	0.939
SS3	2.604	1.615	4.563	0.252	5.193	13.79%	0.743

of the GPLFM to the sensor configuration is examined to determine the minimum sufficient measurement setup for characterizing the dynamic strain response of the OWT. Finally, perturbation to the foundation stiffness profile is introduced to evaluate the ability of the GPLFM to accommodate model uncertainty. The performance of the GPLFM is calculated using the correlation coefficient (CC) and mean relative error (MRE) of the predicted strain time series in relation to the validation measurements, as well as by converting the strain estimate into a Damage Equivalent Load (DEL) [54,59] to measure the impact of estimation error on fatigue assessment.

In this application of the GPLFM, the finite element model of the W27 turbine is translated into a reduced-order model retaining the first $n_m = 3$ modes in each of the FA and SS directions of the structure. The stochastic model of the input assumes the three

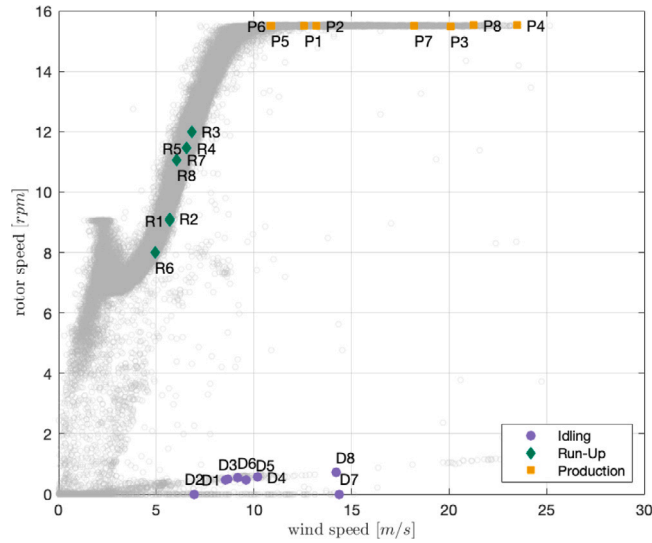


Fig. 6. Wind speed-rotor speed state curve of 10-min records taken over the period of March 2016 to June 2017 for the W27 turbine. The records selected for validation in each of the operating states are labeled.

modal forces are independent and identically characterized by a Matérn kernel with $\nu = 5/2$, such that $\kappa_1(\tau) = \kappa_2(\tau) = \kappa_3(\tau)$. The corresponding augmented model is therefore driven by two hyperparameters, $\theta = [\alpha, l_s]$.

The acceleration-only measurement dataset consists of acceleration time series transformed into FA and SS coordinates at three elevations of the structure; these elevations are referred to as A1 (combination of orthogonal Acc. 1 and Acc. 2 at +4.50 m NAP), A2 (uniaxial Acc. 3 at +42.50 m NAP), and A3 (accelerometer in the RNA at +95.0 m NAP). The validation dataset consists of strain time series recorded at position A for FA motion and position B for SS motion within the rings at six depths below the mudline, referred to by their corresponding ring numbers as Str. 2–7 (refer to Fig. 2 for an illustration). Validation is shown for eight records from each of the operating states of the turbine, with each record identified by a numeral preceded by the letter ‘D’ for the idling state, ‘R’ for the run-up state, and ‘P’ for the production state. The records are selected with the intention to encompass a wide range of wind and rotor conditions in each operating state, as shown on the wind speed-rotor speed state curve in Fig. 6.

4.1. Hyperparameter fitting

For each record in the validation dataset, the hyperparameters of the GPLFM are fit according to the approach of Eq. (34). It is consistently observed that evaluation of the objective function produces a convex surface with a clear optimum θ^* , illustrated in Fig. 7a for sample record P3 with search domain $\alpha = [0, 50]$ and $l_s = [0, 1]$. To visualize the goodness of fit, the empirical distribution of the measurements (in blue) and the predicted prior distribution corresponding to the optimal model hyperparameters θ^* (in orange) are overlaid on the time series plots and histograms of the measured variables in Fig. 7b. The near exact match in the distributions across all three measurement channels indicates an appropriate fit of the model to prior information on the true system dynamics.

The optimal hyperparameters for all validation records are shown in Fig. 8. In the leftmost plot, the hyperparameter sets $\{\alpha, l_s\}$ for each of the operating states are observed to fall in distinct clusters. In particular, the idling state (records D1, ..., D8) corresponds to large values of l_s and small values of α due to the dominance of low frequency oscillations driven by low-powered ambient excitation of the tower. The production state (records P1, ..., P8) corresponds to small values of l_s and large values of α due to the combination of high wind speeds and higher-frequency motion resulting from interaction with the turbine rotor. From these correlations, we may draw physical interpretation of the controlling hyperparameters of the Matérn kernel as they relate to qualities of the unknown input. To substantiate this reasoning, we refer to the center plot showing a positive correlation between α and the mean wind speed of the record, implying the vertical length scale parameter relates to the amplitude of the input as it varies with the magnitude of wind speed. In the rightmost plot, l_s is compared with a metric of relative power in the low frequency range, which is computed as the ratio of the power (integration of area under the power spectral density function) from 0.1 to 1.0 Hz with the power from 1.0 to 5.0 Hz. Larger values of this ratio indicate greater contributions of low frequencies to the response spectrum. From the positive correlation observed, it can be inferred that the horizontal length scale parameter relates to the dominance of low frequency inputs, which varies with both ambient conditions and rotor-induced loads.

The run-up state (records R1, ..., R8) corresponds to hyperparameters falling between the ranges of the idling and production states, as it marks the transition between these two states. Moreover, lower variance in the hyperparameters is observed compared to the other two states because run-up of the OWT always occurs in predefined wind conditions, whereas the OWT may remain stationed in the idling or production state across a wider range of wind speeds (see Fig. 6).

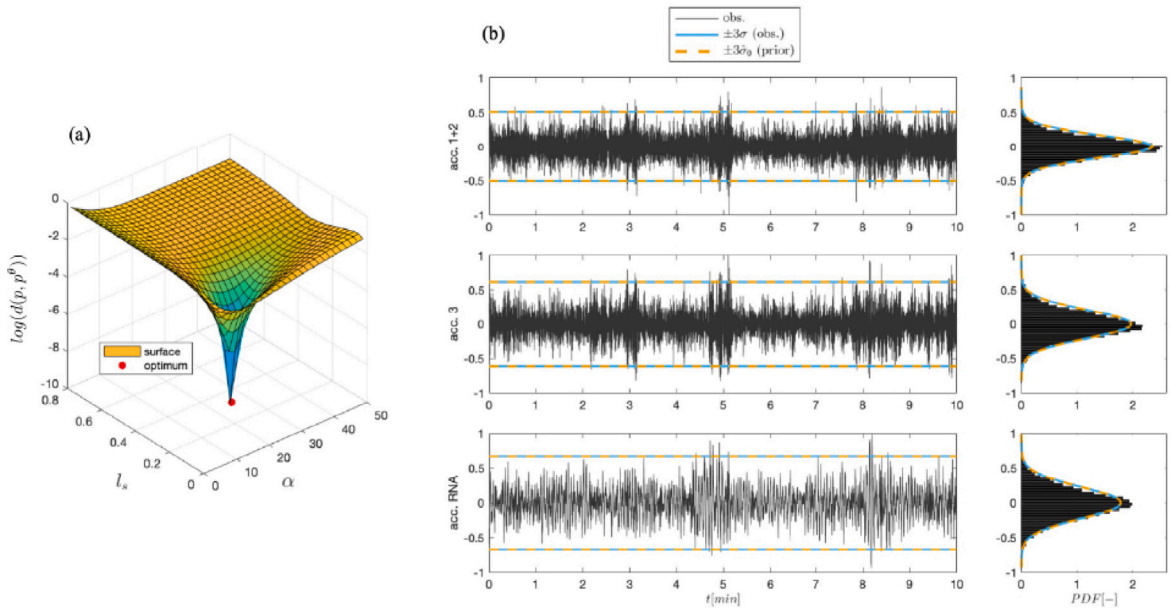


Fig. 7. (a) The optimization surface across the search domain $\alpha = [0, 50]$ and $l_s = [0, 1]$ for hyperparameters of the Matérn $\nu = 5/2$ kernel. (b) Measurement time series and histogram of acceleration data at 3 channels for record P3, with the empirical distribution indicated in blue and modeled distribution indicated in orange.




Fig. 8. Optimal l_s vs. optimal α (left), wind speed vs. optimal α (center), and optimal l_s vs. relative power in low-frequency range 0.1–1.0 Hz (right) for eight records in each operational state.

4.2. Response estimation

The strain estimation results in the idling, run-up, and production states are given in Figs. 9, 10, and 11, respectively. Overall, the GPLFM attains remarkable accuracy in the estimation of the frequency content of sub-soil strains; discrepancy is observed in the high frequency range in Fig. 9, where low signal-to-noise ratio in the idling state results in difficulty with capturing the signal where it falls below the level of sensor noise in the validation data. This discrepancy disappears for the run-up and production states, as the response is no longer obscured by the noise amplitude. Despite the good performance in capturing the dominant frequency components, an overestimation of the strain response is evident from both the time- and frequency domain plots. The error in the amplitude is most apparent in the idling state in Fig. 9 as well as at periodic impulses caused by turbine rotor motion in the run-up and production states, shown in the time series of Figs. 10 and 11 respectively. The discrepancy in the predicted signal amplitude is caused by a slight overestimation in the contribution of low frequencies falling below the first natural frequency of the system, which may be attributed to the general challenge of the estimation of low frequency strains. In particular, low frequency noise in measured accelerations is inevitably amplified over double integration of accelerations to obtain strains [5], leading to the observed inflation of low frequency content in the estimated strains. Moreover, effective characterization of the quasi-static behavior of the turbine generally requires a large number of modes to be included in the modal decomposition of the model, a number which exceeds the 2–3 modes which are able to be accurately identified in the present study [27]. The inflation of low frequencies may

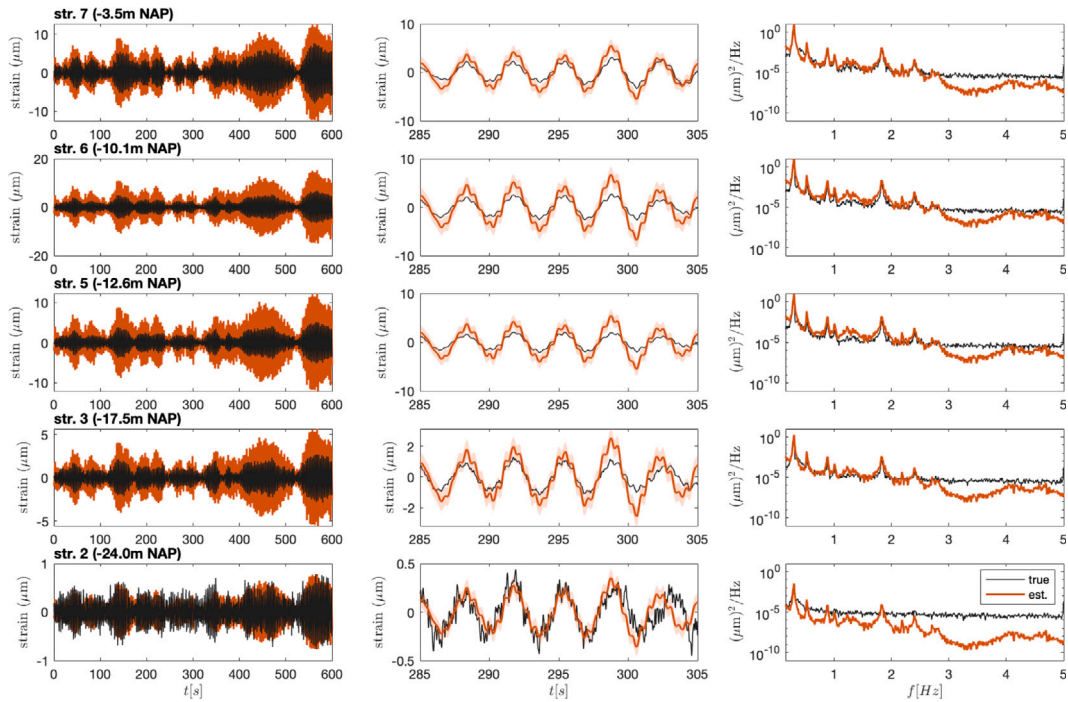


Fig. 9. Strain estimation for a record from the idling state (D7), shown with the full time series (left column), zoomed-in view of the time series (center column), and power spectral density (right column) for five strain channels.

also be indicative of mischaracterization of measurement noise; as shown in the sensitivity study in Appendix A.1, underestimation of measurement noise variance leads to an overestimation of low frequency content in the displacement response. These challenges are encountered in a similar study in [27], where a high-pass filter which removes frequencies below 0.25 Hz is applied to the predicted strain sequence before drawing comparison to the validation data.

Strain estimation results for all validation records is shown in Fig. 12, where the boxplots are computed from accuracy metrics of the strain estimate across eight records in each operating state. In general, high accuracy in the frequency domain estimate is reflected by a very high correlation coefficient (CC) near 1, while fair accuracy in the time-domain estimate is reflected by a mean relative error (MRE) generally falling in the range of 5%–10%. There is a reduction in accuracy in the run-up state due to larger variance in the wind and rotor speed (and therefore aerodynamic damping) among the records. Lower accuracy observed at strain gauge 2 may be a result of sensor noise in the particular strain gauge or extrapolation uncertainty, as the sensor is located at the furthest depth in the foundation where uncertainty in the soilbed properties is highest.

Samples of the modal force estimate from each of the operational states of the OWT are shown in Fig. 13. While it is not possible to obtain direct measurements of the loads for validation of the input estimate, the frequency spectrum of the prediction provides insights to the joint relationship between inputs and states defined in the GPLFM. First, it is clear that the magnitude of the input increases from the idling state to the production state, reflected by the difference in signal amplitude in the time series plots. Higher estimation uncertainty for higher modes is shown by the larger shaded area representing the ± 3 standard deviation bounds on the modal force prediction. Of particular interest is the frequency content estimated at the natural frequencies of the system, noted by dashed lines in the PSD plots. For the run-up state and production state, there is a visible dip in the frequency spectrum at the first natural frequency, which is an indicator of underestimation of the damping ratio in the first mode. In particular, when an underdamped model of the structural system predicts response amplitudes which exceed that of the measured acceleration response, the GPLFM compensates for the mismatch by reducing the power at that natural frequency in the frequency spectrum of the modal force. This effect is also observed in numerical experiments presented in Appendix A.1, where perturbations to modal damping ratios in a synthetic model causes artificial suppression of frequency content about modes where damping is misidentified. The input estimate is likely also affected by errors in mode shapes and natural frequencies of the system. Since congruence in the mode shapes identified from OMA and those of the mechanical model may only be evaluated at three points where accelerations are recorded, there is a high likelihood for deviation in the section of the mode shape profile falling below the mudline where strains are estimated. Moreover, modal properties in the run-up and production states shift under varying operational and environmental conditions [2]. It is observed in the numerical experiments that while the input estimate is particularly sensitive to errors in the modal properties, the strain estimate is left relatively unperturbed. This effect indicates that in the prediction of latent strain states, the inclusion of input estimation acts to accommodate error in modal properties by attributing its source to stochastic noise, reflected by adjustments to the predicted spectrum of external excitation to the system in order to fit the data provided. However, the input estimate has

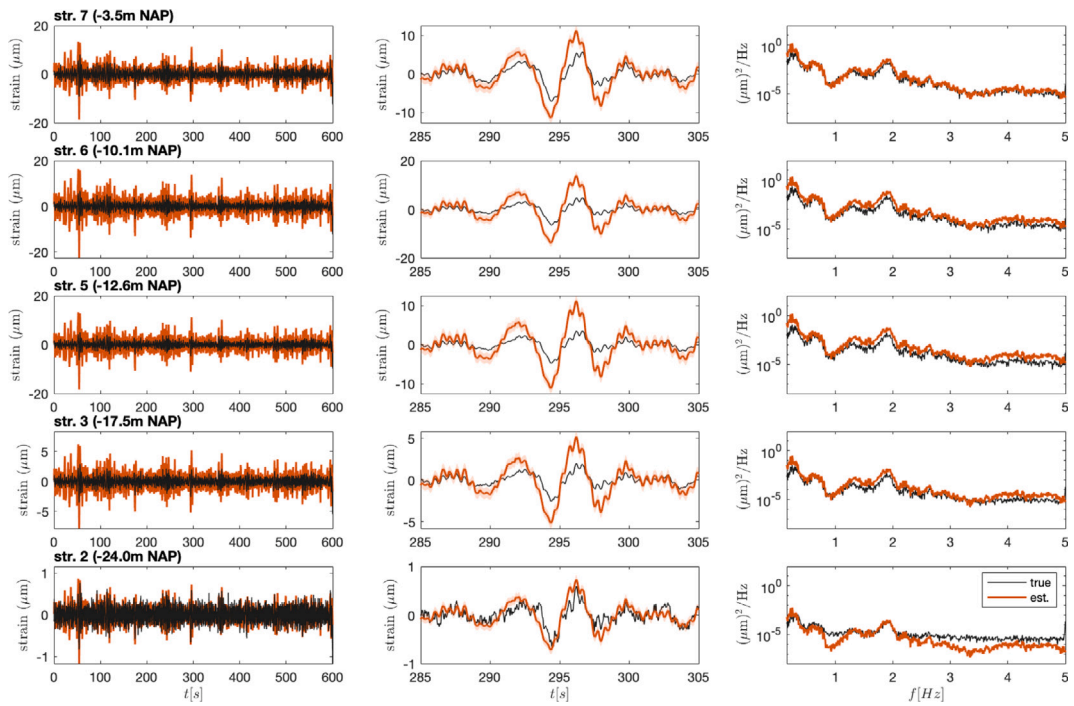


Fig. 10. Strain estimation for a record from the run-up state (R5), shown with the full time series (left column), zoomed-in view of the time series (center column), and power spectral density (right column) for five strain channels.

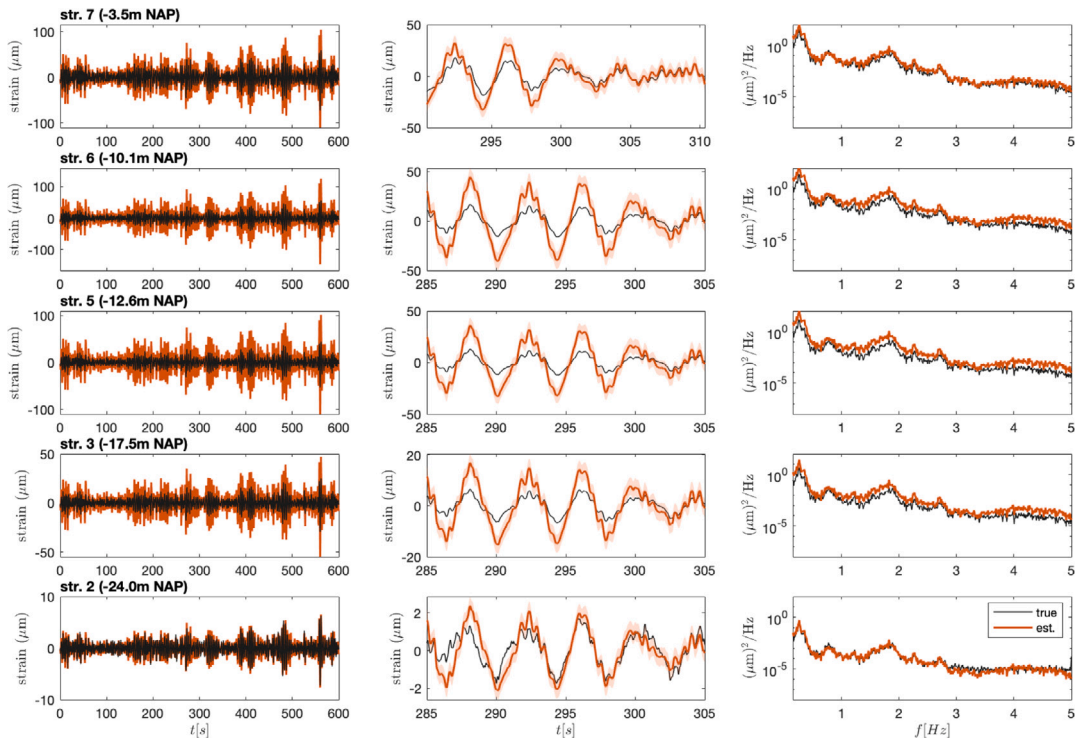


Fig. 11. Strain estimation for a record from the production state (P4), shown with the full time series (left column), zoomed-in view of the time series (center column), and power spectral density (right column) for five strain channels.

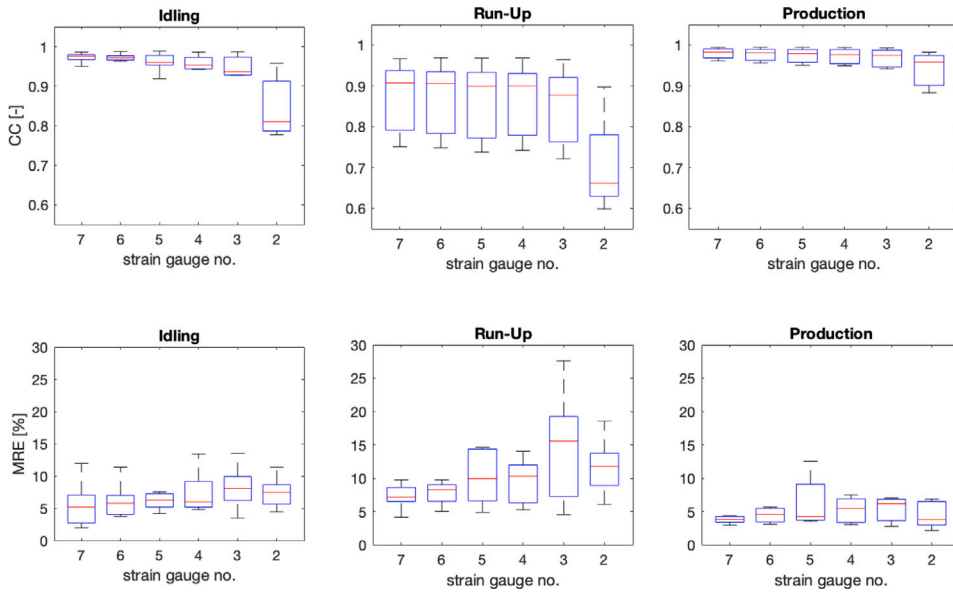


Fig. 12. Correlation coefficient (CC) and mean relative error (MRE) across the 6 sub-soil strain channels in the idling (left), run-up (center), and production (right) states.

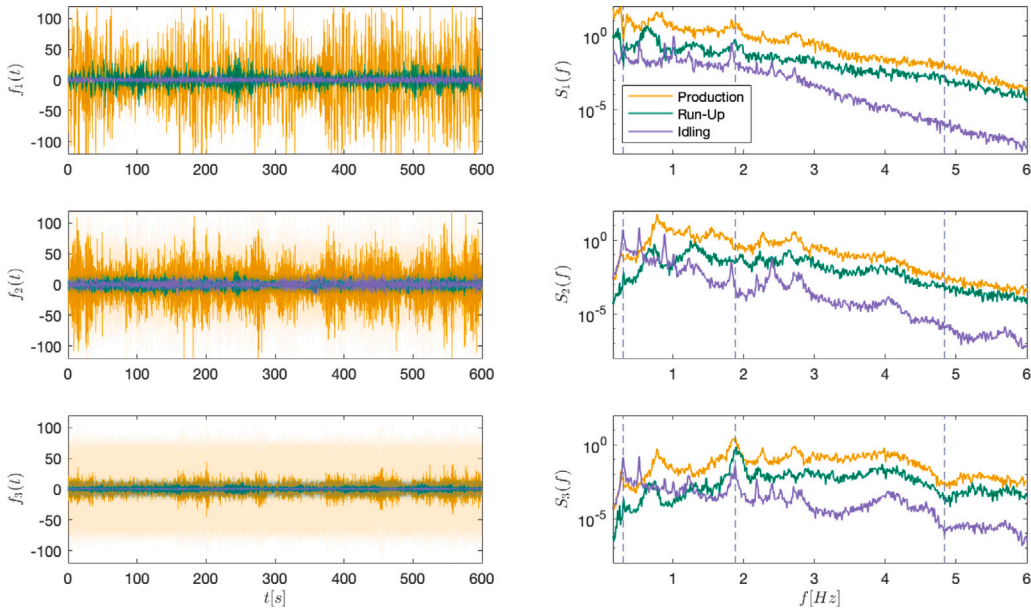


Fig. 13. Estimation of the first three modal forces for records in the idling, run-up, and production states, shown with the full time series on the left and power spectral density on the right.

limited efficacy in accounting for other sources of systematic model error, such as error in the assumed level of measurement noise or mismatch in the relationship defined between displacements and strains. Nevertheless, these results in response estimation provide strong evidence that the GPLFM overcomes several challenges with reconstructing subsoil strain response compared to similar studies conducted using dual-band MD&E for virtual sensing in [3] and the classical Kalman filter in [27], which were not able to achieve the same level of accuracy in the frequency domain estimate and phase synchronization of the time domain estimate.

4.3. Sensitivity to sensor configuration

In this section, the performance of the GPLFM is examined under different assumptions with respect to the availability and configuration of measurement data. The motivation is twofold: to determine an instrumentation scheme which aids in correcting for

the overestimation of strain amplitude, and to determine the minimum instrumentation scheme which provides sufficient evidence to characterize the strain response and inform future measurement campaigns for fatigue monitoring of similarly sized OWTs. Considering the practicality of in-situ implementation, we assume only sensors at accessible locations above the waterline are used for virtual sensing of subsoil response. Strain estimation with the GPLFM is conducted with three sensor configurations, labeled as:

- A3, using only the acceleration channel in the RNA (+95.0 m NAP);
- A3+A1, using two acceleration channels with one located in the RNA (+95.0 m NAP) and the other at +4.50 m NAP; and
- A3+S8, using the acceleration channel in the RNA (+95.0 m NAP) and the strain gauge ring located at +6.95 m NAP.

While the single sensor set-up of the A3 configuration results in the observability of only the first mode, it is included in the following comparison as a point of reference for the added value of the second sensor in the A3+A1 and A3+S8 configurations. Validation of the strain estimate with Record P3 for the three sensor configurations is given in Fig. 14. For brevity, only strain channels 2, 6, and 7 are shown. It is observed the A3 configuration is only capable of capturing the first modal response, with the accuracy at frequencies above the first natural frequency rapidly declining. The introduction of the second acceleration channel in A3+A1 allows for accurate estimation of higher frequency contributions to the response. However, A3 and A3+A1 lead to a comparable level of overestimation of the signal amplitude due to the inability of the model to overcome challenges with low frequency noise amplification using acceleration-only measurements. In contrast, the inclusion of a strain channel in A3+S8 leads to substantial improvements in capturing the strain amplitude across channels, in exchange for reduced accuracy in the estimation of non-governing frequencies falling between natural frequencies of the system. Compared to A3, the single strain channel in A3+S8 adds information on not only the high frequency response (such as the second modal response) in the tower section closer to the mudline, but also the amplitude of strains, thereby overcoming low frequency error.

The quality of the strain estimate for different sensor configurations may be compared using fatigue assessment measures. In this work, the S-N diagram, also referred to as the Woehler curve, is constructed using a rainflow counting algorithm which translates time series data into the number of cycles counted within defined stress ranges [54,59]. The S-N diagram of Fig. 15 shows the stress cycle counts of the true signal and the predicted signals at Str. 6, where the highest bending moments in the tower-foundation system are expected to occur, across four records from the production state. From this illustration, it is clear that configurations A3 and A3+A1 lead to an overestimation of the stress magnitude for lower cycle counts, which corresponds to low frequency oscillations. For the A3 case, the contribution of low-magnitude high frequency oscillations to fatigue stress accumulation is completely miscalculated. Fig. 15 affirms the A3+S8 configuration provides the best estimate of the stress-cycle curve. This result is further reflected in the damage equivalent load (DEL) measures of fatigue stress accumulation, where the deviation of the DEL of the prediction from that of the true strain response is 106.89% for A3, 98.43% for A3+A1, and 16.04% for A3+S8. In other words, the fatigue estimate from the acceleration-only configurations A3 and A3+A1 is on the order of double that of the true fatigue damage accumulation.

4.4. Sensitivity to model uncertainty

Previously, in Section 2.3.2, we showed that the mathematical construction of the GPLFM attributes noise sources to the unknown input. This section evaluates this effect using in-situ validation data, by imposing controlled model error in the mechanical model and evaluating the consequences on the strain estimation. In this case study, as well as in general for offshore wind turbines, one significant source of uncertainty is in properties of the foundation stiffness, which impacts mode shapes corresponding to the mechanical model. Whereas the stiffness profile $k_{eff}(z)$ is derived from soil characterization data which is customarily collected at wind farm sites, the subsoil strain data used to tune the magnitude of the profile γ in [57] is generally not available. Without such data, there is the potential for substantial error in the assumed soil stiffness magnitude. We examine both extremes of introducing error in the soil stiffness scaling γ by an order of magnitude above and below the baseline stiffness magnitude, with all other FE model parameters remaining constant. The following model variants representing different levels of model fidelity are compared:

- $\gamma = 0.078$: limiting case of underestimating the soil stiffness scaling by an order of magnitude
- $\gamma = 0.78$: baseline case of soil stiffness scaling to reflect the best fit to true in-situ conditions, determined in [57]
- $\gamma = 7.8$: limiting case of overestimating the soil stiffness scaling by an order of magnitude

The modal parameters of the three model variants are shown in Fig. 16 and reported in Table 2, where the relative frequency error e_f and MAC values are with respect to the modal parameters identified through OMA. It is observed that the lowest two mode shapes in each of the FA and SS directions, which dominate the OWT response, are relatively unaffected by scale differences in the soil stiffness profile. A larger effect is observed in the third SS mode. In terms of both e_f and MAC, the error is greater in the case of underestimation of γ than in the case of overestimation.

Fig. 17 shows the strain estimate obtained with each model variant using the acceleration-only measurement setup. The difference in the accuracy of the strain estimate varies with the depth of the strain location: there is comparatively lower change in the strain estimate at higher elevations (Str. 6 and 7) than at the lowest elevation (Str. 2), where underestimation of stiffness leads to underdamped response and the overestimation of stiffness leads to overdamped response. The difference is non-uniform across channels due to the fact that there is significantly lower stiffness at the upper section of the profile, where $k_{eff}(z)$ is on the order of 10^7 N/m³, than at lower sections of the profile, where $k_{eff}(z)$ is on the order of 10^{11} N/m³ (see Fig. 4). The scale difference therefore leads to a greater change in subsoil properties at the greatest depths. While the scaling of the predicted signal is affected

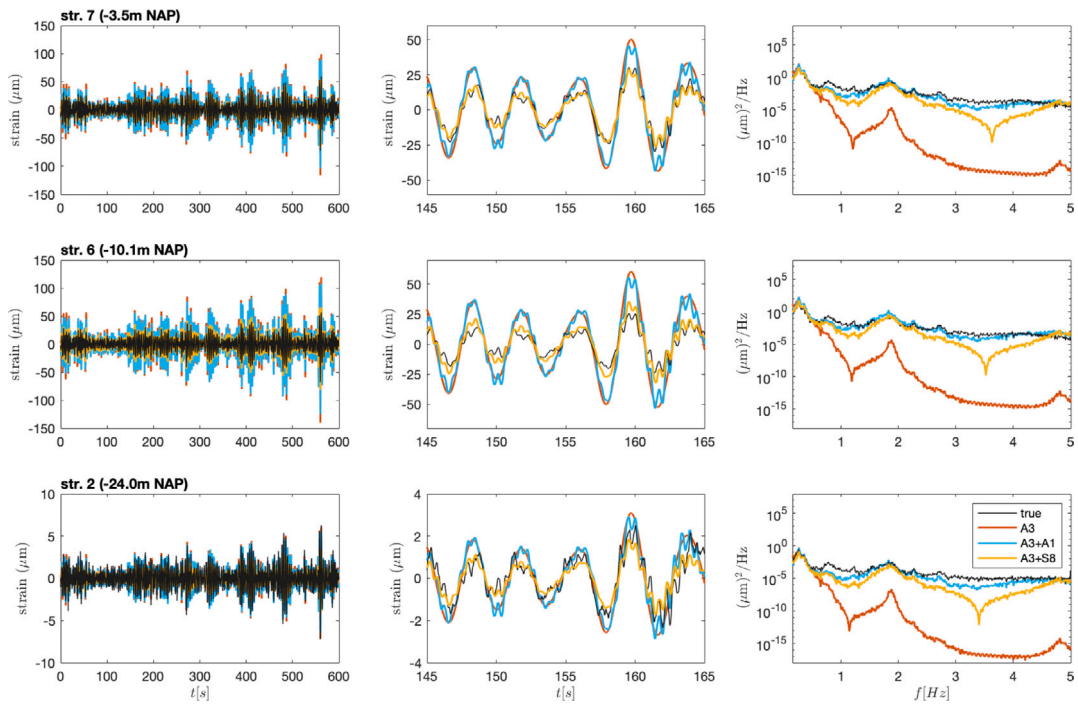


Fig. 14. Strain estimation with different sensor configurations for record P3.




Fig. 15. S-N diagram with different sensor configurations across 4 samples in production state.

by the model error, the shape of the distribution of power over frequencies remains relatively unchanged up to the second natural frequency across all channels.

The difference in prediction by the model variants is more apparent in the modal force estimate, shown in Fig. 18. Here, the discrepancy in the natural frequencies between the model variants is illustrated with dashed lines in the frequency spectrum. What is particularly notable is the dip in the frequency spectrum at the j th natural frequency for the j th modal force, where the natural frequencies differ for each model variant. These dips reflect adjustments to the frequency contributions of the input by the GPLFM in order to compensate for error between the dominant frequencies of the measured response with that of the predicted response

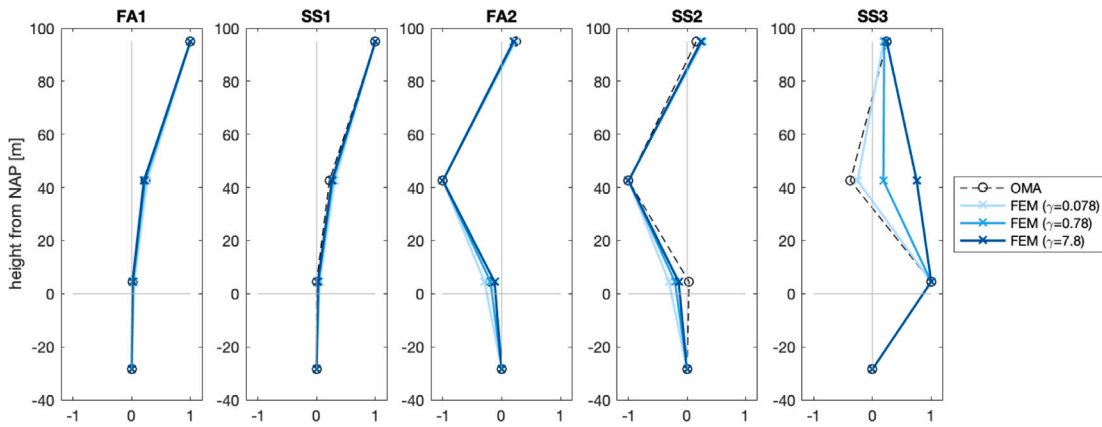


Fig. 16. Modal parameters of model variants.

Table 2

Modal parameters from each model variant, where accuracy metrics e_f and MAC are computed with respect to the modal parameters identified by OMA.

Mode	$\gamma = 0.078$			$\gamma = 0.78$			$\gamma = 7.8$		
	f_{fem} [Hz]	e_f [%]	MAC	f_{fem} [Hz]	e_f [%]	MAC	f_{fem}	e_f [%]	MAC
FA1	0.279	-5.82	0.999	0.296	0.01	0.999	0.307	3.55	0.999
SS1	0.269	6.70	0.990	0.289	-0.12	0.996	0.300	3.74	0.998
FA2	1.636	-12.47	0.991	1.874	0.23	0.999	2.041	9.18	0.994
SS2	1.665	-11.66	0.886	1.883	-0.09	0.939	2.025	7.44	0.964
SS3	4.486	-1.69	0.987	5.193	13.79	0.743	5.745	25.88	0.310

Table 3

Damage equivalent load (DEL) values corresponding to Str. 6 of each model variant for two sensor configurations, A1+A2+A3 and A3+S8. The relative error is computed with respect to the true value of DEL.

	A1+A2+A3		A3+S8	
	DEL	Error [%]	DEL	Error [%]
True	2.946	-	2.946	-
Gamma = 0.078	7.131	142.1%	4.295	45.8%
Gamma = 0.78	6.652	125.8%	3.418	16.0%
Gamma = 7.8	4.994	69.5%	2.227	-24.4%

by the perturbed model. For instance, where the GPLFM expects greater presence of the perturbed second natural frequency of the model than what exists in the measurement data, the contribution of that frequency in the corresponding input estimate is reduced accordingly. In this manner, error in modal properties is accounted for through the stochastic input estimate, thus retaining the relatively robust accuracy observed in the strain estimate.

To evaluate the effect of model fidelity as a function of the instrumentation scheme, Fig. 19 shows the S-N diagram corresponding to four records from the production state, produced from the estimate of Str. 6, for two sensor configurations: the original scheme of using all three acceleration measurement channels (A1+A2+A3) and the best performing scheme from the investigation in Section 4.3 (A3+S8). For both instrumentation schemes, the three model variants result in relatively comparable stress-cycle curves, indicating relative invariance to perturbation in the model. As before, A3+S8 provides the most robust reproduction of the true stress-cycle curve while the acceleration-only scheme in A1+A2+A3 universally overestimates the stress magnitude of low-frequency oscillations. The improvement gained from the use of A3+S8 over the acceleration-only scheme is reflected in the DEL values reported in Table 3, where A3+S8 results in both lower error and lower variance in the DEL estimate across levels of model perturbation.

5. Conclusions

This work is among the first studies to provide in-situ validation of the GPLFM as a virtual sensing technique. In contrast to other Kalman filter-based techniques to virtual sensing, the GPLFM characterizes the unknown input as a Gaussian process which can describe a greater breadth of time-varying behavior than the standard assumption of white noise input. The joint posterior distribution of latent inputs and states is determined by Gaussian process regression of acceleration measurements, which is performed sequentially in augmented state-space using Kalman filtering/RTS smoothing. Rather than using heuristic or ad hoc tuning procedures, the process noise covariance in the dynamics model is solved for in a data-driven fashion, in which the structure

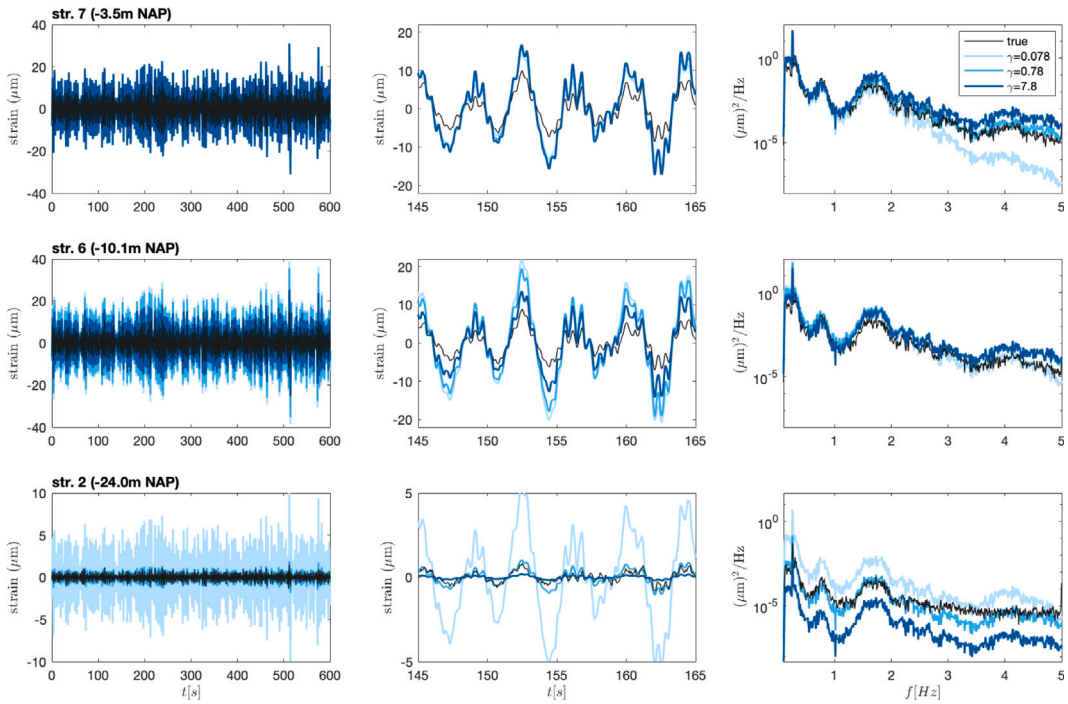


Fig. 17. Strain estimation for varying levels of model fidelity using all acceleration measurement channels for record P2.

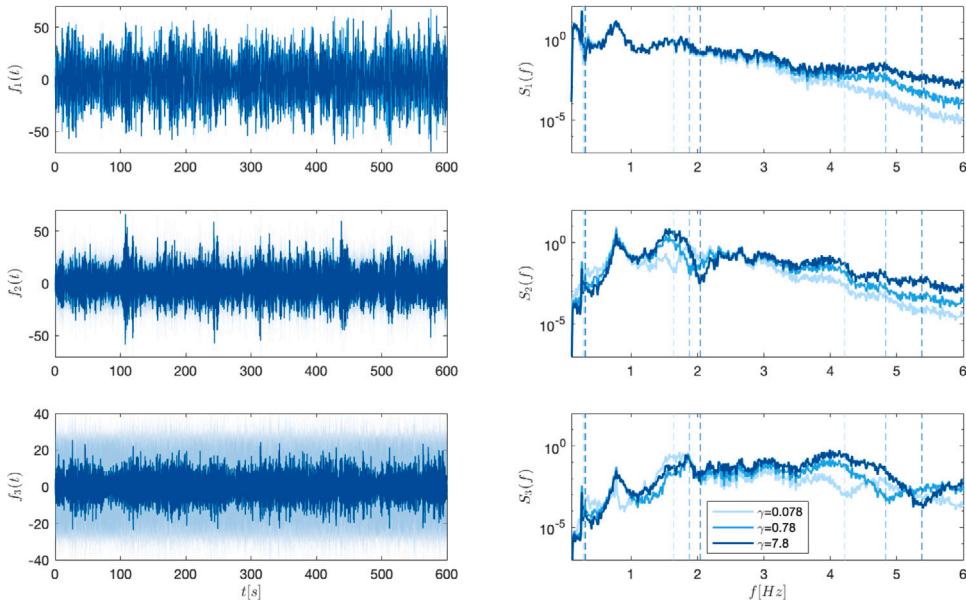


Fig. 18. Modal force estimation for varying levels of model fidelity using all acceleration measurement channels for record P2. The natural frequencies of each model variant are noted by dashed lines in the frequency spectrum.

of the GPLFM is utilized to define the joint relationship between input and state variables and the kernel hyperparameters governing the process noise covariance matrix are fit to acceleration data prior to filtering. Compared to other Kalman filter-based techniques for joint input-state estimation, the GPLFM has several notable theoretical advantages, in enforcing a characterization of covariance between inputs and states consistent with stochastic differential equations theory, as well as practical advantages, in providing an efficient and descriptive parameterization of statistical model parameters.

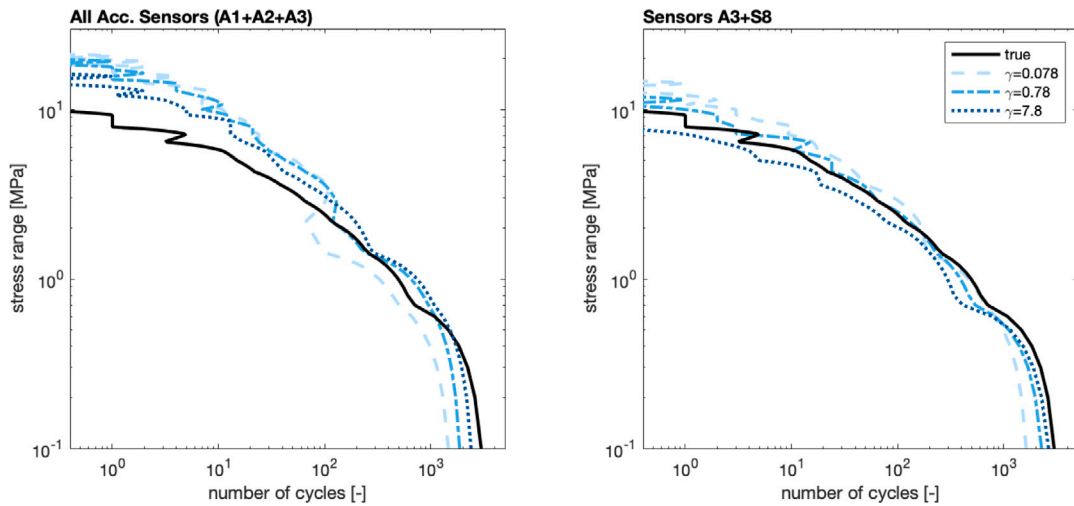


Fig. 19. S-N diagram for varying levels of model fidelity for records from the production state, for sensor configuration A1+A2+A3 on the left and sensor configuration A3+S8 on the right.

In this contribution, the GPLFM is applied to the fatigue monitoring of offshore wind turbines. The validation data is taken from a monitoring campaign of an operating nearshore wind turbine from the Westermeerwind Park in the Netherlands, where six strain gauge rings embedded along the profile of the monopile foundation provide subsoil strain measurements for validation. This dataset presents the opportunity to evaluate the GPLFM for its ability to address the challenge in fatigue monitoring of OWTs with high model uncertainty in the soil-structure system of the monopile foundation. First using acceleration-only data from an offshore wind turbine, the GPLFM implemented with a Matérn covariance kernel is able to reconstruct the frequency spectrum of the subsoil strain response with excellent accuracy across operating conditions of the OWT. In the time domain, overestimation of the strain amplitude is attributed to a small level of low frequency noise amplification. In order to test the sensitivity of the GPLFM to the level of model fidelity, perturbation to the scale of the foundation stiffness model is introduced. It is demonstrated that with induced errors in the natural frequencies and mode shapes of the mechanical model, the strain estimate by the GPLFM remains relatively insensitive, particularly at depths of interest near the mudline where the highest bending moments are expected to occur. This result illustrates that robustness in the strain estimate is attained by attributing sources of error in modal properties to the stochastic input estimate. Moreover, it provides evidence that the input estimate may be used as a type of heuristic indicator of error in modal properties, particularly of error in damping levels in the run-up and production states of OWTs. Through a study of the influence of the sensor configuration on the performance of the GPLFM, it is evident that estimation bias may be systematically accounted for through the inclusion of a strain channel in the measurement data. Not only does the strain data aid in correcting the predicted signal amplitude along all depths of the foundation, but it may be used to supply important information on the quasi-static behavior of the strain response which cannot be captured by high-frequency acceleration measurements alone [4,5]. The best performing instrumentation scheme minimally consists of a single acceleration channel and single strain channel, where the GPLFM achieves a relative error of 16% in the DEL, mean correlation coefficient of 0.968, and mean relative error of 2.28% in the estimated subsoil strain time series across records from the production state of the OWT. The use of a heterogeneous dataset consisting of acceleration and strain measurements is deemed sufficient for curtailing model uncertainty, allowing the GPLFM to attain strong robustness and transferability properties in estimating the subsoil strain response to variable operational and environmental loads.

We believe it is possible to further improve the performance of the GPLFM for subsoil strain estimation by two measures: deliberate estimation of quasi-static response, and explicit consideration of the dynamic evolution of modal properties of the structure. Regarding the former, other than by inclusion of strain data, low frequency error in the predicted strains can be mitigated by estimating measurement noise covariance with greater accuracy or by utilizing accelerometers with a lower noise baseline. Regarding the latter, this work demonstrates that the benefits of the GPLFM in providing a more flexible characterization of the input in the context of filtering are inhibited when this characterization is not extended to the context of modal identification. In other words, traditional OMA techniques, whose formulation relies on the assumption of white noise input, are not appropriate for the modal identification of offshore wind turbines in the run-up and production states, where environmental loads and rotor blade dynamics introduce harmonic components to the excitation [49]. Therefore, applications of the GPLFM for virtual sensing warrants additional developments to online modal identification techniques which allow for real-time modal parameter estimation accounting for non-white excitation in the system. To this end, the authors recommend further evaluation of proposed adjustments and alternatives to damping identification techniques for LTP systems which are reviewed in [49]. The reader is also referred to [34] where modal identification is treated with a Bayesian approach, in which inference of system parameters such as mass, stiffness, and damping is performed jointly with the estimation of the unknown ambient excitation by means of the GPLFM. This approach enables posterior state inference and modal identification to be performed jointly with a fully cohesive model of the input and may be applied across a greater range of contexts compared to traditional OMA methods.

There is potential for further improvements to the strain estimation through model choices in the GPLFM. In particular, a major assumption of uncorrelated modal forces may be overcome by explicitly inferring a covariance parameter between modal force sequences in the state-space matrices of Eq. (26). Secondly, model uncertainty could be addressed through the definition of a process noise variance parameter on response states (displacements and velocities) in the spectral density of the white noise term of the GPLFM. These additional model parameters, together with the hyperparameters of the GP kernel, constitute an augmented vector of hyperparameters which may be inferred with a maximum likelihood estimation approach. In theory, these improvements would lead to more accurately estimated statistics of the latent states, though at the cost of a higher dimensional optimization problem. Finally, there are opportunities to investigate the role of the choice of GP covariance kernel in the GPLFM for response data which is driven by inputs of varying degrees of periodicity and temporal correlation. Beyond model choices, the results of this work promote further studies on efficient and practical sensor placement schemes in monopile-based offshore wind turbines, to select the most informative data for the purpose of virtual sensing-based fatigue monitoring.

CRedit authorship contribution statement

Joanna Zou: Methodology, Data curation, Software, Formal analysis, Investigation, Visualization, Writing – original draft. **Eliz-Mari Lourens:** Conceptualization, Supervision, Writing – review & editing. **Alice Cicirello:** Supervision, Writing – review & editing.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Joanna Zou reports financial support was provided by Stichting Fulbright Commission of the Netherlands. Joanna Zou reports equipment, drugs, or supplies was provided by Siemens Gamesa Renewable Energy.

Data availability

The authors do not have permission to share data

Acknowledgments

This research was funded by a Fulbright Student Research Grant awarded by the Stichting Fulbright Commission of the Netherlands. The data used in this paper is obtained from the DISSTINCT project, a collaboration between Siemens Wind Power, Fugro, Delft University of Technology, and DNV-GL which was in part funded by the Dutch government (project number TKIW02001). The authors gratefully acknowledge the people at Siemens-Gamesa Renewable Energy for providing the data, Joaquin Bilbao-Neiva for sharing code for the implementation of the SSI-Cov/OPTICS algorithm, and Mikk Laanes for sharing data preprocessing code.

Appendix

A.1. Effect of model error on an illustrative example

A synthetic model is used to illustrate the quality of estimation by the GPLFM when model error is introduced, as well as to show the efficacy of the measurement noise covariance tuning procedure described in Section 2.3.2. The synthetic model is a 10-dof vertical mass-spring system with fixed-free end conditions representing a cantilever structure. The system is defined by ten point masses of $m = 200$ kg and a stiffness of $k = 5 \times 10^3$ N/m. Degrees of freedom are horizontal and aligned with the direction of forcing. Mass proportional damping with modal damping ratios of $\xi_j = 0.02$ for all degrees of freedom is assumed. An artificial input is defined as the combination of a sinusoidal force with an amplitude of 100 N and a frequency of 0.2 Hz, applied at the top (10th) level of the model, and a Gaussian process with a Matérn $\nu = 5/2$ kernel with parameters $\alpha = 50, l_s = 1$, applied to all 10 levels of the model. In this example, the number of degrees of freedom are $n_u = 10$ and the number of forces are $n_p = 2$, with the locations of the forces denoted by the selection matrix $S_p \in \mathbb{R}^{n_u \times n_p}$. The synthetic system is constructed to loosely represent environmental and operational load conditions on a wind turbine support structure, where the sinusoidal force represents the harmonic frequency of vibration induced by the rotor motion of the turbine and the GP force represents ambient wind excitation distributed along the height of the tower. Fig. A.20 shows the artificial input acting at the top level of the structure, where the harmonic frequency at 0.2 Hz is clearly visible in the frequency spectrum. The ensemble distribution of the input, shown by the histogram in the top right plot, illustrates how the input is non-Gaussian due to the contribution of the sinusoidal force.

To set up the filtering problem, the ground truth response of the cantilever structure to the artificial load is simulated using the Newmark average acceleration method [60] for numerical integration with a time step of 0.001 s. The simulation model used for generating the ground truth is a full-order model of the system, whereas the prediction model is a modally reduced-order formulation of the GPLFM of the form in Eq. (28) retaining the first $n_m = 3$ modes of the system, such that modal states $\mathbf{x}(t) \in \mathbb{R}^{2n_m}$, modal forces $\mathbf{f}(t) \in \mathbb{R}^{n_m}$, and its higher derivatives up to order β constitute the augmented state vector $\mathbf{z}^a(t)$. For more details on notation, see Section 2. The measurement data consists of the acceleration response at all 10 degrees of freedom, observed at a sampling




Fig. A.20. The ground truth input applied to the top level of the 10-dof cantilever structure, shown in terms of the time series (top left), histogram of realized values (top right), and frequency spectrum (bottom).

frequency of 100 Hz and colored with white noise with covariance $\mathbf{R} = \sigma_R \mathbf{I}$, where $\sigma_R = 10^{-2}$. It is assumed that no process noise exists in the system, such that response states are known exactly, and that measurement noise is known to be independent across channels of measurement data with variance σ_R .

Numerical experiments are conducted to examine the effect of model error. In each experiment, filtering is performed without perturbation, where the prediction model is directly a reduced-order formulation of the simulation model which generated the ground truth. Additionally, filtering is performed where the prediction model is a reduced-order model with perturbation in one of its properties, in order to illustrate the ability of the GPLFM to accommodate for misidentified model parameters in its prediction of latent states. Each version of the prediction model is fit independently using the method of GP hyperparameter inference described in Section 2.3.1; in the following studies, we assess the extent to which the GP hyperparameters corresponding to the Matérn $\nu = 5/2$ kernel for the prediction model with perturbation differ from that without perturbation. In this synthetic example, the true value of the latent states over time, including displacements, velocities, and modal components of the force, are available for comparison with the prediction by the GPLFM.

Error in modal damping ratios

The prediction model without perturbation assumes the modal damping ratios ξ_j are known exactly to be equal to 0.02. Two cases of model perturbation are considered, where (a) the first three modal damping ratios are underestimated by an order of magnitude to be $\xi_j = 0.002$ and (b) the first three modal damping ratios are overestimated by an order of magnitude to be $\xi_j = 0.2$.

Fig. A.21 shows the estimated displacement, velocity, and acceleration response at level 5 of the 10-dof cantilever model for each of the prediction models. We observe good congruence between the true response and predicted response in the time domain for all models, with some error in the predicted displacement time series. Due to the use of model order reduction to 3 modes in the prediction models, the frequency content of the estimate loses accuracy for frequencies above 0.6 Hz, beyond the third natural frequency of the system. In general, there is negligible effect of error in the modal damping ratio on the response states estimated by the GPLFM. This can in part be attributed to compensation in the modal force estimate, shown in Fig. A.22, where the estimated first, second, and third modal forces are compared to the modal components of the true force. In the frequency spectrum, a visible drop in the frequency about the j th natural frequency of the system is observed in the j th modal force estimate by the model with underestimated damping; similarly, a visible inflation of frequency is observed for the model with overestimated damping. The effect illustrates how the input estimate is utilized to account for mismatch between the system model and the measured response: in the case of underdamping, where greater frequency content about the natural frequencies would be expected according to the system model than what is present in the measurement data, the input estimate is adjusted such that the contribution of those particular frequencies from the input are reduced. This experiment lends evidence that Bayesian joint input-state estimators have an advantage over state-only estimators in that the inclusion of input estimation can mitigate certain sources of model error and aid the primary objective of achieving accurate state estimation. It is notable that the presence of the harmonic frequency between the first and second natural frequencies of the structure does not interfere with the accuracy of either the response estimate or the input estimate.

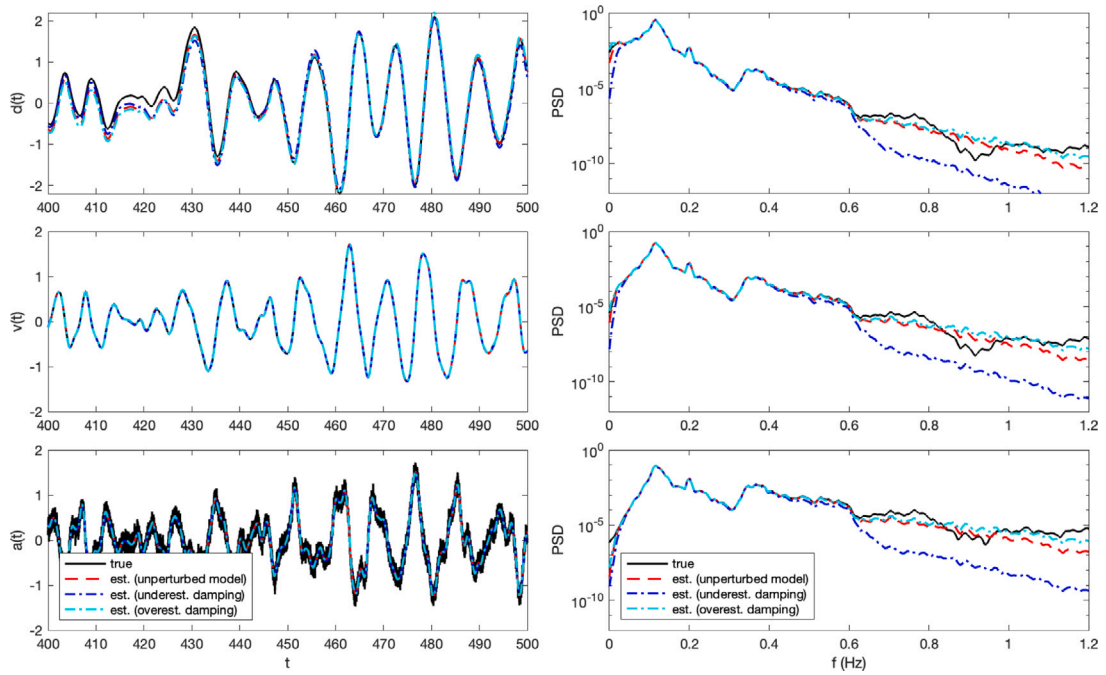


Fig. A.21. Estimates by the GPLFM of displacement (top row), velocity (center row), and acceleration (bottom row) states at level 5 of the 10-dof cantilever model with no perturbation (red), underestimated modal damping (blue), and overestimated modal damping (cyan), compared to the true response (black).

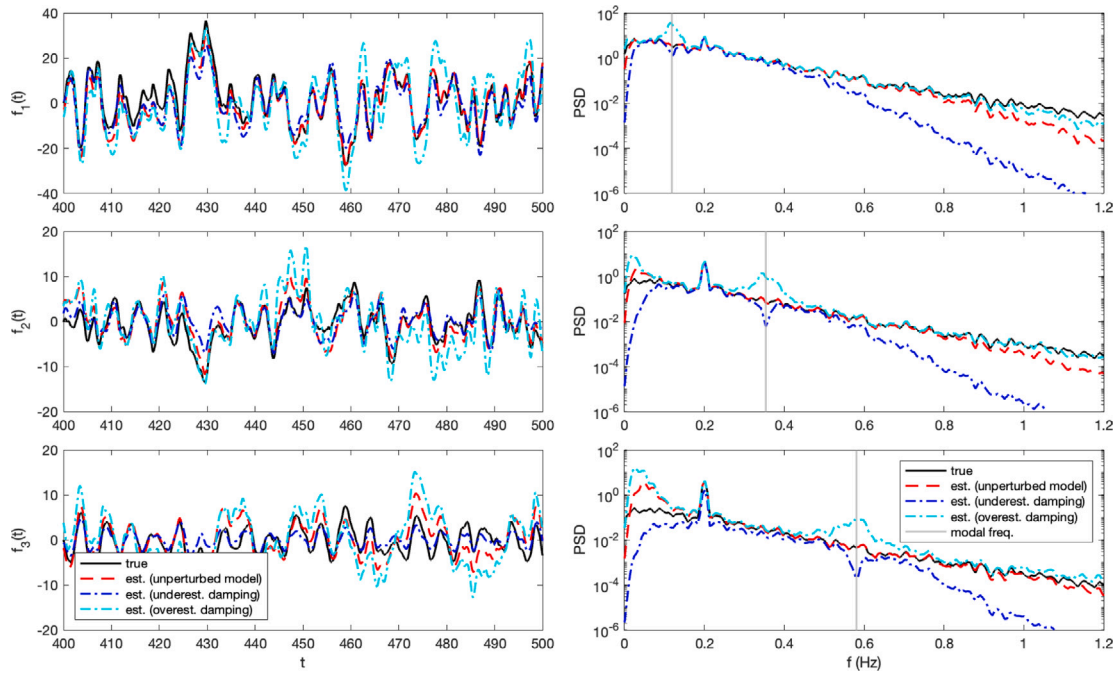


Fig. A.22. Estimates by the GPLFM of the first three modal forces (top, center, bottom rows respectively) using the 10-dof cantilever model with no perturbation (red), underestimated modal damping (blue), and overestimated modal damping (cyan), compared to the true modal forces (black).

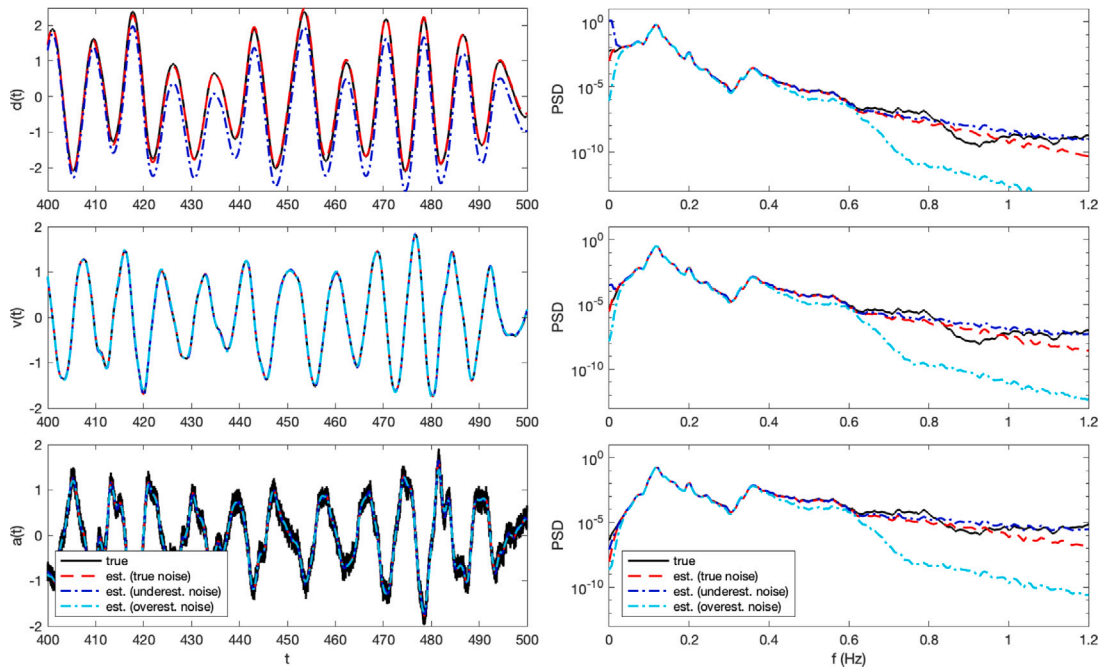


Fig. A.23. Estimates by the GPLFM of displacement (top row), velocity (center row), and acceleration (bottom row) states at level 5 of the 10-dof cantilever model with known measurement noise variance (red), underestimated measurement noise variance (blue), and overestimated measurement noise variance (cyan), compared to the true response (black).

Error in the measurement noise variance

In the prediction model without perturbation, the variance of measurement noise in each channel is known to be $\sigma_R = 10^{-2}$. Two cases of model perturbation are considered, where (a) the measurement noise variance is overestimated by 3 orders of magnitude to be $\sigma_R = 10^1$ and (b) the measurement noise variance is underestimated by 3 orders of magnitude to be $\sigma_R = 10^{-5}$. In this experiment, erroneous values of the measurement noise covariance matrix are provided as input to the Kalman filter with no tuning performed.

In Fig. A.23, it is observed that while high accuracy is retained in the predicted velocity and acceleration response, the introduction of measurement noise results in reduced accuracy in the estimation of low frequency components of the response states, to the greatest degree in the displacement response. This same low frequency drift is observed in the modal force estimate in Fig. A.24, with a larger effect in the second and third modal forces. Measurement noise affects the frequency content of the measurement data by introducing greater contributions of higher frequencies. In the case where measurement noise variance is overestimated, the predicted signal has reduced contributions of high frequencies which are expected by the prediction model but not present in the measured signal. In the case where measurement noise variance is underestimated, the predicted signal captures this presence of high frequencies in the measured signal; however, it is countered by inflation of low frequencies which the prediction model expects to observe in the absence of significant measurement noise.

An additional experiment is conducted to demonstrate the efficacy of the measurement noise covariance tuning procedure in Section 2.3.2, which refines the measurement noise covariance matrix \mathbf{R} in the prediction model using the residual sequence between the predicted and measured acceleration response over k iterations. It is observed that the sequence $\{\mathbf{R}_k\}$ converges to the true value of \mathbf{R} over a number of iterations on the order of $k = 3$. While this convergence is observed regardless of the initial value, this experiment presents the case where the measurement noise covariance is initialized at $\mathbf{R}_0 = \sigma_R \mathbf{I}, \sigma_R = 10^{-5}$. Using the tuning procedure, prediction error is effectively eliminated once the order of magnitude of the measurement noise variance is correctly identified, as shown in the response estimation in Fig. A.25 and modal force estimation in Fig. A.26.

Error in the modal frequencies and mode shapes

The prediction model without perturbation assumes modal properties are known exactly, where the first three natural frequencies of the system are 0.12 Hz, 0.35 Hz, and 0.58 Hz. One case of model perturbation is considered, where the prediction model has reduced stiffness of $0.5k$ and a concentrated mass of 2000 kg at the top of the tower, thus having modal properties which deviate from that of the simulation model. The first three natural frequencies of the perturbed model are 0.05 Hz, 0.19 Hz, and 0.36 Hz. The mode shapes of the two prediction models are shown in Fig. A.27.

In Fig. A.28, it is observed that even with error in the modal properties of the prediction model, the estimation of response states is highly accurate with respect to the true response and on par with the estimate using the error-free prediction model. Fig. A.29

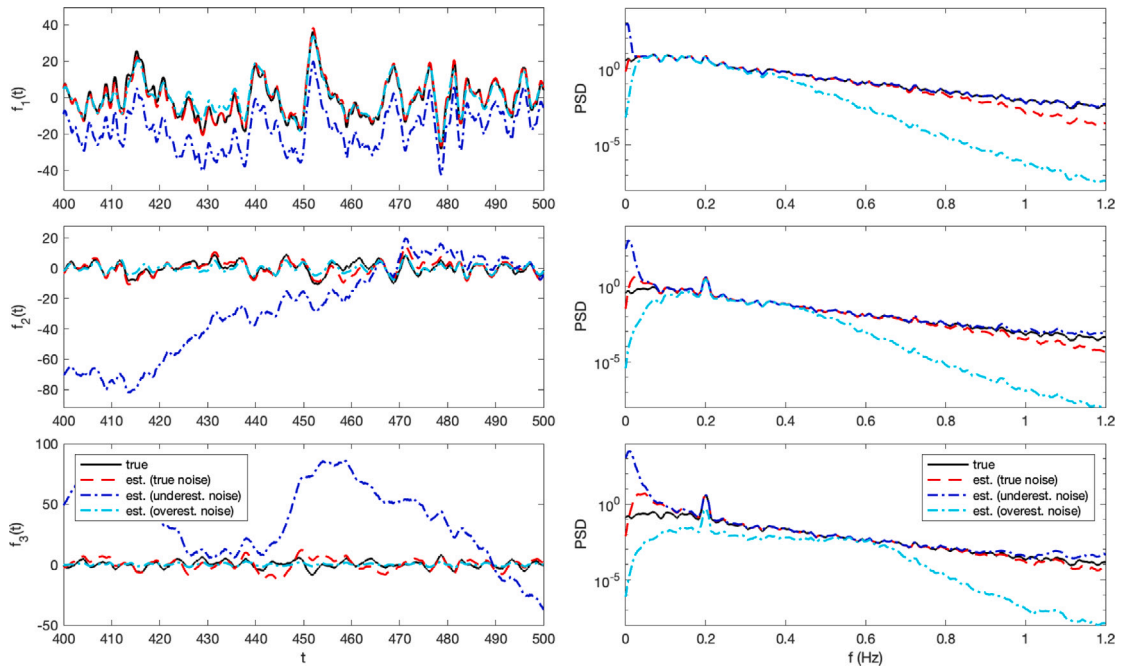


Fig. A.24. Estimates by the GPLFM of the first three modal forces (top, center, bottom rows respectively) using the 10-dof cantilever model with known measurement noise variance (red), underestimated measurement noise variance (blue), and overestimated measurement noise variance (cyan), compared to the true modal forces (black).




Fig. A.25. Estimates by the GPLFM of displacement (top row), velocity (center row), and acceleration (bottom row) states at level 5 of the 10-dof cantilever model with known measurement noise variance (red) and with measurement noise covariance tuning using the procedure in Section 2.3.2 (blue), compared to the true response (black).




Fig. A.26. Estimates by the GPLFM of the first three modal forces (top, center, bottom rows respectively) using the 10-dof cantilever model with known measurement noise variance (red) and with measurement noise covariance tuning using the procedure in Section 2.3.2 (blue), compared to the true response (black).




Fig. A.27. First three mode shapes of the 10-dof cantilever model without perturbation (black) and with perturbation (blue) from the modal properties of the simulation model.

shows that this high accuracy in the response states is attained by a corresponding adjustment to the modal force estimate. In particular, we draw attention to the frequency content of the modal forces about the natural frequencies of both the perturbed and unperturbed prediction models. In the first modal force, the frequency contribution of the first natural frequency of the perturbed model is suppressed, while that of the unperturbed model is inflated. A similar pattern is observed in the second and third modal forces with respect to the second and third natural frequencies, respectively. Even in the second modal force, where the second natural frequency of the unperturbed model nearly coincides with the nonstructural harmonic frequency, the power of frequency in the vicinity of the structural frequency is reduced. This effect illustrates that the input estimate accounts for error in modal properties of the system by identifying frequencies which are dominant in the measurement data as well as frequencies which are less prevalent in the measurement data than what would be expected, such as frequency bands about the natural frequencies of the prediction model.

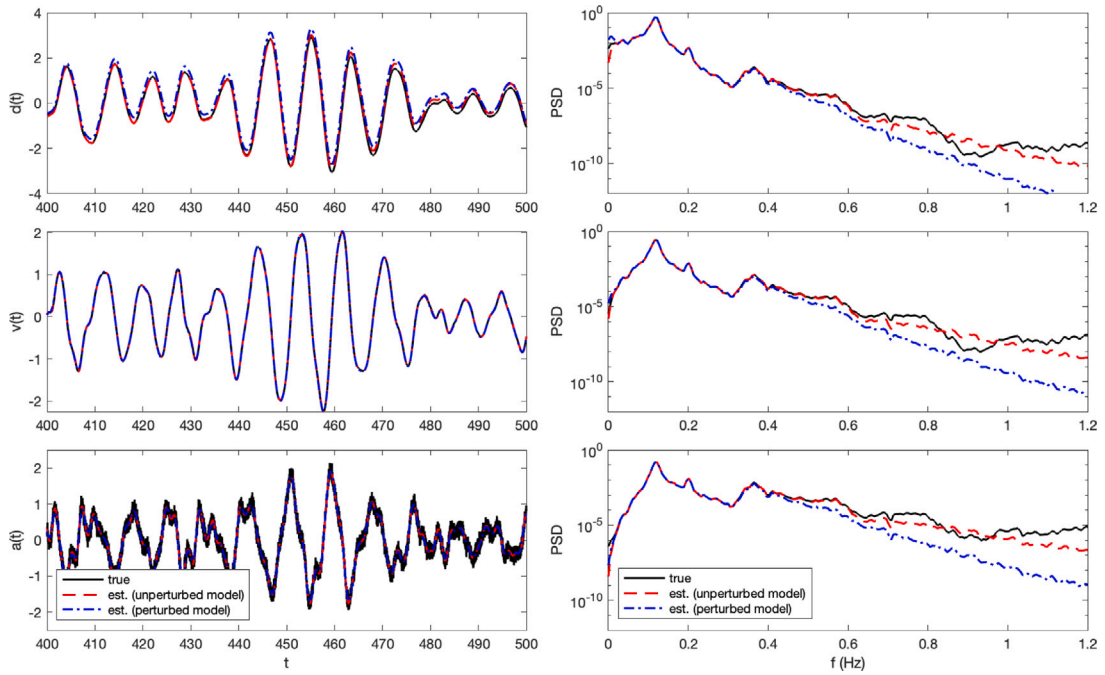


Fig. A.28. Estimates by the GPLFM of displacement (top row), velocity (center row), and acceleration (bottom row) states at level 5 of the 10-dof cantilever model with no perturbation (red) and with perturbation from the true simulation model (blue), compared to the true response (black).

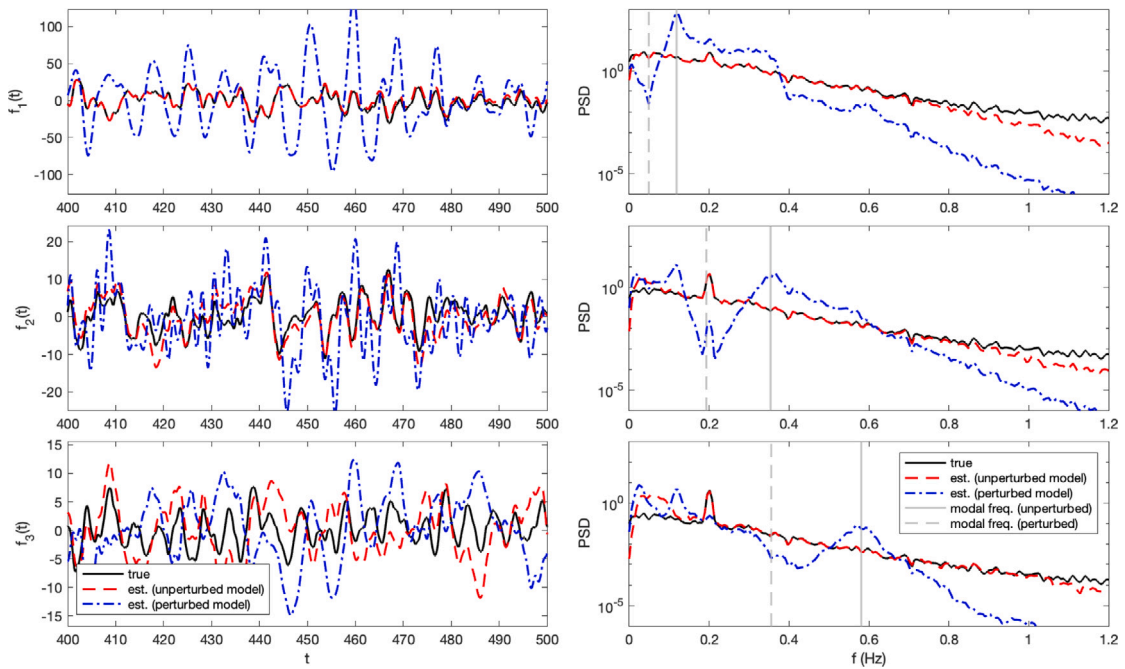


Fig. A.29. Estimates by the GPLFM of the first three modal forces (top, center, bottom rows respectively) using the 10-dof cantilever model with no perturbation (red) and with perturbation from the true simulation model (blue), compared to the true response (black).

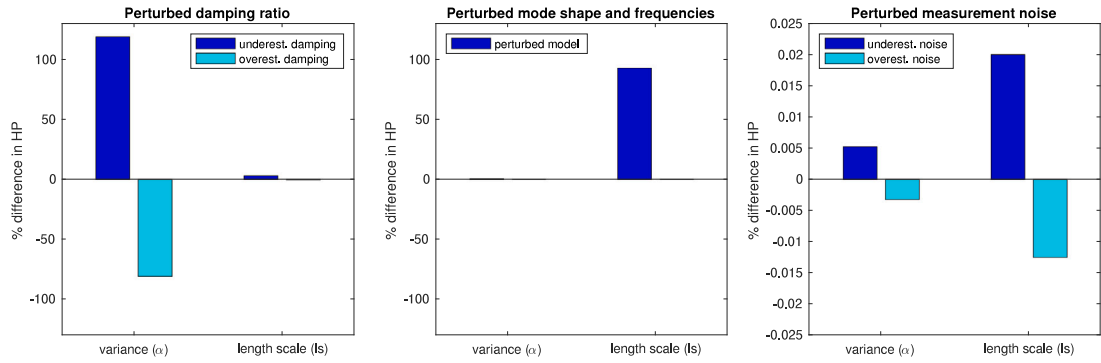


Fig. A.30. Percent difference in the GPLFM hyperparameters corresponding to a Matérn kernel ($\nu = 5/2$) with respect to the value of hyperparameters computed with the unperturbed model. The left plot shows percent change in the two hyperparameters from underestimation and overestimation of the first three modal damping ratios. The right plot shows percent change in the two hyperparameters from underestimation and overestimation of the measurement noise variance.

Changes in GP kernel hyperparameters in response to model error

While this interchange between the state estimate and input estimate is characteristic of the general class of Bayesian joint input-state estimators, the GPLFM can provide a concise parametric description of how the model adjusts to account for mismatch with measured data. For the prediction model both with and without perturbation, the GPLFM is fit by means of the parameter inference procedure in Section 2.3.1, where the parameters in this instance are the GP kernel hyperparameters, $\theta = [\alpha, l_s]$. Fig. A.30 shows the percent change in the parameters θ fit for the perturbed model with respect to the parameters fit for the error-free model for each of the conducted experiments. When error in the modal damping ratio is introduced, a large change in the variance parameter α is observed - a nearly 100% increase in α in the case of overestimated damping and a decrease of a similar magnitude in the case of underestimated damping. In contrast, very small change is observed in the length scale parameter l_s . The opposite is the case in the experiment where error in the mode shapes and natural frequencies is introduced, where large change in the length scale parameter l_s and negligible change in the variance parameter α are observed. The mechanism by which error in modal properties are compensated for in the GPLFM is reflected in the change in GP kernel hyperparameters, where error in damping levels primarily affects the amplitude of the latent states, measured by α , and error in the natural frequencies primarily affects the frequency content of the latent states, measured by l_s . When error in the measurement noise variance is introduced, there is noticeably greater change in l_s compared to α , reflecting the influence of measurement noise on frequency content; however, in comparison to the change in parameters observed from errors in the modal properties of the system, the adjustment is considerably smaller or even negligible. This demonstrates that while the GP kernel hyperparameters of the GPLFM are the statistical model parameters which are most sensitive to error in modal properties such as modal damping ratios, natural frequencies, and mode shapes, these parameters are not sensitive to other sources of error, such as error in qualities of the measurement noise. Rather than use the GP kernel hyperparameter tuning process, error from misidentified measurement noise covariance is more effectively addressed by the aforementioned measurement noise covariance tuning process.

A.2. Transformation of lateral displacements to axial strain

From elastic beam theory, the lateral deformation of a 2D beam element v is represented with a third order polynomial of the form:

$$v(\xi) = a_0 + a_1\xi + a_2\xi^2 + a_3\xi^3 \tag{A.1}$$

$\xi = \frac{x}{L}$ is the normalized axial distance along length L of the element. The unknown coefficients a_0, a_1, a_2, a_3 can be solved for by enforcing boundary conditions which relate the deformation to the vector of degrees of freedom of the element \mathbf{d}_e , one translational and one rotational degree of freedom at each end:

$$v(\xi) = \begin{bmatrix} N_1(\xi) & N_2(\xi) & N_3(\xi) & N_4(\xi) \end{bmatrix} \begin{bmatrix} v_1 \\ \theta_1 \\ v_2 \\ \theta_2 \end{bmatrix} = \mathbf{N}_e \mathbf{d}_e \tag{A.2}$$

The solution provides the element shape functions \mathbf{N}_e , which approximate the continuous spatial distribution of deformation between finite element nodes. The derivations of shape functions for an Euler–Bernoulli beam element and Timoshenko beam element are detailed in [50,61], respectively.

For small deformations, the axial displacement u may be approximated by the rotational deformation θ_z and distance from the element centroidal axis y as:

$$u = -y\theta_z \tag{A.3}$$

Algorithm 2 Rauch–Tung–Striebel (RTS) Smoother

```

for  $k = N - 1, \dots, 1$  do
  Smoothing update equations:
   $\mathbf{G}_k = \mathbf{P}_{k|k} \mathbf{F}^T (\mathbf{P}_{k+1|k}^-)^{-1}$ 
   $\hat{\mathbf{z}}_{k|N} = \hat{\mathbf{z}}_{k|k} + \mathbf{G}_k (\hat{\mathbf{z}}_{k+1|N} - \hat{\mathbf{z}}_{k+1|k}^-)$ 
   $\mathbf{P}_{k|N} = \mathbf{P}_{k|k} + \mathbf{G}_k (\mathbf{P}_{k+1|N} - \mathbf{P}_{k+1|k}^-) \mathbf{G}_k^T$ 
  Moments of measurement states  $\mathbf{y}$ :
   $\hat{\mathbf{y}}_{k|N} = \mathbf{H} \hat{\mathbf{z}}_{k|N}$ 
   $\mathbf{P}_{k|N}^{\mathbf{y}} = \mathbf{H} \mathbf{P}_{k|N} \mathbf{H}^T + \mathbf{R}$ 
end for

```

A.4. SSM for Matérn kernel

The Matérn family is a class of covariance functions of the form

$$\kappa(\tau; \nu, \alpha, l_s) = \alpha^2 \frac{2^{(1-\nu)}}{\Gamma(\nu)} \left(\frac{\sqrt{2\nu}}{l_s} \tau \right)^\nu K_\nu \left(\frac{\sqrt{2\nu}}{l_s} \tau \right) \tag{A.8}$$

where $\Gamma(\nu)$ is the Gamma function and K_ν is the modified Bessel function of the second kind. The Matérn kernel defines the covariance between two points t, t' of a process to be a function of only the distance $\tau = t' - t$ between them, resulting in stationary stochastic processes. The hyperparameters influence qualities of the covariance relationship derived from the distance: the vertical length scale α controls the magnitude of variance at $\tau = 0$, the horizontal length scale l_s controls the distance τ over which covariance diminishes to zero, and the parameter ν controls the smoothness of the resulting process. Note that in its limits, a Matérn kernel with $l_s \rightarrow 0$ obtains the Dirac delta function for the covariance, which corresponds to a white noise kernel, and $\nu \rightarrow \infty$ obtains the squared exponential covariance function corresponding to infinitely differentiable processes [32,39].

This section presents the form of the Matérn kernel with smoothing parameter $\nu = p + 1/2 = 5/2$, where $p = 2$, leading to processes which are twice differentiable. The associated hyperparameters are then $\theta = [\alpha, l_s]$. The covariance function becomes:

$$\kappa_{5/2}(\tau; \alpha, l_s) = \alpha^2 \left(1 + \frac{\sqrt{5}}{l_s} \tau + \frac{5\tau^2}{3l_s^2} \right) \exp \left(-\frac{\sqrt{5}}{l_s} \tau \right) \tag{A.9}$$

A Gaussian process $h(t) \sim GP(0, \kappa_{5/2}(\tau))$ may be expressed as a third-order linear stochastic differential equation of the form:

$$\frac{d^3 h(t)}{dt^3} + 3\lambda \frac{d^2 h(t)}{dt^2} + 3\lambda^2 \frac{dh(t)}{dt} + \lambda^3 h(t) = w(t) \tag{A.10}$$

where $\lambda = \sqrt{5}/l_s$. The system is driven by white noise with spectral density:

$$q_w = S_w(\omega) = \frac{400\sqrt{5}\alpha^2}{3l_s^5} \tag{A.11}$$

Consequently, we can use the form of the Matérn kernel to derive the spectral density of $h(t)$ as $S_h(\omega) = q_w(\lambda^2 + \omega^2)^{-(p+1)}$.

The companion form of (A.10) gives rise to the state-space model of process $h(t)$, with the continuous-time matrices [23]:

$$\mathbf{F}_c = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ -\lambda^3 & -3\lambda^2 & -3\lambda \end{bmatrix}, \quad \mathbf{L}_c = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}, \quad \mathbf{H}_c = [1 \quad 0 \quad 0] \tag{A.12}$$

A.5. Hellinger distance

The Hellinger distance is a type of f-divergence which integrates the error between two probability densities $p(x)$ and $q(x)$ as follows:

$$H^2(p, q) = \frac{1}{2} \int (\sqrt{p(x)} - \sqrt{q(x)})^2 dx = 1 - \int \sqrt{p(x)q(x)} dx \tag{A.13}$$

In the discrete case, the Hellinger distance between the probability distributions $p = \{p_1, \dots, p_m\}$ and $q = \{q_1, \dots, q_m\}$ may be computed as a Riemann sum over m subdivisions:

$$H(p, q) = 1 - \sum_{i=1}^m \sqrt{p_i q_i} \tag{A.14}$$

A.6. Sensor specifications

The instrumentation scheme of W27 wind turbine considered in the present study consists of the following sensor types:

- Strain gauge XY101-3/350: cross strain gauge with two grids in the 0° and 90° directions, with a resistance of 350 Ohm and active gauge factor of 2.01. From stability tests conducted using a Quantum MX1615B amplifier, the strain gauge is estimated to have an accuracy of approximately 3%.
- Accelerometer AAA323-001: accelerometers with an output of ± 5 VDC. The bidirectional accelerometers with a grid of 0° and 90° have a range of ± 2 g and the unidirectional accelerometer has a range of ± 1 g.
- MGCPlus data acquisition system (DAS), which utilizes the HBM Catman data acquisition software to retrieve data at a sampling rate of 50 Hz. A Bessel low-pass filter at 5 Hz is automatically applied as an anti-aliasing feature.

References

- [1] 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.* 76–77 (2016) 592–611, <http://dx.doi.org/10.1016/j.ymssp.2016.01.004>.
- [2] C. Devriendt, W. Weijtjens, M. El-Kafafy, G. De Sitter, Monitoring resonant frequencies and damping values of an offshore wind turbine in parked conditions, *IET Renew. Power Gener.* 8 (4) (2014) 433–441, <http://dx.doi.org/10.1049/iet-rpg.2013.0229>.
- [3] M. Henkel, W. Weijtjens, C. Devriendt, Validation of virtual sensing on subsoil strain data of an offshore wind turbine, in: *8th IOMAC - Int. Oper. Modal Anal. Conf. Proc.*, 2019, pp. 765–774.
- [4] N. Noppe, A. Iliopoulos, W. Weijtjens, C. Devriendt, Full load estimation of an offshore wind turbine based on SCADA and accelerometer data, *J. Phys. Conf. Ser.* 753 (7) (2016) <http://dx.doi.org/10.1088/1742-6596/753/7/072025>.
- [5] A. Smyth, M. Wu, Multi-rate Kalman filtering for the data fusion of displacement and acceleration response measurements in dynamic system monitoring, *Mech. Syst. Signal Process.* 21 (2) (2007) 706–723, <http://dx.doi.org/10.1016/j.ymssp.2006.03.005>.
- [6] A. Iliopoulos, D. Van Hemelrijck, J. Vlassenbroeck, D. Aggelis, Assessment of grouted samples from monopile wind turbine foundations using combined non-destructive techniques, *Constr. Build. Mater.* 122 (2016) 855–862, <http://dx.doi.org/10.1016/j.conbuildmat.2015.11.047>.
- [7] A. Iliopoulos, W. Weijtjens, D. Van Hemelrijck, C. Devriendt, Fatigue assessment of offshore wind turbines on monopile foundations using multi-band modal expansion, *Wind Energy* 20 (2017) 1463–1479, <http://dx.doi.org/10.1002/we.2104>.
- [8] 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: *Proc. ISMA 2018 - Int. Conf. Noise Vib. Eng. USD 2018 - Int. Conf. Uncertain. Struct. Dyn.*, 2018, pp. 4847–4855.
- [9] M. Henkel, N. Noppe, W. Weijtjens, C. Devriendt, Sub-soil strain measurements on an operational wind turbine for design validation and fatigue assessment, *J. Phys. Conf. Ser.* 1037 (5) (2018) <http://dx.doi.org/10.1088/1742-6596/1037/5/052032>.
- [10] R. Kalman, A new approach to linear filtering and prediction problems, *J. Basic Eng.* 82 (1) (1960) 35–45, <http://dx.doi.org/10.1115/1.3662552>.
- [11] P. Kitandis, Unbiased-minimum variance linear state estimation, *Automatica* 23 (6) (1987) 775–778.
- [12] M. Hou, R. Patton, Optimal filtering for systems with unknown inputs, *IEEE Trans. Automat. Control* 43 (3) (1998) 445–449, <http://dx.doi.org/10.1109/9.661621>.
- [13] C. Hsieh, Robust two-stage Kalman filters for systems with unknown inputs, *IEEE Trans. Automat. Control* 45 (12) (2000) 2374–2378, <http://dx.doi.org/10.1109/9.895577>.
- [14] M. Darouach, M. Zasadzinski, M. Boutayeb, Extension of minimum variance estimation for systems with unknown inputs, *Automatica* 39 (2003) 867–876.
- [15] S. Gillijns, B. De Moor, Unbiased minimum-variance input and state estimation for linear discrete-time systems, *Automatica* 43 (1) (2007) 111–116, <http://dx.doi.org/10.1016/j.automatica.2006.08.002>.
- [16] S. Gillijns, B. De Moor, Unbiased minimum-variance input and state estimation for linear discrete-time systems with direct feedthrough, *Automatica* 43 (5) (2007) 934–937, <http://dx.doi.org/10.1016/j.automatica.2006.11.016>.
- [17] H. Rauch, F. Tung, C. Striebel, Maximum likelihood estimates of linear dynamic systems, *AIAA J.* 3 (8) (1965) 1445–1450, <http://dx.doi.org/10.2514/3.3166>.
- [18] K. Maes, S. Gillijns, G. Lombaert, A smoothing algorithm for joint input-state estimation in structural dynamics, *Mech. Syst. Signal Process.* 98 (2018) 292–309, <http://dx.doi.org/10.1016/j.ymssp.2017.04.047>.
- [19] 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, <http://dx.doi.org/10.1016/j.ymssp.2012.01.011>.
- [20] K. Maes, A.W. Smyth, G. De Roeck, G. Lombaert, Joint input-state estimation in structural dynamics, *Mech. Syst. Signal Process.* 70–71 (2016) 445–466, <http://dx.doi.org/10.1016/j.ymssp.2015.07.025>.
- [21] S.E. Azam, E. Chatzi, C. Papadimitriou, A dual Kalman filter approach for state estimation via output-only acceleration measurements, *Mech. Syst. Signal Process.* 60 (2015) 866–886, <http://dx.doi.org/10.1016/j.ymssp.2015.02.001>.
- [22] K. Tatsis, E. Lourens, A comparison of two Kalman-type filters for robust extrapolation of offshore wind turbine support structure response, in: *Life-Cycle Eng. Syst. Emphas. Sustain. Civ. Infrastruct. - 5th Int. Symp. Life-Cycle Eng. IALCCE 2016*, 2017, pp. 209–216, <http://dx.doi.org/10.1201/9781315375175-25>.
- [23] R. Nayek, S. Chakraborty, S. Narasimhan, A Gaussian process latent force model for joint input-state estimation in linear structural systems, *Mech. Syst. Signal Process.* (2019) <http://dx.doi.org/10.1016/j.ymssp.2019.03.048>.
- [24] S. Azam, E. Chatzi, C. Papadimitriou, A. Smyth, Experimental validation of the Kalman-type filters for online and real-time state and input estimation, *JVC/J. Vib. Control* 23 (15) (2017) 2494–2519, <http://dx.doi.org/10.1177/1077546315617672>.
- [25] K. Maes, K. Van Nimmen, E. Lourens, A. Rezayat, P. Guillaume, G. De Roeck, G. Lombaert, Verification of joint input-state estimation for force identification by means of in situ measurements on a footbridge, *Mech. Syst. Signal Process.* 75 (2016) 245–260.
- [26] O. Petersen, O. Oiseth, E. Lourens, Estimation of the dynamic response of a slender suspension bridge using measured acceleration data, *Procedia Eng.* 199 (2017) 3047–3052, <http://dx.doi.org/10.1016/j.proeng.2017.09.547>.
- [27] K. Maes, G. De Roeck, A. Iliopoulos, W. Weijtjens, C. Devriendt, G. Lombaert, Kalman filter based strain estimation for fatigue assessment of an offshore monopile wind turbine, in: *Proc. ISMA 2016 - Int. Conf. Noise Vib. Eng. USD2016 - Int. Conf. Uncertain. Struct. Dyn.*, 2016, pp. 1649–1661.
- [28] E. Lourens, E. Reynders, G. De Roeck, G. Degrande, G. Lombaert, An augmented Kalman filter for force identification in structural dynamics, *Mech. Syst. Signal Process.* 27 (1) (2012) 446–460, <http://dx.doi.org/10.1016/j.ymssp.2011.09.025>.
- [29] F. Naets, J. Cuadrado, W. Desmet, Stable force identification in structural dynamics using Kalman filtering and dummy-measurements, *Mech. Syst. Signal Process.* 50–51 (2015) 235–248, <http://dx.doi.org/10.1016/j.ymssp.2014.05.042>.
- [30] K. Yuen, K. Hoi, K. Mok, Selection of noise parameters for Kalman filter, *Earthq. Eng. Eng. Vib.* 6 (49) (2007) 445–466, <http://dx.doi.org/10.1007/s11803-007-0659-9>.

- [31] M. Alvarez, D. Luengo, N.D. Lawrence, Latent force models, *J. Mach. Learn. Res.* 5 (2009) 9–16.
- [32] J. Hartikainen, S. Särkkä, Kalman filtering and smoothing solutions to temporal Gaussian process regression models, in: *Proc. 2010 IEEE Int. Work. Mach. Learn. Signal Process. MLSP 2010*, 2010, pp. 379–384, <http://dx.doi.org/10.1109/MLSP.2010.5589113>.
- [33] J. Hartikainen, S. Särkkä, Sequential inference for latent force models, in: *Proc. 27th Conf. Uncertain. Artif. Intell. UAI 2011*, 2011, pp. 311–318.
- [34] T. Rogers, K. Worden, E. Cross, On the application of Gaussian process latent force models for joint input-state-parameter estimation: With a view to Bayesian operational identification, *Mech. Syst. Signal Process.* 140 (2020) 106580, <http://dx.doi.org/10.1016/j.ymssp.2019.106580>.
- [35] O. Petersen, O. Oiseth, E. Lourens, Wind load estimation and virtual sensing in long-span suspension bridges using physics-informed Gaussian process latent force models, *Mech. Syst. Signal Process.* 170 (2022) 108742, <http://dx.doi.org/10.1016/j.ymssp.2021.108742>.
- [36] J. Bilbao Neiva, E. Lourens, A. Schulze, L. Ziegler, Virtual sensing in an onshore wind turbine tower using a Gaussian Process Latent Force Model, *Data-Centric Engineering* 3 (3) (2022) <http://dx.doi.org/10.1017/dce.2022.38>.
- [37] K. Maes, E. Lourens, K. Van Nimmen, E. Reynders, G. De Roeck, G. Lombaert, Design of sensor networks for instantaneous inversion of modally reduced order models in structural dynamics, *Mech. Syst. Signal Process.* 52–53 (1) (2015) 628–644, <http://dx.doi.org/10.1016/j.ymssp.2014.07.018>.
- [38] R. Ferrari, F. Pioldi, E. Rizzi, C. Gentile, E. Chatzi, R. Klis, E. Serantoni, A. Wieser, Heterogeneous sensor fusion for reducing uncertainty in structural health monitoring, in: *UNCECOMP 2015 - 1st ECCOMAS Themat. Conf. Uncertain. Quantif. Comput. Sci. Eng.*, 2015, pp. 511–528, <http://dx.doi.org/10.7712/120215.4289.821>.
- [39] C. Rasmussen, C. Williams, *Gaussian Processes for Machine Learning*, MIT Press, 2006.
- [40] S. Särkkä, *Applied Stochastic Differential Equations*, Helsinki University of Technology, 2012.
- [41] M. Grewal, A. Andrews, *Kalman Filtering: Theory and Practice with MATLAB*, Wiley-IEEE Press, 2001.
- [42] S. Bangert, *Lyap Function: Continuous Lyapunov Equation Solution in MATLAB*, MathWorks Inc., 1986.
- [43] S. Särkkä, M. Alvarez, N. Lawrence, Gaussian process latent force models for learning and stochastic control of physical systems, *IEEE Trans. Automat. Control* 64 (7) (2019) 2953–2960, <http://dx.doi.org/10.1109/TAC.2018.2874749>, arXiv:1709.05409.
- [44] J. Kaipio, E. Somersalo, *Statistical and computational inverse problems*, in: *Applied Mathematical Sciences*, Vol. 160, Springer, 2005.
- [45] J. Bilbao Neiva, *Fatigue load monitoring of an onshore wind turbine [MS thesis]*, in: *TU Delft Repository*, 2021.
- [46] A. Solin, S. Särkkä, Explicit link between periodic covariance functions and state space models, *J. Mach. Learn. Res.* 33 (2014) 904–912.
- [47] W. Stephenson, S. Ghosh, T. Nguyen, M. Yurochkin, S. Deshpande, T. Broderick, Measuring the sensitivity of Gaussian processes to kernel choice, 2021, ArXiv arXiv:2106.06510.
- [48] B. Peeters, G. De Roeck, Reference-based stochastic subspace identification for output-only modal analysis, *Mech. Syst. Signal Process.* 13 (6) (1999) 855–878, <http://dx.doi.org/10.1006/mssp.1999.1249>.
- [49] M. van Vondelen, *Damping identification of operational offshore wind turbines: an operational modal analysis approach in the presence of harmonics [MS thesis]*, in: *TU Delft Repository*, 2021.
- [50] G. Liu, S. Quek, *The Finite Element Method: A Practical Course*, Elsevier Science, 2003.
- [51] E. Lourens, D. Fallais, On the use of equivalent forces for structural health monitoring based on joint input-state estimation algorithms, *Procedia Eng.* 199 (2017) 2140–2145, <http://dx.doi.org/10.1016/j.proeng.2017.09.152>.
- [52] E. Lourens, D. Fallais, Full-field response monitoring in structural systems driven by a set of identified equivalent forces, *Mech. Syst. Signal Process.* 114 (2019) <http://dx.doi.org/10.1016/j.ymssp.2018.05.014>.
- [53] Q. Sun, C. Liu, C. Zhen, Abnormal detection of wind turbine operating conditions based on state curves, *J. Energy Eng.* 145 (5) (2019) 06019001, [http://dx.doi.org/10.1061/\(asce\)ey.1943-7897.0000612](http://dx.doi.org/10.1061/(asce)ey.1943-7897.0000612).
- [54] M. Laanes, *Identification of fatigue and resonance frequency influencing parameters of an offshore wind turbine via data analysis of real measurements [MS thesis]*, in: *TU Delft Repository*, 2021.
- [55] M. Ankerst, M. Breunig, H. Kriegel, J. Sander, OPTICS: Ordering points to identify the clustering structure, in: *Int. Conf. on Management of Data*, 1999.
- [56] R. Boroschek, J. Bilbao Neiva, Interpretation of stabilization diagrams using density-based clustering algorithm, *Eng. Struct.* (2019) 245–257, <http://dx.doi.org/10.1016/j.engstruct.2018.09.091>.
- [57] W. Versteijlen, F. Renting, P. van der Valk, J. Bongers, K. van Dalen, A. Metrikine, Effective soil-stiffness validation: Shaker excitation of an in-situ monopile foundation, *Soil Dyn. Earthq. Eng.* 102 (June) (2017) 241–262, <http://dx.doi.org/10.1016/j.soildyn.2017.08.003>.
- [58] W. Versteijlen, F. Renting, P. van der Valk, J. Bongers, K. van Dalen, A. Metrikine, Effective stiffness method for rigid monopile foundations of offshore wind turbines and in-situ validation, *Procedia Eng.* 199 (2017) 3248–3253, <http://dx.doi.org/10.1016/j.proeng.2017.09.349>.
- [59] ASTM-E1049-85, *Standard Practices for Cycle Counting in Fatigue Analysis*, ASTM International, West Conshohocken, PA, 2017.
- [60] A. Chopra, *Dynamics of Structures: Theory and Application To Earthquake Engineering*, Prentice Hall, 1995.
- [61] A. Bazoune, Y. Khulief, Shape functions of three-dimensional timoshenko beam element, *J. Sound Vib.* 259 (2) (2003) 473–480, <http://dx.doi.org/10.1006/jsvi.2002.5122>.