



Delft University of Technology

Document Version

Final published version

Licence

CC BY-NC

Citation (APA)

Patryniak, K., Collu, M., Yeter, B., Minisi, S., Oneto, L., & Coraddu, A. (2026). Hybrid frequency–time domain surrogate-augmented model of dynamics of floating offshore wind turbines. *Ocean Engineering*, 365(P1), Article 127227. <https://doi.org/10.1016/j.oceaneng.2026.127227>

Important note

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

Copyright

In case the licence states “Dutch Copyright Act (Article 25fa)”, this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

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.

This work is downloaded from Delft University of Technology.



Research paper

Hybrid frequency–time domain surrogate-augmented model of dynamics of floating offshore wind turbines

Katarzyna Patryniak ^a,*, Maurizio Collu ^a, Baran Yeter ^b, Simone Minisi ^c, Luca Oneto ^c, Andrea Coraddu ^d^a Department of Naval Architecture, Ocean and Marine Engineering, University of Strathclyde, Glasgow, UK^b Department of Energy, Aalborg University, Esbjerg, Denmark^c DITEN, University of Genoa, Genoa, Italy^d Faculty of Mechanical, Maritime and Materials Engineering, Delft University of Technology, Delft, Netherlands

ARTICLE INFO

Keywords:

Floating offshore wind turbine

Fatigue analysis

Design optimisation

Mooring design

Offshore engineering

Surrogate model

Machine learning

ABSTRACT

Floating offshore wind turbines can harness the more consistent and powerful winds in deep-water offshore areas. However, they face challenges due to complex loading mechanisms across various load cases. Fully coupled nonlinear time-domain response and load analyses are computationally expensive, making them unsuitable for highly iterative design optimisation under all limit states' constraints. This research presents a hybrid domain (frequency and time) and hybrid paradigm (physics-based, surrogate-augmented) approach that strikes a balance between accuracy and computational efficiency. The motion response is obtained in the frequency domain, while structural loads are computed from the inverse fast Fourier transform of the motion spectra. This approach combines the efficiency of the frequency-domain approach with the ability to model some load nonlinearities in the time domain. To further enhance efficiency, two dedicated surrogate models have been developed and embedded into the workflow. The first surrogate predicts hydrodynamic potential coefficients over the relevant range of wave frequencies, thereby replacing the need for repeated potential-flow computations in the design loop. The second surrogate emulates the time-domain mooring line loads as a function of mooring design and fairlead motion, acting as a fast proxy for more computationally demanding dynamic mooring analyses. The mooring surrogate is agnostic to platform design and environmental conditions, offering broader applicability than existing approaches. Both surrogates are constructed as data-driven, multi-output regression models trained on large datasets generated from higher-fidelity numerical simulations, tailored to provide accurate predictions in milliseconds. This novel approach provides accurate design trend predictions while significantly reducing solution time. It enables fast and comprehensive FOWT design optimisation, accelerating early-stage development and allowing greater complexity from the outset.

1. Introduction and previous work

Floating offshore wind turbines (FOWTs) have the potential to harness the abundant wind resources in deep-water offshore conditions, where bottom-fixed structures are no longer feasible. With taller towers and larger rotors, they can access stronger wind at higher altitudes, thus operating at higher capacity factors (Johnston et al., 2020). Moreover, FOWTs benefit from more uniform and consistent wind, with minimal terrain roughness. Installed further away from the coast, they minimise interference with other coastal activities as well as audio-visual, societal, and environmental impacts (IRENA, 2019). However, they are also exposed to a harsher, unsheltered marine environment. Due to the lack of rigid connection to the seabed, these structures may experience

large rigid-body motions and structural loads due to the wind, wave, and current loads, with strong interdependence between the inertial, wind, and wave-driven responses (Bachynski and Moan, 2012; Wang et al., 2023). Therefore, integrated design methods are essential for balancing performance and cost across all components of the complex, coupled FOWT system. Efficient and accurate numerical modelling is critical for design optimisation, as high computational demand can make this highly iterative process too time-consuming to be practical, while insufficient accuracy risks over-engineered or unrealistic design solutions.

FOWT modelling is challenging due to the interactions between the multivariate environmental loading and the system's interaction with the environment. As shown in multiple references including (Liu

* Corresponding author.

E-mail address: katarzyna.patryniak@strath.ac.uk (K. Patryniak).<https://doi.org/10.1016/j.oceaneng.2026.127227>

Received 28 February 2026; Received in revised form 22 July 2026; Accepted 23 July 2026

Available online 30 July 2026

0029-8018/© 2026 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY-NC license (<http://creativecommons.org/licenses/by-nc/4.0/>).

et al., 2017; Lienard et al., 2019; Huang and Wan, 2020), wave-induced rigid motion of the platform substantially affects the tower-top motion, hence influencing the magnitude and variance of the rotor thrust and torque, which in turn lead to fluctuating structural loads on the system's components (Veldman, 2020; Zhu et al., 2024). The structural loads are further augmented by the gravitational and inertial loads related to the angular displacements and accelerations of the rigid body (Robertson et al., 2017). Moreover, large mean displacements of the platform due to aerodynamic loads lead to a nonlinear force–displacement relationship of the mooring lines (Brown, 2005; Liu et al., 2017; Tracy, 2007), which, together with the variable displacements of the platform, leads to significant variability of the mooring tension (Robertson et al., 2017).

The mean and fluctuating wind and wave loads contribute to distinct structural failure mechanisms. Therefore, conducting analysis across various scenarios throughout the lifespan of FOWTs is crucial to mitigate risks associated with extreme environmental forces (ultimate limit states), operational conditions (serviceability limit states), and progressive structural degradation over time (fatigue limit states). In the context of design optimisation, fatigue assessment presents a significant challenge, as it requires the evaluation of internal structural loads under a variety of environmental loading conditions (at least 20, according to DNV-RP-C203 (DNVGL, 2016)). The applicability and accuracy of spectral-based fatigue assessment methods are limited by the broad-band loading in offshore conditions, and time-domain approaches may be required, further increasing computational effort.

A full stochastic analysis of wind- and wave-induced loading on FOWTs over a 25–30-year lifetime using state-of-the-art, fully coupled nonlinear time-domain solvers requires substantial computational resources, even for a single analysis (Butterfield et al., 2005). Long simulation windows are necessary to reduce statistical uncertainty, and the computational burden is further amplified by the iterative nature of design optimisation (Moan et al., 2020). As fully coupled simulations are therefore impractical for optimisation — particularly in early-stage concept design — numerous alternative modelling approaches have been proposed for predicting motions, loads, and fatigue, including both physics-based and data-driven methods.

1.1. Efficient physics-based models

The following sections introduce the key challenges associated with modelling motions and loads, as well as assessing fatigue damage.

1.1.1. Motion and load response

Simplified dynamic models have long been used to capture the essential coupled dynamics of FOWT systems (Moan et al., 2020), and linearised formulations in particular offer clear advantages for iterative optimisation. A comprehensive review of such methods is given in Patryniak et al. (2022). In what follows, we draw on a small number of representative models to frame the context of existing frequency-domain approaches, noting that this discussion is not intended to be exhaustive or systematic.

Early frequency-domain (FD) approaches primarily treated wind- and wave-induced responses separately. For example, Bachynski (2014) modelled wave-driven motions by multiplying the hydrodynamic response amplitude operators (RAOs) with the incident wave spectrum, and obtained wind-driven motions by combining wind RAOs with a thrust spectrum precomputed in the time domain. More integrated FD formulations have since been developed. Hegseth et al. (2020) assembled a closed-loop aero–hydro–servo–elastic model by expressing structural and control components in state–space form. Hydrodynamic excitation was obtained from MacCamy–Fuchs theory, damping from linearised Morison's equation, and added mass from analytical 2D formulations. Aerodynamic loads were computed using a linearised quasi-steady blade element momentum method, and mooring forces

were represented using the simplified FD model of Larsen and Sandvik (1990). While the aerodynamic predictions compared well with a nonlinear reference, full-system verification was not demonstrated. Recent frameworks incorporate more accurate disciplinary models while retaining computational efficiency. RAFT (Hall et al., 2022), for example, couples potential-flow and strip-theory hydrodynamics with BEM aerodynamics, quasi-static moorings, and linear turbine control within a unified FD solver. Verification showed generally good agreement with OpenFAST for global motions, though noticeable discrepancies remained in predicted tower-base loads. Overall, FD models have progressed from highly simplified spectral formulations to more comprehensive multi-physics frameworks, improving accuracy while maintaining tractability for design-space exploration. Nevertheless, most published models still lack systematic verification or show notable inaccuracies in load prediction.

1.1.2. Fatigue assessment

The simplest spectral-based approaches for fatigue analysis are based on the assumption of a narrow-band stress process. The wind and wave-driven loads exhibit distinct statistical properties, as they act in different (although not completely disjoint) frequency regimes. Depending on the location, the first-order wave loads are concentrated in a specific, higher frequency band of the order of 0.1 Hz (or 10 s period), while the difference frequency second-order wave loads and turbulent wind loads typically act over a wider range of frequencies at the lower end of the spectrum of the order of 0.001 to 0.01 Hz (or a period of 100 to 1000 s). While each load source can be considered narrow-band, the total load follows a more complex distribution. Consequently, narrow-band spectral methods are not suitable when both wind and wave loads are applied simultaneously (Martinez-Puente et al., 2024), and, as shown in Mao et al. (2025), require corrections and calibration factors.

Dual-narrow-band or wide-band frequency-domain methods are more often applied, as they can account for the complex nature of the combined wind-wave loading. Several studies highlighted the importance of incorporating wind and wave-induced loading in fatigue and reliability analyses of FOWTs. The work of Sørensen and Sørensen (2010), Burton et al. (2011) and Song et al. (2021) showed that the variable aerodynamic loads can significantly contribute to fatigue damage of FOWT. Xu et al. (2019) and Wang et al. (2023) showed that while wind-induced loading is often the primary contributor to fatigue during operational conditions, high sea states with large wave heights can also lead to substantial fatigue damage. In agreement with this, Zhu et al. (2024) showed that while the relative contributions of the wind and waves to fatigue damage depend on the control strategy, wind-induced fatigue tends to dominate in lower wind speed operational conditions, and wave-induced fatigue becomes relatively more important in harsh sea states in the high wind speed (reduced thrust) regime.

Simply summing the fatigue damage from low- and high-frequency responses can lead to underestimation when both sources contribute significantly (DNVGL, 2016; Kai-Tung Ma et al., 2019). Kai-Tung Ma et al. (2019) propose two practical solutions: the combined-spectrum method, which computes total damage from the root-sum-of-squares stress standard deviation; and the dual-narrow-band method, which applies a correction factor reflecting the relative contributions of the two frequency bands. Appendix D of DNV-RP-C203 (DNVGL, 2016) also provides a simplified combination rule based on mean zero-upcrossing periods. Beyond these, wide-band spectral methods such as the Dirlik (1985) and Tovo and Benasciutti (2004) formulations are widely regarded as the state-of-the-art in frequency-domain fatigue estimation. A comprehensive overview was given in Dirlik and Benasciutti (2021), which highlighted that these methods were initially derived using Monte Carlo sampling of stress processes, followed by inverse FFT reconstruction of synthetic time series from which empirical cycle distributions were obtained. Their parameters are therefore tuned to

reproduce rainflow statistics, but they remain approximations of the underlying time-domain process. Accordingly, DNV-E301 recommends using the dual narrow-band method and checking a few cases with a more accurate rainflow counting method.

Although computationally efficient, particularly valuable in high-resolution multibody or FEA structural assessments, the spectral methods are generally conservative (Dirlik and Benasciutti, 2021; Kai-Tung Ma et al., 2019) and lead to over-engineered structures. In contrast, the rainflow counting method operates directly on the stress time series. It does not require assumptions about the underlying spectral shape, making it the most accurate engineering approach for fatigue assessment (Kai-Tung Ma et al., 2019). When dynamic responses can be efficiently obtained for selected structural details/locations under the relevant metocean and operational conditions, time-domain rainflow counting offers a balance between accuracy and computational cost, making it highly suitable for fatigue evaluation during early-stage design.

1.2. Data-driven models

Alternatively, data-driven surrogate models offer accuracy comparable to fully nonlinear dynamic models, but at a fraction of the computational cost (Queipo et al., 2005; Forrester et al., 2008). In a typical supervised-learning setting, a surrogate is trained offline on a database of input-output pairs generated using higher-fidelity numerical simulations or experiments, thereby learning an approximation of the underlying input-response mapping. Once trained, the surrogate can be evaluated in milliseconds, enabling its use within gradient-based or gradient-free optimisation, uncertainty quantification, and reliability analyses that require thousands to millions of model evaluations (Forrester et al., 2008; Rasmussen and Williams, 2006). In this way, the computational burden is shifted from the optimisation loop to an upfront pre-computation phase devoted to dataset generation and model training.

From a methodological standpoint, surrogate modelling encompasses a broad family of regression techniques. Classical approaches include low- and high-order polynomial response surfaces, Gaussian process/Kriging models, and radial-basis-function networks, which provide smooth global approximations and, in the case of Gaussian processes, principled uncertainty quantification (Queipo et al., 2005; Rasmussen and Williams, 2006). More recently, machine-learning models such as support vector regression, random forests, gradient-boosting ensembles, and deep neural networks have gained prominence due to their ability to handle high-dimensional, nonlinear mappings (Bishop and Bishop, 2023; Oneto, 2020). Their practical scalability to very large datasets and to coupled multi-output problems, however, varies considerably across methods and implementations. Each of these techniques offers different trade-offs between expressiveness, data efficiency, robustness to noise, and interpretability, and their selection is typically guided by problem dimensionality, available data, and computational budget (Oneto, 2020).

In the maritime and offshore domain, data-driven surrogates have already been successfully deployed for ship performance prediction and fouling assessment via digital twins (Coraddu et al., 2019; Oneto et al., 2016), for the estimation of ultimate strength reduction in stiffened panels under welding residual stresses (Li et al., 2022; Coraddu et al., 2023), and for hybrid physical/data-driven modelling of marine engines (Coraddu et al., 2022). These studies show that, when coupled with rigorous model selection, regularisation, and error-estimation procedures (Oneto, 2020), surrogate models can achieve high predictive accuracy while maintaining a controlled generalisation error. In particular, recent works have combined advanced learning pipelines (feature engineering, hyperparameter optimisation, cross-validation, and uncertainty quantification) with physically informed inputs and constraints, yielding surrogates that are not only accurate but also robust and physically plausible within their domain of applicability (Coraddu et al.,

2019, 2023, 2022). This line of research emphasises the importance of aligning model complexity with the quantity and quality of data, as well as systematically validating surrogates before incorporating them into design or operational decision-making frameworks.

A typical example directly relevant to FOWTs is the construction of a surrogate model that emulates a high-fidelity aero-hydro-servo-elastic simulator. In this setting, the inputs consist of both design variables (e.g. platform geometry, mooring-system properties, and controller parameters) and environmental descriptors (e.g. wave spectrum, wind speed, and directionality), while the outputs are key performance indicators such as platform motions, tower-base bending moments, and mooring-line tensions. The training dataset is generated by sampling the joint design-environmental space and running the high-fidelity simulator for each sample, typically at a computational cost on the order of an hour per case. This database is then used to train a data-driven surrogate model — such as a Gaussian process or a shallow neural network — that learns the mapping from inputs to system responses. Once trained and carefully validated, the surrogate can be evaluated within its domain of applicability in milliseconds, acting as an efficient emulator of the full nonlinear time-domain model. This enables extensive design-space exploration, sensitivity analysis, and scenario screening that would otherwise be computationally prohibitive, particularly in early-stage FOWT design where thousands of evaluations are required before committing resources to a limited number of detailed simulations.

The potential of these models has been recognised in several recent publications on FOWT design. For example, Zhang et al. (2022) conducted a multi-objective optimisation of a semisubmersible platform using a Standard Response Surface, which was constructed from a database of results obtained from OpenFAST. In this study, the platform parameters served as inputs, while the outputs were the motions and natural frequencies, which also acted as design constraints. Trubat et al. (2022) performed FOWT mooring optimisation using a Differential Evolution algorithm, with an Artificial Neural Network (ANN) model predicting static and dynamic responses (line forces and platform motions) based on input parameters related to line geometry. Ye et al. (2024) optimised the FOWT mooring system using the NSGA-II algorithm and a Kriging model, which was trained on results from OpenFAST. As in the previous studies, the mooring system parameters were input to the model, while platform motions (which also served as design constraints) were the outputs. Wang et al. (2025) trained a Deep Neural Network (DNN) using the geometric parameters of the mooring system of a semisubmersible FOWT to predict optimisation objectives – mooring tension and platform motions – for three design load cases. A key innovation in the approach was the use of initial static calculations to filter out infeasible samples from the dataset, which were then excluded from the training of the surrogate model.

Unlike the other studies, a recent semisubmersible FOWT optimisation by Bessone et al. (2024) employed a surrogate model for potential hydrodynamic coefficients (rather than directly modelling the design constraints). In this case, a data-driven replacement of the boundary element method code was utilised within a highly efficient frequency-domain framework to compute system responses, which were then considered as design constraints. The authors trained an XGBoost model based on results from the HAMS (NREL, 2024b) hydrodynamic solver. A similar approach was independently demonstrated by Coraddu et al. (2020) and Rohrer and Bachynski-Polic (2024), though these studies did not directly apply surrogate models to design optimisation.

Most of the studies discussed employed surrogate models to accelerate design evaluations within stochastic optimisation frameworks. With the exception of one study, these studies employed surrogate models to predict design objectives and constraints, rather than replacing specific bottleneck models in the coupled solution process. Furthermore, each study trained its models for specific environmental conditions, mooring studies focused on particular platform designs, and platform studies were based on specific mooring configurations, limiting their applicability.

Model	C1		C2	
	Paradigm	Domain	Paradigm	Domain
Mean motion	PM		PM	
Potential coefficients		FD	DDM	FD
Dynamic motions		FD → TD		FD → TD
Tower loads		TD	PM	TD
Mooring loads			DDM	
Section	3		4.1, 4.2	

Fig. 1. Dynamic models overview. PM—physical model; DDM—data-driven model; FD—frequency domain; TD—time domain.

1.3. Paper outline

This paper presents an efficient approach to the design evaluation of FOWTs for use in optimisation studies. It provides a method for assessing the metrics across the ultimate, fatigue, and serviceability limit states, as required by a holistic design approach at the conceptual design stage. In Section 3, we introduce a physics-based hybrid frequency–time domain dynamic model and evaluate its accuracy and efficiency. The limitations identified in this evaluation motivate the development of a surrogate-augmented model, presented in Section 4, in which surrogates for the potential-flow coefficients and mooring loads address the remaining computational bottlenecks. A demonstration of design optimisation performed with each model is presented in Section 5, with the article concluding in Section 6.

2. Methodology

The two models presented in this paper differ in their modelling paradigm: physics-based (PM) or data-driven (DDM), and modelling domain: frequency domain (FD) or time domain (TD), as summarised in Fig. 1. Typical frequency-domain models rely on frequency-domain solutions for both motion and load response, which limits their accuracy, particularly in the prediction of mooring tension and fatigue, and tower base bending moment (Patryniak, 2026).

The hybrid-domain model developed in this paper (Section 3) retains the efficient computation of the motion response in the frequency domain, but transforms the resulting motion to the time domain, where the structural loads are calculated directly. This allows some motion-load nonlinearities to be captured, following an approach inspired by the work of Pegalajar-Jurado et al. (2018). The second model (Section 4) augments the hybrid-domain model with two surrogates: one replacing the boundary element method for computation of the hydrodynamic coefficients, and another for the computation of the mooring loads.

2.1. Verification study setup

The models developed are tested on the example of the OC3 Hywind spar floating platform (Jonkman, 2010) with the 5 MW reference rotor (Jonkman et al., 2009). The mooring system was sized to ensure the mean offset below 10% of the water depth and no vertical load on the anchors in the rated wind (maximum thrust) condition.¹ The system is illustrated in Fig. 2 and detailed in Tables 1–2.

¹ This was achieved by modifying the original line diameter and anchor radius, as one of the possible ways of achieving the desired response. Note that this is not an optimal design in any sense.

Table 1

Reference FOWT design characteristics. CG — centre of gravity; CB — centre of buoyancy.

Property	Unit	Value
Platform, including ballast:		
Draft	m	120
Diameter	m	6.5–9.7 (tapered)
Offset column diameter	m	–
Base column diameter	m	–
Mass	kg	$7.47 \cdot 10^6$
Vertical CG	m	–89.92
Pitch inertia about CG	kg m ²	$4.23 \cdot 10^9$
Hydrostatics:		
Vertical CB	m	–61.17
Metacentric height	m	28.7
Stiffness coefficient:		
Heave	N m ^{–1}	$3.33 \cdot 10^5$
Pitch	N m rad ^{–1}	$–4.99 \cdot 10^9$
Infinite-frequency added mass:		
Surge-surge	kg	$7.569 \cdot 10^6$
Heave-heave	kg	$2.35 \cdot 10^5$
Pitch-pitch	kg m ²	$3.7 \cdot 10^{10}$
Surge-pitch	kg m	$–4.71 \cdot 10^8$
Mooring system:		
Fairlead height	m	–70
Fairlead radius	m	5.2
Water depth (anchor position)	m	–320
Anchor radius	m	845.2
Unstretched line length	m	902.2
Chain diameter (nominal)	m	0.07
Pretension	kN	859.0
Natural frequency:		
Surge	Hz	0.007
Heave	Hz	0.032
Pitch	Hz	0.034

Table 2

NREL 5 MW reference rotor and tower characteristics.

Property	Unit	Value
Rotor-Nacelle Assembly:		
RNA mass	kg	$3.5 \cdot 10^5$
Hub height	m	90.0
Cut-in, rated, cut-out wind speed	m s ^{–1}	3, 11.4, 25
Rated wind speed thrust	kN	805.11
Tower:		
Elevation to tower base	m	10.0
Elevation to tower top	m	87.6
Tower mass	kg	249 718
Tower vertical CG	m	43.4

Each model is verified against a state-of-the-art time-domain coupled model of dynamics OpenFAST (National Renewable Energy Laboratory, 2020). The simulation setup includes only rigid degrees of freedom, no controller (operational point set according to the control look-up table), first-order wave loads (both potential and viscous) integrated up to the mean waterline, no current, and turbulent wind generated in accordance with IEC-recommended spectrum (IEC 61400-1), aerodynamic loads computation with a blade element momentum model without tower influence but with tip and hub loss corrections, and a dynamic lumped-mass mooring model.

To verify the FOWT motions and loads, three series of verification cases are performed: with only the wave load (cases 1.x), with only the wind load (cases 2.x), and with both wind and wave loads (cases 3.x), to provide insight into the mechanisms driving the discrepancies. The analyses rely on the metocean data from the Scottish sectoral marine plan site “NE8” (Scottish Government, 2020). As summarised in Table 3, the wave-only test cases are run at varying peak periods and significant wave heights, to investigate the response in different frequency and amplitude regimes. Wind-only and wind-wave cases are run in below-rated, rated, and above-rated conditions. The significant wave height H_s and peak period T_p were selected based on the joint




Fig. 2. Reference spar FOWT system illustration.




Fig. 3. Representative wind speed, significant wave height and peak period selection based on conditional probabilities.

Table 3

Dynamic model verification cases.

Case	Wave-only cases					Wind-only cases				Wind and wave cases		
	1.1	1.2	1.3	1.4	1.5	2.1	2.2	2.3	2.4	3.1	3.2	3.3
H_s (m)	1.0	2.0	1.0	2.0	1.0	–	–	–	–	1.44	1.87	2.68
T_p (s)	8.0	8.0	16.0	16.0	32.0	–	–	–	–	8.14	7.48	7.63
U_∞ (m s ⁻¹)	–	–	–	–	–	4.0	8.5	11.4	15.1	8.5	11.4	15.1

wind–wave statistics, as illustrated in Fig. 3. The metocean dataset was first filtered to include only records with wind speeds within a ± 1.0 m s⁻¹ interval around the selected wind speed. Within this subset, H_s was defined as the expected value of the empirical H_s distribution, and a range within ± 0.25 m of this value was selected. From within this range of records, the T_p was then taken as the expected value of the corresponding empirical T_p distribution.

The fatigue assessment in the verification studies is carried out for a single sea state using a 1-hour simulated response, which is then extrapolated to the full lifetime of the structure (25 years). This approach contrasts with a fully probabilistic fatigue assessment, where contributions from all sea states expected over the structure's lifetime are accounted for. By focusing on a single sea state, direct comparisons of fatigue damage across models can be made without averaging over multiple environmental loading realisations. When required, multiple representative sea states are analysed—for example, to investigate fatigue behaviour under different loading regimes.

A single wind and waves direction of 0° is modelled, and a fatigue reduction factor R_μ is applied to approximately account for the directionality of the waves (distribution of the mean wave direction), following the approach proposed in Vugts (2005). The metocean data is

used to obtain the probability of each mean wave direction $p(\mu_i)$ (shown in Fig. 4), and the correction factor is obtained according to Eq. (1):

$$R_\mu = \sum_{i=1}^I \left[p(\mu_i) |\cos \mu_i|^3 \right] \quad (1)$$

The output time signals from OpenFAST (tower base bending moment and mooring tension at the fairlead) are post-processed to compute the stress and fatigue damage, which are not direct outputs of OpenFAST (the method will be outlined in Section 3).

3. Physics-based model (C1)

The subsequent sections introduce and verify the hybrid-domain model (C1). The model relies on the prediction of the steady-state position of FOWT as excited by mean aero- and mooring loads, as well as the dynamic, harmonic response due to the fluctuating wave and wind loads, both the deterministic component related to the variation of the wind load due to the rotor rotation, as well as the stochastic component, related to the turbulent wind inflow.




Fig. 4. Wave rose for the NE8 site.

3.1. Static equilibrium

The floating support structure is modelled as a rigid body with three degrees of freedom (DOFs): surge, heave, and pitch. Only in-plane loads and motion response are considered in the current study. The static response (equilibrium position) of the floating system is obtained by iteratively solving for the mean aerodynamic and mooring load on the platform ($\bar{F}_{aero}$, $\bar{F}_{moor}$) and the offset and attitude of the system ($\bar{\eta}$), considering the hydrodynamic and gravitational restoring (C_h , C_g), using an efficient Broyden solver, as recommended in Patryniak and Collu (2023):

$$(C_h + C_g)\bar{\eta} = \bar{F}_{aero}(\bar{\eta}) + \bar{F}_{moor}(\bar{\eta}) \quad (2)$$

The load on the platform from the mooring system is modelled using a quasi-static approach implemented in MoorPy (Hall et al., 2021). This method utilises standard catenary equations to calculate the instantaneous position and tension in the mooring lines, based on the net weight and axial stiffness. The mean aerodynamic load and its derivatives are solved by a steady blade element momentum model implemented in CCBlade (Ning, 2013), including the Prandtl tip and hub loss corrections and wake rotation (non-zero tangential factors).

3.2. Dynamic motion response

The response of the platform to turbulent wind and wave excitation is computed based on the solution of the three DOF forced oscillation problem. The complex amplitude of motion $\hat{\eta}(\omega)$ is obtained by solving the equation of motion in the frequency domain:

$$\hat{\eta}(\omega)[- \omega^2(\mathbf{M} + \mathbf{A}(\omega)) + i\omega\mathbf{B}(\omega) + \mathbf{C}] = \hat{\mathbf{F}}_{hydro}(\omega) + \hat{\mathbf{F}}_{aero}(\omega) \quad (3)$$

$\mathbf{M}$ and $\mathbf{C}$ are the mass/inertia and stiffness matrices, respectively, and $\mathbf{A}$, $\mathbf{B}$ are the frequency-dependent added mass and damping matrices, respectively. Inertia and hydrostatic stiffness matrices are calculated analytically in the initial undisplaced position. The contribution to total stiffness from the mooring system is computed by a quasi-static mooring model MoorPy (Hall et al., 2021) (linear stiffness), in the equilibrium position (as it varies significantly from the stiffness in the initial undisplaced position). The derivatives of the aerodynamic load, i.e., aerodynamic damping, are computed by a steady blade element momentum model (CCBlade (Ning, 2013)).

The frequency-dependent potential hydrodynamic coefficients (added mass, radiation damping, and excitation) are computed using a 3D linear frequency-domain BEM solver, Capytaine (Ancellin and Dias, 2019), for the floating platform in the equilibrium position (i.e., from the solution of the radiation and diffraction problems). The excitation forces are required to obtain the loads on the right-hand side of Eq. (3), as will be detailed shortly.




Fig. 5. Boundary element method mesh (panels) over submerged part of the floating platform. The mesh represents a sample spar design, not the original OC3 design.

The platform's underwater surface is discretised using quadrilateral panels on lateral and bottom faces. Mesh refinement followed the Nyquist-Shannon sampling criterion (Shannon, 1949), with at least four panels per wavelength. A sample mesh is shown in Fig. 5.

Viscous hydrodynamic loads are introduced by including a viscous damping coefficient in the total damping term in Eq. (3). The coefficient is obtained from a stochastic linearisation of the drag term in Morison's equation, as proposed by Borgman in Borgman (1967). For each horizontal slice of the floater (dL), viscous drag (dF_d) is approximated by:

$$dF_d/dL = 1/2 \cdot c_D \rho D \sqrt{8/\pi} \sigma_u u, \quad (4)$$

Here, c_D is the drag coefficient, D is the local diameter of the floater, u and σ_u are the local relative normal velocity and its standard deviation. The drag coefficient was assumed to be 0.6, as characteristic for a cylinder in high Reynolds number oscillatory flow in moderate and severe conditions (Jonkman, 2010). The standard deviation of velocity is calculated from the velocity response spectrum:

$$\sigma_u = \left(\int_0^\infty S_u(\omega) d\omega \right)^{1/2} \quad (5)$$

which, in turn, is obtained as a product of the square of the velocity RAO and the JONSWAP sea spectrum:

$$S_u(\omega) = |RAO_u(\omega)|^2 S_\zeta(\omega) \quad (6)$$

The velocity RAO is obtained from the displacement RAO:

$$RAO_u = \omega \cdot RAO_\eta \quad (7)$$

The displacement RAO is calculated from Eq. (3) by dividing the motion response amplitude by the wave amplitude:

$$RAO_{\eta} = \eta/a \quad (8)$$

The water particle velocity contribution to the relative velocity is ignored following the assumption of linear wave kinematics. Viscous damping is then computed as:

$$B_v = dF_d/du \quad (9)$$

Since viscous damping depends on the platform's response statistics, and the response calculation requires viscous drag as input, viscous damping and motion response (RAOs) are solved iteratively (full solution is outlined in Patryniak and Collu (2023)).

In Eq. (3), $\hat{F}_{hydro}(\omega)$ and $\hat{F}_{aero}(\omega)$ are the hydrodynamic and aerodynamic fluctuating loads, respectively. More specifically, $\hat{F}_{hydro}(\omega)$ is the frequency- and direction-dependent complex wave excitation matrix, obtained as a product of the first order wave excitation coefficient $X(\omega)$ from the diffraction problem solution, and the wave amplitude in the frequency domain $\hat{\xi}(\omega)$:

$$\hat{F}_{hydro}(\omega) = \hat{\xi}(\omega)X(\omega) \quad (10)$$

Second-order hydrodynamics are not considered, supported by Jonkman (2010), Roald et al. (2013) and Duarte et al. (2014). The complex wave amplitude is derived with amplitudes from the two-parameter JONSWAP spectrum $S_J(\omega)$ and random phases $\phi(\omega)$:

$$\hat{\xi}(\omega) = \sqrt{2S_J(\omega)d\omega} \cdot \exp(i\phi(\omega)) \quad (11)$$

$\hat{F}_{aero}(\omega)$ in Eq. (3) is the frequency-dependent complex turbulent wind excitation. It is derived based on the fast Fourier transform (FFT) of the fluctuating aerodynamic load time series precomputed using OpenFAST (National Renewable Energy Laboratory, 2020) for a range of mean wind speeds, turbulence intensity factors, and rotor plane inclination angles, with all rigid and flexible DOFs disabled (therefore specific to a given rotor, but agnostic to the support structure design). In the simulations, the controller was disabled, and the rotor speed and blade pitch angle were set to the values from the control table, interpolated based on the mean wind speed. Directly precomputing the aerodynamic load and not only the wind inflow velocities was possible, as this study concerns a fixed rotor design. One of the benefits of this approach is that some of the nonlinearities in the wind load could be accounted for in the otherwise linear approach.

The complex motions in the frequency domain were transformed to the time domain by performing the inverse Fourier transform (iFFT) of $\hat{\eta}(\omega)$:

$$\eta(t) = \mathcal{F}^{-1}(\hat{\eta}(\omega)) \quad (12)$$

3.3. Loads assessment

The motion time series obtained from the dynamic model is used to directly compute the system loads, specifically the tower base bending moment and the most loaded mooring line tension at the fairlead, as shown in Fig. 6. These two locations are considered critical from the perspective of ultimate and fatigue limit states assessment.

The calculation of the tower base bending moment involves summing up the contributions of the aerodynamic load on the rotor and tower, transformed to the tower base-fixed coordinate system, as well as the contributions from the inertial effects of the acceleration and inclination of the structures above the tower base, which include the tower itself and the rotor-nacelle assembly. The computation of the mean bending moment relies on the input of the mean inflow velocity, aerodynamic load, and pitch displacement. In contrast, the dynamic tower load calculation requires the input of the time series of the aerodynamic load, pitch displacement, and acceleration, all computed at the mean position of the system. The mean moment due to aerodynamic

drag on the tower is computed based on the integration of the drag along the tower with the drag coefficient $c_d = 1$ (Det Norske Veritas (DNV), 2014) and wind shear exponent $\alpha = 0.2^2$:

$$dM_d(z) = \frac{1}{2} \rho c_d V_{ref}^2 \left(\frac{z}{z_{ref}} \right)^{2\alpha} dA(z) z \quad (13)$$

V_{ref} is the reference wind speed at hub height z_{ref} . The instantaneous moment due to the acceleration of the i th mass item located at the height z_i above the tower base is formulated as:

$$M_{accel,i}(t) = m_i a_i(t) z_i \quad (14)$$

The instantaneous moment due to the inclination of the weight item N_i over the tower base is formulated as:

$$M_{incl,i}(t) = z_i (\sin(\theta + \delta\theta(t)) - \sin(\theta)) N_i \quad (15)$$

accounting for the reduction of the effective lever of the weight with the mean pitch angle.

The mooring loads were computed by feeding the motion time series to the quasi-static mooring model (MoorPy). This method computes the mooring loads at each time step, accounting for the mooring tension nonlinearity with respect to the platform motion (although it does not model the two-way motion-load coupling that a purely time-domain model would).

3.4. Fatigue assessment

The fatigue assessment follows the S-N approach suggested by DNV (2013), in which the number of stress cycles n occurring at critical areas of the structure at a given hotspot stress range is compared with the number of cycles to failure N . The number of cycles to failure for each stress range $\Delta\sigma$ is obtained from fatigue test results reported in DNVGL (2016) as curves characterised by the intercept of the $\log(N)$ axis ($\log(a)$) and the negative inverse slope (m):

$$\log(N) = \log(a) - m \log(\Delta\sigma) \quad (16)$$

Using the available load time signals, the rainflow counting method is applied to directly count the stress cycles. A stress range histogram is generated, and the fatigue damage accumulated over a single sea state (1-hour simulation) is computed by summing the contributions from k histogram bins (linear accumulation), following the Palmgren-Miner rule, augmented by the wave directionality correction factor R_μ :

$$D_{ss} = \sum_{i=1}^k \frac{n_i}{N_i} R_\mu = \frac{1}{a} \sum_{i=1}^k n_i \Delta\sigma_i^m R_\mu \quad (17)$$

The cumulative fatigue damage over the design life of the structure is then obtained by summing the products of the one-hour damage, the design life in hours (T_d), and the probability of occurrence of each sea state (p_{ss}):

$$D = T_d \sum_{ss=1}^l D_{ss} p_{ss} \quad (18)$$

Note that the verification presented in this paper considers one environmental condition at a time ($l = 1$, $p_{ss} = 1$). The parameters of the adopted S-N curves are listed in Table 4. The C1 curve from DNVGL (2016) is applied for the tower, assuming good-quality welded joints in air, while the curve for the mooring chain is taken from DNV-OS-E301 (DNV, 2013) for a studless chain.

² The results presented in this paper were obtained assuming $\alpha = 0.2$, as recommended by the IEC 61400-1:2019 Section 6.3.2.2 (International Electrotechnical Commission (IEC), 2019) for the normal wind model. However, lower values, 0.11–0.14, are typically assumed and recommended in more recent editions of standards.




Fig. 6. The process behind model C2: from frequency-domain motion to time-domain loads.

Table 4

Relevant S-N curves parameters.

Location	Material	m	a
Tower base	Welded joint	3.0	2.8119×10^{12}
Fairlead	Studless chain	3.0	6.0×10^{10}

Mooring lines. This analysis focuses on tension–tension fatigue, excluding the out-of-plane bending effects. The nominal stress range (double amplitude) in the i th environmental condition $\sigma_{S,i}$ is given as the tension range ($\sigma_{T,i}$) divided by the nominal cross-section area of the chain (corrected by corrosion):

$$d_{corr} = d - 0.1 \frac{mm}{year} \cdot 25 \text{ year} \quad (19)$$

$$\sigma_{S,i} = \sigma_{T,i} / (2\pi d_{corr}^2 / 4) \quad (20)$$

Tower base. The nominal stress at the tower base is calculated based on DNV-GL-RP-C202 recommendation as the difference between the axial compressive and tensile stress:

$$\sigma_{S_i} = \sigma_{S,i,bending} - \sigma_{S,i,comp} \quad (21)$$

$$\sigma_{S,i,bending} = \frac{r_i M_{y,i}}{I_{y,i}} \quad (22)$$

$$\sigma_{S,comp} = M a_z \cos(\bar{\theta}) / A_i \quad (23)$$

Then, the hotspot stress range $\Delta\sigma$ is calculated as the product of the nominal stress and the appropriate stress concentration factor (SCF):

$$\Delta\sigma = \sigma_{S,i} SCF \quad (24)$$

The SCF is applied to account for geometric discontinuity (e.g., slope between two sections of varying diameters), complexity of connections, and manufacturing tolerances, in accordance with DNVGL (2016). For the tower base of the current design $SCF \approx 1.22$.

3.5. Verification

The environmental conditions for the verification of the model are outlined in Table 3 in Section 2.1. To reiterate, cases 1.x relate to wave-only loading, cases 2.x relate to wind-only cases, and cases 3.x relate to combined wind-wave loading of increasing severity. The detailed numerical results for all quantities discussed in this section are presented in Tables A.18–A.19 in Appendix.

3.5.1. Motion response

As illustrated in Fig. 7, the mean motion responses from the current model (C1) closely matched the OpenFAST results (OF), with a slight overprediction of the offsets (errors below 5.0%), attributed to the overprediction of the mean aerodynamic load.

As illustrated in Fig. 8, the prediction of the wave-driven motion response (cases 1.x) matched the benchmark results well, with discrepancies generally below 10%. The discrepancies remained similar for both significant wave height cases considered, suggesting that the nonlinearity between the hydrodynamic loads and platform motion does not play a major role.

In wind-only (low-frequency) loading scenarios (cases 2.x), the misprediction of the motions remained below 10.0%, with the highest errors in the near-rated wind conditions.³ This is primarily due to the invalidity of the linearised approach at large surge offsets. Unlike wave-only cases, wind exerts a significant mean load on the floating system, leading to a notable platform offset. Additionally, wind loads concentrate most of their energy in the very low-frequency range, which partly coincides with the surge natural frequency. This overlap amplifies the response near that frequency, causing significant dynamic surge motions around the mean. As illustrated by the force–displacement relationship in Fig. 9, larger mean offsets are associated with larger variations in mooring tension due to the nonlinear characteristics of the mooring system. Consequently, the validity of representing the mooring restoring behaviour using a constant linearised stiffness diminishes as the operating point moves further into the nonlinear region of the force–displacement curve, leading to increased discrepancies between the benchmark nonlinear model and the present linear formulation.

3.5.2. Loads

As summarised in Fig. 10, the mean mooring tension closely matches the benchmark results, owing to the model's ability to account for the nonlinearity of the loads in the motion response. The mean tower base bending moment is overestimated, which is attributed to the overprediction of the mean aerodynamic load (seen in Fig. 7) and the underprediction of damping. While both the aerodynamic precomputation for the current model and the OpenFAST solution utilised the same aerodynamic model, the precomputation was performed for a fixed platform; therefore, the coupling between platform motion and aerodynamic loads has not been fully accounted for. While OpenFAST resolves this coupling at each time step, the current model only partly accounts for it by including the precomputed aerodynamic damping term in the frequency-domain equation of motion.

³ Although the heave motion amplitude relative error reached 13.1%, the absolute error was insignificant at 0.01 m.




Fig. 7. Verification of the motion response in operational cases. Mean values. See Table 3 for reference.




Fig. 8. Verification of the motion response in wave-only, wind-only, and operational cases. Standard deviation. See Table 3 for reference.




Fig. 9. Mooring force-displacement curve obtained from the response to turbulent wind at different mean wind speeds and fixed turbulence intensity of 14%.

Regarding the dynamic response, Fig. 11 shows mooring tension overestimation, with a relative error of 13.1% in the rated wind condition (case 3.2). The tower base bending moment, in contrast, is predicted accurately: the transformation of the motions into the time domain allows the nonlinear relationship between the load and the pitch motion to be accounted for, and the error remains below 4% for all cases. The remaining discrepancies are attributed to the overprediction of the low-frequency, wind-driven loads.

To support the above discussion, Fig. 12 presents a comparison of the motion and load time series, as well as their corresponding Fourier amplitudes, for the $8.5 m/s$ wind case (case 3.1). All motion time signals generally agree well with those from OpenFAST, with a slight




Fig. 10. Verification of the loads in operational cases. Mean values.

overestimation of low-frequency surge motion and an underestimation of heave motion in the 0.025–0.05 Hz range, which corresponds to the most energetic part of the turbulent wind spectrum. It is worth noting that the absolute discrepancies in heave motion are very small in absolute terms and have a negligible influence on the resulting loads.

As expected from the quasi-static mooring model, the mooring tension solution fails to capture the wave frequency load, although it is significantly lower than the low-frequency contribution, which is predicted accurately. The main discrepancy in the tower base bending moment is the mean value, while the dynamic amplitudes match very well, as seen in both the time series and frequency amplitudes (also confirmed by comparison to Figs. 10 and 11).

3.5.3. Fatigue

The small errors in the time series of loads had an effect on the fatigue assessment. Figs. 13–15 compare the 1 h mooring tension time series and the rainflow counting results (stress versus number of cycles




Fig. 11. Verification of the loads in operational cases. Standard deviation.




Fig. 12. Verification of the responses — time series (left) and Fourier amplitudes (right) for in-plane platform displacements, nacelle horizontal acceleration, line 2 tension at the fairlead, and tower base bending moment.




Fig. 13. Tension and stress cycles counting for the mooring line (fairlead). Case 3.1 (below rated conditions).

curves) between the current model and OpenFAST, for three operational cases (3.1, 3.2, 3.3).

The over- and underprediction of the tension in the low and wave frequency ranges had opposite effects on the fatigue analysis: the signal of the current model included more cycles of high stress and fewer cycles of low stress. These two effects balanced each other out, resulting in a good general prediction of fatigue (relative errors of -3.4% , 31.3% , and -18.4% in cases 3.1, 3.2, and 3.3). Depending on the level of over- or under-estimation in a particular condition, and the probability distribution of the conditions over the lifetime of the structure, the net effect can either be an over- or under-estimation of the total fatigue.

Figs. 16–18 compare the 1 h tower base bending stress time series and the rainflow counting results between the current model and OpenFAST, for the same three operational cases.

Although there are some differences in predicting the tower base bending moment time history (as shown in the left-hand subplots), these are less significant than in the case of the mooring tension, and so the model counts the load cycles more accurately (relative errors in fatigue of 9.1% , 8.7% , and 6.8% in cases 3.1, 3.2, and 3.3). The model consistently overestimates the fatigue, which is a desired, conservative behaviour. Unlike for the mooring lines, a higher cycle count is observed at most stress levels (as shown in the right-hand subplots).

Overall, the hybrid frequency/time-domain model predicts the motions and tower base loads accurately, without significant computational cost. The mooring tension, however, exhibits considerable error, driven mainly by underprediction at the wave frequency, which propagates into the fatigue analysis and markedly reduces its accuracy. This underprediction could be mitigated by applying a dynamic mooring





Fig. 14. Tension and stress cycles counting for the mooring line (fairlead). Case 3.2 (rated conditions).





Fig. 15. Tension and stress cycles counting for the mooring line (fairlead). Case 3.3 (above rated conditions).





Fig. 16. Bending stress and stress cycles counting for the tower base section. Case 3.1 (below rated conditions).





Fig. 17. Bending stress and stress cycles counting for the tower base section. Case 3.2 (rated conditions).

model (for example, a lumped-mass model) to the motion time series, instead of the quasi-static model; however, the associated computational cost would become a bottleneck in the optimisation context. The next section discusses the development of a data-driven model to address these shortcomings without increasing the solution time.

4. Surrogate models (C2)

The limitations of the physics-based model, including excessive computational effort related to the computation of hydrodynamic coefficients and the misprediction of mooring tension, are addressed by




Fig. 18. Bending stress and stress cycles counting for the tower base section. Case 3.3 (above rated conditions).




Fig. 19. Physics-based and data-driven models integration.

introducing two surrogate models into the workflow depicted in Fig. 19. The development and verification of the surrogate models are detailed in the next two subsections.

4.1. Potential hydrodynamics surrogate

The computationally expensive boundary element method solver was replaced with a surrogate model that computes the first-order hydrodynamic potential coefficients in the frequency domain based on the input of the spar floating platform geometry features: draft, nine radii along the draft, centre of buoyancy, and system mass. Note that this surrogate was developed specifically for spar-type geometries; it is naturally agnostic to the environmental conditions but not to the floater type. Therefore, it is applicable to spar-type platforms only.

4.1.1. Dataset generation

This section presents the inputs and outputs of the dataset, as outlined in Fig. 20, as well as the dynamic model used to map the inputs to the outputs.

Dataset input. The independent variables, consisting of the floating platform draft and nine radii, were defined over the ranges listed in Table 5. For each design, the dependent parameters: mass and centre of buoyancy, were computed. Each design was subsequently evaluated across a non-uniform frequency vector: frequencies from 0.05 rad s^{-1} to 0.18 rad s^{-1} were sampled at an interval of 0.05 rad s^{-1} , frequencies from 0.18 rad s^{-1} to 1.25 rad s^{-1} were sampled at an interval of 0.01 rad s^{-1} , while the range from 0.18 rad s^{-1} to 2.5 rad s^{-1} , where hydrodynamic coefficients gradients are smaller, was sampled at an interval of 0.5 rad s^{-1} . A full factorial design of these parameters would yield $10 \times 7^9 = 403\,536\,070$ potential configurations, each to be evaluated at 113 frequencies. From this large design space, 7 million samples were randomly selected to ensure broad coverage while maintaining computational feasibility.

The platform draft is capped at the reference value of 120 m, as deeper-drafted structures would be challenging to manufacture and

Table 5

Ranges of the inputs of the dataset for potential coefficients surrogate training.

Input	Min	Max	Interval
Independent parameters:			
Draft (m)	75	120	5
9 radii (m)	3	9	1
Frequency (rad s^{-1})	0.05	2.5	Non-uniform (113 items)
Dependent parameters:			
Mass (kg)	$4.90\text{E}+06$	$23.19\text{E}+06$	–
Centre of buoyancy (m)	–80.85	–26.28	–

deploy given port restrictions. Drafts below approximately 68 m would result in a negative total pitch stiffness C_{55} (the sum of hydrostatic and gravitational contributions), leading to a loss of initial stability. For context, the shallowest spar-type FOWT currently installed, Hywind Scotland, has a draft of about 78 m, while typical designs range between 90 m and 120 m (e.g., OC3-Hywind reference design, Hywind Demo, Hywind Tampen, UMaine spar). Consequently, the draft range is broadly defined as 75 m–120 m. Platform radii are limited to slightly smaller than that of the reference design (OC3 spar) and up to approximately 50 % larger than its maximum diameter, which remains within the range commonly considered feasible for offshore steel structures. As summarised in Table 5, this results in a diverse set of platform geometries spanning a wide range of mass and centre of buoyancy values.

Dataset output. The output consisted of the diagonal (uncoupled) and the surge-pitch coupled components of the added mass, radiation damping, and wave excitation coefficients, as summarised in Table 6.

Dataset characterisation. The 13 input features and 14 target coefficients were normalised to zero mean and unit variance prior to training using a standard scaler. Fig. 21 presents the marginal distributions of the input and output variables, based on a representative 5 % subset




Fig. 20. Database generation process.

Table 6

BEM surrogate dataset outputs (hydrodynamic coefficients).

Output	Components	Units	Description
$Ah_{i,j}$	$(i, j) \in \{(1, 1), (1, 5), (3, 3), (5, 5)\}$	kg, kg m ²	Added mass
$Bh_{i,j}$	$(i, j) \in \{(1, 1), (1, 5), (3, 3), (5, 5)\}$	kg s ⁻¹ , kg m ² s ⁻¹	Radiation damping
$\Re(X_i)$	$i \in \{1, 3, 5\}$	N m ⁻¹	Potential excitation (real)
$\Im(X_i)$	$i \in \{1, 3, 5\}$	N m ⁻¹	Potential excitation (imaginary)

of the dataset. Only two radii are shown, as all radius variables exhibit similar distributions. The geometric input parameters (radii and draft) display approximately uniform coverage across their respective ranges. The derived parameters (mass and vertical centre of buoyancy) exhibit approximately normal distributions, whereas the hydrodynamic coefficients display distinctly non-Gaussian distributions, reflecting the underlying nonlinear hydrodynamic dependencies.

To visualise the sampling uniformity and input–output relationships, pairwise scatter plots of selected inputs and outputs are shown in Fig. 22. The input variables exhibit broad and uniform coverage across the sampled space, confirming that the random sampling strategy effectively spans the design domain. The outputs display strongly nonlinear relationships with the inputs. The absence of significant clustering or empty regions indicates that the dataset provides adequate coverage for surrogate model training.

Dynamic simulations. For each sampled configuration, frequency-domain radiation–diffraction simulations were performed using the BEM solver Capytaine (Ancellin and Dias, 2019), as detailed in Section 3. The procedure involved meshing according to the defined parameters, computing the dependent variables (mass and centre of buoyancy), running the simulations, and post-processing the results to extract the required hydrodynamic coefficients. Evaluating 7 million samples, where each design was discretised with $\mathcal{O}(10^3)$ panels, required about 800 h on 60 cores of an Intel(R) Xeon(R) Silver 4216 CPU at 2.1 GHz. Since the hydrodynamic linear coefficients are independent of the platform motion, the derived surrogate is broadly applicable to projects with similar platform topologies (independent of environmental conditions, mooring design, etc.). Thus, this computational effort represents a one-time cost.

4.1.2. Model development

This section describes the data-driven, machine-learning-based model developed to predict the variables reported in Table 6 (14 scalar outputs) from the variables listed in Table 5 (13 scalar inputs). The model will be constructed using the dataset introduced in the previous section, which provides approximately 7 million examples of the input–output relation. This problem can be formulated as a multi-output regression task in machine learning, for which many algorithms with associated hyperparameters are available (Shalev-Shwartz and

Ben-David, 2014; Bishop and Bishop, 2023). According to the no-free-lunch theorem, no algorithm can be chosen *a priori* without empirical evaluation on the data (Adam et al., 2019). Nevertheless, experience and knowledge of the properties of different algorithms can guide some initial *a priori* choices (Shalev-Shwartz and Ben-David, 2014; Bishop and Bishop, 2023).

Given the tabular nature of the data under consideration, shallow machine-learning models (Shalev-Shwartz and Ben-David, 2014), such as shallow neural networks (Bishop and Bishop, 2023) and tree-based ensemble methods (Zhou, 2025), are widely recognised for their strong inductive bias and competitive performance on tabular problems, and were therefore considered as candidate approaches.

The suitability of the candidate methods was assessed empirically in terms of predictive performance and the computational feasibility of the complete model-development procedure. The experiments were performed on a Google Cloud *c4-standard-16* virtual machine equipped with 16 CPU cores, 60 GiB of RAM, and an NVIDIA L40 GPU. The complete neural-network development workflow, including the prescribed hyperparameter search, final architecture selection, and training on approximately seven million examples, required approximately 78 h.

By comparison, XGBoost required approximately 72 h to train one hyperparameter configuration on a subset of 400,000 examples, corresponding to approximately 5.7% of the complete dataset. The resulting mean absolute errors were approximately three to five times larger than those of the selected neural-network surrogate. This accuracy comparison should be interpreted cautiously because the XGBoost experiment used fewer training data and only one hyperparameter configuration. It nevertheless showed an unfavourable accuracy–computational-cost trade-off for the tested configuration.

When XGBoost was trained on the complete dataset, one hyperparameter configuration had not completed after 168 h and the run was terminated. A corresponding full-scale LightGBM configuration likewise did not complete within the seven-day computational window. Consequently, a systematic hyperparameter search using either tree-based implementation was not computationally feasible within the resources available for this study. Since the full-scale models did not complete training, their full-dataset predictive performance and inference times could not be measured without unsupported extrapolation. As no comparison is possible, the inference-time test results for




Fig. 21. BEM surrogate dataset input and output parameters distributions.

the neural network are therefore presented in the model verification (Section 4.1.3).

The reported results should not be interpreted as a general claim that neural networks are superior to gradient-boosting methods for tabular data. They establish that, for the present seven-million-example, 14-output regression problem and computational platform, the neural-network approach was the only evaluated method that completed the required full-scale hyperparameter-selection and training procedure within a practicable time.

In addition to the above, it is worth noting that shallow neural networks can handle large datasets by fully exploiting parallel architectures and hardware accelerators, such as state-of-the-art GPUs (Silvano et al., 2025), and can predict all 14 output quantities simultaneously through a shared representation. In contrast, the tree-based methods do not natively support the coupled, 14-dimensional output of

the present problem: they would require either training 14 independent models or resorting to multi-output extensions with more limited functionality, whereas the hydrodynamic coefficients considered here are physically correlated and should vary jointly and smoothly with platform parameters and frequency.

All considered, for this work, we adopted a shallow neural network with a fully connected architecture composed of: an input layer of dimension 13; a tunable number $\{1, 2, 4\}$ of hidden fully connected layers, each with a tunable width in $\{1, 2, 4, 8, 16\} \times 10^3$ neurons and activation function chosen from $\{\text{ReLU}, \text{sigmoid}\}$; and a final fully connected layer with linear activation producing the 14-dimensional output. The network was trained using the Adam optimiser with a tunable learning rate in $\{0.001, 0.002, 0.005, 0.01\}$ for a tunable number of epochs in $\{5, 6, \dots, 100\}$, minimising the mean squared error between predicted




Fig. 22. BEM surrogate dataset output versus input parameters scatters.

and actual values. It is important to note that, given the extremely large size of the dataset, which required several days of training on an NVIDIA A100 GPU equipped with 40 GB of RAM, we could not explore a larger set of architectural or training hyperparameters.

In order to tune the free hyperparameters, namely the number of hidden layers, the number of neurons in each hidden layer, the activation functions, and the learning rate, and to estimate the performance of the final model, we employ a state-of-the-art hold-out resampling

strategy (Oneto, 2020). Given the very large size of the dataset (approximately 7 million examples), more advanced approaches such as cross-validation are unnecessary and, in any case, computationally impractical. The approach is based on a simple yet robust idea (Oneto, 2020): for each iteration, the original dataset is split without replacement into three independent subsets. These subsets are the training set (approximately 6 million samples), used to learn the model; the validation set (approximately 0.5 million samples), used to tune the

Table 7

Performance metrics used for surrogate model evaluation. y_i and $\hat{y}_i$ – true and predicted values of the output, $\bar{y}$ – mean of the true values, N – number of samples.

Metric	Description	Formulation
MSE	Average squared deviation between predicted and true values; sensitive to large errors.	$MSE = \frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2$
RMSE	Square root of MSE; expressed in the same units as the output variable.	$RMSE = \sqrt{MSE}$
MAE	Average absolute deviation; less sensitive to outliers than RMSE.	$MAE = \frac{1}{N} \sum_{i=1}^N y_i - \hat{y}_i $
R^2	Indicates how much of the variance in the true values is explained by the model; values closer to 1 denote better agreement.	$R^2 = 1 - \frac{\sum_{i=1}^N (y_i - \hat{y}_i)^2}{\sum_{i=1}^N (y_i - \bar{y})^2}$
NRMSE	RMSE normalised by the range or mean of the true values; provides a scale-independent accuracy measure.	$NRMSE_{\text{range}} = \frac{RMSE}{y_{\text{max}} - y_{\text{min}}}$ $NRMSE_{\text{mean}} = \frac{RMSE}{\bar{y}}$

Table 8

BEM coefficients surrogate aggregate error metrics grouped by mode.

Metric	Surge	Heave	Pitch	Surge-pitch
MSE	3.01e+09	2.71e+07	1.39e+17	3.3e+13
RMSE	3.46e+04	4.88e+03	1.87e+08	4.12e+06
MAE	3.08e+04	3.08e+03	1.71e+08	3.79e+06
R^2	0.999	0.999	0.999	0.999
NRMSE (range)	0.00417	0.00409	0.00456	0.00568
NRMSE (mean)	0.0161	0.0448	0.0188	0.0168

hyperparameters; and the test set (approximately 0.5 million samples), used to assess the performance of the final model.

In this work, it is not only important to achieve high accuracy but also to ensure fast prediction times. This implies that smaller networks yield faster inference, although often at the cost of reduced accuracy (Bishop and Bishop, 2023). For this reason, we search for the smallest (i.e., fastest) network whose mean squared error increases by no more than 5 % with respect to the best-performing network in the validation search space. This 5 % threshold was chosen empirically, as it provides a good trade-off between accuracy and prediction speed.

4.1.3. Model verification

On the test set, performance can be evaluated using both quantitative metrics (Naser and Alavi, 2023) (e.g., prediction time, mean squared error, mean absolute error, or coefficient of determination), as defined in Table 7 and qualitative (Sainani, 2016) (e.g., scatter plots, histograms of the errors, or trends with respect to individual inputs).

The metrics values are reported in Table 8, grouped by the mode (surge, heave, pitch, and surge-pitch). The surrogate model achieved near-perfect agreement with the reference hydrodynamic data, with coefficients of determination (R^2) exceeding 0.999 across all motion modes (Table 8). The range-normalised errors ($NRMSE_{\text{range}} < 0.6\%$) and mean-normalised errors ($NRMSE_{\text{mean}} < 5\%$) indicate excellent consistency throughout the training domain.

To visualise local variations in model performance, scatter plots of true versus predicted values were generated across different modes: surge, heave, pitch, and surge-pitch coupled mode (Fig. 23). Perfect agreement would correspond to points lying on the 45° line. These plots provide an intuitive indication of the global goodness of fit and reveal any systematic trends or bias. To complement this view, histograms of the absolute relative errors are shown alongside the scatter plots. The plots were generated based on 5 % randomly selected samples. Frequency was included as one of the model inputs, alongside the platform parameters, meaning each point in the plot corresponds to the prediction for a specific spar design at a specific frequency.

On average, the predictions align closely with the true values, except for the excitation coefficient (X) and radiation damping (B_h) in the heave direction, where significant deviations from the ideal fit are observed. $X_{im,3}$ and $B_{h,33}$ were therefore further investigated through

additional diagnostics, including a set of plots that examine the residual behaviour and distributional agreement between the true and predicted values.

For B_{33} , Fig. 24 shows the residual distribution (top-right panel) with a small systematic bias, indicating a mild tendency of the surrogate to underpredict the radiation damping coefficient $B_{h,33}$. A subtle skewness is also visible, with a heavier left tail and lighter right tail, suggesting that negative residuals (underpredictions) occur more frequently or with larger magnitude. This behaviour is consistent between the residual histogram and the QQ plot, where deviations from the reference line at the lower quantiles further confirm the asymmetry of the residual distribution. Nevertheless, these effects are small in magnitude and do not significantly compromise the overall fidelity of the surrogate.

The diagnostics for $X_{im,3}$ are similar (Fig. 25), except that the bias runs in the opposite direction: a slight positive shift in the residual mean and a modest right-tail skew indicate a minor overprediction tendency, rather than the underprediction seen for $B_{h,33}$. The deviations are again small and within acceptable limits (see Fig. 25).

Additionally, the performance of the surrogate model across the full frequency range for each hydrodynamic coefficient was evaluated using the NRMSE metric. To illustrate the model's ability to reproduce design trends, the NRMSE was plotted against one design variable (e.g., draft) while keeping the other (e.g., waterline radius) fixed, as shown in Figs. 26 and 27.

An acceptably low NRMSE was observed for most coefficients across the design space, except for the heave components of radiation damping and the real and imaginary parts of wave excitation (B_{33} , $X_{re,3}$, $X_{im,3}$), particularly in designs with a small waterplane area (radius). However, these discrepancies occur in coefficients that contribute minimally to the total response of the FOWT and therefore do not compromise the model's overall usefulness. Specifically, for the slender spar-type platform, potential damping is small compared to additional linear⁴ and viscous damping, and wave excitation is small compared to aerodynamic excitation.

Nonetheless, to assess the impact of these discrepancies, the motion and load responses of the FOWT were evaluated in three operational conditions, as shown in Fig. 28 and Table 9, by comparing the baseline physics-based frequency-domain model with the version using the BEM surrogate. The resulting errors in total response were generally very small, confirming the validity of the surrogate implementation. Notably, the total solution time was significantly reduced—an essential factor for effective design optimisation.

Finally, to characterise the computational cost of the BEM surrogate at inference time, we benchmarked it on a workstation with an Intel

⁴ Additional linear damping is added to account for the contributions to hydrodynamic damping other than the linear radiation damping and nonlinear viscous drag, as observed in free-decay tests (Jonkman, 2010)




Fig. 23. BEM surrogate true-predicted scatter plots with absolute relative error histograms.

Xeon Silver 4216 CPU (2×16 physical cores, 2.10 GHz). No data was loaded from disk during the timed calls, and no CPU-to-GPU transfer occurred, since inference was performed entirely on CPU. Python 3.11.13 and TensorFlow 2.20.0 are used. Model and scaler loading were performed once and excluded from the timing; each timed call included input scaling, forward prediction, and output inverse scaling. Two scenarios were tested, following 5 warm-up runs each: (i) a single

query (one design, one frequency), (ii) a batch of 75 frequencies—as normally used within the coupled model. Results are summarised in Table 10.

4.2. Mooring surrogate

As illustrated in the flow diagram of Fig. 19, the mooring dynamics are solved at three key stages of the coupled solution process:



Fig. 24. Diagnostic plots for $B_{h,33}$ prediction.Fig. 25. Diagnostic plots for $X_{im,3}$ prediction.

Table 9

Relative errors due to incorporating the BEM surrogate model in the frequency domain baseline model (C1).

Case	Platform motion (%)				Loads (%)	
	Surge	Heave	Pitch	Nac. Accel.	T2	My
3.1	1.00	0.52	0.08	-2.25	1.24	-0.38
3.2	1.37	-2.77	-0.15	-4.27	1.71	-0.39
3.3	1.39	-0.99	0.03	-3.23	2.00	-0.38

Table 10

BEM surrogate inference-time benchmark. Batch size 1: 4000 timed repetitions; batched call over 75 frequencies: 3000 repetitions.

	Batch size 1	Batched call (75 freq.)
Median	78.7 ms	86.3 ms
Mean	79.0 ms	86.7 ms
95th percentile	80.9 ms	88.5 ms
Per example (median)	78.7 ms	1.15 ms

1. Static response: solution of the mean motions of the platform requires calculating the mooring load on the floating platform. This is achieved using the quasi-static model MoorPy (see Section 3.1).

2. Dynamic response: solution of the dynamic motions of the platform around its mean position relies on the mooring stiffness computed at the mean position, which is also done using MoorPy.
3. Loads: given the platform motion, line tensions are determined using either a dynamic model (e.g., MoorDyn) or a quasi-static model (e.g., MoorPy).

In the following, the application of a machine learning method to predict mooring tension at the final stage is presented, addressing the inaccuracies of the physics-based model while avoiding the high computational cost of the time-domain dynamic simulations. The surrogate model maps catenary mooring system design parameters and fairlead motion time series to tension time series. The remainder of this section describes the database generation, model training, and verification process.

4.2.1. Dataset generation

A state-of-the-art numerical model is employed to obtain the FOWT motion responses and mooring load time series, thereby generating the synthetic database necessary for surrogate model training in the absence of experimental data. The database generation process consists of the following stages (also illustrated in Fig. 29):




Fig. 26. NRMSE between the physics-based and data-driven BEM model results (hydrodynamic coefficients), computed over all frequency components. Trends against waterline radius for different values of draft.




Fig. 27. NRMSE between the physics-based and data-driven BEM model results (hydrodynamic coefficients), computed over all frequency components. Trends against draft for different values of waterline radius.

1. Design of experiments (DOE): define a parameter space covering floating platform and mooring design variables, as well as environmental parameters, using a full factorial design of N samples.
2. Sampling: randomly select a subset of $M \ll N$ DOE samples,
3. Simulation: perform dynamic simulations for the selected cases,
4. Post-processing: conduct convergence checks and sanity checks on the simulation results,




Fig. 28. Comparison of the baseline FD model and the FD model with BEM surrogate in terms of the standard deviation of motions and loads. Performed for a reference design redefined with a reduced number of radii. See Table 9 for numeric results.




Fig. 29. Mooring surrogate database generation process.

Table 11

Mooring chain properties.

Property	Value	Unit
Grade	R3	—
Mass density per diameter squared (m_2)	2e4	$\text{kg m}^{-1} \text{m}^{-2}$
Stiffness per diameter squared (EA_2)	85.6e9	N m^{-2}
Stiffness per diameter cubed (EA_3)	-3.93e7	N m^{-3}
Normal drag coefficient	1.333	—
Axial drag coefficient	0.639	—
Normal added mass coefficient	1.0	—
Axial added mass coefficient	0.5	—

5. Feature and target selection: identify the parameters and responses to be used as features and targets, drawing from both DOE variables and dynamic simulation outputs.

Dataset input. Every mooring system in this database consists of three catenary lines of equal length, each made entirely of studless chain. The material and hydrodynamic properties are summarised in Table 11.

Based on these properties and the nominal chain diameter d , the linear mass density of the lines m and the extensional stiffness EA are computed based on the empirical expressions from NREL (2025):

$$m = m_2 d^2 \quad (25)$$

$$EA = EA_2 d^2 + EA_3 d^3 \quad (26)$$

Keeping the material and hydrodynamic chain properties fixed for all designs, six parameters are needed to uniquely define each variant of the mooring design: nominal diameter of the chain d_{moor} (i.e., the diameter of the rod used to produce the chain), line length l_{moor} (measured along the line from fairlead to the anchor), fairlead z-coordinate z_F (point of attachment of a mooring line to the platform), fairlead radius r_F , anchor z-coordinate z_A (equal to the water depth), and anchor radius r_A , all illustrated in Fig. 30.

Design of experiments (DOE). The design space is sampled based on seven scalar variables representing the geometric properties of the




Fig. 30. Mooring line parametrisation — parameters of the design of experiments and ML features.

platform and the mooring design, as well as four scalar variables representing the environmental loading conditions, as listed in Table 12. Full factorial DOE is used to sample the space evenly, resulting in 25 600 design configurations and 108 environmental conditions (2 764 800 combinations in total). Out of these, a random selection of 300 000 samples is performed to limit the computational effort of database generation.⁵ Initially, the parameter ranges were intentionally broad to maximise the model's general applicability. However, broad ranges demand a prohibitively large number of samples; simulating only a small subset leaves the dataset too coarse to capture the nonlinear dependence of line tension on the inputs, and thus insufficient to train

⁵ Further to that, some simulations will fail to converge due to an ill-defined system, conflicting variables, or too harsh an environmental loading; therefore, the set of valid samples for training purposes will be smaller than 300 000. In the current case, 170 000 samples were used.

Table 12

Primary platform, mooring system, and environmental parameters (design of experiments).

	Min	Max	Units	Count
Platform and mooring design variables				
Nominal chain diameter (d_{moor})	0.056	0.084	m	5
Fairlead radius (r_F)	4.16	6.24	m	4
Fairlead z-coordinate (z_F)	-84.0	-56.0	m	4
Anchor radius (r_A)	676.16	929.72	m	4
Anchor z-coordinate (z_A)	-330.0	-260.0	m	4
Line length (l_{moor})	721.76	1082.64	m	5
Platform vertical centre of mass (z_{CM})	-98.91	-71.93	m	4
Environmental conditions variables				
Mean wind speed (U)	3.0	25.0	m s^{-1}	4
Turbulence intensity (TI)	12	16	–	3
Significant wave height (H_s)	0.5	6.0	m	3
Peak period (T_p)	6.0	10.0	s	3

Table 13

Mooring load surrogate model features.

Feature	Min (m)	Max (m)	Type
Nominal chain diameter (d)	0.056	0.084	Scalar
Line length (l)	721.76	1082.64	Scalar
Fairlead-to-anchor height (z)	204.0	246.0	Scalar
Fairlead-to-anchor radius (R)	672.0	923.48	Scalar
Horizontal motion at fairlead (x_F)	(–)	(–)	Time series 3600 s at 0.05 s
Vertical motion at fairlead (z_F)	(–)	(–)	Time series 3600 s at 0.05 s

a surrogate model with the required precision. To avoid this curse of dimensionality, the present study employs narrower ranges. While this limits general applicability,⁶ it enables a valid proof-of-concept. Ranges were set at $\pm 20\%$ of the original design's values, except for the anchor radius ($\pm 10\%$).

Features extraction. The central design choice of the mooring surrogate is that it is formulated from the fairlead point of view: although the dataset is generated using a specific platform, RNA, and controller, only the fairlead motion time series and the catenary mooring parameters are retained as features. The model is therefore independent of the system and environment that produced the motion, since inference likewise requires only a fairlead motion time series. The DOE was made broad enough to span a wide diversity of motion signals, keeping its application to other systems within the interpolation regime. However, the model does only remain valid for catenary mooring systems.

In line with this choice, the features used in the machine learning model are not identical to the design variables mentioned earlier. Instead, we select a specific set of features that directly influence the mooring line loads. Given the fixed properties listed in Table 11, the catenary mooring design can be uniquely represented by four features: line diameter, line length, radial fairlead-anchor distance, and vertical fairlead-anchor distance, which render the model independent of the platform design. Additionally, two features represent the time series of fairlead motion in the horizontal and vertical directions, making the model independent of environmental conditions. All six features are summarised in Table 13.

Dataset output. For the selected samples, numerical simulations are performed to evaluate the dynamic behaviour of the FOWT system under combined wind and wave loading. The analysis focuses on the horizontal and vertical motion of the fairlead, as well as on the corresponding horizontal and vertical loads in the most heavily loaded mooring line, as summarised in Table 14. The loads are evaluated at the first chain link connected to the fairlead.

Table 14

Mooring load surrogate model targets.

Target	Units	Type
Upwind line F_x at fairlead	kN	Time series 3600 s at 0.05 s
Upwind line F_z at fairlead	kN	Time series 3600 s at 0.05 s

Dataset characterisation. Fig. 31 shows the marginal distributions of the input and output variables, based on a representative 5 % subset of the full dataset. The shown inputs include four geometric parameters (chain diameter, length, and the fairlead-to-anchor vertical and radial distances) together with the mean and standard deviation of the fairlead motions. The shown outputs comprise the mean and standard deviation of the line load components. Overall, the dataset spans a wide range of mooring designs and motion conditions. The multimodal or skewed distributions reflect the discrete sampling of design parameters and the nonlinear nature of the mooring responses. All outputs show broad, asymmetric distributions. The relationships between the input and output parameters are illustrated in Fig. 32. For clarity, only one geometric input variable is shown, as the discrete sampling of the design space results in clustered input–output patterns that are not particularly informative. Importantly, the plots highlight the nonlinear relationships between motion statistics and the resulting tension and fatigue responses.

Dynamic simulations. All simulations are performed using a model of the spar-supported 5MW wind turbine (Jonkman, 2010; Jonkman et al., 2009), as detailed in Section 2.1. However, the features selected for the surrogate training are platform-design independent. Therefore, the choice of the floating platform concept and specific rotor does not affect the applicability of the surrogate to other designs. The samples are evaluated using OpenFAST: an extensively validated nonlinear aero-hydro-servo-elastic coupled time-domain model of dynamics (NREL, 2024a). The degrees of freedom include six rigid body motion modes. The NREL 5MW baseline variable-pitch variable-speed controller allows for blade pitch and rotor speed variations depending on the mean wind speed and slight rotor speed variations due to the system motion and turbulent inflow (Jonkman et al., 2009). The aerodynamic solution relies on the blade-element momentum theory, combined with the Beddoes–Leishman unsteady blade airfoil aerodynamics model, tip

⁶ The surrogate model is not meant to extrapolate well.




Fig. 31. Mooring surrogate dataset input and output parameters distributions.




Fig. 32. Mooring surrogate dataset output versus input parameters scatters.

and hub loss corrections, and tower influence, all based on potential flow (details in NREL, 2026). The first-order potential hydrodynamics of the spar platform are modelled through frequency-to-time-domain transforms based on the potential coefficients obtained from the boundary element method (BEM) code WAMIT (Lee and Newman, 2006). Viscous loads are computed from Morison's theory. Second-order hydrodynamics are not considered, supported by Jonkman (2010), Roald et al. (2013) and Duarte et al. (2014). The nonlinear mooring loads are computed using a lumped-mass mooring line model, MoorDyn (Hall and Goupee, 2015). The model incorporates the effects of axial stiffness and damping, as well as weight and buoyancy, hydrodynamic viscous forces, and vertical spring-damper forces resulting from contact with the seabed.

4.2.2. Model development

This section describes the data-driven, machine-learning-based model developed to predict the variables reported in Table 14 (two time series of 72 000 elements) from the variables listed in Table 13 (four scalars and two time series of 72 000 elements). The model will be constructed using the dataset introduced in the previous section, which provides approximately 170 000 examples of the input–output relation.

In this section, we do not repeat the concepts presented in Section 4.1.2, as most of them apply here in exactly the same way. The only difference here lies in the machine learning model that we chose to adopt based on the available data. In this case, we have a huge input (144 000 elements) and output (144 000 elements) space, which, moreover, are not simply scalar (four scalar inputs) but sequences (two input sequences and two output sequences), meaning that there is structure both in the input and in the output. Moreover, we have a large number of examples of this input–output relation (170 000 examples), representing almost 100 GB of data.

To handle this amount of data, analogous to what was done in Section 4.1.2, the only viable option is to use a neural network (Bishop and Bishop, 2023). Because of the nature of the input and output, the most reasonable choice would be to use architectures specifically developed for sequence data, such as Gated Recurrent Units (Chung et al., 2014), Long Short-Term Memory networks (Hochreiter and Schmidhuber, 1997), Temporal Convolutional Networks (Bai et al., 2018), Transformers (Vaswani et al., 2017), or Time Series Foundation Models (Das et al., 2024). Nevertheless, while attempting to rely on these types of networks, it soon became apparent that either their accuracy was too low or the size of the network (with associated memory and computational requirements) was too large to be handled, even with our hardware accelerators, while still achieving a prediction time suitable for subsequent use in optimisation.

As a consequence, we had to adopt a different approach, one that relied on the physical meaning of these inputs and outputs. The input series and a few scalar parameters represent the system and the excitation applied to it, which produce the output sequences. For this reason, given that the structure cannot have infinite memory, the output at time t can only be related to the scalar inputs plus the input series within a window around t , namely $[t - \Delta, t + \Delta]$, with Δ measured in seconds. Moreover, since the structure (FOWT characteristics) does not change over time, the response should not depend on t itself. Thanks to this observation, we only need to construct a neural network applied at time t whose inputs are the four scalars plus the elements of the input series in the window $[t - \Delta, t + \Delta]$, where Δ is tunable in $\{0.1, 0.5, 1, 2, 4, 8\}$ seconds. We then use a neural network exactly like the one adopted in Section 4.1.2, and as output, we predict the values at time t of the two output time series. Despite this reduction in dimensionality, the resulting dataset remained extremely large and required several days of training on an NVIDIA A100 GPU with 40 GB of RAM. Consequently, we were unable to explore a broader range of architectural or training hyperparameters, nor larger values of Δ .

In order to tune the free hyperparameters, Δ together with the same hyperparameters of the neural network described in Section 4.1.2,

Table 15

Mooring model surrogate aggregate error metrics.

Metric	F_x mean	F_x std	F_z mean	F_z std
MSE	6.3774	2.0043	2.5979	3.1995
RMSE	2.5253	1.4157	1.6118	1.7887
MAE	1.3005	0.8445	0.9065	1.0096
MAPE (%)	0.4045	9.9621	0.1699	14.5947
R^2	0.9998	0.9767	0.9999	0.9002
NRMSE (range)	0.0025	0.0251	0.0015	0.0431
NRMSE (mean)	0.0088	0.1157	0.0032	0.2116

namely the number of hidden layers, the number of neurons in each hidden layer, the activation functions, and the learning rate, and to estimate the performance of the final model, we follow the same procedure described in Section 4.1.2 where 80% of the data were used for training, 10% for validation, and 10% for test.

4.2.3. Model verification

Aggregate error measures across the full dataset are reported in Table 15 as concise numerical benchmarks of the surrogate's accuracy.

The surrogate model demonstrated excellent predictive performance for all response components, with $R^2 > 0.90$ across metrics. Errors in the mean force predictions were negligible (NRMSE < 0.003), indicating near-perfect agreement with the reference model. While the standard deviation predictions showed higher relative errors (up to about 15% in MAPE), the overall accuracy remains high, confirming that the surrogate effectively captures both the mean and variability of the mooring line forces.

The scatter plots of Fig. 33 confirmed the near-perfect agreement between predicted and true values for the mean components of the mooring loads. In contrast, the standard deviation predictions exhibit a broader scatter, particularly for the vertical load component, indicating comparatively lower accuracy in capturing the variability of the response. Nevertheless, the overall distributions of absolute relative errors remain concentrated near zero, demonstrating that large deviations are rare and that the surrogate model performs consistently well across all components.

Fig. 34 presents the residual-based diagnostic plots for the least accurate prediction case, namely the standard deviation of the vertical mooring load component, as introduced in Section 4.1.3.

The residuals exhibit a funnel-shaped pattern with increasing variance at higher predicted values, indicating mild heteroscedasticity and slightly reduced accuracy for larger responses. The residual distribution is approximately centred around zero but shows heavier tails, suggesting a few larger deviations in both directions. The kernel density comparison shows that the surrogate accurately reproduces the bimodal nature of the true response distribution, albeit with minor discrepancies around both peaks. The QQ plot reveals noticeable departures from normality in the tails, consistent with the presence of outliers, but the majority of the residuals align well with the theoretical quantiles. Overall, the model accurately captures the general response behaviour, with only limited bias and acceptable error dispersion across the prediction range.

The preceding analysis focused on the statistical characteristics of the mooring load components. The final verification examines the time-domain signals directly, complemented by Welch-smoothed power spectral density (PSD) comparisons. Figs. 35 and 36 present the true and predicted horizontal and vertical load components, respectively, for three example samples from the dataset.

Both time-domain and frequency-domain comparisons show that the surrogate accurately reproduces the dynamic behaviour of the mooring loads. For the low-frequency dominated case (sample 129), the predicted signals closely follow the true signals, with minimal phase lag and amplitude deviation. The corresponding PSDs demonstrate excellent alignment across the dominant frequency range. In the other




Fig. 33. Mooring surrogate true-predicted scatter plots with absolute relative error histograms.

cases (samples 2026 and 8888), characterised by a broader spectral content, the agreement remains good, although the surrogate tends to overpredict the range 0.02 Hz to 0.07 Hz, and underpredict the higher-frequency load. Overall, the surrogate model preserves both the mean trends and the dynamic characteristics of the mooring load components in diverse loading conditions.

Fig. 37 compares the mooring line tension time series (left) and the corresponding rainflow cycle process (right) obtained using the mooring surrogate-augmented model (C2) and OpenFAST for the original design under the rated wind speed condition. As previously observed for the samples within the dataset, the surrogate reproduces the general temporal trends from OpenFAST reasonably well. However, some underestimation of high-frequency tension fluctuations is noted. As shown

in the rainflow diagram, these small high-frequency discrepancies result in fewer low-amplitude cycles being counted. Moreover, the peak tensions are underestimated, despite the overestimation of the peak motions in the underlying motion solution (see Fig. 12). These effects together result in an overall underestimation of cumulative fatigue damage.

It should be noted that these findings pertain to a specific design under a specific environmental condition, and a perfect match on any single sample is neither expected nor required. The surrogate is assessed on its aggregate accuracy across the design space, as demonstrated earlier in this section. For example, the wave-frequency underprediction visible in Fig. 37 is a per-sample residual that does not compromise the intended use of the surrogate in optimisation.




Fig. 34. Diagnostic plots for standard deviation of F_z .




Fig. 35. True versus predicted mooring load component F_x .

Notably, the implementation of the surrogate models within the coupled framework enabled the level of simulation efficiency required for the design optimisation studies. Table 16 summarises the average computational time for a single design evaluation, including 12 fatigue and one ultimate strength environmental condition. All simulations were executed on a single CPU (Intel® Xeon® Silver 4216 CPU at 2.10 GHz). The results highlight the substantial reduction in computational cost achieved by the surrogate-augmented model (C2) compared with the physics-based model (C1). Application of the BEM surrogate halved the solution time of the hybrid frequency–time domain model (C1). The application of the mooring surrogate model further reduced the simulation time by a factor of two. Overall, the surrogate augmented model only takes about 5 min to provide a comprehensive evaluation of the FOWT performance in 13 loading conditions.

5. Case study

To demonstrate the impact of the modelling approaches, a series of support-structure optimisation exercises is presented, each employing one of the two models: C1 (physics-based hybrid-domain), or C2 (surrogate-augmented hybrid-domain). Rather than demonstrating a fully realistic design workflow, the aim of this exercise is to isolate and evaluate the influence of the modelling approach on the optimisation process and its outcomes. For this purpose, a deliberately minimal optimisation problem is formulated.

The objective is to minimise the material cost of the mooring lines⁷ by varying the mooring line shape parameter (α_{moor}), chain diameter

⁷ The values presented refer to the cost of a single line.




Fig. 36. True versus predicted mooring load component F_z .




Fig. 37. Tension and stress cycles counting for the mooring line (fairlead), comparison between OpenFAST and the surrogate model. Case 3.2 (rated conditions).

Table 16

Average simulation time per design evaluation for models C1–C2 under 12 fatigue and 1 ultimate strength condition.

Model	BEM model	Mooring model	Computation time [s]
C1	Capytaine	MoorPy on motion time series	1232.23
C2	Surrogate	MoorPy on motion time series	586.83
C2	Surrogate	Surrogate	289.87

(d), anchor radius (r_{anchor}), and the fairlead vertical position normalised by the draft (z_F/T), subject to the constraints of the maximum offset and dynamic tension, all evaluated in the rated-wind speed condition. The line-shape parameter varies between 0 and 1, corresponding respectively to the shortest and longest feasible line lengths for the given anchor and fairlead positions. The constraint limits are set to those of the reference design (first iteration), ensuring that any cost reduction is accompanied by maintained or improved overall performance. The optimisation problems are solved using the IPOPT optimiser (MDO Lab, 2025) with finite-difference derivatives.

As shown in Fig. 38, the two optimisation cases — despite following different optimisation paths — ultimately produced significantly more economical designs. In both cases, the optimisation converged to the lowest allowable values of α_{moor} and the chain diameter, relying primarily on adjustments to the fairlead and anchor positions as the main cost levers. The optimisation using model C2 converged to a distinct and cheaper solution, characterised by a smaller anchor radius, a lower fairlead position, and a shorter line length.

The dynamic tension constraint was active at the final solution in both optimisation cases, indicating that it governed the optimal designs. However, the optimisation trajectories differed significantly between the physics-based and surrogate-augmented models.

For C1, the optimiser initially prioritised objective reduction, moving into the infeasible region with respect to the dynamic tension constraint. As the optimisation progressed, the constraint violation penalties became increasingly influential (consistent with IPOPT's interior-point strategy), steering the solution back towards the feasibility boundary and increasing the cost slightly. Consequently, the final design was reached by approaching the constraint from the infeasible side, with objective improvements ultimately limited by the need to satisfy the constraint.

In contrast, for C2, after an initial moderate reduction in objective value, further improvements occurred while remaining on the feasible side of the constraint boundary. The optimiser approached the limiting tension value from within the feasible region, allowing




Fig. 38. Comparison of optimisation process with models C1 and C2. The last panel compares the original and optimised mooring line profiles.

a smoother search trajectory and more effective exploration of feasible descent directions. This behaviour is reflected in the markedly smoother optimisation history.

The differing trajectories can be attributed to discrepancies in the dynamic tension and tension gradients predictions between the models. Since gradient information determines the search direction at each iteration, small differences in derivative estimates can lead to substantially different optimisation paths. In a non-convex design space with multiple local minima, such path dependency is expected. If the problem were convex, all optimisations would converge to the same solution regardless of initial search direction; however, the presence of multiple local minima makes the final design sensitive to the early gradient information and subsequent trajectory. In the present case, the surrogate-augmented optimisation converged to a solution with a lower objective value than that obtained using the physics-based model. However, this outcome should not be interpreted as systematically superior performance. Rather, in a non-convex setting, differences in gradient information are expected to yield different local optima, which may be better or worse depending on the specific landscape encountered during the search.

Because no constraint was imposed on the direction of the load at the anchors,⁸ both optimal designs resulted in mooring lines that do not rest on the seabed in the equilibrium position, which is consistent with the objective of minimising cost. In practice, however, the savings associated with shorter lines could be offset by higher anchor costs, as traditional drag-embedded anchors might no longer be suitable for such configurations.

Note that the oscillations visible in the optimisation histories are normal and characteristic of gradient-based interior-point optimisation, where early iterations take large steps that prioritise reducing the objective before the constraint term becomes more influential. This effect is amplified by the non-convexity of the design space and by the use of finite-difference gradients.

The initial design and the two optimised designs were subsequently analysed in OpenFAST to verify the improvements using a single consistent model. As summarised in Table 17, both optimisations produced stiffer mooring systems, achieving more than a 50 % reduction in maximum offset, albeit with an increase in pretension of approximately 30 %. When evaluated with the higher-fidelity model, the optimised configurations also exhibited higher dynamic tensions: around a 10 % increase for the surrogate-based optimisation and 45 % increase for the

Table 17

Initial and optimised designs comparison and verification. Values based on OpenFAST simulation results. Tension relates to the upwind line tension.

	Initial design	C1 best	C2 best
Mooring line cost (\$)	139 745	127 140	115 281
Mean + standard dev. offset (m)	9.64	12.84	13.23
Mean tension (kN)	1660	1600	1560
Standard dev. of tension (kN)	88.7	101.0	76.3
Total optimisation time (min)	N/A	106.9	15.7

physics-based optimisation. This occurred despite the tension standard deviation constraint being active in the final iteration, indicating that modelling inaccuracies influenced the optimisation outcomes. The surrogate model resulted in a smaller increase in dynamic tension in the final design. Importantly, while both optimisations were relatively fast, the surrogate model delivered the final result in only 15 min, making it a valuable tool for supporting conceptual design with high potential for solving more complex optimisation problems.

6. Conclusion

This paper addressed the need for an efficient coupled dynamic modelling framework suitable for the conceptual design optimisation of FOWT support structures. The goal was to balance computational efficiency with sufficient accuracy and fidelity to enable realistic design solutions. The developed model evaluates key performance metrics across the ultimate, fatigue, and serviceability limit states, supporting a holistic and integrated design approach.

The coupled model was formulated by first establishing a fully converged static equilibrium solution, followed by dynamic simulations around this equilibrium. External loads and system responses were then mapped to internal loads in the tower and mooring lines, and the model was systematically verified against a fully coupled nonlinear reference.

The hybrid-domain model computes the motions in the frequency domain and transforms them into the time domain, where the structural loads are calculated directly, allowing motion-load nonlinearities to be partially captured and the fatigue to be assessed by rainflow counting. The verification showed that this approach predicts the motions and the tower base loads accurately, at a low computational cost. Two limitations remained, however. First, the quasi-static mooring model did not capture the wave-frequency tension, and the resulting errors in the mooring tension propagated into the fatigue analysis. Second, the reliance on the boundary element method for solving the potential-flow problem remained a computational bottleneck.

⁸ Purposefully, to keep the problem minimal as explained at the beginning of the section.

Table A.18

Verification results in operational cases for the hybrid-domain model (C1) against OpenFAST (OF). Mean of the motions and loads.

Case	Surge (m)		Heave (m)		Pitch (deg)		T2 (kN)		My (MNm)		Fxaero (kN)	
	OF	C1	OF	C1	OF	C1	OF	C1	OF	C1	OF	C1
3.1	15.54	15.98	0.65	0.66	2.86	2.91	1093.16	1092.44	49.9	56.0	401.14	427.55
3.2	26.84	27.24	1.00	1.01	4.96	4.87	1255.06	1253.43	86.3	94.2	679.69	718.08
3.3	15.27	15.98	0.63	0.66	2.82	2.91	1099.87	1083.84	49.2	56.6	364.64	417.80

Table A.19

Verification results in wave-only, wind-only, and operational cases for the hybrid-domain model (C1) against OpenFAST (OF). Standard deviation of the motions and loads.

Case	Surge (m)		Heave (m)		Pitch (deg)		Nacelle Accel. (m s^{-2})		T2 (kN)		My (MNm)	
	OF	C1	OF	C1	OF	C1	OF	C1	OF	C1	OF	C1
1.1	0.08	0.07	0.01	0.01	0.04	0.04	0.0936	0.0910	2.56	0.32	3.52	3.54
1.2	0.15	0.15	0.02	0.02	0.08	0.08	0.1872	0.1830	4.58	0.62	7.03	7.08
1.3	0.22	0.21	0.05	0.05	0.11	0.10	0.0786	0.0748	1.28	1.19	3.36	3.31
1.4	0.45	0.42	0.09	0.10	0.22	0.20	0.1572	0.1495	2.49	2.39	6.73	6.63
1.5	0.55	0.56	0.29	0.31	0.42	0.45	0.0666	0.0650	3.24	4.56	4.76	5.34
2.1	0.87	0.84	0.01	0.01	0.13	0.12	0.0050	0.0045	9.48	9.80	2.22	2.26
2.2	2.71	2.90	0.07	0.06	0.41	0.39	0.0216	0.0211	39.15	45.54	6.98	7.20
2.3	3.94	4.29	0.16	0.14	0.68	0.63	0.0386	0.0384	70.48	85.33	11.0	11.5
2.4	4.50	4.84	0.13	0.12	0.81	0.76	0.0558	0.0522	66.55	76.85	12.9	13.4
3.1	2.69	2.74	0.07	0.06	0.42	0.39	0.1374	0.1327	39.07	42.47	8.62	8.80
3.2	3.91	4.06	0.16	0.13	0.68	0.63	0.1766	0.1734	70.25	79.99	12.8	13.2
3.3	4.46	4.48	0.13	0.11	0.81	0.75	0.2574	0.2494	66.23	70.03	16.0	16.3

These limitations were addressed through the development of two surrogate models, for the prediction of the mooring loads and the hydrodynamic coefficients, respectively. In particular, the data-driven surrogate for the mooring loads in the time domain represents a significant advancement over state-of-the-art methods, as its dataset construction strategy is agnostic to both the environmental conditions and the design of the floating platform. It relies solely on the fairlead motion time series and the mooring design parameters, making it broadly applicable across different projects while delivering near-instantaneous predictions with accuracy comparable to time-domain dynamic models.

The case study demonstrated that, while both the hybrid-domain and surrogate-augmented models enabled substantial reductions in the mooring line cost, the surrogate-augmented model identified a distinct and notably cheaper configuration. This difference was attributed to variations in the gradient information, which led to different optimisation trajectories within the non-convex design space. When evaluated in OpenFAST, both optimised designs showed increased stiffness and reduced offset, but also higher dynamic tensions, reflecting the influence of modelling inaccuracies on the active constraints. The surrogate-augmented optimisation produced the design whose dynamic-tension predictions aligned most closely with OpenFAST, and delivered it in only 15 min, substantially faster than the hybrid-domain model.

Overall, these findings highlight the potential of the surrogate-augmented approach to support early-stage design by enabling the rapid and reliable exploration of the design space.

7. Future work

Future work can extend this approach in several directions. First, alternative spectral-based fatigue analysis methods could be investigated to better represent long-term damage accumulation and further reduce reliance on computationally expensive time-domain simulations. The accuracy of the surrogate model in predicting the long-term fatigue accumulation should also be validated, together with an evaluation of the practical impact on the design process (this has partly been achieved in the thesis of Patryniak (2026)). The surrogate model's dataset could be extended to extreme and survival sea states (e.g., 50-year return period conditions) to establish its applicability beyond the turbine's normal operational envelope. In addition, the boundary element method and mooring load surrogate models could be expanded

to span a wider range of design variables, enabling greater flexibility in early-stage design exploration and extending applicability beyond spar-type platforms and catenary mooring configurations. Finally, more complex optimisation case studies should be performed to demonstrate the full capabilities and limitations of the developed models.

CRediT authorship contribution statement

Katarzyna Patryniak: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Maurizio Collu:** Writing – review & editing, Supervision, Funding acquisition, Conceptualization. **Baran Yeter:** Writing – review & editing, Supervision, Conceptualisation. **Simone Minisi:** Writing – original draft, Methodology, Data curation. **Luca Oneto:** Writing – review & editing, Methodology, Data curation. **Andrea Coraddu:** Writing – review & editing, Writing – original draft, Supervision, Methodology, Data curation, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Andrea Coraddu serves as a Deputy Editor in the journal. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix. Verification numerical results

This appendix presents the numerical results obtained during the dynamic model verification process, as introduced in Section 3.5, in terms of the mean responses and loads (Table A.18) and the standard deviation of the responses and loads (Table A.19).

References

- Adam, S.P., Alexandropoulos, S.A.N., Pardalos, P.M., Vrahatis, M.N., 2019. No free lunch theorem: A review. *Approx. Optim.* 57–82.
- Accellin, M., Dias, F., 2019. Cappytaine: a Python-based linear potential flow solver. *J. Open Source Softw.* 4.

- Bachynski, E.E., 2014. Design and Dynamic Analysis of Tension Leg Platform Wind Turbines. Norwegian University of Science and Technology, Trondheim.
- Bachynski, E.E., Moan, T., 2012. Design considerations for tension leg platform wind turbines. *Mar. Struct.* 29 (1), <http://dx.doi.org/10.1016/j.marstruc.2012.09.001>.
- Bai, S., Kolter, J.Z., Koltun, V., 2018. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv Preprint arXiv:1803.01271*.
- Bessone, M.B., Singh, D., Kalimeris, T., Bachynski-Polić, E., Viré, A., 2024. Surrogate-assisted optimization of floating wind turbine substructure. *J. Phys.: Conf. Ser.* 2767, 062032.
- Bishop, C.M., Bishop, H., 2023. *Deep Learning: Foundations and Concepts*. Springer, Cham, <http://dx.doi.org/10.1007/978-3-031-45468-4>.
- Borgman, L.E., 1967. The Spectral Density for Ocean Wave Forces. Department of Mathematics University of California, Davis.
- Brown, D.T., 2005. Chapter 8 - Mooring systems. In: Chakrabarti, S.K. (Ed.), *Handbook of Offshore Engineering*. Elsevier, London, England, pp. 663–708.
- Burton, T., Jenkins, N., Sharpe, D., Bossanyi, E., 2011. *Wind Energy Handbook*, second ed. John Wiley & Sons, <http://dx.doi.org/10.1002/9781119992714>.
- Butterfield, S., Musial, W., Jonkman, J., Sclavounos, P., 2005. Engineering challenges for floating offshore wind turbines. In: *Proceedings of the Copenhagen Offshore Wind Conference*.
- Chung, J., Gulcehre, C., Cho, K., Bengio, Y., 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*.
- Coraddu, A., Kalikatzarakis, M., Theotokatos, G., Geertsma, R., Oneto, L., 2022. Physical and data-driven models hybridisation for modelling the dynamic state of a four-stroke marine diesel engine. In: Agarwal, A.K., Kumar, D., Sharma, N., Sonawane, U. (Eds.), *Engine Modeling and Simulation*. In: *Energy, Environment, and Sustainability*, Springer, Singapore, pp. 145–193. http://dx.doi.org/10.1007/978-981-16-8618-4_6.
- Coraddu, A., Oneto, L., Baldi, F., Cipollini, F., Atlas, M., Savio, S., 2019. Data-driven ship digital twin for estimating the speed loss caused by the marine fouling. *Ocean Eng.* 186, 106063. <http://dx.doi.org/10.1016/j.oceaneng.2019.05.045>.
- Coraddu, A., Oneto, L., Kalikatzarakis, M., Ilardi, D., Collu, M., 2020. Floating Spar-Type Offshore Wind Turbine Hydrodynamic Response Characterisation: a Computational Cost Aware Approach. In: *Glob. Ocean. 2020 Singapore – U.S. Gulf Coast*. pp. 1–8. <http://dx.doi.org/10.1109/IEEECONF38699.2020.9389074>.
- Coraddu, A., Oneto, L., Li, S., Karpenko, O., 2023. Surrogate models to unlock the optimal design of stiffened panels accounting for ultimate strength reduction due to welding residual stress. *Eng. Struct.* 293, 116645. <http://dx.doi.org/10.1016/j.engstruct.2023.116645>.
- Das, A., Kong, W., Sen, R., Zhou, Y., 2024. A decoder-only foundation model for time-series forecasting. In: *International Conference on Machine Learning*.
- Det Norske Veritas (DNV), 2014. Recommended Practice DNV-RP-C205.
- Dirlik, T., 1985. *Application of Computers in Fatigue Analysis* (Ph.D. thesis). University of Warwick, Coventry, UK.
- Dirlik, T., Benasciutti, D., 2021. Dirlik and tovo-benasciutti spectral methods in vibration fatigue: A review with a historical perspective. *Metals* 11 (9), <http://dx.doi.org/10.3390/met11091333>.
- DNV, 2013. Offshore standard DNV-OS-E301. Position mooring. DNV.
- DNVGL, 2016. DNVGL-RP-C203: Fatigue Design of Offshore Steel Structures. Tech. Rep., DNV GL, Recommended Practice.
- Duarte, T., Sarmiento, A., Jonkman, J., 2014. Effects of second-order hydrodynamic forces on floating offshore wind turbines. In: *AIAA SciTech Forum*. <http://dx.doi.org/10.2514/6.2014-0361>, no. AIAA 2014-0361.
- Forrester, A.J.J., Sobester, A., Keane, A.J., 2008. *Engineering Design via Surrogate Modelling: A Practical Guide*. John Wiley & Sons, Chichester, <http://dx.doi.org/10.1002/9780470770801>.
- Hall, M., Goupee, A., 2015. Validation of a lumped-mass mooring line model with DeepCwind semisubmersible model test data. *Ocean Eng.* 104.
- Hall, M., Housner, S., Srinivas, S., Wilson, S., 2021. MoorPy: Quasi-Static Mooring Analysis in Python. National Renewable Energy Laboratory.
- Hall, M., Housner, S., Zalkind, D., Bortolotti, P., Ogdén, D., Barter, G.E., 2022. An open-source frequency-domain model for floating wind turbine design optimization. *J. Phys. Conf. Ser.* 2265 (4).
- Hegseth, J.M., Bachynski, E.E., Martins, J.R.R.A., 2020. Integrated design optimization of spar floating wind turbines. *Mar. Struct.* 72, <http://dx.doi.org/10.1016/j.marstruc.2020.102771>.
- Hochreiter, S., Schmidhuber, J., 1997. Long short-term memory. *Neural Comput.* 9 (8), 1735–1780.
- Huang, Y., Wan, D., 2020. Investigation of interference effects between wind turbine and spar-type floating platform under combined wind-wave excitation. *Sustainability* 12.
- International Electrotechnical Commission (IEC), 2019. IEC 61400-1 International Standard. Wind energy generation systems – Part 1: Design requirements.
- IRENA, 2019. *Future of Wind: Deployment, Investment, Technology, Grid Integration and Socio-Economic Aspects* (A Global Energy Transformation Paper). International Renewable Energy Agency, Abu Dhabi, IRENA.
- Johnston, B., Foley, A., Doran, J., Littler, T., 2020. Levelised cost of energy, A challenge for offshore wind. *Renew. Energy* 160, 876–885.
- Jonkman, J., 2010. Definition of the Floating System for Phase IV of OC3. Tech. Rep., National Renewable Energy Laboratory (NREL).
- Jonkman, J., Butterfield, S., Musial, W., Scott, G., 2009. Definition of a 5-MW Reference Wind Turbine for Offshore System Development. National Renewable Energy Laboratory (NREL), Golden, CO.
- Kai-Tung Ma, Luo, Y., Kwan, T., Wu, Y., 2019. Chapter 6 - Fatigue analysis. In: *Mooring System Engineering for Offshore Structures*. Gulf Professional Publishing, pp. 115–137.
- Larsen, K., Sandvik, P.C., 1990. Efficient methods for the calculation of dynamic mooring line tension. In: *European Offshore Mechanics Symposium*, vol. All Days, pp. ISOPE-E-90-058.
- Lee, C.H., Newman, J.N., 2006. WAMIT User Manual, Versions 6.4. Tech. Rep., Massachusetts Institute of Technology.
- Li, S., Coraddu, A., Oneto, L., 2022. Computationally aware estimation of ultimate strength reduction of stiffened panels caused by welding residual stress: From finite element to data-driven methods. *Eng. Struct.* 264, 114423. <http://dx.doi.org/10.1016/j.engstruct.2022.114423>.
- Lienard, C., Boisard, R., Daudin, C., 2019. Aerodynamic behavior of a floating offshore wind turbine. In: *AIAA Scitech*. San Diego, United States.
- Liu, Y., Xiao, Q., Incecik, A., Peyrard, C., Wan, D., 2017. Establishing a fully coupled CFD analysis tool for floating offshore wind turbines. *Renew. Energy* 112, 280–301.
- Mao, W., Nembach, B., Tian, W., Wu, D., 2025. A spectral fatigue assessment framework for floating offshore wind turbines. *Ships Offshore Struct.* 20 (3), 306–315.
- Martinez-Puente, E., Zarketa-Astigarraga, A., Martinez-Agirre, M., Zabala, A., Esnaola, J.A., Muñoz-Calvente, M., Llavori, I., Penalba, M., 2024. Benchmarking of spectral methods for fatigue assessment of mooring systems and dynamic cables in offshore renewable energy technologies. *Ocean Eng.* 308, 118311.
- MDO Lab, 2025. IPOPT — pyOptSparse documentation. <https://mdolab-pyoptsparse.readthedocs-hosted.com/en/latest/optimizers/IPOPT.html>. (Online; Accessed 7 November 2025).
- Moan, T., Gao, Z., Bachynski, E.E., Nejad, A.R., 2020. Recent advances in integrated response analysis of floating wind turbines in a reliability perspective. *J. Offshore Mech. Arct. Eng.* 142 (5), 052002.
- Naser, M.Z., Alavi, A.H., 2023. Error metrics and performance fitness indicators for artificial intelligence and machine learning in engineering and sciences. *Archit. Struct. Constr.* 3 (4), 499–517.
- National Renewable Energy Laboratory, 2020. OpenFAST v2.6. URL <https://github.com/OpenFAST/openfast>.
- Ning, S.A., 2013. CCBlade Documentation. Release 0.1.0.
- NREL, 2024a. OpenFAST, v3.5. GitHub [code], <https://github.com/openfast/openfast/>.
- NREL, 2024b. pyHAMS: A python-based hybrid analysis model for simulation of offshore wind turbines. URL <https://github.com/WISDEM/pyHAMS>. (Accessed: 02 December 2024).
- NREL, 2025. MoorProps.py. URL <https://github.com/NREL/MoorPy/blob/master/moorpy/MoorProps.py>. (Accessed: 09 April 2025).
- NREL, 2026. AeroDyn users guide and theory manual. <https://openfast.readthedocs.io/en/main/source/user/aerodyn/index.html>.
- Oneto, L., 2020. Model Selection and Error Estimation in a Nutshell. In: *Modeling and Optimization in Science and Technologies*, vol. 15, Springer, Cham, <http://dx.doi.org/10.1007/978-3-030-24359-3>.
- Oneto, L., Anguita, D., Coraddu, A., Cleophas, T., Xepapa, K., 2016. Vessel monitoring and design in industry 4.0: A data driven perspective. In: *2016 IEEE 2nd International Forum on Research and Technologies for Society and Industry. RTSI, IEEE*, pp. 1–6. <http://dx.doi.org/10.1109/RTSI.2016.7740594>.
- Patryniak, K., 2026. Multidisciplinary Design Analysis and Optimisation for Conceptual Design of Floating Offshore Wind Turbine Support Structures (Ph.D. thesis). University of Strathclyde, Glasgow, UK, <http://dx.doi.org/10.48730/9ghq-jc45>, URL <https://stax.strath.ac.uk/concern/theses/xp68kg944>, Thesis identifier T18031.
- Patryniak, K., Collu, M., 2023. Multidisciplinary design analysis and optimisation of floating offshore wind turbine support structures - Coupled model solution strategies. *J. Phys. Conf. Ser.* 2626.
- Patryniak, K., Collu, M., Coraddu, A., 2022. Multidisciplinary design analysis and optimisation frameworks for floating offshore wind turbines: State of the art. *Ocean Eng.* 251, 111002. <http://dx.doi.org/10.1016/j.oceaneng.2022.111002>.
- Pegalajar-Jurado, A., Borg, M., Bredmose, H., 2018. An efficient frequency-domain model for quick load analysis of floating offshore wind turbines. *Wind. Energy Sci.* 3 (2), 693–712.
- Queipo, N.V., Haftka, R.T., Shyy, W., Goel, T., Vaidyanathan, R., Tucker, P.K., 2005. Surrogate-based analysis and optimization. *Prog. Aerosp. Sci.* 41 (1), 1–28. <http://dx.doi.org/10.1016/j.paerosci.2005.02.001>.
- Rasmussen, C.E., Williams, C.K.I., 2006. *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, MA, <http://dx.doi.org/10.7551/mitpress/3206.001.0001>.
- Roald, L., Jonkman, J., Robertson, A., Chokani, N., 2013. The effect of second-order hydrodynamics on floating offshore wind turbines. *Energy Procedia* 35, <http://dx.doi.org/10.1016/j.egypro.2013.07.178>, URL <https://www.sciencedirect.com/science/article/pii/S1876610213012642>. DeepWind'2013 – Selected papers from 10th Deep Sea Offshore Wind R&D Conference, Trondheim, Norway, 24 – 25 January 2013.
- Robertson, A.N., Wendt, F., Jonkman, J.M., Popko, W., Dagher, H., Gueydon, S., Qvist, J., Vittori, F., 2017. OC5 Project Phase II: Validation of Global Loads of the DeepCwind Floating Semisubmersible Wind Turbine. In: *Energy Procedia*, vol. 137.

- Rohrer, P.J., Bachynski-Polic, E.E., 2024. Analytical gradients of first-order diffraction and radiation forces for design optimization of floating structures. *Appl. Ocean Res.* 152.
- Sainani, K.L., 2016. The value of scatter plots. *PM&R* 8 (12), 1213–1217.
- Scottish Government, 2020. Sectoral Marine Plan for Offshore Wind Energy. Tech. Rep., URL <https://www.gov.scot/publications/sectoral-marine-plan-offshore-wind-energy/>.
- Shalev-Shwartz, S., Ben-David, S., 2014. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press.
- Shannon, C.E., 1949. Communication in the presence of noise. *Proc. the IRE* 37 (1), 10–21.
- Silvano, C., Ielmini, D., Ferrandi, F., Fiorin, L., et al., 2025. A survey on deep learning hardware accelerators for heterogeneous hpc platforms. *ACM Comput. Surv.* 57 (11), 1–39.
- Song, Y., Basu, B., Zhang, Z., Sørensen, J.D., Li, J., Chen, J., 2021. Dynamic reliability analysis of a floating offshore wind turbine under wind-wave joint excitations via probability density evolution method. *Renew. Energy* 168, 991–1014.
- Sørensen, J.D., Sørensen, J.N., 2010. *Wind Energy Systems: Optimising Design and Construction for Safe and Reliable Operation*. Woodhead Publishing, <http://dx.doi.org/10.1533/9780857090638>.
- Tovo, R., Benasciutti, D., 2004. A simplified approach to evaluate fatigue damage under wide-band stationary random processes. *Int. J. Fatigue* 26 (9), 831–839. <http://dx.doi.org/10.1016/j.ijfatigue.2004.02.009>.
- Tracy, C., 2007. *Parametric Design of Floating Wind Turbines*. Massachusetts Institute of Technology.
- Trubat, P., Herrera, A., Molins, C., 2022. Mooring optimization using ML techniques. In: *Proceedings of the ASME 2022 International Offshore Wind Technical Conference. IOWTC 2022, American Society of Mechanical Engineers (ASME)*.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., et al., 2017. Attention is all you need. *Neural Inf. Process. Syst.*
- Veldman, P.N., 2020. *Essentials in Coupled Dynamics of Floating Offshore Wind Turbines*. Delft University of Technology.
- Vugts, J.H., 2005. Fatigue damage assessments and the influence of wave directionality. *Appl. Ocean Res.* 27 (3), 173–185.
- Wang, S., Moan, T., Gao, Z., 2023. Methodology for global structural load effect analysis of the semi-submersible hull of floating wind turbines under still water, wind, and wave loads. *Mar. Struct.* 91.
- Wang, H., Ran, Q., Ma, G., Wen, J., Zhang, J., Zhou, S., 2025. Optimization design of floating offshore wind turbine mooring system based on DNN and NSGA-III. *Ocean Eng.* 316, 119915.
- Xu, K., Zhang, M., Shao, Y., Gao, Z., Moan, T., 2019. Effect of wave nonlinearity on fatigue damage and extreme responses of a semi-submersible floating wind turbine. *Appl. Ocean Res.* 91, 101879.
- Ye, X., Zheng, P., Qiao, D., Zhao, X., Zhou, Y., Wang, L., 2024. Multi-objective optimization design of a mooring system based on the surrogate model. *J. Mar. Sci. Eng.* 12 (10).
- Zhang, H., Wang, H., Cai, X., Xie, J., Wang, Y., Zhang, N., 2022. Novel method for designing and optimising the floating platforms of offshore wind turbines. *Ocean Eng.* 266.
- Zhou, Z.H., 2025. *Ensemble Methods: Foundations and Algorithms*. CRC Press.
- Zhu, F., Yeter, B., Brennan, F., Collu, M., 2024. Time-domain fatigue reliability analysis for floating offshore wind turbine substructures using coupled nonlinear aero-hydro-servo-elastic simulations. *Eng. Struct.* 318, 118759. <http://dx.doi.org/10.1016/j.engstruct.2024.118759>.