https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

Modeling wind farm response: a modular, integrated, and multi-stakeholder approach

Andreas Vad[1], Adrien Guilloré[1], Abhinav Anand[1], Vasilis Pettas[2], Anik H. Shah[1], Ion Lizarraga-Saenz[3], Maria Aparicio-Sanchez[3], Irene Eguinoa[3], Nicolau Conti Gost[2], Iasonas Tsaklis[2], Ariane Frère[4], Koen W. Hermans[5], Joseph K. Kamau[6], Nikolay Dimitrov[2], Tuhfe Göçmen[2], and Carlo L. Bottasso[1]

[1]Wind Energy Institute, Technical University of Munich, Boltzmannstraße 15, 85748 Garching b. München, Germany
[2]Department of Wind and Energy Systems, Technical University of Denmark, Roskilde, Denmark
[3]Wind Energy Department, National Renewable Energy Centre (CENER), Spain, Ciudad de la Innovación, 7, 31621 Sarriguren (Navarra), Spain
[4]ENGIE Research & Innovation, Laborelec, 1630 Linkebeek, Belgium
[5]Ramboll Netherlands, Hooikade 13, 2627 AB Delft, Netherlands
[6]Ramboll Danmark A/S, DK REG. NO. 35128417, Hannemanns Allé 53, 2300 Copenhagen S, Denmark

Correspondence: Carlo L. Bottasso (carlo.bottasso@tum.de)

Abstract. Accurate and computationally efficient modeling of wind farm response is essential to support a wide range of stakeholders, including research institutes, wind turbine and wind farm designers, operators, and control algorithm developers. This paper presents a modular and integrated framework for modeling wind farm response, enabling consistent and multi-purpose predictions across turbine- and farm-level applications. The proposed approach combines computationally efficient and site-specific wind farm flow modeling, high-fidelity aero-servo-elastic simulations, wake-resolved inflow characterization, and data-driven response surrogates within a flexible architecture that allows individual components to be independently developed, validated, and exchanged.

Within this framework, key novelties are introduced such as a modular and holistic wind farm model, as well as a wake-slice methodology to represent local waked inflow conditions in a compact and physically meaningful form, enabling efficient training of response surrogate models using single-turbine simulations. Artificial neural network surrogates are developed to predict individual turbine responses based on a reduced set of local inflow and control descriptors, allowing the effects of wakes, turbulence, and operational strategies to be captured without resorting to full farm-level aeroelastic simulations. Another key feature of the proposed framework is its ability to consistently model multiple turbine types as well as a wide range of operational modes (power production, start-up, shut-down and parking) combined with several control modes (normal operation, yaw-steering, derating, down-regulation and noise-curtailment) within a single formulation. To this end, the methodology employs location-agnostic load surrogates, applicable to a given turbine type irrespective of its position within a farm and at any site. The overall framework is wind farm agnostic, with a modular structure that enables application to arbitrary farm layouts, environmental conditions, and operating modes without structural modification.

The framework is tested using one open-source reference turbine and two anonymized commercial turbines. For each turbine type, surrogates were developed using a single holistic library of inflow profiles representing clean and waked conditions. The performances are evaluated through an exemplary wind farm configuration composed of six turbines, demonstrating the

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

location agnosticism of the proposed approach. Furthermore, the framework is systematically evaluated through surrogate validation and analysis across different turbine types, environmental conditions, and operational and control modes. The results demonstrate that the proposed toolchain accurately reproduces the load variations induced by wake interactions, operational modes and control modes, while maintaining a low computational cost. By combining modular physics-based modeling with scalable data-driven surrogates, the framework provides a multi-stakeholder solution for wind farm response modeling, supporting applications ranging from design analysis to operational assessment and wind farm control studies.

1 Introduction

Wind farms operate over lifetimes exceeding 20 years, during which they are subject to evolving meteorological conditions, electricity markets, grid requirements, regulatory and environmental constraints, and interactions with neighboring wind farms. These changes directly affect power production, structural loading, and operational flexibility, requiring control strategies that can adapt to dynamic operating conditions.

Wind farm control (WFC) has traditionally focused on maximizing power and energy production, primarily through wake mitigation and load-aware operation (van Wingerden et al., 2020; Göçmen et al., 2022). However, the evolving economic and political context of wind energy demands a shift toward value-oriented control objectives. In this broader perspective, WFC must balance economic performance, structural integrity, environmental impact, and contributions to power system stability, requiring modeling frameworks capable of quantifying and trading off these diverse value dimensions (Veers et al., 2019, 2022; Meyers et al., 2022). Central to such an assessment is a robust understanding of how individual turbines within a wind farm respond to the local flow field, which evolves over a wide range of environmental and operational conditions throughout the asset lifetime. These responses – encompassing both power production and structural loading – arise from the combined influence of ambient inflow conditions and wake interactions among neighboring turbines. In the following, this collective behavior is referred to as the wind farm response.

Modeling wind farm response is challenging due to tightly coupled aerodynamic, structural, electrical, and control processes. Physics-based models offer generality and interpretability but are limited by simplifying assumptions and high computational cost, while data-driven approaches offer efficient representations of complex non-linear behaviors but may lack physical consistency and perform poorly outside the training domain. Data-augmented, physics-based hybrid modeling approaches combine these strengths to achieve improved fidelity and adaptability under varying operating conditions. Rather than relying on a single approach to represent the complex behavior of wind farms, a combination of complementary approaches is required to capture the different sub-aspects of wind farm response. Furthermore, developing integrated response models requires the consistent combination of heterogeneous data sources and modeling components, as well as coordination among multiple stakeholders, typically including turbine manufacturers, wind farm operators, research and development (R&D) departments and academic research institutions. The resulting challenges include:

- Data availability and accessibility: Diverse stakeholder interests and proprietary constraints limit access to measurement data and detailed wind turbine information.

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

- Modularity and integration: Models must be modular and independently developed while remaining compatible, requiring advanced data fusion to integrate multi-model approaches.
- Model output evaluation: Limited access to absolute outputs shifts the focus to relative changes ("deltas") driven by input variations to support WFC and optimization. Operators with site-specific and regulatory knowledge are able to validate relative model outputs and translate them into absolute values.

This work presents a multi-model framework that addresses these key challenges. The framework models wind farm response for both power production forecasting and component-level load prediction across key turbine subsystems (e.g., blades, tower, foundations, drivetrain, and pitch and yaw systems), as functions of environmental conditions, wind farm characteristics, and control actions. The framework prioritizes power and loads as fundamental variables governing wind farm dynamics, as they serve as essential inputs for deriving higher-level value metrics that represent operational efficiency, structural integrity, and economic performance. By focusing the framework on these key variables, it creates a simple and flexible basis for adding other performance indicators. This allows wind farm performance to be evaluated, compared, and optimized in a consistent way for different goals. Additionally, the modular structure of the framework not only facilitates but actively promotes productive collaboration among multiple stakeholders, supporting both the joint development of individual modules and their application across a wide range of use cases.

Current wind farm modeling is primarily focused on power estimation in industrial wind farm control, relying on physics-based tools that are highly parameterized and sensitive to calibration (Kölle et al., 2022). While these steady-state analytical models can accurately predict farm-wide energy yield, accurate and fast estimation of component-level fatigue loads remains challenging (Meyers et al., 2022). This is primarily due to the greater relevance of turbulence modeling for fatigue loading and the pronounced sensitivity to any inflow heterogeneity across the rotor, which can induce substantial load variations (Muñoz-Simón et al., 2022). In wind farms, the number of distinct inflow conditions can become very large due to the many possible combinations of turbine–turbine wake interactions. The effects of wind turbine wakes on structural loads can be modeled either by increasing ambient turbulence using the effective turbulence method or by performing dynamic wake meandering (DWM) simulations that resolve wake behavior which can be coupled with structural load models across the entire wind farm (Larsen, 2007; Doubrawa et al., 2023). While the effective turbulence method is simple and computationally inexpensive, it does not provide turbine-level load evaluation and is therefore unsuitable for applications such as WFC, which require load- and state-resolved information at the individual-turbine level.

The estimation of turbine structural loads in an optimization-friendly manner remains a significant research gap, particularly for a practical implementation of coordinated control within a wind farm. While machine learning has been increasingly applied in research to predict power and load, these approaches are generally limited to isolated, non-generalizable case studies, with transferability across turbines, sites, and operating conditions not yet consistently demonstrated, pointing to the need for methods with stronger generalization (He et al., 2022; Moss et al., 2023; Witter et al., 2025; Mönnig et al., 2025).

Turbine component loads are typically obtained from time-domain aero-servo-elastic simulations coupled with a DWM model (Larsen, 2007), which together better account for wake effects and turbine responses. Additionally, the effect of different

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

control modes on individual turbines has been shown to significantly impact turbine loading and should be considered in the context of wind farm control evaluation (Pettas and Cheng, 2024). Therefore, recent research focuses on fast and data-driven load surrogate models and shows that such surrogates achieve high accuracy while being both location-agnostic (independent of turbine position within a farm) and control-oriented (capturing the influence of control setpoints) (Shaler et al., 2022; Liew et al., 2024; Guilloré et al., 2024). However, these models remain turbine-specific due to variations in aeroelastic behaviour and controller dynamics. As a result, introducing a new turbine type requires generating an entirely new training dataset of aero-servo-elastic simulations driven by inflow conditions that reflect a wide range of dynamic wake interactions. Efficient and generalizable approaches capable of representing turbine–turbine flow interactions for different turbine types and associated dynamic wake effects remain lacking (Guilloré et al., 2026b).

To this end, the proposed framework for modeling wind farm power and load responses introduces several key contributions that advance the current state of the art.

First, the fully modular architecture of the framework allows individual stakeholders, such as turbine manufacturers, wind farm operators, developers and researchers, to independently enhance and update specific model components while preserving proprietary knowledge. As a result, model improvements can be continuously incorporated without compromising data ownership or intellectual property.

Second, it allows physics-based wake models to be complemented with correction terms learned from operational data, which strengthens the modeling of intra-farm flow interactions and thereby improves predictions of the turbine aero-servo-elastic behavior. This hybrid formulation leverages the strengths of both modeling paradigms to improve predictive accuracy under realistic operating conditions.

Third, the proposed approach provides a general methodology for building comprehensive simulation databases covering a wide range of turbine operating modes, including variations in environmental conditions, control strategies, and wake interactions. It introduces a standardized representation of partial and full wakes for aeroelastic simulations, based on two hypotheses: in the self-similar far-wake regime, turbines of similar size and operating under comparable thrust and yaw conditions exhibit similar normalized velocity-deficit and wake-meandering behavior; and complex flow fields with multiple interacting wakes can be represented by equivalent local inflow states. This systematic coverage enables robust training, testing, and validation of surrogate models across realistic and diverse operational regimes and ensures consistent and comparable results for different turbine types.

Finally, the framework establishes an accurate and transferable mapping between farm inflow conditions and turbine control actions on the one hand, and the resulting power and load responses on the other. This capability enables the development of farm-agnostic response models that retain physical relevance while being applicable across different sites and layouts, thereby facilitating scalable control, optimization, and decision-support applications.

In the present approach, the turbine load surrogates are location-agnostic, meaning that a single surrogate model is applicable to a given turbine type irrespective of its position in a farm, installation site, or inflow environment (e.g., free-stream, partially waked, fully waked, or terrain-influenced effects). This location agnosticism is achieved by parameterizing the surrogate exclusively through local inflow descriptors at the rotor disk, independent of the physical mechanisms that generated the

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

inflow conditions. Consequently, the overall framework is wind farm agnostic, in the sense that the complete framework can be applied to arbitrary wind farm layouts and environmental conditions without structural modification. This is enabled by a modular formulation with independent blocks that can be combined for different farm configurations, turbine types, and operating modes. The library of synthetic clean and waked inflow conditions is reusable across turbines, and the trained location-agnostic turbine surrogates can be deployed in any wind farm configuration. Optional data-driven tuning using site-specific information may further adapt the wind farm model to a particular installation. While individual surrogates remain turbine-specific and tuned farm models remain site-specific, the proposed framework itself is readily and efficiently transferable across application cases.

The remainder of the paper is structured as follows. Section 2 introduces the proposed approach for modeling the wind farm power and load response. It begins with the wind farm flow modeling in Sect. 2.1, followed by the synthetic inflow generation described in Sect. 2.2, which provides the input for the aero-servo-elastic simulations outlined in Sect. 2.3. These simulations form the training dataset used to develop the surrogate models presented in Sect. 2.4. In Sect. 3, first the specifications of the case study setup are described in Sect. 3.1. This includes the specifications of the considered turbine types, farm modeling, inflow profiles, and the trained surrogate model. The individual surrogate models are validated in Sect. 3.2. Section 3.3 evaluates the generalizability of the synthetic wake inflow database by comparing framework-based load predictions to independent farm-level dynamic simulations. The turbine responses across different turbine types and a range of environmental and operational conditions are examined in Sect. 3.4. Finally, Sect. 3.5 demonstrates the application of the wind farm response framework to an exemplary wind farm. Section 4 concludes the paper with a summary of the main results and a discussion of limitations and future directions.

2 Methodology: wind farm response framework

The wind farm response framework (Vad et al., 2026b) is illustrated in Fig. 1. It outlines the process of modeling the wind farm response – expressed in terms of power production and fatigue metrics – for given environmental conditions and turbine control setpoints, which form the basis for further response quantities and applications.

The environmental conditions comprise the wind speed (WS), wind direction (WD), and turbulence intensity (TI) at the wind farm level. Four operational modes are considered: power production, start-up, shut-down, and parked. During each of the operation modes, applicable turbine-specific control inputs can additionally be provided, namely the yaw setpoint and the power mode for each turbine. The power mode is defined through turbine-specific performance curves (e.g., power and thrust as functions of WS).

The flow within the wind farm is modeled using an engineering flow modeling tool, as described in Sect. 2.1. The flow model yields the turbine-specific inflow. This inflow can be used either directly to compute the power output of each turbine or further processed into sector-averaged quantities. In the latter case, the rotor is divided into four equally sized sectors, and the inflow (characterized using WS and TI) is averaged within each sector, as illustrated in Fig. 3 (e). These sector-averaged inflow quantities (SAIQs) serve as input to a surrogate model trained on a dataset of aeroelastic simulations. While the aeroelastic

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

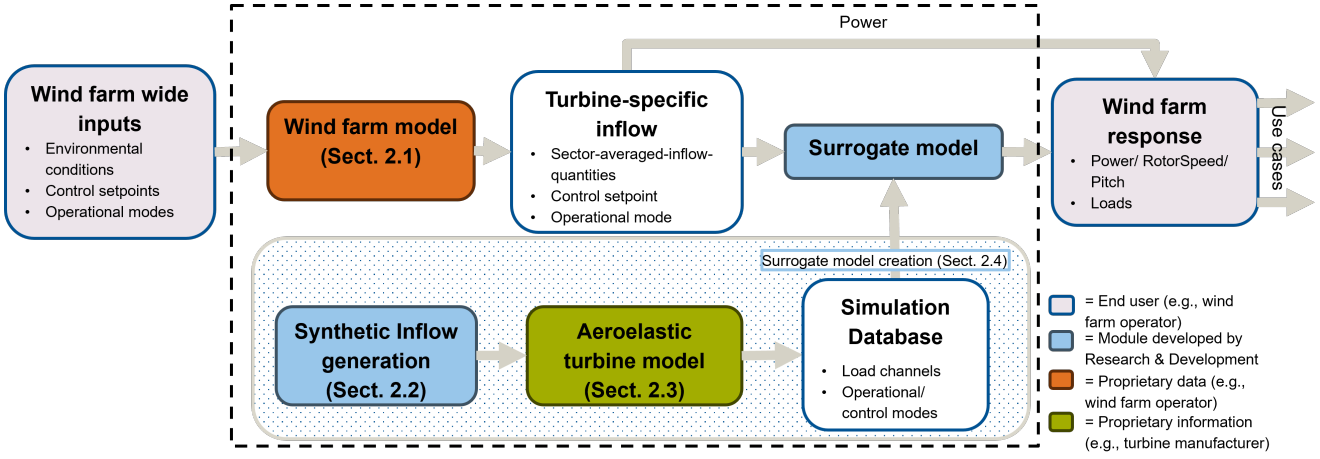

Figure 1. Flowchart of the multi-module framework used to model the wind-farm response. Colored elements indicate typical contributions from different stakeholder groups. The dashed black box marks the scope of the present code framework (Vad et al., 2026b). The gray-patterned box summarizes the steps required to build the simulation database for surrogate-model training. A typical use case follows the upper branch, relying solely on computationally efficient modules.

simulations can be carried out with any commercial or open-source software – potentially involving proprietary data and thus supporting collaboration among different stakeholders – the inflow fields used as simulation input are generated in a standardized manner.

The procedure for generating synthetic inflow is detailed in Sect. 2.2, and the details regarding the aeroelastic simulations are presented in Sect. 2.3. The response surrogates are described in Sect. 2.4 as machine learning models trained to predict turbine responses based on four sector-averaged wind speeds (SAWSs) and four sector-averaged turbulence intensities (SATIs) across the rotor disk (see Fig. 3 (e)), along with the applicable control setpoints. It was demonstrated in Guilloré et al. (2024) that these reduced-order descriptors are sufficient to accurately capture the behavior of fatigue loadings regardless of the position of a turbine in a farm.

Finally, when following the upper branch of the flowchart, the framework efficiently produces wind-farm agnostic quantities —- such as power and loads -— that serve as the foundation for deriving additional wind-farm-level response metrics. These, in turn, enable a wide range of subsequent case studies, including wind farm control, lifetime extension, repowering, wildlife impact assessments, and investigations of wind turbine noise. The specific response metrics employed in this study are defined in Sect. 2.5.

2.1 Wind farm modeling

The flow model represents wake interactions within the wind farm and relates the ambient site inflow to turbine-level inflow conditions and corresponding power output. Steady-state engineering wake models provide an effective balance between predictive accuracy and computational cost, leading to their broad adoption in both industry and research (Porté-Agel et al.,

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

2019). In addition, their modular structure and adjustable parameters enable calibration to better reflect site-specific conditions. In practice, wind farm flow fields are also affected by factors such as terrain and farm-scale boundary layer interactions. Modeling approaches for these effects are under active development (Stipa et al., 2024), or can alternatively be inferred from operational data (Braunbehrens et al., 2023).

Figure 2 provides a flowchart summarizing the interfaces between the input and outputs of the wind farm model.

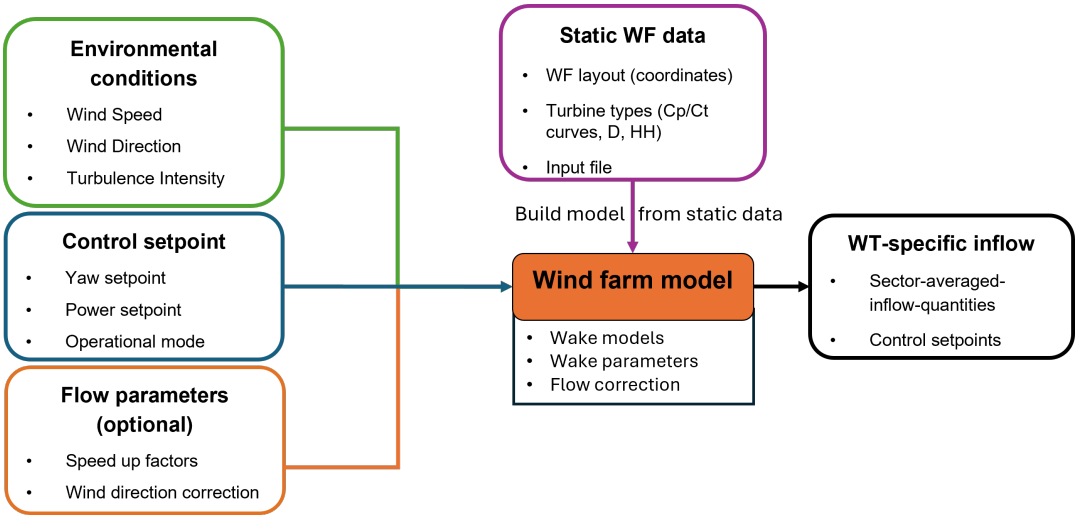

Figure 2. Wind farm model. Static wind farm data is used to build up the model. Environmental conditions, control setpoints, and optional flow parameters drive the flowfield simulation and turbine-specific inflow extraction.

Before the wind farm model can be used, it must be initialized with static information describing the wind farm layout and turbine types. This information is provided in structured data tables that adhere to standardized nomenclature conventions for column names (González-Salcedo et al., 2025), developed in accordance with guidelines from multiple ontology standards. Once initialized, the model can compute the resulting flow field for any specified environmental conditions and control setpoints. In this step, wakes are modeled using either default wake parameters or those provided in the static wind farm data, while the background flow is represented by flow parameters that capture WS variations induced by local terrain features. These flow parameters are jointly calibrated using site-specific historical data collected at the wind farm, following the wind farm-as-a-sensor methodology described in Braunbehrens et al. (2023). The calibration minimizes the discrepancy between simulated and measured wind farm power output, thereby learning site-specific flow patterns and refining the wake parameters.

The calibrated wind farm model yields turbine-specific inflows that account for changing velocity and turbulence distributions induced by upstream turbine wakes and terrain-induced flow speed-ups. These turbine-specific inflows can be used to compute the power output of each turbine. Additionally, SAIQs can be derived for WS and TI, which – along with the control inputs – serve as inputs to the surrogate models. The underlying assumption is that steady-state engineering wake models yield $10\,\mathrm{min}$ averaged wake characteristics comparable to those produced by the dynamic wake model applied in the aeroelastic simulations (Ardillon et al., 2023).

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

For this study, the engineering flow-modeling tool `Floris` (NREL, 2025) is employed for wind farm representation. SAIQs at the turbine locations are computed using a 10×10 discretization of the rotor plane. The model is accessible via the `TWAIN` environment (Vad, 2026a) or on `Zenodo` (Vad et al., 2026b).

2.2 Synthetic inflow generation

Turbulent wind fields used as inputs for single-turbine aero-servo-elastic simulations, which are employed to evaluate aerodynamic and structural responses, can be generalized as they are not subject to proprietary constraints. A shared library of such turbulent wind fields, spanning a wide range of inflow conditions including wind farm effects, can therefore form the basis for multi-stakeholder collaboration. Consistent with the overall framework, this module is itself modular, ensuring flexibility for different stakeholders.

In this study, turbulent flow fields are generated in two categories: clean inflow and waked inflow profiles. Clean inflow profiles represent ambient, free-stream turbulence, whereas waked inflow profiles account for the influence of upstream turbines and their operating conditions.

2.2.1 Clean inflow profiles

Free-stream, time-resolved turbulent wind fields can be generated using any stochastic turbulence field model commonly employed in wind energy applications. In this study, the turbulence formulation consists of two main components: the temporal variability at a fixed spatial location, represented by a velocity spectral model, and the spatial correlation of the wind field across the rotor plane, described by a spatial coherence model.

The Kaimal spectral model specified in the IEC 61400-1 for wind turbine design is selected to model the velocity spectra (International Electrotechnical Commission, 2019). It describes the one-sided power spectral density of the three wind velocity components (longitudinal u, lateral v, and vertical w) and is widely applied in aeroelastic simulations for wind turbines. The spatial correlation of the longitudinal wind component u between two points separated by distance r in the rotor plane is modeled using the IEC coherence model (International Electrotechnical Commission, 2019). This model captures the frequency-dependent decay of the turbulence correlation along the rotor plane using an exponential function.

Furthermore, TI levels can be specified using different IEC classes, including normal and extreme conditions, depending on the application. Apart from the mean WS and TI at a reference point (typically the wind turbine hub-height), additional meteorological boundary conditions – such as vertical and lateral profiles of WS and WD – can be superimposed a posteriori. This approach allows different stakeholders to adapt a common set of turbulent inflow files to their specific aero-servo-elastic simulation requirements.

2.2.2 Waked inflow profiles

Wake interactions have a pronounced impact on the structural fatigue of turbines within a wind farm due to the combined effects of local velocity deficits, wake meandering, and wake-added turbulence (Larsen, 2007; Doubrawa et al., 2023). Ac-

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

curate assessment of the structural response of individual turbines, therefore, requires that dynamic wake effects be properly represented during surrogate model development. This work builds upon the location-agnostic and control-oriented load surrogate model introduced by Guilloré et al. (2024), which demonstrated that knowledge of local inflow quantities at the rotor disk – characterized using SAWSs and SATIs – together with turbine control setpoints, is sufficient to predict the structural response of a turbine independent of its position within a farm. However, that study was limited to a single turbine model, and the surrogate models were trained using wind-farm-scale dynamic simulations. The present work extends this framework to multiple turbine models, encompassing turbines of different ratings and structural designs, as well as multiple control strategies and operating modes for each turbine. In addition, it addresses various stakeholder use cases, including improvements in inflow and response modeling techniques (R&D and research institutes perspective), wind turbine structural and control design (turbine manufacturer perspective) and wind farm operation under varying environmental conditions and control optimization strategies (wind farm operator perspective). For these various stakeholder use cases, performing iterative farm-level simulations that account for varying turbine designs and operating conditions becomes impractical and computationally prohibitive. To address this limitation, this work introduces a novel approach that employs generic wake slices as inputs to single-turbine simulations, yielding a substantially more computationally efficient and generic modeling strategy.

Figure 3 illustrates the workflow for efficiently generating and using a large number of wake slices to train location-agnostic turbine response surrogate models. First, a DWM model is employed to generate a time-resolved field of a single wake (step (a)). Cross-stream slices are then extracted across a wide domain at the free-stream and multiple downstream locations (step (b)). In step (c), reduced turbulent fields are obtained at various streamwise and lateral positions, capturing different levels of wake recovery and impingement efficiently. Finally, these turbulent fields serve as inflow for aero-servo-elastic simulations of any new (similar-sized) turbine (step (d)). Based on the rotor diameter of the turbine of interest, and for each inflow cases considered, the SAWS and SATI are derived from the spatiotemporally varying local inflow, which is partitioned into sectors as illustrated in step (e), to serve as inputs to the response surrogates. Steps (a), (b) and (c) are performed once using a single turbine type to build a comprehensive library of representative dynamic wake fields, whereas steps (d) and (e) can be efficiently repeated for any (similar-sized) turbine model or control strategy.

The present approach relies on two underlying hypotheses. First, beyond a sufficiently large downstream distance at which the self-similar far-wake regime is reached, two turbine models of comparable size operating under similar conditions (i.e., similar inflow, thrust regime, and yaw misalignment) are assumed to exhibit similar normalized velocity-deficit and meandering characteristics. Second, complex flow situations involving multiple interacting wakes are assumed to be representable by an equivalent local inflow state, described through the reduced SAIQ descriptors. Under this assumption, even when multiple wakes overlap, the resulting local inflow can be mapped onto the training domain defined by single-wake simulations, provided that the corresponding SAIQ descriptors fall within the range covered by the surrogate model library and that these descriptors are sufficient to predict fatigue loadings (see Guilloré et al. (2024)). For more details on the methodology and validation of this efficient wake slices approach, please refer to Guilloré et al. (2026b). The generated wake slices should be applied to wind turbine models of comparable sizes. The method was validated for multiple turbine types with rotor diameters between 90 m and 120 m and hub heights between 60 m and 100 m

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

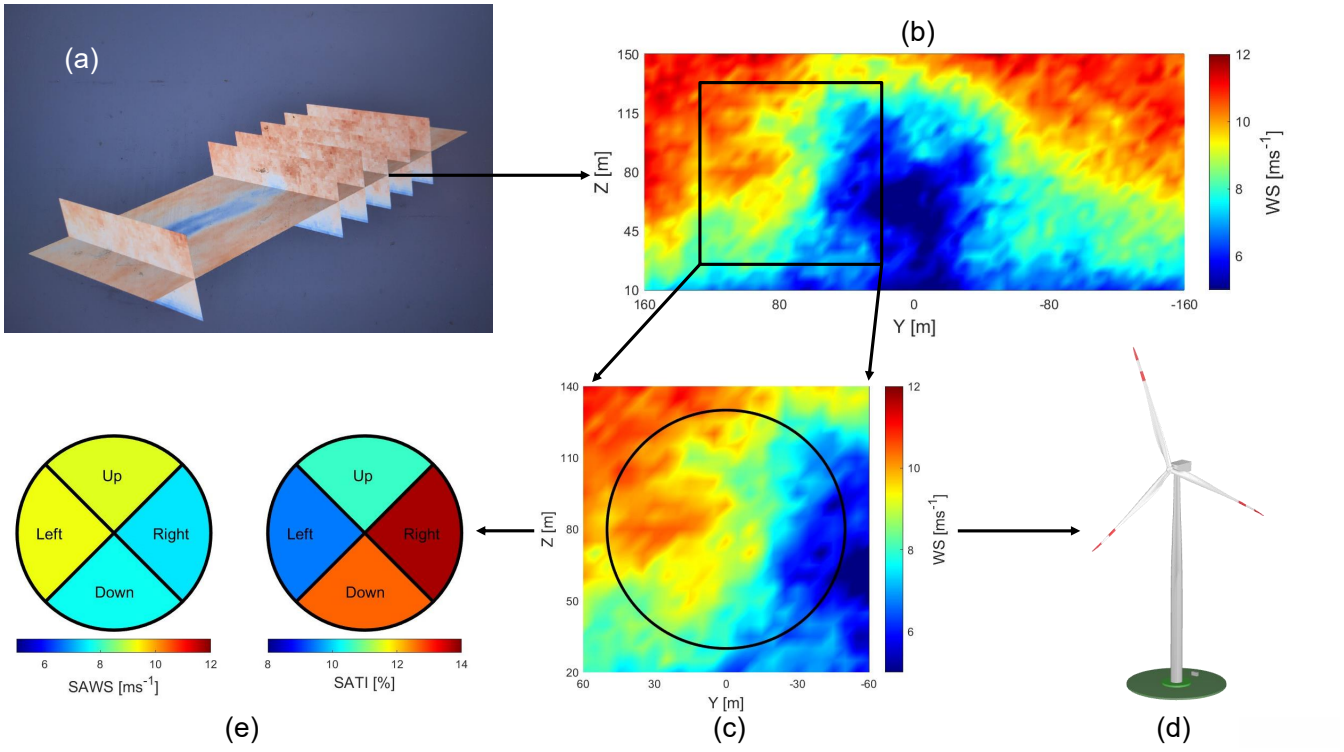

Figure 3. The approach to generate and apply to a wind turbine a library of wake slices. **(a)** simulations of a wake based on DWM for various ambient inflow conditions on a farm scale. **(b)** sampling of lateral slices at different streamwise distances. **(c)** extraction of rotor-scale turbulent grids for various streamwise and lateral locations. **(d)** computationally efficient aero-servo-elastic simulations of a single turbine including waked inflow fields. **(e)** processing of the spatiotemporally varying local inflow into four SAWSs and four SATIs as inputs to the location-agnostic surrogates.

2.2.3 Wind condition selection

Following the multi-stakeholder approach, it is important to create wind fields that span the full range of relevant environmental conditions. However, the computation effort for both generating the turbulent wind fields and building the subsequent turbine-response database for surrogate model training can be reduced while maintaining accuracy. SAIQs are used to describe the local inflow at the turbine level, providing the necessary detail for both free-stream and waked conditions (Guilloré et al., 2024). This results in an 8-dimensional (four SAWSs and four SATIs, as shown in Fig. 3 (e)) representation of local inflow characteristics.

For clean inflow, due to the a posteriori superimposition of WS and WD profiles, the 8-dimensional grid combination of local inflow values collapses into a 2-dimensional grid of ambient WS and TI values for generation of turbulent wind fields. However, for waked inflow profiles, superimposition is not possible due to interacting wake effects. Hence, generating waked turbulent

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

wind fields spanning the entire 8-dimensional grid would require a significant number of farm-level dynamic simulations, which involve a relatively higher computational effort. To address this, a smart sampling approach is developed to identify the relevant environmental conditions for running dynamic farm-level simulations.

This approach employs Latin hypercube sampling (LHS) (McKay et al., 1979) to select a representative subset of wind farm ambient conditions, while ensuring that the resulting turbine-specific inflow conditions – characterized by the eight-dimensional distribution of inflow parameters – remains representative of the full parameter space. Additionally, edge cases are included to prevent surrogate extrapolation errors. The Bhattacharyya distance (Theodoridis and Koutroumbas, 2011) is used as a metric to compare the original 8-dimensional distribution of sector average quantities for the full set of ambient conditions with the 8-dimensional distribution obtained from the Latin-hypercube sampled subset of ambient conditions. A value closer to 0 indicates a stronger overlap between the two distributions.

Generating the full parameter space using dynamic wake models is computationally prohibitive. However, Ardillon et al. (2023) demonstrated that steady-state engineering wake models reproduce the same key trends in wake characteristics as dynamic wake models. Therefore, the wake model introduced in Sect. 2.1, `Floris` (NREL, 2025), is employed as a reliable proxy to generate the 8-dimensional distribution of sector-averaged quantities.

2.3 Aeroelastic simulations

The objective of the aeroelastic simulations is to characterize the response of a single turbine to a wide range of operational conditions experienced within a wind farm over its operational life. These operational conditions include different inflow conditions, operational modes, as well as control modes. Section 2.2 presented detailed discussions on identifying and generating different inflow profiles. This section extends the discussion to consider the different operational and control modes required to build a comprehensive simulation database that fully characterizes turbine response and enables future analysis of the impact of control actions on wind farm response. Furthermore, these results are then used to train, validate, and test the surrogate models of the wind farm response framework as described in Sect. 2.4.

To this extent, simulation models of two industry-standard turbines are considered here. For data confidentiality reasons, the turbines will henceforth be referred to as Turbine-A and Turbine-B. Additionally, an open-source reference turbine has also been included within this study, the NREL-1.79-100 (NREL, 2025). Each turbine can be simulated using a different aeroelastic solver, in accordance with the multi-stakeholder approach considered here, where stakeholders may have proprietary in-house simulation tools.

To capture the main operating conditions of a turbine within a wind farm, this study considers the following operation modes:

- **Power-production (PP)** mode, denoting the wind turbine operation when it is producing power.
- **Start-up (SU)** mode, denoting the starting up phase of a wind turbine.
- **Shut-down (SD)** mode, denoting the shutting down phase of a wind turbine.
- **Parked (PK)** mode, denoting the idling phase of a wind turbine.

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

All operating modes are simulated over 800 seconds using the synthetic inflow winds described in Sect. 2.2. The first 200 seconds are discarded as transients to allow the turbine to reach its nominal point and the wake inflow to develop fully, while the remaining 600 seconds are retained for evaluation. For each simulation case, multiple turbulence realizations are simulated to ensure statistical convergence.

In the power-production mode, the turbine generates power continuously for the entire simulation, while in the parked mode, it remains idle, allowing the rotor to rotate freely without producing power. The start-up and shut-down modes capture transient transitions between idling and power production, with each simulation representing either a single start or stop event. To ensure balanced load representation, production and idling phases are allocated roughly equally within the 600 seconds of evaluation.

In standard load assessment practice, start-up and shut-down load cases are typically simulated over the duration of the maneuver itself, which is on the order of 20–40 seconds, depending on the controller design. In the present work, however, load responses for all operating modes – including start-up and shut-down – are computed over a fixed $10\,\mathrm{min}$ time window, in order to ensure direct comparability with normal power-production conditions. To this end, a standardized start-up or shut-down maneuver is embedded within the 600-second evaluation interval. In the shut-down case, the turbine operates in power-production mode for the first 150 seconds, after which the shut-down maneuver is initiated, and the turbine transitions into idling mode. Conversely, in the start-up case, the turbine remains parked until the maneuver is initiated, after which it enters power-production mode and operates for 150 seconds. This approach was adopted here as it reflects realistic operational conditions, in which individual $10\,\mathrm{min}$ intervals may include multiple operating states with varying duration of normal operation, while ensuring that load responses are consistently evaluated over identical time windows across all operating modes.

Furthermore, for the power-production mode, the full range of input wind conditions characterized by both clean and waked inflow profiles is considered. In contrast, for the start-up, shut-down, and parked modes, the full range of wind conditions is characterized using only clean inflow profiles, with the additional constraint that, for the start-up mode, wind conditions are limited to values up to the rated WS of the turbine. Waked inflow conditions are not simulated for these modes in order to reduce computational effort and because prior analysis showed that the maneuver itself dominates the loading response, making the influence of wake effects comparatively negligible in these cases.

Within each operation mode, for a given inflow condition, the turbine can be controlled in various ways, resulting in different responses. From the perspective of WFC, the turbine control setpoints can be specified as a yaw-offset, to steer the wake away from the downstream turbines, or as a power setpoint (further translated into pitch and torque setpoints), to curtail the turbines in order to reduce the thrust to reduce the intensity of the turbine wakes, or as a combination of the two. Hence, this study considers various turbine control modes relevant to wind farm operation, control, and decision support, as listed below.

- **Normal-operation (NO)**: Wind turbine normal operation using baseline controller (also referred to as greedy operation in literature).
- **Yaw-steering (YS)**: Static yaw-offset for wake steering flow control.

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

- **Derating (DE)**: Wind turbine curtailed operation by clipping off the performance curves by pitching before the rated WS. The goal of this mode is to limit the maximum power, which is commonly used for grid-curtailed operation in case of electricity grid congestion.

- **Down-regulation (Iso-TSR DR)**: Wind turbine curtailed operation by scaling down the performance curves using a constant tip speed ratio (TSR)-based down-regulation approach. Although different down-regulation approaches have been presented in the literature, this study considered the Iso-TSR approach due to its ease in controller adaptation logic, which prioritizes reducing generator torque over rotor speed to achieve the desired curtailment, thereby leading to less fatigue loading and pitch activity while also not changing the rotational frequencies (Pettas and Cheng, 2018, 2024).

- **Noise-curtailment (NC)**: Wind turbine curtailed operation according to performance curves denoting different noise modes. The noise emission level is reduced by decreasing the wind turbine rotational speed. Consequently, power is also reduced in the operating ranges where the rotational speed is decreased due to noise. The goal of this control mode is to limit the turbine noise emission to fulfill various levels of regulation requirements.

2.4 Surrogate model creation

The surrogate model creation process can be categorized into defining model inputs and targets, post-processing raw simulation data to estimate the input and target variables, and segregating the post-processed data into meaningful training, validation, and testing sets. The following subsections discuss the methodologies employed for each of these aspects.

2.4.1 Model characteristics

For a given turbine type, a dedicated family of surrogate models is developed. This family comprises distinct surrogate sets associated with the applicable operational modes, as discussed in Sect. 2.3. Within each operational-mode-specific set, individual surrogate models are trained for each target channel. The target channels encompass a comprehensive range of system responses. These may include operational metrics such as rotor speed, blade pitch angle, generator torque, power yield, and load responses of relevant structural components, as well as derived performance indicators, including actuator usage. This structured surrogate formulation enables a comprehensive and modular representation of turbine behavior across operating conditions and performance dimensions.

Each individual surrogate model is formulated using a multiple-input–single-output (MISO) configuration. In this setup, a single target quantity is predicted from a shared set of input features. The input feature space consists of eight inflow variables that describe the local inflow conditions, represented using the SAWS and SATI (see Fig. 3 (e)). Additionally, where applicable, two control-related variables are also included in the input feature space: the power setpoint and the yaw setpoint, encompassing different control modes discussed in Sect. 2.3. This combination of environmental and control inputs allows the surrogate models to account for both external operating conditions and turbine control actions.

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

2.4.2 Data preparation

This section offers insights into the data preparation steps necessary to convert the raw simulation data to surrogate input and target features set. The data preparation step is organized into three stages.

The first stage calculates the input feature set by extracting the $10\,\mathrm{min}$ equivalent SAIQs from the full-field clean inflow profiles (refer to Sect. 2.2.1) and the waked inflow profiles (refer to Sect. 2.2.2). This task is carried out using a dedicated extraction script, which is designed to process collections of wind data in a standardized manner. The script operates on a user-specified set of wind files and generates corresponding output in a designated directory. During execution, key options allow the temporal range of the extracted data to be defined, control the amount of wind shear to be added, and enable the processing to be adapted to different wind conditions.

The second stage focuses on processing the individual simulation results to calculate the target features. Here, the outputs from multiple simulation seeds are processed and aggregated into a representative statistical dataset. The processed data includes metadata describing the load cases and operating conditions, as well as statistical descriptors of the relevant variables. In addition, fatigue-related metrics in the form of damage equivalent loads (DEL) are evaluated across multiple material exponents, and actuator usage is quantified through dedicated duty-cycle metrics. Together, these outputs provide a consolidated and statistically robust basis for assessing turbine performance and structural response across the simulated load cases.

For a given time-step denoted as k, representing a $10\,\mathrm{min}$ simulation duration, the corresponding DEL is computed using the conventional rainflow counting as

$$M_{\mathrm{eq},k} = \left(\frac{1}{N_{\mathrm{eq}}} \sum_i (c_i \cdot L_i^m) \right)^{1/m}, \tag{1}$$

where c_i is the number of load cycles of amplitude L_i from the rainflow counting, N_{eq} is the equivalent number of total cycles (taken as per convention at 600 cycles over $10\,\mathrm{min}$ to represent $1\,\mathrm{Hz}$ DELs), and m is the Wöhler exponent of the S-N curve of the material (Miner, 1945).

The third stage applies a generalized pre-processing pipeline for developing machine learning surrogates. It ingests structured datasets from multiple sources, selects relevant inputs and outputs, and applies standard scaling to ensure variable comparability. The data are then split into stratified training, validation, and testing subsets. All processed data and transformation parameters are saved in a standardized format for reproducible inference, and the pipeline generates visualizations that summarize data distributions and reveal structure in the input space, supporting interpretability and model development.

Python-based scripts containing the utilized code base for the data preparation pipeline have been shared as a supplement to this work and can be accessed on `Zenodo` (Vad et al., 2026b).

2.4.3 Model training

Although several supervised machine learning model architectures and training approaches are available in the literature, classic artificial neural network (ANN) models have been found to give an acceptable degree of performance (Shaler et al., 2022; Liew et al., 2024; Guilloré et al., 2024). Within this study, a deep neural network with a simple grid search-based hyperpa-

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

rameter optimization is considered. Additionally, more sophisticated machine learning models are also being investigated and developed. The performance of these advanced models will be compared with the baseline model considered and will be part of a future benchmark study.

2.5 Aggregation into wind farm response

The presented wind farm response framework yields power and damage equivalent loads for a large number of load channels and components. To study the wind farm response in terms of fatigue, the damage equivalent loads must be aggregated to calculate relevant fatigue metrics.

This study employs a relative assessment methodology to estimate fatigue damage and the remaining useful lifetime (RUL) following IEC 61400-28 (International Electrotechnical Commission, 2025). The method compares the loads experienced during the operational life of the turbine with the design loads associated with its certified IEC class. To support this process, CENER has developed a dedicated tool that integrates the fatigue-life calculation into the wind-farm response framework (Aparicio-Sanchez et al., 2026), implementing the IEC guidelines for fatigue computation (International Electrotechnical Commission, 2019, 2025). The tool is accessible via the `TWAIN` environment (Aparicio-Sanchez et al., 2026).

First, the fatigue-lifetime limit must be defined for each turbine and component. These reference values are pre-computed and provided as inputs to the tool, forming the basis for damage and RUL estimation. The reference lifetime damage $\mathbb{D}_{\mathrm{ref}}$ is determined according to IEC 61400-1 (International Electrotechnical Commission, 2019). Using damage-equivalent loads representative of the operational and meteorological conditions of the respective IEC class, the limit is computed as

$$\mathbb{D}_{\mathrm{ref}} = \sum_{j=1}^{N_{\mathrm{IEC}}} n_j \cdot M_{\mathrm{eq},j}^{m}\,, \tag{2}$$

where $M_{\mathrm{eq},j}$ denotes the damage equivalent load for condition j, N_{IEC} the number of considered conditions, n_j their occurrences over the design lifetime.

Consequently, the cumulative damage $\mathbb{D}$ associated with a time series of N damage-equivalent loads can be expressed as

$$\mathbb{D} = \frac{\sum_{k=1}^{N} M_{\mathrm{eq},k}^{m}}{\mathbb{D}_{\mathrm{ref}}}\,. \tag{3}$$

Assuming linear damage accumulation, the remaining useful life can be estimated by extrapolating the observed damage progression until a cumulative damage level of 1 is reached:

$$\mathrm{RUL} = T\left(\frac{1-\mathbb{D}_{\mathrm{init}}}{\mathbb{D}} - 1\right), \tag{4}$$

where T denotes the considered time interval, $\mathbb{D}$ the damage accumulated during this interval, and $\mathbb{D}_{\mathrm{init}}$ the initial damage state. RUL is computed for all turbine–load combinations, and a single representative value for the wind farm is obtained by selecting the most restrictive case, i.e., the minimum RUL across all turbines and channels.

The present version of this framework does not explicitly include low-frequency fatigue cycles (LFFC). While LFFCs from operating-mode transitions such start-up and shut-down events are included, transitions between different power-producing

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

modes are ignored. Furthermore, cycles associated with wind variations longer than the standard IEC prescribed 10-minute intervals (International Electrotechnical Commission, 2019) – such as low-frequency turbulence, diurnal cycles, and seasonal changes – are also neglected. LFFCs have been shown to have non-negligible effect on the long-term fatigue of wind turbines (Sadeghi et al., 2023; Faria et al., 2024), especially at tower base. Recent work on developing efficient surrogates to include these low-frequency fatigue effects (Shah et al., 2026; Vad et al., 2026c) can be implemented with the present holistic framework in the near future.

3 Results and discussions

3.1 Case study setup

3.1.1 Inflow profile generation setup

The inflow fields used as simulation input are generated in a standardized manner using the open-source tools `TurbSim` (Jonkman, 2009) and `FAST.Farm` (Jonkman et al., 2018), where the former serves as a time-resolved stochastic turbulence box generator, and the latter serves as the DWM model for capturing the dynamic wake behavior.

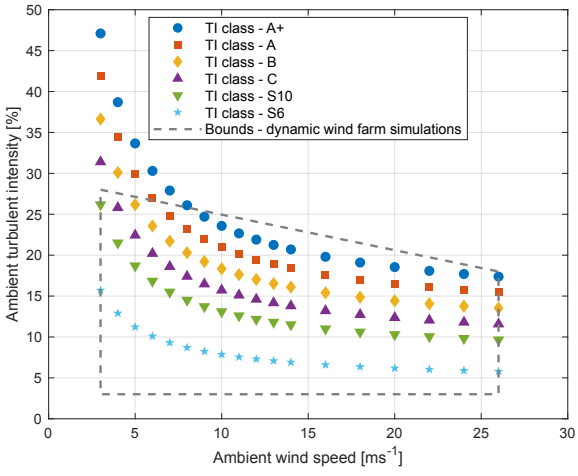

Figure 4. Clean inflow ambient cases based on IEC normal turbulence model and bounds for case selection for dynamic farm simulations.

For clean inflow, turbulent flow fields are generated with a size of $120\,\mathrm{m} \times 120\,\mathrm{m}$, centered at the hub-height of the NREL-1.79-100 turbine, $80\,\mathrm{m}$. Wind speeds ranging from 3 to 26 ms^{-1} are considered with six different turbulent classes (A+, A, B, C, S10, S6) based on the IEC normal turbulence model guidelines International Electrotechnical Commission (2019). This results in a total of 108 combinations that cover the relevant ambient conditions instead of a rectangular grid, as shown in Fig. 4. The spatial resolution is 5 m and the sampling frequency is 10 Hz. To ensure statistical convergence, 6 turbulent seeds

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

are generated for each inflow combination. Wind profiles for the shear exponent of 0.2 are superimposed onto the generated flow fields.

For waked inflow conditions, DWM simulations are performed using `FAST.Farm` with the NREL-1.79-100 turbine. The latest version of `FAST.Farm` (NREL, 2025) is used, including the latest wake-added turbulence (Kretschmer et al., 2021; Branlard et al., 2024) and curled wake (Branlard et al., 2023) models, as they have been shown to be key drivers of fatigue loads in a farm. The simulations considered a range of ambient wind speeds and turbulence intensities, with turbulent inflow fields generated using `TurbSim` (Jonkman, 2009). Inflow slices are extracted at multiple downstream and lateral positions, as described in Sect. 2.2.2. To efficiently select combinations of ambient conditions that span the full parameter space representative of the inflow conditions experienced by turbines within a wind farm, the smart sampling approach outlined in Sect. 2.2.3 is applied.

For each combination of ambient wind speed and turbulence intensity, SAIQs are extracted at streamwise locations of $[-0.5$, 4, 5, 6, 7, and $8\,D]$ and lateral locations of $[-1 : 0.25 : 1\,D]$ relative to the wake-generating turbine (with $D = 100\,\mathrm{m}$, referencing the NREL-1.79-100), capturing different levels of wake recovery and impingement. Here, D denotes the rotor diameter. This procedure yields 54 distinct combinations of local SAWS and SATI for a given ambient wind speed and turbulence intensity. The values at streamwise location of $-0.5\,D$ provide additional clean inflow cases.

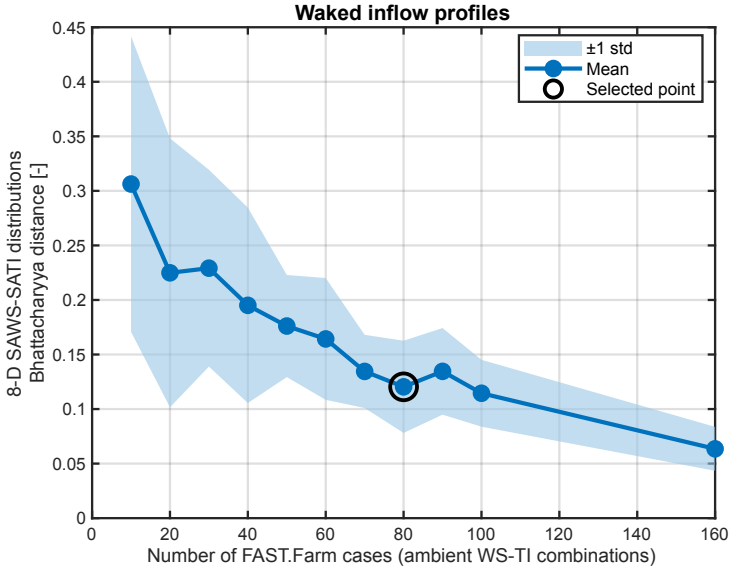

Figure 5. Convergence plot of the farm-level cases selection using the Bhattacharyya distance to generate waked inflow profiles.

A convergence study is performed to determine the minimum number of dynamic wake simulations required, with candidate cases identified using Latin-hypercube sampling (McKay et al., 1979) applied within a trapezoidal grid of ambient WS (between 3 and 26 ms^{-1}) and TI (between 3 and 28 %), as shown by the grey dashed line in Fig. 4. Lower turbulence intensity

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

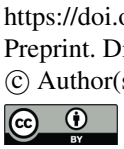

eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

levels are considered here to expand the range of inflow conditions. Finally, edge cases are always included to prevent surrogate extrapolation errors. The Bhattacharyya distance (Theodoridis and Koutroumbas, 2011) values closer to zero indicate greater overlap between the original and downsampled 8-dimensional distribution of sector-averaged quantities.

Based on the results shown in Fig. 5, a total of 83 combinations (80 + 3 edge cases) of ambient wind speeds and turbulence intensities are selected to run dynamic farm level simulations. Figure 6 compares the individual histograms of sector-averaged quantities after downsampling the ambient cases for dynamic wake simulations, highlighting a good match and showcasing the wide range of local inflow conditions. Eventually, the turbine response surrogates performance for inflow cases not present in the training dataset are shown in Sect. 3.2.

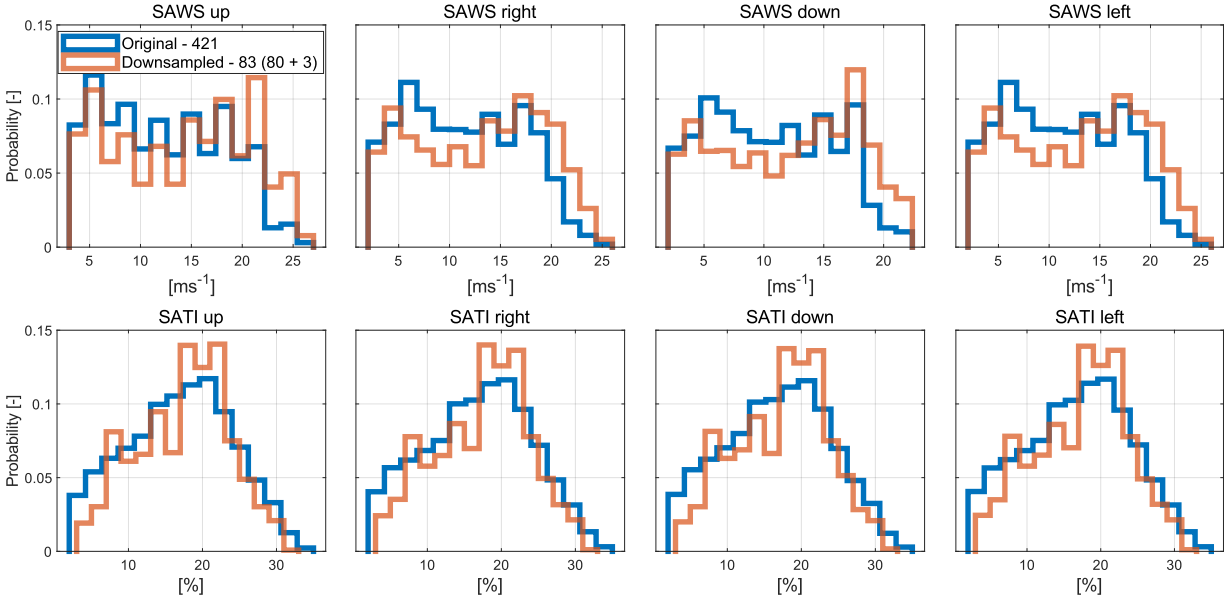

Figure 6. Comparison of probability distributions of sector-averaged quantities after downsampling ambient cases for waked inflow profiles.

For waked inflow profiles, step (a) from Sect. 2.2.2 (Fig. 3) is performed using the DWM implementation in `FAST.Farm` (Jonkman et al., 2018). The NREL-1.79-100 turbine (rotor diameter $100\,m$) is used to generate the wake (NREL, 2025). The turbine is operated at its full power curve (no curtailment) and without any yaw misalignment. Consequently, the resulting library of waked inflow does not include cases with deflected or curled wakes. The assumption is made that partial-wake cases are sufficient for the load surrogate to remain applicable in steered-wake scenarios. This simplification avoids an exponential increase in the number of dynamic wake simulations required for varying yaw misalignments. While the results in Sect. 3.3 show a satisfactory match of the surrogates on the wind farm level, future work will expand the wake-slice library to include misaligned conditions.

Once the 3-dimensional DWM simulations of the single wake are completed, turbulent fields of $120\,\mathrm{m} \times 120\,\mathrm{m}$ are extracted at streamwise locations of $[-0.5,\ 4,\ 5,\ 6,\ 7,\ \text{and}\ 8\,D]$ and lateral locations of $[-1 : 0.25 : 1\,D]$ relative to the wake-generating

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

turbine (with $D = 100\,\mathrm{m}$, referencing the NREL-1.79-100). This procedure yielded a total of 4482 local wake slices. The turbulent slices consist of 25×25 points (i.e., a 5 m spatial resolution) and are sampled at 10 Hz. To ensure statistical convergence of both the inputs (SAWS and SATI) and the outputs (DELs) of the surrogates, 6 turbulent seeds are simulated for each inflow case.

The library of selected synthetic inflow profiles for training and testing dataset generation, including clean and waked cases, is available at (Guilloré et al., 2026a).

3.1.2 Turbine simulation setup

This study considers three utility-scale wind turbine types to capture a representative range of rotor sizes, rated power levels, and hub heights commonly encountered in modern onshore wind farms. The first turbine, NREL-1.79-100, is an open-source reference turbine model (NREL, 2025), with rated power of 1.79 MW, rotor diameter of 100 m and hub height of 80 m. This NREL-1.79-100 is used both to generate the library of wake slices and to demonstrate the suitability of the proposed approach by comparison with farm-level simulations in Sect. 3.3. The two other turbines of interest (Turbine-A and Turbine-B) are commercially deployed turbines from two different manufacturers. These turbines are modern onshore turbines, with rated power around 2 MW and their dimensions are in a similar range as the NREL-1.79-100. For confidentiality reasons, the actual names of these turbines are anonymized.

Each of the turbines is simulated using a different aeroelastic solver, in accordance with the multi-stakeholder approach considered here, where the stakeholders might have proprietary in-house simulation tools. In this study, state-of-the-art aeroelastic codes have been used to simulate the single turbine response, where the simulation code of the Turbine-A is implemented in OpenFAST (NREL, 2025) and the simulation code of Turbine-B is implemented in HAWC2 (Larsen and Hansen, 2015). For the Turbine-A, ENGIE developed and shared the first version of the aeroelastic model, which was subsequently calibrated by CENER for use within the wind farm response framework. For the Turbine-B, the aeroelastic model was developed and calibrated by DTU to be used within the wind farm response framework. Additionally, an open-source reference turbine NREL-1.79-100, developed by NREL and implemented in OpenFAST, has also been included within this study (NREL, 2025).

For the Turbine-A and Turbine-B, dedicated dynamic controllers have been tuned for each operating mode by CENER and DTU, respectively. For Turbine-B, the controller builds on the DTUWEC open-source implementation (Meng et al., 2020). This includes a normal operation baseline controller, as well as controllers for start-up and shut-down transient events. The controller tuning is based on on-site measurements from the wind farms where these turbines are installed. For the NREL-1.79-100 turbine, the ROSCO open-source controller (Abbas et al., 2022) has been used, without modification from the set-up shared at NREL (2025).

Table 1 summarizes the information related to different turbine types, aeroelastic simulation code, controllers, and the corresponding operation and control modes considered within this study. For each turbine and operation mode, multiple control mode combinations are possible; however, this study considers only those that are most relevant from a WFC perspective.

The post-processed simulation output database for the NREL-1.79-100 that is used for surrogate creation is available at (Guilloré et al., 2026a).

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

Table 1. Wind turbine models, operation modes, and considered control modes.

Wind turbine name	Simulation code	Operation mode	Controller	Control mode
Turbine-A	OpenFAST	PP	CENER in-house	NO + NC + YS
		SU		NO + NC
		SD		NO + NC
		PK		–
Turbine-B	HAWC2	PP	DTU in-house	NO + Iso-TSR DR + YS
		SU		NO + Iso-TSR DR
		SD		NO + Iso-TSR DR
		PK		–
NREL-1.79-100	OpenFAST	PP	ROSCO open-source	NO + DE + YS

3.1.3 Surrogate model setup

Although the method is applicable to any load and actuator measurements available, the results in this paper focus on four of the most commonly used DEL channels, denoted as:

- Blade IP: blade bearing in-plane bending moment (non-rotating with pitch).
- Blade OoP: blade bearing out-of-plane bending moment (non-rotating with pitch).
- Tb SS: tower base side-side bending moment.
- Tb FA: tower base fore-aft bending moment.

All of these output channels use a fatigue Wöhler exponent of $m = 4$ for the steel material (Thedin et al., 2025). A summary of all available surrogate models, together with their description and selected Wöhler exponents, is provided in appendix B. The MISO response surrogates trained and evaluated in this work are predicting the mean power and $10\,\mathrm{min}$ DELs at various structural components of different turbine models, for given local inflow conditions (characterized as SAWS and SATI) and control setpoints.

To identify the optimal ANN architecture, a grid search is conducted on two hyperparameters: the number of layers and the number of neurons per layer. In total, 16 different configurations are evaluated across all target outputs and operation modes. Figure 7 presents the scaled mean-squared-error distributions, averaged over all target channels for the power-production and start-up modes. For the PP and PK operation modes, the best-performing architecture consists of three layers with 64 neurons each, whereas for SU and SD with smaller sample sizes per operation mode, a simpler architecture with a single layer and 8 neurons yielded the lowest errors.

Table 2 summarizes the selected hyperparameters for each operation mode. All models are trained using a $80\,\%$–$20\,\%$ split for training and validation. Optimization is carried out in `TensorFlow` (Abadi et al., 2016) using the Adam algorithm (Kingma and Ba, 2014) and the `tanh` activation function.

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

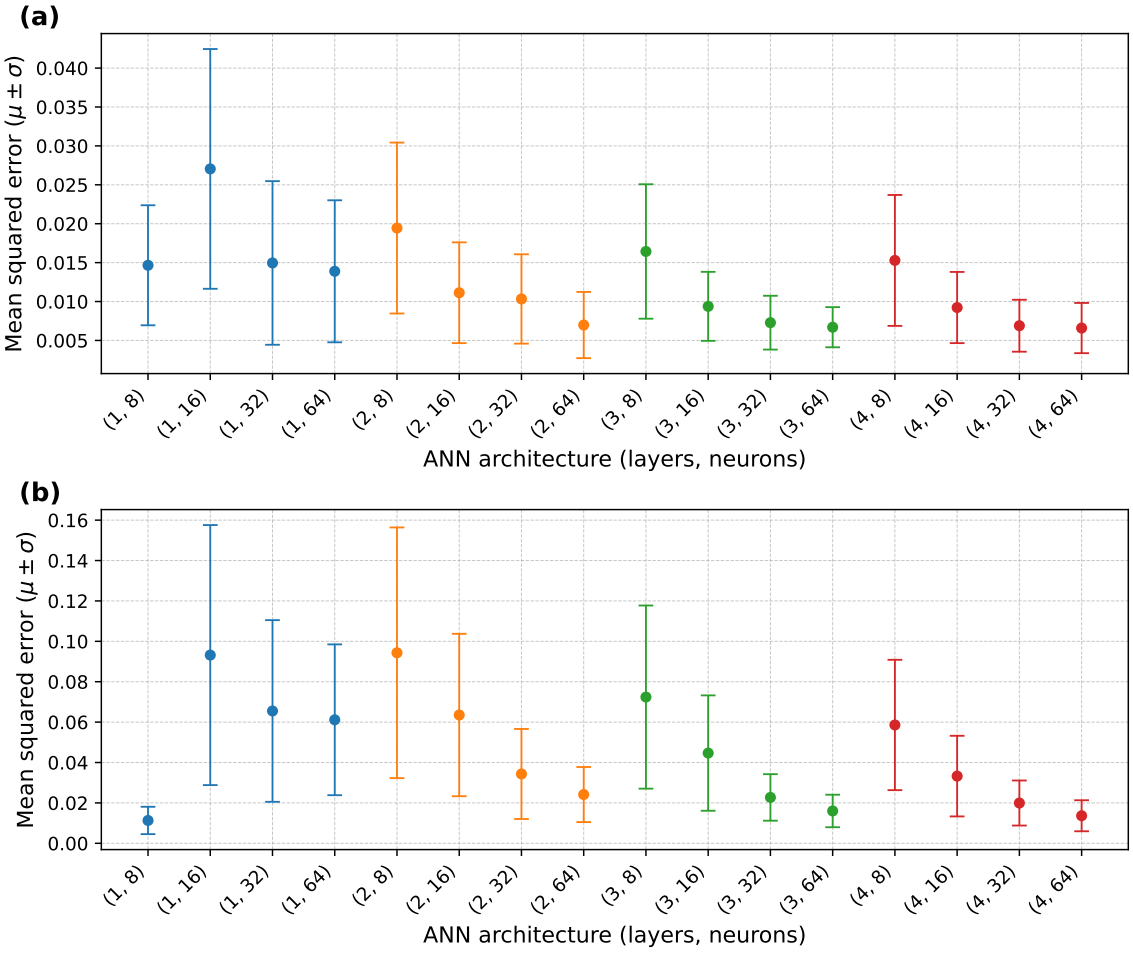

Figure 7. Grid-search results of the ANN architecture for the Turbine-A for the two operating modes: (a) power production, and (b) start-up. Shown are the error bars for the scaled mean-squared errors averaged across all target channels for different hyperparameter combinations (number of layers and number of neurons).

3.2 Performance of the turbine response surrogates

This section presents the performance of the trained ANN models on the training, validation, and testing datasets. The predictions of the DELs by the load surrogates are compared against the values obtained from the aeroelastic simulations to assess predictive capability and quantify the associated errors.

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

Table 2. Neural network hyperparameters for four operation modes: PP, SU, SD, and PK.

Operation mode	# layers	# neurons	activation	optimizer
PP	3	64	tanh	Adam
SU	1	8	tanh	Adam
SD	1	8	tanh	Adam
PK	3	64	tanh	Adam

To quantify the performances of the surrogates the normalized root mean square error (NRMSE) is used, defined as

$$\mathrm{NRMSE} = \frac{\sqrt{\frac{1}{N}\sum_{l=1}^{N}\left(y_l - \hat{y}_l\right)^2}}{y_{\max} - y_{\min}}, \tag{5}$$

where y_l denotes the true (measured) value, $\hat{y}_l$ denotes the corresponding predicted value for the point l, N is the number of samples, and $y_{\max}$ and $y_{\min}$ are the maximum and minimum values of the true data, respectively. The NRMSE is evaluated over the training, validation, and testing sets, as presented in Tables 3, 4, 5 and 6 (values expressed in $\%$). The normalization range ($y_{\max} - y_{\min}$) is fixed on the training set, for each load channel and turbine model.

In addition to NRMSE, the normalized error for each point, defined as

$$\epsilon_M = \frac{y_l - \hat{y}_l}{y_{\max} - y_{\min}}, \tag{6}$$

is also included. The probability density functions of these errors (expressed in $\%$), quantified using the mean (μ) and standard deviation (σ) of the distribution, are considered as an additional indicator of the performances of the surrogates across various datasets, load channels, and turbine type.

Table 3 reports the NRMSE of the load surrogate models for all four DEL channels and both turbine models in the PP mode across the three datasets. For each turbine type, the root mean sqaure error (RMSE) values are normalized by the range of the corresponding channel in the training dataset, following Eq. 5, and expressed as a percentage

Table 3. NRMSE of the trained surrogates for various sets, DEL channels, and turbine types for the PP mode. The range of each channel in the *training* set is used for normalizing the RMSE, expressed in $\%$.

	Blade IP M_{eq}	Blade OoP M_{eq}	Tb SS M_{eq}	Tb FA M_{eq}
Turbine-A– *training*	0.30 %	0.30 %	0.54 %	0.37 %
Turbine-A– *validation*	0.36 %	0.37 %	0.60 %	0.42 %
Turbine-A– *testing*	0.66 %	0.65 %	1.26 %	1.44 %
Turbine-B– *training*	0.47 %	0.26 %	0.58 %	0.35 %
Turbine-B– *validation*	0.50 %	0.30 %	0.59 %	0.37 %
Turbine-B– *testing*	1.45 %	1.55 %	1.14 %	1.00 %

The performance metrics indicate that the surrogates trained on the extensive datasets perform very well, with all NRMSEs below 2 % of the corresponding prediction range. In particular, the NRMSEs for the validation datasets are comparable to those

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

of the training datasets, demonstrating good generalization and indicating that overfitting is avoided. For the testing datasets, which include unseen inflow conditions and control set-points, the NRMSEs are higher across all load channels but remain within an acceptable accuracy range. Notably, the NRMSEs are consistent across different load channels and turbine models, confirming that the surrogate methodology is broadly applicable and that the proposed wake-slice approach achieves high predictive accuracy across diverse datasets.

To examine the surrogate performance in the PP operation mode in more detail, Fig. 8 presents probability density functions of the surrogate errors normalized by the range of each channel for the training, validation, and testing datasets of Turbine-A and Turbine-B. The error distributions are well centered around $\mu = 0$, indicating no systematic bias, and all errors remain small, with standard deviations σ below 1.5 % of the prediction range. It should be noted that the PP mode encompasses a wide range of inflow conditions (free-stream, partially and fully waked) as well as control strategies (NO, YS, Iso-TSR DR, and NC). As expected, the error distributions are broader for the testing datasets, which include unseen cases that require interpolation; however, the surrogate performance remains more than adequate for the intended applications.

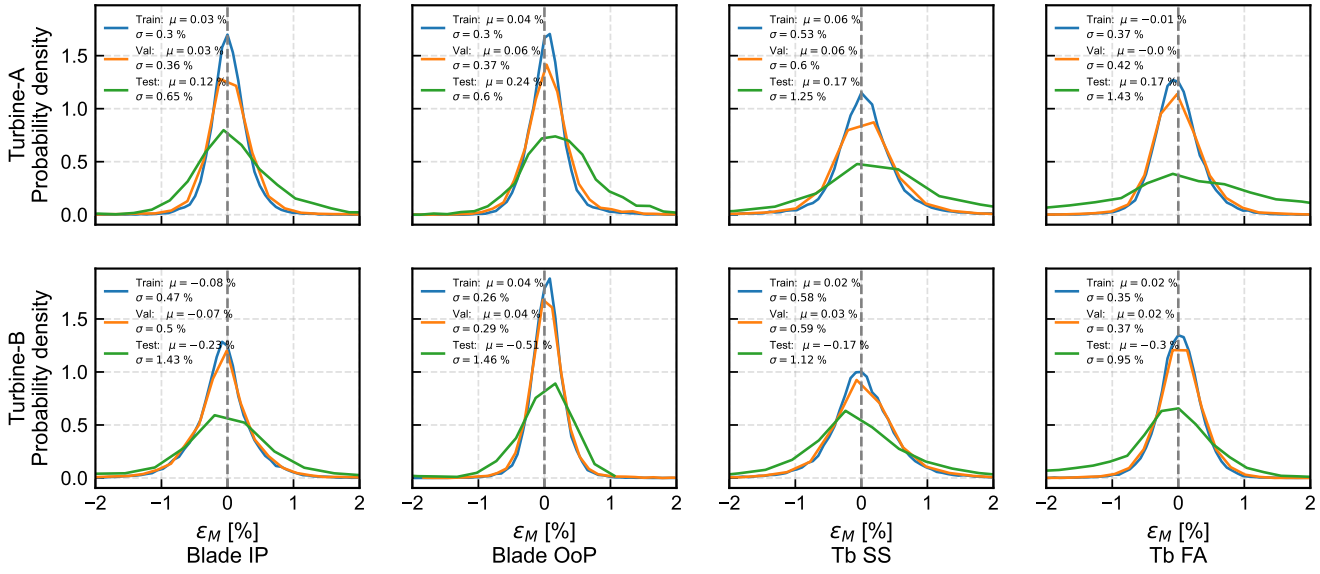

Figure 8. Probability density function of the error distributions (expressed in %) between the surrogate prediction and the target data from the aeroelastic simulations for the PP mode, normalized to the respective target value ranges, for training, validation and testing datasets. The rows of the subplots represent the two turbine types and the columns the selected load channels.

All extended NRMSE results for the SU, SD and PK operation modes are reported in Tables 4, 5, and 6, respectively. The performance in SU and SD modes is notably lower than in the PP mode, due to the smaller number of training and testing cases (i.e., lower number of samples for surrogate model training) and the inherently higher complexity of these maneuvers. The corresponding loads are more stochastic because SU and SD occur over short time windows, during which the local inflow

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

experienced by the turbines is highly variable. Nevertheless, the NRMSEs for SU and SD remains below 10 %. In contrast, the surrogates exhibit strong predictive capability for the PK mode, with errors remaining below 4 %, reflecting the reduced complexity of the idling condition and the lower sensitivity to inflow turbulence during these events. While the NRMSE values are very similar for Turbine-A and Turbine-B under PP, SD, and SU operation, Turbine-A exhibits significantly higher NRMSEs in the PK mode. This difference is likely related to disparities in the representation of structural design and idling behavior between the two turbines, which may lead to different sensitivities to stochastic inflow realizations.

Table 4. NRMSE of the trained surrogates for various sets, DEL channels, and turbine types for the SU mode. The range of each channel in the *training* set is used for normalizing the RMSE, expressed in %.

	Blade IP M_{eq}	Blade OoP M_{eq}	Tb SS M_{eq}	Tb FA M_{eq}
Turbine-A– *training*	3.28 %	2.51 %	2.95 %	2.11 %
Turbine-A– *validation*	3.01 %	2.55 %	3.02 %	2.64 %
Turbine-A– *testing*	6.24 %	3.12 %	6.16 %	2.78 %
Turbine-B– *training*	1.80 %	1.16 %	5.91 %	1.71 %
Turbine-B– *validation*	2.14 %	1.55 %	6.97 %	2.09 %
Turbine-B– *testing*	4.69 %	3.70 %	9.46 %	5.27 %

Table 5. NRMSE of the trained surrogates for various sets, DEL channels, and turbine types for the SD mode. The range of each channel in the *training* set is used for normalizing the RMSE, expressed in %.

	Blade IP M_{eq}	Blade OoP M_{eq}	Tb SS M_{eq}	Tb FA M_{eq}
Turbine-A– *training*	2.98 %	2.61 %	1.16 %	2.64 %
Turbine-A– *validation*	3.62 %	2.85 %	1.24 %	3.52 %
Turbine-A– *testing*	5.95 %	6.53 %	2.63 %	6.92 %
Turbine-B– *training*	2.18 %	1.04 %	1.39 %	1.52 %
Turbine-B– *validation*	1.89 %	0.98 %	1.46 %	1.62 %
Turbine-B– *testing*	7.10 %	8.81 %	2.69 %	2.97 %

3.3 Testing of the surrogates at the wind farm level

To evaluate its suitability, the wake-slices approach is assessed using independent wind-farm-level simulations. The DWM implementation in `FAST.Farm` (Jonkman et al., 2018) is employed, as described in 3.1.1. The farm-level simulations comprised six instances of the NREL-1.79-100 turbine (NREL, 2025), arranged to capture a range of wake impingement and wake overlap scenarios. The selected farm layout is shown in Fig. 9. Ambient wind speeds of 8 and 10 ms^{-1} with a turbulence intensity of 6 % are simulated. The farm is rotated to represent inflow wind directions from 225° to 315° in 5° increments, as illustrated in Fig. 9. This results in 38 `FAST.Farm` simulation cases. Since the farm includes six turbines, this farm-level

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

Table 6. NRMSE of the trained surrogates for various sets, DEL channels, and turbine types for the PK mode. The range of each channel in the *training* set is used for normalizing the RMSE, expressed in %.

	Blade IP M_{eq}	Blade OoP M_{eq}	Tb SS M_{eq}	Tb FA M_{eq}
Turbine-A– *training*	3.39 %	0.67 %	1.02 %	0.97 %
Turbine-A– *validation*	2.86 %	0.62 %	0.46 %	1.31 %
Turbine-A– *testing*	3.87 %	1.08 %	3.05 %	3.28 %
Turbine-B– *training*	0.19 %	0.10 %	0.57 %	0.37 %
Turbine-B– *validation*	0.24 %	0.13 %	0.54 %	0.43 %
Turbine-B– *testing*	1.07 %	0.16 %	0.85 %	0.80 %

dataset comprises 228 testing cases for the location-agnostic load surrogates. In these cases, the turbines are operating in full-power normal operation mode (PP+NO), as the focus is on testing multiple wake interactions. Each case is evaluated using six independent turbulent inflow realizations, which are averaged during post-processing to ensure convergence of both the inputs (SAWS and SATI) and the outputs (DELs). The SAWS and SATI are derived from the `FAST.Farm` inflow field at each rotor disk. The loads surrogates of the NREL-1.79-100, trained using single-turbine simulations driven by wake-derived inflow slices as presented in Sect. 2.2.2, are evaluated against this dataset, which captures realistic wind farm operating conditions, including different levels of wake impingement and multiple wake interactions.

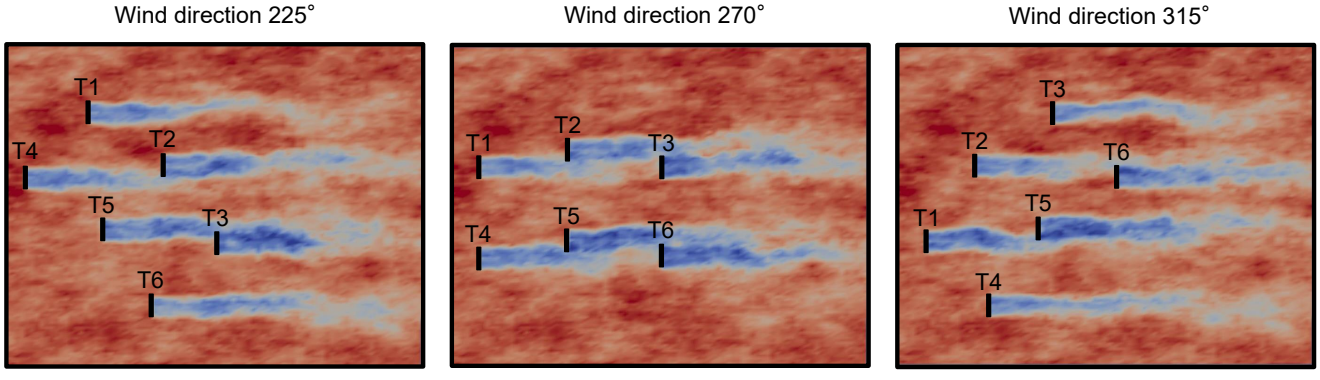

Figure 9. Wind farm layout (with three exemplary wind directions) used for the farm-level testing of the approach. The background illustrate instantaneous view of the flow field at hub height from the dynamic wind farm simulation based on DWM.

Figure 10 illustrates the performance of the different load surrogates in predicting DELs for the farm-level testing cases. The surrogate models are shown by columns, while the individual turbine instances – corresponding to the layout in Fig. 9 – are shown by rows. Results are presented for ambient wind speeds of 8 and 10 ms^{-1}, with the wind direction serving as the primary varying parameter that induces different levels of wake impingement across the farm. For relative comparison across

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

595 cases and load channels, all results in Fig. 10 are normalized by the corresponding value of the free-stream turbine T1 at a wind direction of 270°.

Overall, the results demonstrate that surrogates trained exclusively on single-turbine simulations – using the proposed approach of wake-slices – are able to accurately capture the detailed variations of all four DEL channels under changing wind directions in the farm-level testing. The observed DEL variations arise from the combined effects of changing ambient con-
600 ditions and locally varying inflow conditions caused by wakes. Despite the complex behaviour associated with multiple wake interactions, the surrogates successfully reproduce the detailed variations of the DELs using only the local SAWS, SATI, and control set-points as inputs.

Among the four load channels, the blade root in-plane DELs are predicted with the highest accuracy, as these loads are primarily governed by gravity and relatively simple partial-wake effects. In contrast, the blade root out-of-plane and the tower
605 base side–side and fore–aft DELs are more challenging to predict, owing to their sensitivity to local turbulence levels resulting from a complex interplay between wake-added turbulence – which predominantly increases local turbulence intensity under full-wake conditions – and wake meandering, which mainly affects partial-wake cases. While the blade root out-of-plane and tower base fore–aft DELs are still predicted with very high accuracy by the surrogates, the tower base side–side channel proves to be the most challenging, reflecting its increased sensitivity to flow asymmetry across the rotor disk. Nevertheless,
610 these results demonstrate that the wake-slices approach enables efficient and very accurate training of location-agnostic load surrogates based solely on single-turbine simulations and characterization of the local inflow into SAWS and SATI.

3.4 Comparison of turbine response

Section 3.2 demonstrates that the overall trends predicted by the surrogate models are consistent with the results of the full aeroelastic simulations. Nevertheless, it must be emphasized that extrapolation beyond the training domain can lead to un-
615 physical behaviour, and the surrogates cannot be expected to reproduce the exact turbine response for previously unseen input conditions. The accuracy of the interpolation is highly dependent on the specific surrogate architecture, and commonly quantified via an extensive benchmark. Such an investigation lies outside the scope of the present work and will be addressed in a dedicated follow-up study. Furthermore, Sect. 3.3 showed that the waked behaviour of both single and multiple upstream turbines can be captured with acceptable accuracy for wind farm control applications using the methodology introduced in
620 Sect. 2.2. In summary, the general trends across the investigated input space are validated and found to be in good agreement with the reference simulations. Nonetheless, localized nonphysical artifacts may still arise under certain specific input conditions. These artifacts typically stem from interpolation limitations of the surrogate models in regions sparsely represented in the training data and can manifest as wavy or oscillatory behaviour.

A central contribution of this work lies in demonstrating that, by training surrogate models with an identical architecture
625 and using a common dataset of aeroelastic simulations driven by the same wind inflow conditions, it becomes feasible to perform a consistent comparison of the structural responses of different wind turbine types across a wide range of environmental and operational scenarios. This unified surrogate-modeling framework enables turbines originating from different stakehold-

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

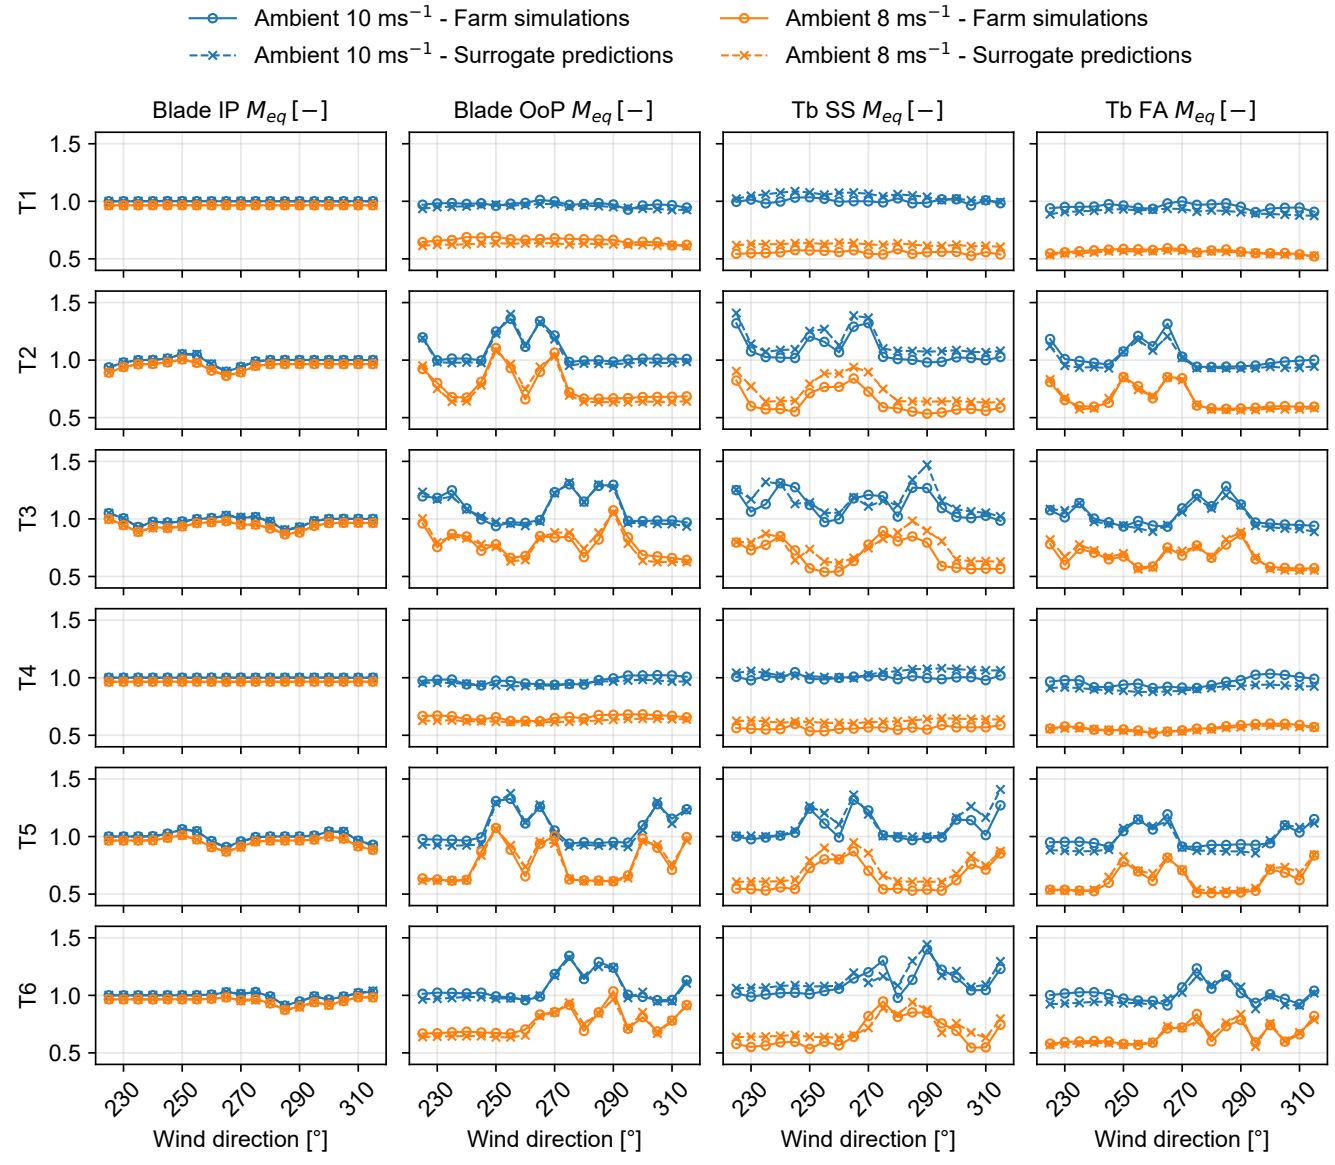

Figure 10. Surrogate predictions (dashed lines and crosses) versus reference values (solid lines and circles) obtained from DWM simulations for the wind-farm-level test cases (six instances of the NREL-1.79-100 arranged as illustrated in Fig. 9), shown as a function of wind direction. All quantities are normalized by the corresponding free-stream reference value of turbine T1 at a wind speed of 10 ms^{-1} and a wind direction of 270°, for each load channel.

ers, simulated with different aeroelastic solvers, to be evaluated on a comparable basis under both clean and waked inflow conditions.

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

In the following, the overall framework is applied to systematically compare Turbine-A and Turbine-B regarding their structural response under varying environmental conditions, waked inflow scenarios, and control setpoints. The rotor inflow used as input to the surrogate models is obtained with the wind-farm modeling framework `Floris` (NREL, 2025), introduced in Sect. 2.1, employing the default parameters of the widely used Gaussian–Curl–Hybrid model (Martínez-Tossas et al., 2021) as well as the Crespo Hernandez turbulence model (Crespo and Hernández, 1996). To evaluate the sector-averaged quantities at the rotor, a discretization of the rotor using a 10×10 grid is chosen. Ambient WS and TI are specified at hub height. A constant vertical shear with a power-law exponent of 0.2 referenced to hub height is assumed. Figures 11, 12, and 14 present results for an isolated turbine operating in free stream, whereas Fig. 13 considers a two-turbine configuration of identical turbine type with varying streamwise and lateral spacings expressed in rotor diameters.

Figure 11 illustrates the trends of five surrogate output channels as a function of the ambient wind speed, with different colors indicating the respective ambient turbulence intensity levels. The first row presents the power normalized to the rated power, as it provides the most intuitive basis for comparison. The subsequent four rows show the Blade IP, Blade OoP, Tb SS and Tb FA DELs.

The power output closely follows the nominal power curve and exhibits noticeable deviations across TI levels primarily in the vicinity of the rated wind speed in Fig. 11. The Blade IP loading remains at a comparable level for both turbine types, with slight increases between 5 and 10 ms^{-1}. At higher wind speeds, the in-plane loading displays a broader spread across TI levels, with higher turbulence intensities leading to increased IP loading. This behavior is consistent for both turbines and can be attributed to the onset of pitch activity at rated wind speed. The increasing pitch angle introduces a component of the flapwise blade response into the in-plane moment, making it more sensitive to wind fluctuations. Additionally, intensified pitch actuation required for power regulation at elevated wind speeds induces stronger aerodynamic responses in the flapwise direction (Zheng et al., 2023). In the out-of-plane direction, the spread of DELs with relation to turbulence intensity becomes noticeable already at low wind speeds. A distinct change in the slope of the bending moments is observed at rated wind speed, which can be attributed to variations in the airfoil properties in the out-of-plane direction as the blade undergoes pitch rotation. The tower-base side–side loads increase primarily with higher wind speeds and elevated turbulence intensities. This trend arises because side–side loading is strongly influenced by the lateral components of turbulent inflow as well as by control-induced effects, which intensify above rated wind speed due to increased pitch activity (Pamososuryo et al., 2024). For the fore–aft loading, turbine thrust and turbulence intensity are the dominant contributors. Consequently, an increase in loading is observed up to rated wind speed. At higher wind speeds, turbulence intensity becomes the more influential factor. Although the blade loads of the Turbine-A and Turbine-B are similar in magnitude, the Turbine-A exhibits higher tower-base loads. This occurs despite its generally lower thrust force across the investigated wind-speed range and its smaller hub height, indicating that additional factors – such as differences in mass distribution, control strategies, and tower eigenfrequencies – play a significant role.

Figure 12 shows the selected target channels as a function of wind speed, but now differentiated by operation mode. This representation makes it possible to clearly identify the wind-speed ranges in which specific operation modes impose higher loads than power production for a given channel. As expected, parked conditions consistently result in the lowest loading, whereas

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

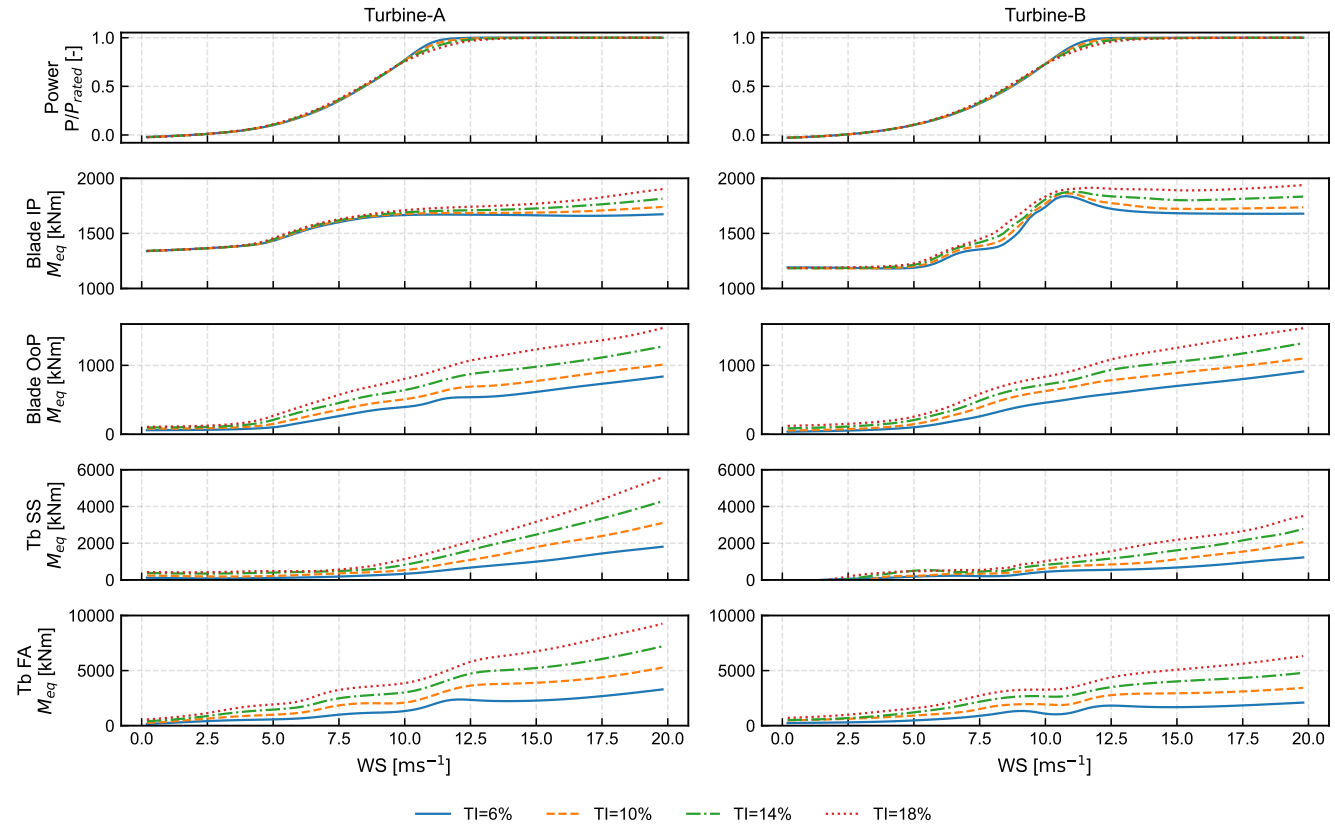

Figure 11. Surrogate-model outputs as functions of wind speed for varying turbulence intensities. The left column shows the response of the Turbine-A, the right column that of the Turbine-B.

start-up and shut-down events substantially increase loads across a broad wind-speed range. For the tower-base fore–aft and blade out-of-plane channels in particular, a pronounced peak occurs around 10 ms^{-1}. At this wind speed and for the given turbulence intensity, a shut-down event can produce approximately twice the damage equivalent load observed during normal power production.

The influence of wake effects on the structural response of wind turbines remains an active field of research. The novel methodology for generating an extensive database of synthetic waked inflows, as described in Sect. 2.2, enables systematic comparisons of wake-induced responses across different turbine types. This capability is illustrated in Fig. 13, which presents results for various partial and full wake configurations, evaluated for five output channels at a wind speed of 10 ms^{-1} and a turbulence intensity of 10 %. The variable ΔY on the x-axis denotes the lateral offset between two turbines. A value of ΔY=0 corresponds to perfect alignment, whereas negative offsets indicate that the downstream turbine is shifted to the right relative to the inflow direction, causing the wake to impinge on the left side of its rotor. The three curves in each subplot represent

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

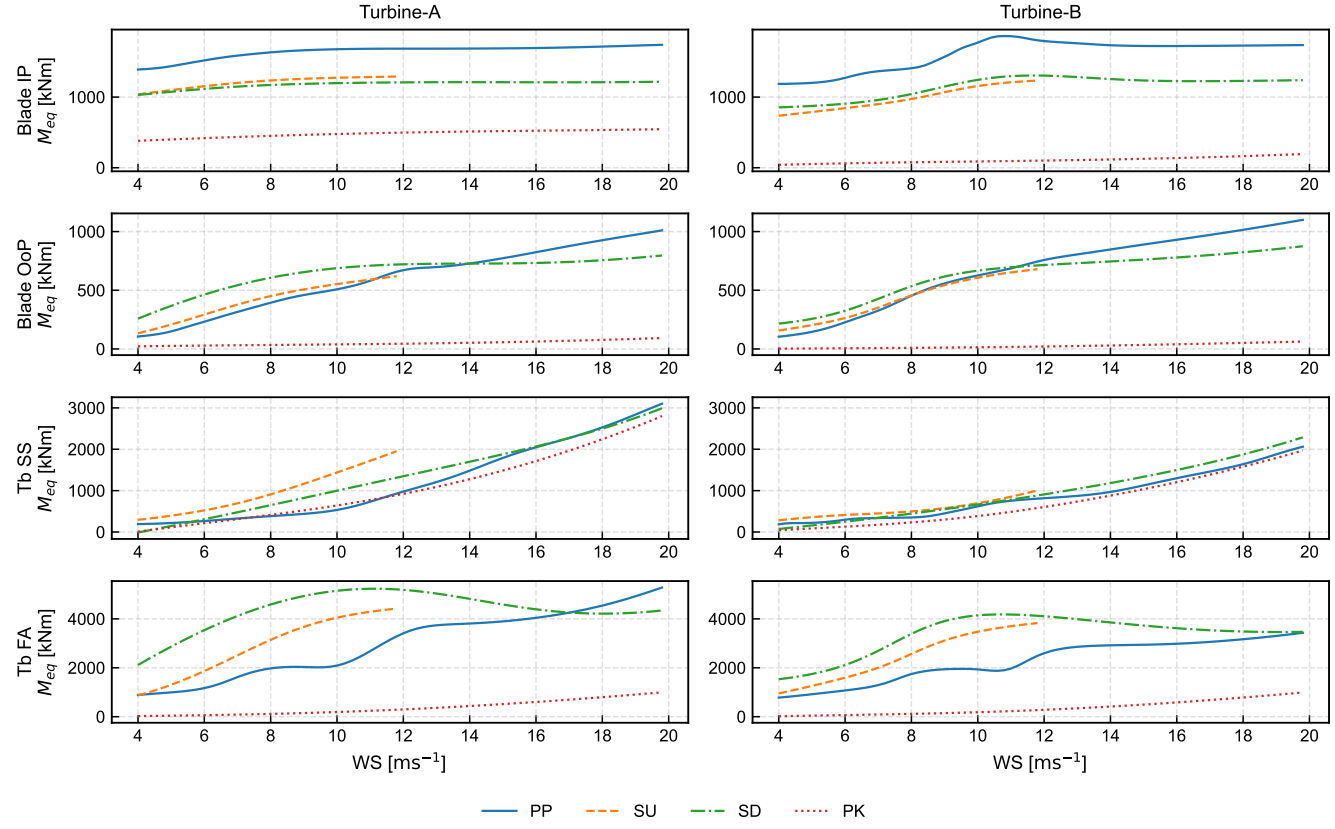

Figure 12. Surrogate-model outputs as a function of wind speed for various operating modes at a turbulence intensity of 10 %.

streamwise spacings of four, five, and six rotor diameters. For the power output, the characteristic Gaussian distribution of the wake-induced velocity deficit is clearly visible. As the wake recovers with increasing downstream distance, the magnitude of the velocity deficit decreases accordingly. For the blade in-plane loads, the asymmetry between left-impinging (ΔY<0) and right-impinging wakes is governed by the direction of rotation of the rotor. When the wake impinges on the left side of the rotor (lateral position Y<0 relative to the wake center), the upward-moving blade encounters reduced aerodynamic momentum while working against gravity, resulting in increased cyclic load amplitudes. Conversely, a wake impinging on the right side of the rotor (Y>0) leads to reduced DELs, in some cases even lower than those observed under full-wake or free-stream conditions. This is well aligned with previously reported results in the literature (Stanley et al., 2022). This trend is consistent across both turbine types. Notably, the influence of turbulence intensity is less pronounced for left-impinging wake conditions. Similarly, the blade out-of-plane loads increase for left-impinging wakes, although both left- and right-side impingement lead to higher loads compared to a fully centered wake. The influence on the side–side loading at the tower base shows less consistent behavior between the two turbine types. Whereas the Turbine-A exhibits the highest loading under full-wake conditions, the

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

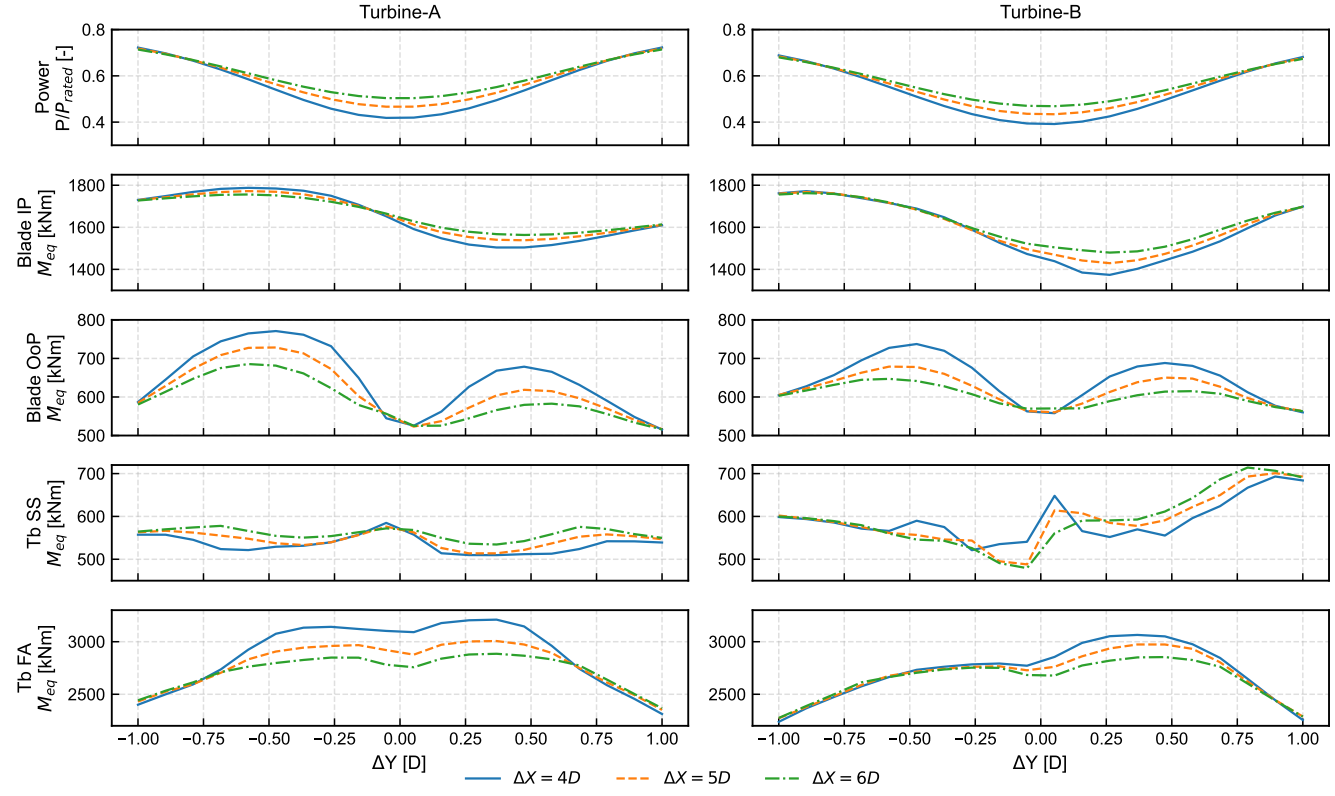

Figure 13. Surrogate-model predictions for partial wake overlaps at a wind speed of 10 ms^{-1} and a turbulence intensity of 10 %.

Turbine-B experiences its maximum loading during right partial-wake conditions. Given the comparatively high training errors for Tb SS, these findings should be interpreted with caution. Differences between the turbine models are also evident in the fore–aft loading. The Turbine-B shows the highest loading for right partial wakes, while the Turbine-A displays a moderate increase in loading for both partial-wake directions.

Figure 14 illustrates the turbine response to varying control setpoints. The x-axis shows different yaw misalignment angles, while the plotted lines represent the respective noise-reduction modes for the Turbine-A and the power down-regulation levels for the Turbine-B. Because the underlying controller strategies for noise modes and power down-regulation differ substantially, a direct one-to-one comparison between the two turbine types is not feasible. Nevertheless, certain qualitative trends with respect to yaw misalignment can still be contrasted. For the Turbine-B, the power output exhibits a trend that aligns well with a cosine-law-based model (Tamaro et al., 2024). In contrast, the Turbine-A shows a more complex and less easily parameterizable power response. A notable feature in both turbines is the approximately linear decrease in the Blade OoP load for yaw angles between -30° and +30°. The tower base loads, however, differ markedly between the two turbine types. While the side-to-side

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

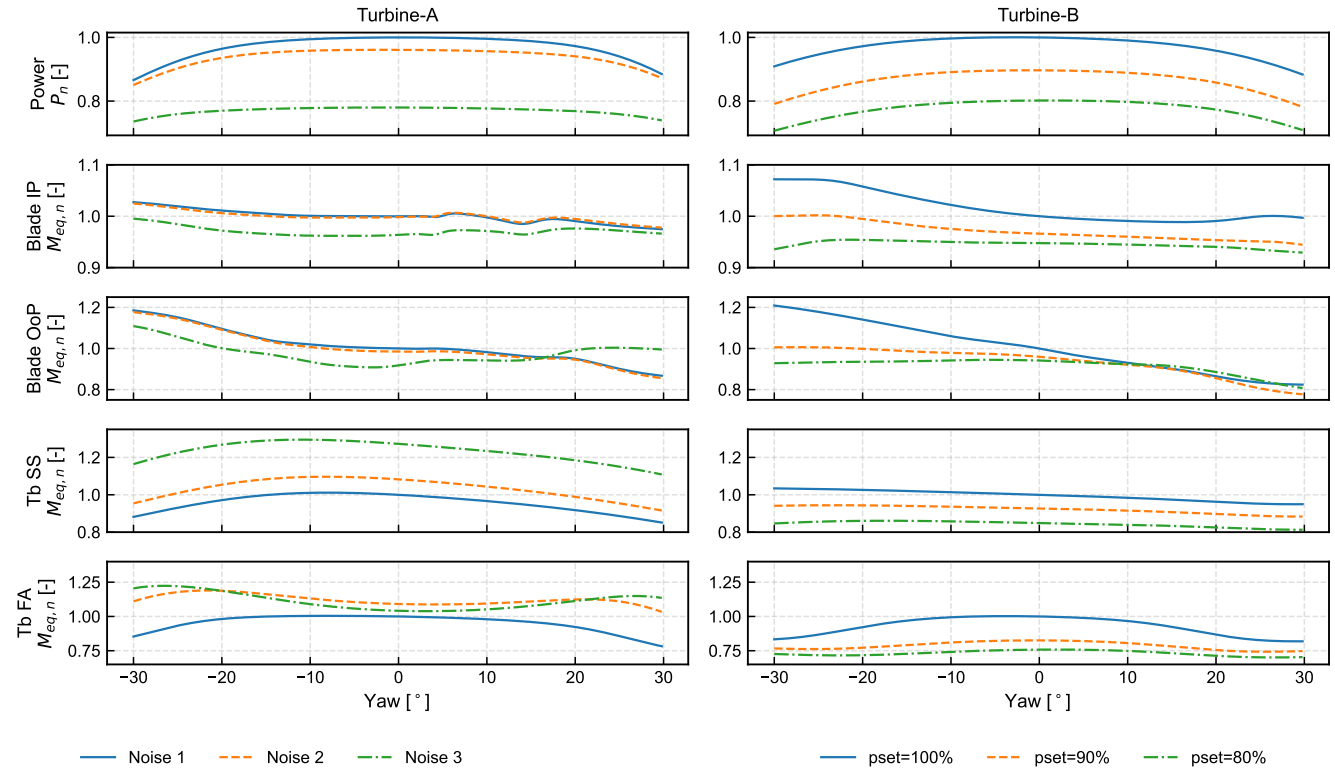

Figure 14. Surrogate-model outputs for various yaw angles and power/noise setpoints, normalized to the 0 degree-yaw output of the normal operating mode (Noise 1, 100 %) at a wind speed of 12 ms^{-1} and a turbulence intensity of 10 %.

loading of the Turbine-A decreases for both positive and negative yaw misalignments, the Turbine-B shows no clear sensitivity of side-to-side loading to yaw angle. Similarly, no consistent or shared trends are observed in the fore-aft loading direction.

3.5 Demonstration on an example wind farm

The primary objective of the developed framework is to model the response of a wind farm to various control strategies. By leveraging an extensive database of partial-wake conditions, a broad spectrum of intra-farm flow scenarios can be reproduced. This capability enables the assessment of lifetime-related metrics, such as the RUL at the wind-farm level. The applicability of the framework is demonstrated through a case study involving a wind farm composed of six wind turbines, either of type Turbine-A or Turbine-B. The corresponding layout has been introduced previously in Sect. 3.3 and illustrated in Fig. 9. The same wake model setup as in Sect. 3.4 is used.

A synthetic time series of the wind speeds, wind directions and turbulence intensities is generated for one year with the following distribution parameters:

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

- Wind speed: Weibull distribution with shape parameter $a = 4$ and scale factor $c = 10$ ($U \sim \mathrm{Weibull}(a, c)$)
- Wind direction: von Mises distribution with the dominant direction $\Gamma_0 = 270$ and concentration parameter $\kappa = 4$ ($\Gamma \sim \mathrm{vonMises}(\Gamma_0, \kappa)$)
- Turbulence intensity: log-normal distribution with mean $\mu = 0$, standard deviation $\sigma = 6\,\%$, and amplitude factor I_0 ($I \sim I_0 \cdot \mathrm{Lognormal}(\mu, \sigma)$)

The resulting distributions of the environmental conditions are presented in Fig. 15. These lead to heterogeneous turbine responses across the wind farm in terms of structural loading and power production, driven by each turbine position and its corresponding exposure to waked inflow.

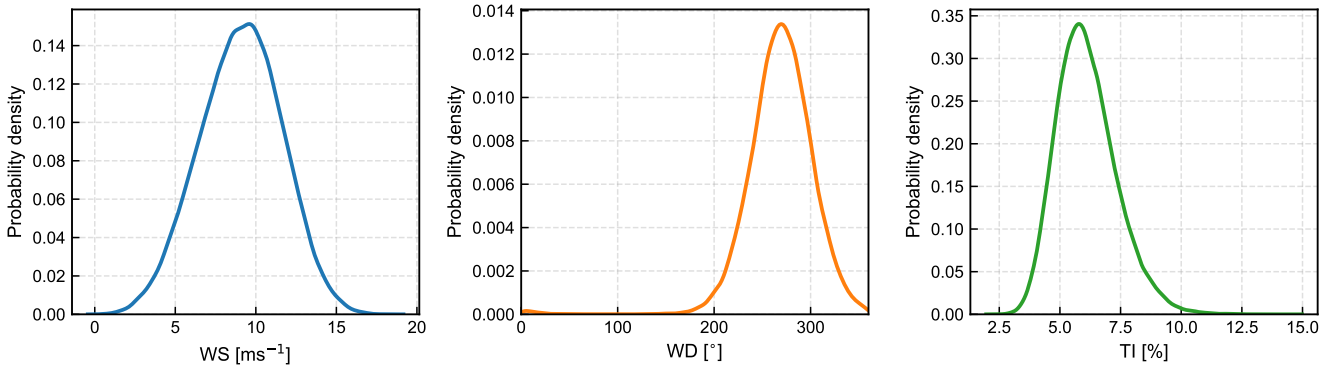

Figure 15. Distributions of wind speed (left), wind direction (center), and turbulence intensity (right).

Furthermore, a yaw optimization was carried out for each turbine type using the `YawOptimization` class integrated in `Floris` (NREL, 2025). The resulting optimal yaw settings are compiled into a look-up table based on a wind-direction discretization of 1 degree over the full 0–360 degree range and a wind-speed discretization of 1 ms^{-1} between 4 and 10 ms^{-1}. The resulting look-up-table is depicted in the Appendix C in Fig. C1 for a wind speed of 8 ms^{-1}. Since the wind direction distribution is dominated by westerly winds, turbines 1, 2, 4, and 5 experience the most frequent yaw misalignment. Figure 16 shows the resulting yaw distributions for all six turbines over the one-year analysis period.

By supplying the defined environmental conditions and yaw setpoints to the framework, the resulting response of each turbine in the farm can be evaluated. Figure 17 depicts scatter plots of the corresponding DELs for turbines 1 and 3 as a function of the rotor-averaged wind speed. The response of the Turbine-A is depicted in the first row, while the response of the Turbine-B is shown in the second row. The color scale indicates the rotor-averaged turbulence intensity. For the given wind conditions, with a predominant wind direction from the west, turbine 1 operates mostly in free-stream conditions, whereas turbine 3 is frequently located in the wake of turbines 1 and 2. Consequently, the four illustrated load channels exhibit markedly different response characteristics. The DEL magnitudes are governed primarily by the rotor-averaged wind speed and turbulence intensity, which

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

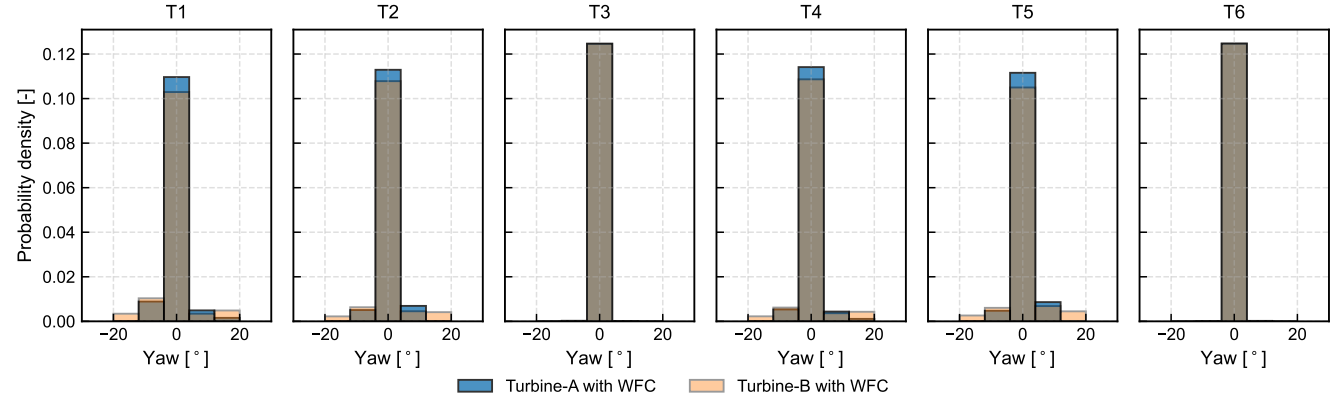

Figure 16. Distribution of yaw setpoints for all six turbines over the one-year analysis period for each turbine type.

explains the elevated DELs observed for many operating points of the downstream turbine (right column of Fig. 17) compared to the free-stream turbine (left column). For waked turbines, the wake position – and thus the SAWS and SATI – introduces an additional layer of sensitivity. In partial-wake conditions, identical rotor-averaged wind speeds and turbulence intensities can still produce substantially higher DELs than for the free-stream turbine when the corresponding sector-averaged quantities differ. This effect is particularly pronounced for blade out-of-plane loads, for which partial-wake overlaps are especially damaging (cf. Fig. 13).

To quantify the variability in the response of individual turbines within the farm, the 10-minute DEL time series is aggregated into a long-term DEL according to

$$M_{\mathrm{eq,LT}} = \left(\frac{1}{N} \sum_{k=1}^{N} M_{\mathrm{eq},k}^{m} \right)^{1/m}, \tag{7}$$

where $M_{\mathrm{eq},k}$ denotes the damage equivalent load in time step k of the one-year time series of length N, $M_{\mathrm{eq,LT}}$ is the resulting long-term damage equivalent moment, and m is the Wöhler exponent.

Figure 18 shows the relative deviations in long-term damage equivalent loads $\Delta M_{\mathrm{eq},WT}$ and power ΔP_{WT} for a given turbine WT with respect to the mean values of all six turbines, within the wind farm of a given turbine-type and control case. These deviations are defined as

$$\Delta M_{\mathrm{eq},WT} = \frac{M_{\mathrm{eq,LT},WT} - \overline{M_{\mathrm{eq,LT}}}}{\overline{M_{\mathrm{eq,LT}}}}, \tag{8a}$$

$$\Delta P_{WT} = \frac{P_{WT} - \overline{P}}{\overline{P}}. \tag{8b}$$

The turbine-wise deviations are presented for both turbine types and for the greedy as well as the yaw-steering control strategies. Across turbine types, the trends in power change, blade out-of-plane loading, and tower-base fore–aft loading are broadly consistent. In contrast, notable differences between turbine types arise for blade in-plane loading and tower-base

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

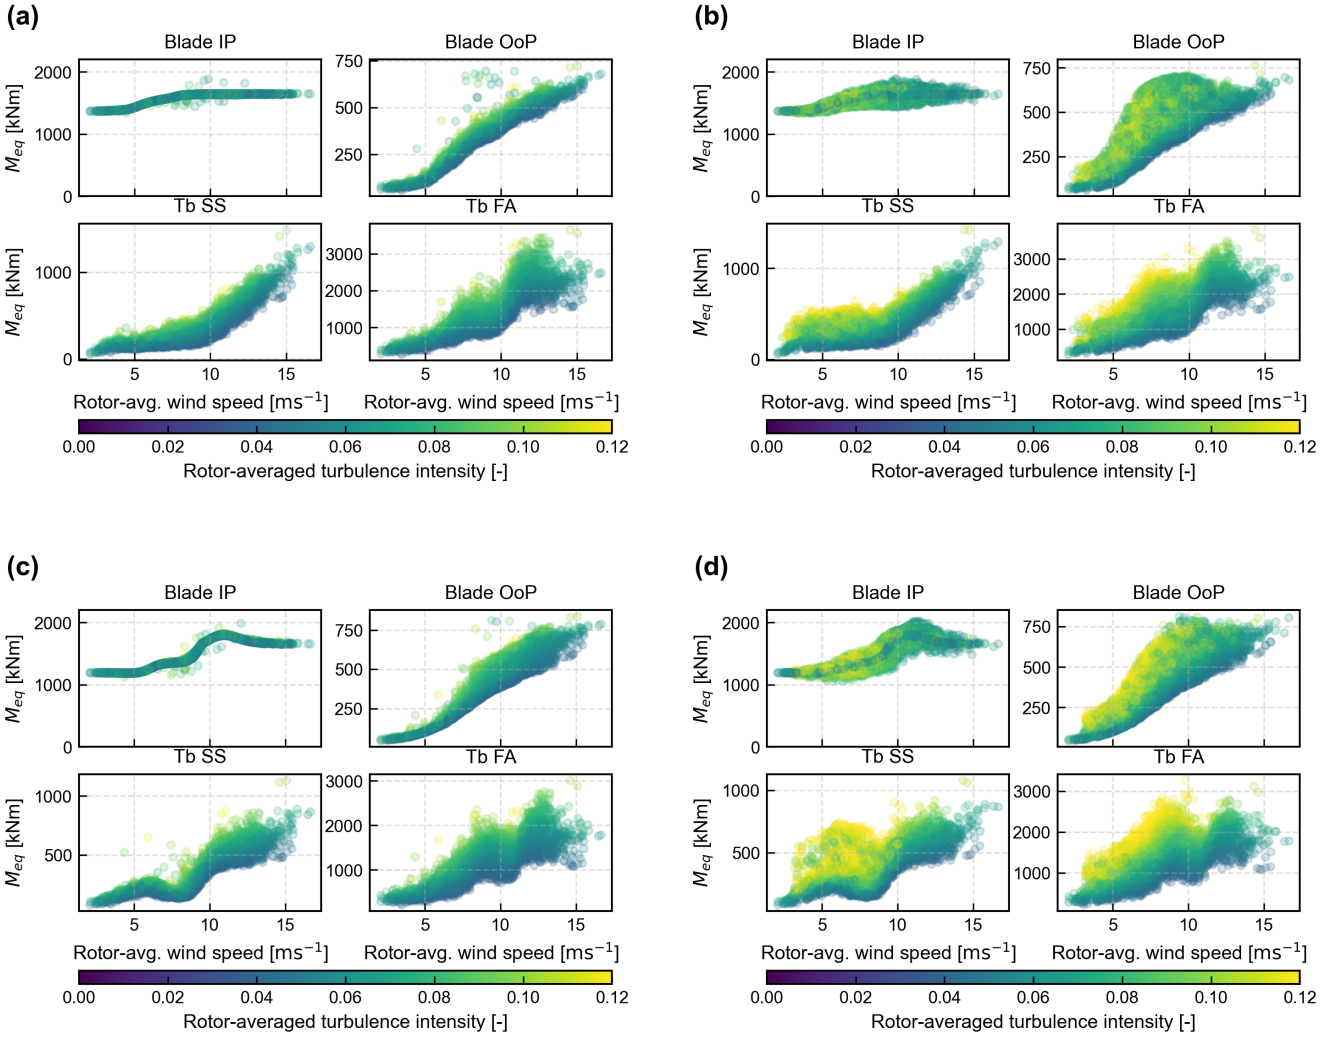

Figure 17. Scatter plots of the DELs for Turbines 1 and 3 (see wind-farm layout in Fig. 9) over the first simulation month without WFC. Panels show (a) T1, Turbine-A; (b) T3, Turbine-A; (c) T1, Turbine-B; and (d) T3, Turbine-B.

side–side loading. These load channels also exhibited the strongest turbine-specific variations in Sec. 3.2 (cf. Figs. 13 and 14). Wake steering influences turbines in distinct ways: upstream turbines experience more frequent intentional misalignment, whereas downstream turbines encounter fewer full-wake overlaps. Nevertheless, the net change relative to greedy control remains subtle.

755 The RUL estimation is limited to the analyzed one-year period and assumes a 20-year design lifetime, IEC Class IIA loading conditions, and no initial damage. The farm-level RUL is obtained by selecting the most restrictive case, i.e., the minimum RUL across all turbines and channels, as defined in 2.5. Blade IP loading proved to be the most restrictive, and the corresponding

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

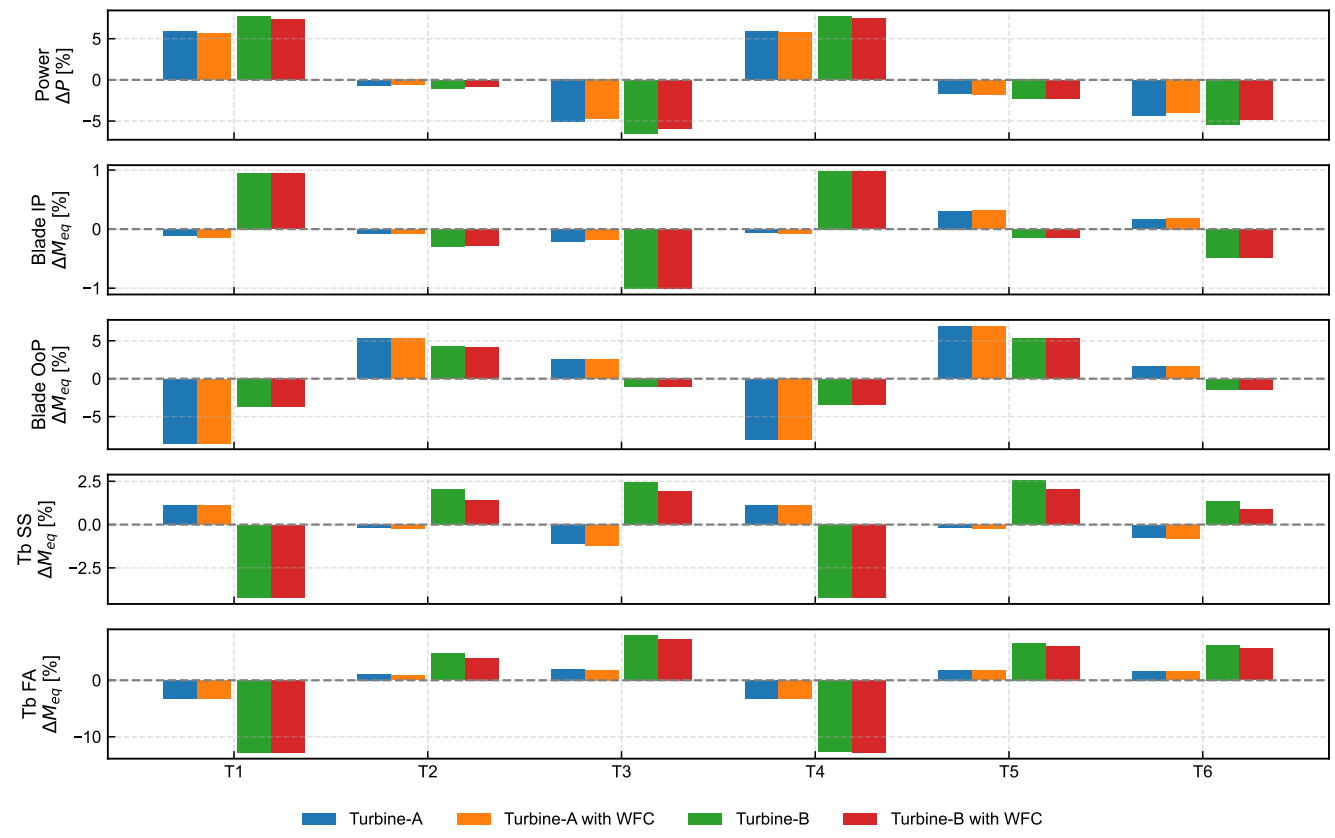

Figure 18. Deviation of power (top) and long-term damage equivalent loads (bottom four subplots) from the farm-wide mean across all six turbines.

Table 7. Remaining useful life calculated for the wind farm setup and the environmental conditions and control inputs presented in Fig. 15

	Turbine-A	Turbine-A + WFC	Turbine-B	Turbine-B + WFC
RUL [years]	21.26	21.29	21.89	21.92

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

4 Conclusions

This paper has introduced a modular and integrated framework for wind farm response modeling that combines computationally efficient steady-state flow modeling, high-fidelity aero-servo-elastic simulations, wake-resolved inflow characterization, and data-driven response surrogates within a single, coherent, and modular architecture. The framework is designed to provide physically consistent and computationally scalable predictions of turbine- and farm-level responses, supporting a wide range of applications spanning research, design, operation, and control of wind farms.

A key contribution of this work lies in the modular architecture of the framework, which allows individual modules to be developed, validated, and replaced independently while remaining fully interoperable. This flexibility supports a broad spectrum of modeling approaches – from physics-based to data-driven and hybrid approaches – and makes the framework well suited for multi-stakeholder use. Researchers, developers, and operators can adopt the level of model fidelity that matches their needs, thereby strengthening collaboration across stakeholder groups. Furthermore, by introducing a standardized wake-slice methodology to represent local waked inflow conditions, it becomes possible to generate a comprehensive library of synthetic inflows that spans both waked and clean states. This, in turn, enables the creation of an aeroelastic simulation database covering a broad range of wind conditions, operating states, and control setpoints. Such a database supports the efficient and practical training of surrogate models using single-turbine simulations, while retaining sensitivity to wake-induced flow features, turbulence characteristics, and control actions. The methodology thus effectively links high-fidelity aeroelastic simulations with large-scale wind farm analyses, avoiding the need for computationally expensive farm-level aeroelastic modeling in many practical applications.

The applicability of the proposed toolchain has been demonstrated across multiple turbine types, including one open-source reference turbine and two commercially deployed turbines. This confirms that the framework can be efficiently applied to turbines with different levels of data availability, structural layouts, and control philosophies, highlighting its suitability for both academic and industrial stakeholders. By employing a consistent simulation-database structure for surrogate training across all turbine types, case-specific customization is eliminated, and the resulting predictions become directly comparable.

The framework consistently captures the effects of diverse operating strategies, including power maximization, yaw-based wake steering, derating, iso-TSR down-regulation, transient shutdowns and startups events, idling, as well as multiple noise curtailment strategies. This consistency across turbine models and control modes represents a key advancement compared to conventional approaches that rely on case-specific tuning or simplified assumptions. The ability to capture both operational and control-induced load variations in a consistent manner is essential for realistic wind farm response assessments and for evaluating trade-offs between energy production, structural loading, and operational constraints.

While the framework can be applied to many end-use applications, this study focuses on a specific set of output channels: electrical power and DELs at the blade root (in- and out-of-plane) and at the tower base (fore–aft and side–side). Overall, the load surrogate models exhibit very satisfactory performance across various datasets, turbine models, and control modes. In power production mode, the NRMSE values of the test sets remain well below 2% for all considered output channels. For start-up and shut-down events, surrogate accuracy reflects the inherently higher stochasticity of these transient maneuvers and

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

their stronger sensitivity to inflow turbulence realizations, resulting in testing errors below 10%. In contrast, under the less stochastic parked conditions, the surrogate models again achieve low testing errors, remaining below 4%.

Validation against full dynamic wind farm simulations shows that the wake-slice methodology provides an accurate and computationally efficient description of local waked inflow conditions for training location-agnostic surrogates. The validation demonstrated that the approach is suitable for applications involving complex multi-wake interactions. Despite its compact formulation, the wake-slice representation enables the construction of response surrogates that are largely independent of turbine location within the farm, while still capturing the load variations induced by wake interactions. This result demonstrates that detailed wake effects can be embedded into turbine-level surrogates without the need for farm-level aeroelastic simulations during surrogate training.

A detailed comparison between the two commercial turbine types highlighted that, while their overall responses may appear similar under certain conditions, significant differences arise due to variations in structural design and control implementation. These differences manifest in both load levels and trends across operation and control modes, illustrating that turbine-specific behavior can strongly influence wind farm response. Such findings emphasize the limitations of simplified or generic modeling approaches and underline the need for a holistic, multi-modal, and computationally efficient framework such as the one presented here, capable of accommodating turbine-dependent characteristics.

The developed wind-farm response framework enables the systematic assessment of alternative farm-level operating strategies for different turbine types. Its capabilities are demonstrated in a case study analyzing changes in power production and DELs for an exemplary six-turbine wind farm operated under two control strategies: greedy control and wake steering. In addition to power and DEL time series, the framework provides computationally efficient estimates of the remaining useful life. This enables a detailed investigation of turbine-specific responses to varying wind and operational conditions, complemented by intuitive lifetime metrics applicable across diverse use cases and stakeholder groups.

Although substantial progress toward holistic modeling has been achieved, the proposed framework still exhibits limitations. The turbines in the intended applications must be in a similar size range as the turbine that was used to generated the library of synthetic waked inflows (for instance, for the inflow profiles generated for the case-study in this work, rotor diameter between 90 m and 120 m and hub height between 60 m and 100 m are acceptable). At present, only a single ambient wind shear value of 0.2 is included, although the inflow dataset could be expanded to incorporate additional shear conditions or other more complex inflow profiles such as veer or low-level jets. Despite covering a broad range of environmental conditions, operational modes, and control modes, the full input space of a specific application study could still not be covered, potentially necessitating reasonable extrapolation strategies. Likewise, interpolation of the surrogate models may lead to non-physical behavior in regions with sparse training data. Furthermore, the surrogate models are limited to capturing structural dynamics within the standard 10-minute simulation window. As a result, low-frequency fatigue cycles driven by wind variability over longer timescales are not represented in the current formulation. Flow phenomena – particularly terrain-induced effects – can be introduced through appropriate flow parameters, yet they remain fundamentally site-dependent. As a result, for sites requiring terrain modeling via speed-up factors, the framework can no longer be considered wind-farm agnostic. Nevertheless,

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

the modular architecture of the framework allows individual components to be replaced with more advanced models, mitigating current limitations as improved capabilities become available.

Future work is needed to extend the synthetic inflow library to further improve surrogate generalization under highly unsteady inflow and extreme events. This would allow to quantify uncertainties to support risk-informed decision-making. The framework should also be applied to a broader range of turbine models, enabling systematic investigation of turbine-specific structural and control design effects on wind farm response. The modular architecture of the framework enables continuous refinement of individual components. For example, the steady-state wind-farm flow model can be advanced by integrating state-of-the-art physics-based formulations with data-driven methods, thereby improving accuracy and robustness across a broader range of operating conditions. Finally, the framework should be deployed in diverse application case studies spanning grid compliance, structural health monitoring, turbine lifetime extension, repowering, and the evaluation of noise emissions and ecological impacts (e.g., wildlife). Addressing these use cases requires the development of additional metrics based on the wind farm response quantities established in this framework. The framework also provides a natural backbone for further studies on wind farm layout and control optimization, including advanced dynamic actuation strategies such as time-varying yaw control and helical wake excitation approaches. Coupling the framework more tightly with wind farm control strategies and field measurement data represents a promising direction for various stakeholders, including research institutes, wind turbine and farm designers, and operators. Overall, the proposed framework provides a scalable and versatile foundation for wind farm response modeling and is intended as an open and extensible platform, encouraging contributions from the wider wind energy community to further develop, validate, and expand its capabilities toward more reliable, efficient, and comprehensive analysis tools for modern wind energy systems.

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

Appendix A: Nomenclature and abbreviations

ADC	Actuator duty cycle
ANN	Artificial neural network
DEL	Damage equivalent load
DE	Derating
DR	Down-regulation
DWM	Dynamic wake meandering
FA	Fore-aft
IEC	International electrotechnical commission
IP	In-plane
LFFC	Low-frequency fatigue cycle
LHS	Latin hypercube sampling
MISO	Multiple-input–single-output
NC	Noise-curtailment
NO	Normal-operation
NRMSE	Normalized root mean square error
OoP	Out-of-plane
PK	Parked
PP	Power-production
RMSE	Root mean square error
R&D	Research and development
RUL	Remaining useful life
SAIQ	Sector-averaged inflow quantity
SATI	Sector-averaged turbulence intensity
SAWS	Sector-averaged wind speed
SD	Shut-down
SS	Side-side
SU	Start-up
Tb	Tower base
TI	Turbulence intensity
TSR	Tip speed ratio
WD	Wind direction
WFC	Wind farm control
WS	Wind speed

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

	YS	Yaw-steering
	Cp	Power coefficient
	Ct	Thrust coefficient
	D	Rotor diameter
	$\mathbb{D}$	Damage
	HH	Hub height
	L	Load cycle amplitude
	M	Load moment
	N	Total number of samples/time-steps
	P	Power
	T	Considered time interval for RUL calculation
	$T1$, $T2$, etc.	Turbine number in the farm layout
	Y	Spatial lateral coordinate (cross-stream), pointing left when looking downwind
	Z	Spatial horizontal coordinate, pointing up against gravity
	c	Number of load cycles
	i	Index variable for ranges of load cycle amplitude
	j	Index variable for meteorological conditions
	k	Index variable for simulation time-step
	l	Index variable for point in the datasets
	m	Wöhler exponent for the material S-N curve
	n	Frequency of occurrence
	r	Radial incremental distance along the rotor blade
	u	longitudinal wind velocity component
	v	lateral wind velocity component
	w	vertical wind velocity component
	y	Variable denoting surrogate output feature
	$\square_{eq}$	Equivalent
	$\square_{IEC}$	Meteorological and operational conditions as per IEC standard
	$\square_{init}$	Initial value
	$\square_{LT}$	Long-term accumulated value
	$\square_{max}$	Maximum value over a sample in a dataset
	$\square_{min}$	Minimum value over a sample in a dataset
	$\square_{ref}$	Reference

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

	$\square_{WT}$	Value for a wind turbine (compared to wind farm average)
	$\Delta\square$	Relative variation compared to a mean
920	$\hat{\square}$	Predicted value obtained from the surrogate
	$\overline{\square}$	Mean value over a sample of points
	ϵ	Normalized error for each data point
	μ	Mean error
925	σ	Standard deviation of error

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

Appendix B: Target channels considered during surrogate creation

Table B1 summarizes the aeroelastic simulation outputs used to train the surrogate models. The target-channel names follow the standardized convention developed in Task 1.4 of TWAIN, which is documented in the published technical report deliverable D1.3 (González-Salcedo et al., 2025).

Table B1. Overview of the load surrogate target channels available within the framework (Vad, 2026a).

Component	Standardized name	Description	Wöhler exp. m
Blade/hub bearing	wrot_Bl1Rad0InPlnMnt	Blade bearing in plane moment (fixed)	4
	wrot_Bl1Rad0OutPlnMnt	Blade bearing out of plane moment (fixed)	4
	wrot_Bl1Rad0ProjInPlnOutPlnMnt	Blade bearing combined moment (fixed)	4
Blade	wrot_Bl1Rad0EdgMnt	Blade root edgewise moment (rotates with pitch)	10
	wrot_Bl1Rad0FlpMnt	Blade root flapwise moment (rotates with pitch)	10
	wrot_Bl1Rad0ProjFlpEdgMnt	Blade root combined moment (rotates with pitch)	10
Shaft	wtrm_LoSpdShaftRotXMt	Shaft torsion moment (same for stationary and rotating)	4
	wtrm_LoSpdShaftRotYMnt	Shaft bending moment (same for Y and Z rotation)	4
	wtrm_LoSpdShaftStaProjYZMnt	Shaft combined bending moment (does not rotate with shaft)	4
Tower	wtow_H0SSMnt	Tower base side-side moment	4
	wtow_H0FAMnt	Tower base fore-aft moment	4
	wtow_H0ProjFASSMnt	Tower base combined moment	4
	wtow_H100SSMnt	Tower top side-side moment (does not rotate with yaw)	4
	wtow_H100FAMnt	Tower top fore-aft moment (does not rotate with yaw)	4
	wtow_H100ProjFJASSMnt	Tower top combined moment (does not rotate with yaw)	4
Power	wtur_W	Wind turbine electrical power	–
Rotor speed	wrot_RotSpd	Rotor speed	–
Blade pitch angle	wrot_Bl1PthAngVal	Blade pitch angle	–
Blade pitch angle ADC	wrot_Bl1PthAngVal_adc	Blade pitch angle actuator duty cycle	–

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

Appendix C: Look-up table of optimal yaw angles used for the application case study

Figure C1 presents the optimal yaw angles as a function of wind direction and turbine for a wind speed of 8 ms^{-1}, which underpins the results discussed in Sect. 3.5. The color scale indicates the yaw magnitude, while the stacked bars show the contributions of the individual turbines.

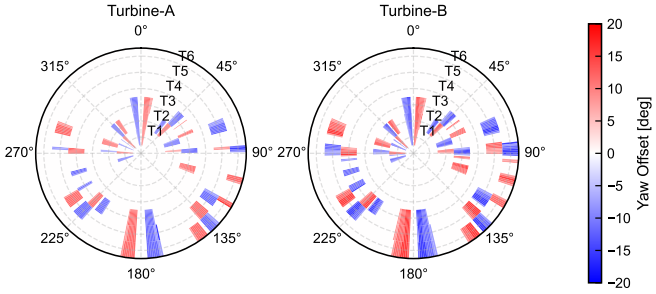

Figure C1. Optimal yaw angles as a function of wind direction and turbine. The color scale encodes yaw magnitude, and stacked bars represent the values for the individual turbines.

Code and data availability. A Python-based implementation of the proposed wind farm response framework, along with the associated data preparation scripts, is available on Zenodo at https://doi.org/10.5281/zenodo.18633504 (Vad et al., 2026b). The database containing a subset of the synthetic inflow slices used for the aeroelastic simulations of the different turbine types considered in this study, as well as the database of simulation outputs used in the surrogate model creation of the NREL-1.79-100, are also available on Zenodo at https://doi.org/10.5281/zenodo.18634032 (Guilloré et al., 2026a). (NOTE: The final DOIs and complete datasets will be provided upon acceptance of the paper).

Author contributions. AV and AG contributed equally to this work, both for the framework implementation and the writing of the manuscript, and share first-coauthorship of this paper. AV led the implementation of the wind farm response framework, the link of steady-state wake models to the surrogates, and generated the results applying the whole toolchain. AG led the implementation of the wake slices approach, the method and training of location-agnostic surrogates, and generated the dynamic farm-level validation results. AA contributed to methodological development and the overall framework implementation, led the standardization of all data sharing and preparation, and wrote a significant part of the manuscript. CLB developed the concepts for a farm-agnostic response framework, a data-augmented wind farm model, and location-agnostic load surrogates, supervised the research, and reviewed the manuscript. TG contributed to the methodological development and provided guidelines on the conceptualization and interpretation of results. VP contributed to conceptualisation and method development, aeroelastic model development and controller tuning, data analysis and visualization, and reviewing the manuscript. AHS generated the library of clean inflow profiles, implemented the selection of all inflow profiles via LHS, and wrote the corresponding parts of the

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

manuscript. ILS contributed to framework and database conceptualization, aeroelastic model development, standardization and simulation, post-processing framework, data analysis, and wrote corresponding parts of the manuscript. MAS contributed to framework conceptualization, standardization of data preparation, implemented the fatigue damage aggregation, wrote the corresponding parts of the manuscript and contributed to the wind farm application example. IE contributed to framework conceptualization, control tuning of the aeroelastic model, operating and control modes, interpretation of results, and reviewing the manuscript. NCG contributed to aeroelastic model development, the development of a post-processing framework, and the creation of an aeroelastic simulation database. IT contributed to data analysis and visualization, code development, and manuscript review. AF provided guidelines on R&D needs and practices, gave feedback on realistic use cases and data availability from the perspective of wind farm operators, and participated in the development of the manuscript through several reviews. KWH and JKK provided inputs on practical aspects and application perspectives for a wind farm operator, and reviewed the manuscript. ND contributed to the methodological development and provided guidelines on the conceptualization and interpretation of results. All authors provided valuable input to this research work through discussions, feedback, and improvement of the manuscript.

Competing interests. The authors have the following competing interests: At least one of the (co-)authors is a member of the editorial board of Wind Energy Science.

Financial support. This work has been supported by the TWAIN project, which receives funding from the European Union's Horizon Europe Programme under the grant agreement number 101122194.

Acknowledgements. The authors would like to thank Antoine Vergaerde (ENGIE Laborelec) for his early technical contributions, as well as Thomas Duc (ENGIE Green) for his valuable inputs reviewing the manuscript.

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

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

References

Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., Kudlur, M., Levenberg, J., Monga, R., Moore, S., Murray, D. G., Steiner, B., Tucker, P., Vasudevan, V., Warden, P., Wicke, M., Yu, Y., and Zheng, X.: TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems, https://www.tensorflow.org, software available from tensorflow.org, 2016.

Abbas, N. J., Zalkind, D. S., Pao, L., and Wright, A.: A reference open-source controller for fixed and floating offshore wind turbines, Wind Energy Science, 7, 53–73, https://doi.org/10.5194/wes-7-53-2022, 2022.

Aparicio-Sanchez, M., Martínez-Merino, I., and Lizarraga-Saenz, I.: FatigueDamage-v1.0.0, https://jupyterhub.twainproject.app/, online code environment for fatigue calculations. Accessed: 2026-01-23, 2026.

Ardillon, E., Bakhoday Paskyabi, M., Cousin, A., Dimitrov, N., Dupoiron, M., et al.: Turbine Loading and Wake Model Uncertainty, D3-2, European Union, https://hal.science/hal-04096504, 2023.

Branlard, E., Martínez-Tossas, L. A., and Jonkman, J.: A time-varying formulation of the curled wake model within the FAST.Farm framework, Wind Energy, 26, 44–63, https://doi.org/10.1002/we.2785, 2023.

Branlard, E., Jonkman, J., Platt, A., Thedin, R., Martínez-Tossas, L. A., and Kretschmer, M.: Development and verification of an improved wake-added turbulence model in FAST.Farm, Journal of Physics: Conference Series, 2767, 092 036, https://doi.org/10.1088/1742-6596/2767/9/092036, 2024.

Braunbehrens, R., Vad, A., and Bottasso, C. L.: The wind farm as a sensor: learning and explaining orographic and plant-induced flow heterogeneities from operational data, Wind Energy Science, 8, 691–723, https://doi.org/10.5194/wes-8-691-2023, 2023.

Crespo, A. and Hernández, J.: Turbulence characteristics in wind-turbine wakes, Journal of Wind Engineering and Industrial Aerodynamics, 61, 71–85, https://doi.org/10.1016/0167-6105(95)00033-X, 1996.

Doubrawa, P., Shaler, K., and Jonkman, J.: Difference in load predictions obtained with effective turbulence vs. a dynamic wake meandering modeling approach, Wind Energy Science, 8, 1475–1493, https://doi.org/10.5194/wes-8-1475-2023, 2023.

Faria, B. R., Sadeghi, N., Dimitrov, N., Kolios, A., and Abrahamsen, A. B.: Inclusion of low-frequency cycles on tower fatigue lifetime assessment through relevant environmental and operational conditions, Journal of Physics: Conference Series, 2767, 042 021, https://doi.org/10.1088/1742-6596/2767/4/042021, 2024.

Göçmen, T., Campagnolo, F., Duc, T., Eguinoa, I., Andersen, S. J., Petrović, V., Imširović, L., Braunbehrens, R., Liew, J., Baungaard, M., van der Laan, M. P., Qian, G., Aparicio-Sanchez, M., González-Lope, R., Dighe, V. V., Becker, M., van den Broek, M. J., van Wingerden, J.-W., Stock, A., Cole, M., Ruisi, R., Bossanyi, E., Requate, N., Strnad, S., Schmidt, J., Vollmer, L., Sood, I., and Meyers, J.: FarmConners wind farm flow control benchmark – Part 1: Blind test results, Wind Energy Science, 7, 1791–1825, https://doi.org/10.5194/wes-7-1791-2022, 2022.

González-Salcedo, A., Eguinoa, I., Nicuesa, X., Duc, T., Sood, I., Kimani, J., Anand, A., Vukobrat, A., Zheng, F., Kersta, P., Slys, P., Davis, N., Dimitrov, N., Dirik, D. G., Göçmen, T., Pierre, S., and Sommer, A.: TWAIN Deliverable 1.3 on Data Standardization Methodology, Technical Report Deliverable 1.3, Version v0.3, CENER; ENG-G; ENG-L; RAM-DK; TUM; SSP; DTU; EDF; VATT, https://doi.org/10.3030/101122194, tWAIN project, funded by the European Union's Horizon Europe Programme under Grant Agreement No. 101122194., 2025.

Guilloré, A., Shah, A. H., and Anand, A.: Datasets as supplement for the journal publication "Modeling wind farm response: a modular, integrated, and multi-stakeholder approach", https://doi.org/10.5281/zenodo.18634032, data set, 2026a.

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

Guilloré, A., Campagnolo, F., and Bottasso, C. L.: A control-oriented load surrogate model based on sector-averaged inflow quantities: capturing damage for unwaked, waked, wake-steering and curtailed wind turbines, Journal of Physics: Conference Series, 2767, 032 019, https://doi.org/10.1088/1742-6596/2767/3/032019, 2024.

Guilloré, A., Chaudhary, A., Anand, A., Vad, A., Shah, A. H., Pettas, V., Göçmen, T., and Bottasso, C. L.: Efficient generation of location-agnostic wind turbine load surrogate models using wake slices, Journal of Physics: Conference Series (under review), 2026b.

He, R., Yang, H., Sun, S., Lu, L., Sun, H., and Gao, X.: A machine learning-based fatigue loads and power prediction method for wind turbines under yaw control, Applied Energy, 326, 120 013, https://doi.org/10.1016/j.apenergy.2022.120013, 2022.

International Electrotechnical Commission: IEC 61400–1: Wind energy generation systems – Part 1: Design requirements, Tech. rep., International Electrotechnical Commission, 2019.

International Electrotechnical Commission: IEC TS 61400–28:2025: Wind energy generation systems – Part 28: Through-life management and life extension of wind power assets, Tech. rep., International Electrotechnical Commission, 2025.

Jonkman, B. J.: TurbSim User's Guide: Version 1.50, Tech. Rep. NREL/TP-500-46198, National Renewable Energy Laboratory (NREL), https://doi.org/10.2172/965520, 2009.

Jonkman, J., Doubrawa, P., Hamilton, N., Annoni, J., and Fleming, P.: Validation of FAST.Farm Against Large-Eddy Simulations, Journal of Physics: Conference Series, 1037, 062 005, https://doi.org/10.1088/1742-6596/1037/6/062005, 2018.

Kingma, D. P. and Ba, J.: Adam: A Method for Stochastic Optimization, arXiv preprint arXiv:1412.6980, 2014.

Kölle, K., Göçmen, T., Garcia-Rosa, P. B., Petrović, V., Eguinoa, I., Vrana, T. K., Long, Q., Pettas, V., Anand, A., Barlas, T. K., Cutululis, N., Manjock, A., Tande, J. O., Ruisi, R., and Bossanyi, E.: Towards integrated wind farm control: Interfacing farm flow and power plant controls, Advanced Control for Applications, 4, https://doi.org/10.1002/adc2.105, 2022.

Kretschmer, M., Jonkman, J., Pettas, V., and Cheng, P. W.: FAST.Farm load validation for single wake situations at alpha ventus, Wind Energy Science, 6, 1247–1262, https://doi.org/10.5194/wes-6-1247-2021, 2021.

Larsen, G. C.: Dynamic Wake Meandering Modeling, Tech. Rep. Risø-R-1607(EN), Risø National Laboratory for Sustainable Energy, Technical University of Denmark, Roskilde, Denmark, https://orbit.dtu.dk/en/publications/dynamic-wake-meandering-modeling, 2007.

Larsen, T. J. and Hansen, A. M.: How 2 HAWC2: The User's Manual, Tech. rep., DTU Wind Energy, Technical University of Denmark, Roskilde, Denmark, https://orbit.dtu.dk/en/publications/how-2-hawc2-the-users-manual, 2015.

Liew, J., Riva, R., Friis-Møller, M., and Göçmen, T.: Wind Farm Control Optimisation Under Load Constraints Via Surrogate Modelling, Journal of Physics: Conference Series, 2767, 092 039, https://doi.org/10.1088/1742-6596/2767/9/092039, 2024.

Martínez-Tossas, L. A., King, J., Quon, E., Bay, C. J., Mudafort, R., Hamilton, N., Howland, M. F., and Fleming, P. A.: The curled wake model: a three-dimensional and extremely fast steady-state wake solver for wind plant flows, Wind Energy Science, 6, 555–570, https://doi.org/10.5194/wes-6-555-2021, 2021.

McKay, M. D., Beckman, R. J., and Conover, W. J.: A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code, Technometrics, 21, 239–245, https://doi.org/10.1080/00401706.1979.10489755, 1979.

Meng, F., Lio, W. H., and Barlas, T.: DTUWEC: an open-source DTU Wind Energy Controller with advanced industrial features, Journal of Physics: Conference Series, 1618, 022 009, https://doi.org/10.1088/1742-6596/1618/2/022009, 2020.

Meyers, J., Bottasso, C., Dykes, K., Fleming, P., Gebraad, P., Giebel, G., Göçmen, T., and van Wingerden, J.-W.: Wind farm flow control: prospects and challenges, Wind Energy Science, 7, 2271–2306, https://doi.org/10.5194/wes-7-2271-2022, 2022.

Miner, M. A.: Cumulative Damage in Fatigue, Journal of Applied Mechanics, Transactions ASME, 12, A159–A164, https://doi.org/10.1115/1.4009458, 1945.

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

Mönnig, A., Hahn, A., Lampert, A., and Römer, U.: Load Estimation in Onshore Wind Farms Using Surrogate Modeling and Generic Turbine Models, Wind Energy Science Discussions, 2025, 1–24, https://doi.org/10.5194/wes-2025-112, 2025.

Moss, C., Maulik, R., Moriarty, P., and Iungo, G. V.: Predicting wind farm operations with machine learning and the P2D-RANS model: A case study for an AWAKEN site, Wind Energy, https://doi.org/10.1002/we.2874, 2023.

Muñoz-Simón, A., Palacios, R., and Wynn, A.: Some modelling improvements for prediction of wind turbine rotor loads in turbulent wind, Wind Energy, 25, 333–353, https://doi.org/10.1002/we.2675, 2022.

NREL: FLORIS v4.5, https://doi.org/10.5281/zenodo.17186844, 2025.

NREL: IEA-scaled wind turbine models, https://github.com/NREL/openfast-turbine-models/tree/main/IEA-scaled, gitHub repository, accessed 12 March 2025, 2025.

NREL: OpenFAST v4.0.2, https://doi.org/10.5281/zenodo.14847846, 2025.

Pamososuryo, A. K., Mulders, S. P., Ferrari, R., and van Wingerden, J.-W.: On the Analysis and Synthesis of Wind Turbine Side–Side Tower Load Control via Demodulation, IEEE Transactions on Control Systems Technology, 32, 1865–1880, https://doi.org/10.1109/TCST.2024.3377508, 2024.

Pettas, V. and Cheng, P. W.: Down-regulation and individual blade control as lifetime extension enablers, Journal of Physics: Conference Series, 1102, 012 026, https://doi.org/10.1088/1742-6596/1102/1/012026, 2018.

Pettas, V. and Cheng, P. W.: Surrogate modeling and aeroelastic analysis of a wind turbine with down-regulation, power boosting, and IBC capabilities, Energies, 17, 1284, https://doi.org/10.3390/en17061284, 2024.

Porté-Agel, F., Bastankhah, M., and Shamsoddin, S.: Wind-turbine and wind-farm flows: A review, Boundary-Layer Meteorology, 174, 1–59, https://doi.org/10.1007/s10546-019-00473-0, 2019.

Sadeghi, N., D'Antuono, P., Noppe, N., Robbelein, K., Weijtjens, W., and Devriendt, C.: Quantifying the effect of low-frequency fatigue dynamics on offshore wind turbine foundations: a comparative study, Wind Energy Science, 8, 1839–1852, https://doi.org/10.5194/wes-8-1839-2023, 2023.

Shah, A. H., Vad, A., Guilloré, A., and Bottasso, C. L.: Towards a load surrogate model for low-frequency fatigue cycles in wind turbines, Journal of Physics: Conference Series (under review), 2026.

Shaler, K., Jasa, J., and Barter, G. E.: Efficient loads surrogates for waked turbines in an array, Journal of Physics: Conference Series, 2265, 032 095, https://doi.org/10.1088/1742-6596/2265/3/032095, 2022.

Stanley, A. P. J., King, J., Bay, C., and Ning, A.: A model to calculate fatigue damage caused by partial waking during wind farm optimization, Wind Energy Science, 7, 433–454, https://doi.org/10.5194/wes-7-433-2022, 2022.

Stipa, S., Ajay, A., Allaerts, D., and Brinkerhoff, J.: The multi-scale coupled model: a new framework capturing wind farm–atmosphere interaction and global blockage effects, Wind Energy Science, 9, 1123–1152, https://doi.org/10.5194/wes-9-1123-2024, 2024.

Tamaro, S., Campagnolo, F., and Bottasso, C. L.: On the power and control of a misaligned rotor – beyond the cosine law, Wind Energy Science, 9, 1547–1575, https://doi.org/10.5194/wes-9-1547-2024, 2024.

Thedin, R., Barter, G., Jonkman, J., Mudafort, R., Bay, C. J., Shaler, K., and Kreeft, J.: Load assessment of a wind farm considering negative and positive yaw misalignment for wake steering, Wind Energy Science, 10, 1033–1053, https://doi.org/10.5194/wes-10-1033-2025, 2025.

Theodoridis, S. and Koutroumbas, K.: Pattern Recognition, Elsevier, Boston, MA, USA, 4 edn., ISBN 9781597492720, https://www.sciencedirect.com/book/9781597492720/pattern-recognition, 2011.

Vad, A.: twain_wind_farm_model: Wind Farm Response Simulation Tool, https://jupyterhub.twainproject.app/, online code environment for wind farm response simulation. Accessed: 2026-01-23, 2026a.

https://doi.org/10.5194/wes-2026-45
Preprint. Discussion started: 9 March 2026

Vad, A., Anand, A., Guilloré, A., Lizarraga-Saenz, I., Gonzalez Lope, R., Aparicio-Sanchez, M., Conti Gost, N., Tsaklis, I., and Göçmen, T.: Code repositories of the wind farm response framework and data preparation scripts as supplement for the journal publication "Modeling wind farm response: a modular, integrated, and multi-stakeholder approach", https://doi.org/10.5281/zenodo.18633504, data set, 2026b.

Vad, A., Shah, H. A., Aktan, D. H., and Bottasso, C. L.: Identifying harmful stress cycles in wind turbine operations, Journal of Physics: Conference Series (under review), 2026c.

van Wingerden, J. W., Fleming, P. A., Göçmen, T., Eguinoa, I., Doekemeijer, B. M., Dykes, K., Lawson, M., Simley, E., King, J., Astrain, D., Iribas, M., Bottasso, C. L., Meyers, J., Raach, S., Kölle, K., and Giebel, G.: Expert Elicitation on Wind Farm Control, Journal of Physics: Conference Series, 1618, 022 025, https://doi.org/10.1088/1742-6596/1618/2/022025, 2020.

Veers, P., Dykes, K., Lantz, E., Barth, S., Bottasso, C. L., Carlson, O., Clifton, A., Green, J., Green, P., Holttinen, H., Laird, D., Lehtomäki, V., Lundquist, J. K., Manwell, J., Marquis, M., Meneveau, C., Moriarty, P., Munduate, X., Muskulus, M., Naughton, J., Pao, L., Paquette, J., Peinke, J., Robertson, A., Sanz Rodrigo, J., Sempreviva, A. M., Smith, J. C., Tuohy, A., and Wiser, R.: Grand challenges in the science of wind energy, Science, 366, https://doi.org/10.1126/science.aau2027, 2019.

Veers, P., Dykes, K., Basu, S., Bianchini, A., Clifton, A., Green, P., Holttinen, H., Kitzing, L., Kosovic, B., Lundquist, J. K., Meyers, J., O'Malley, M., Shaw, W. J., and Straw, B.: Grand Challenges: wind energy research needs for a global energy transition, Wind Energy Science, 7, 2491–2496, https://doi.org/10.5194/wes-7-2491-2022, 2022.

Witter, S., Roth, J., Schelenz, R., and Jacobs, G.: Creating a real-time capable surrogate model of a wind turbine using machine learning, Wind Energy, 28, https://doi.org/10.1002/we.70070, 2025.

Zheng, X., Yao, Y., Hu, Z., Yu, Z., and Hu, S.: Influence of Turbulence Intensity on the Aerodynamic Performance of Wind Turbines Based on the Fluid-Structure Coupling Method, Applied Sciences, 13, 250, https://doi.org/10.3390/app13010250, 2023.