# WES-2026-45 – Answers to Reviewers

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

Authored by: Vad, A., Guilloré, A., Anand, A., Pettas, V., Shah, A. H., Lizarraga-Saenz, I., Aparicio-Sanchez, M., Eguinoa, I., Conti Gost, N., Tsaklis, I., Frère, A., Hermans, K. W., Kamau, J. K., Dimitrov, N., Göçmen, T., and Bottasso, C. L.

Wind Energ. Sci. Discuss. [preprint], https://doi.org/10.5194/wes-2026-45, in review, 2026.

Dear Reviewers,

We sincerely thank you for the time and effort devoted to reviewing our manuscript and for providing thoughtful and constructive feedback. We believe that your comments have helped us strengthen and improve the quality of the paper considerably.

In response to the suggestions and concerns raised during the review process, we have carefully revised the manuscript. In the present document, we provide a detailed point-by-point response to all reviewer comments. Our responses are highlighted in blue.

In addition, we have included a PDF diff highlighting all revisions introduced in the updated version of the manuscript.

Yours sincerely,

The authors.

Black: reviewers comments

Blue: authors responses

## Reviewer 1

### RC1: 'Comment on wes-2026-45', Anonymous Referee #1, 06 Apr 2026

The manuscript presents a framework for surrogate modeling of wind farms. While not providing breakthroughs compared to previous loads surrogate work, it is a good overview and can be a useful standard reference paper.

Dear Reviewer, thanks for your comment. Indeed, the novelty of the paper does not lie in the individual modeling blocks but more on the framework itself, which allows different stakeholders to use it even with proprietary constraints while ensuring accuracy and standardization. The goal is to provide a consistent workflow that is useful to the wind energy community. A paragraph was added to the Introduction (lines 142-146) to highlight this.

The paper lacks evidence for some of the claims in my opinion:

1. The toolchain is claimed to have a low computational cost, but I cannot find that numbers were given to support the claim directly (although of course we can trust it is relatively quick compared to the models providing the training data). How quick/low computational cost are the simulations with the surrogate?

   This is a highly relevant aspect, and we have therefore introduced a dedicated section, "3.1.4 Computational Cost", providing a detailed breakdown of the computational demands associated with synthetic inflow generation, aeroelastic simulations, and the subsequent application of the toolchain.

2. The toolchain is claimed to be fully modular and it is said that blocks can be independently enhanced / updated (line 100-105). There is no clear evidence given imo. Further it is mentioned that wake parameters are to be calibrated in line 186-187, referring to Braunbehrens, 2023. Indeed, the assumption mentioned in line 192-194 is very critical. My own experience is that the turbine loads surrogate models need to be retrained when changing the flow model to get reliable results. Hence maybe the first statement on independent modular model building is too broad/optimistic and we need to be very careful with changing modeling blocks (retraining/calibration may be needed)?

   We appreciate this thoughtful comment regarding the limits of modularity and the circumstances under which surrogate models may need to be retrained. We agree that, in practice, the coupling between the flow model and the load surrogate requires careful consideration, particularly because surrogate models can be trained to compensate for systematic deficiencies of low-fidelity steady-state flow models relative to the mid-fidelity dynamic aeroelastic simulations.

We have expanded both the Introduction and Sect. 2 to address the limitations associated with the modularity of the framework. In addition, we have added a dedicated subsection, 2.1 Framework design principles, which provides a more detailed discussion of the underlying assumptions of the framework. This subsection is organized into three sub-subsections: “Choice of inflow representation”, “Local agnosticity and limitations of surrogate models”, and “Confidentiality barriers.”

Section 2.1.1 “Choice of surrogate inflow representation” specifically examines the two possible approaches for training the surrogate model. The first approach trains the surrogate to represent the mapping between time-averaged inflows of the dynamic flow model (DFM) and the resulting loads. Under this option, the steady-state model must accurately reproduce these inflows to ensure that the framework yields correct DELs.

The second approach trains the surrogate to represent the mapping between steady-state flow model (SFM) inflows and loads. In this case, any modification to the SFM requires retraining the surrogate, an additional step that is not necessary when adopting the first approach.

3. Furthermore in relation with the previous point, the methodology is not quite clear to me in Section 3.4. The wake model used now is a FLORIS model instead of FAST.FARM DWM as used for training the surrogate. Was any recalibration of the FLORIS wake model parameters applied? How do the FLORIS results of Figure 13 compare to FAST.FARM model results? That to me seems a very relevant verification of the framework?

   We agree that this point was not fully addressed in the original manuscript. The surrogate models were trained on time-averaged inflows from FAST.Farm; this has now been clarified in Sect. 3.4. In addition, we added a comparison in Appendix C, where the outputs of the framework using FLORIS are compared with aeroelastic simulations under both clean and waked inflow conditions. As shown in Fig. C1, the surrogate models demonstrate robust predictive capability when coupled with FLORIS, despite having been trained on FAST.Farm inflows.

4. Further, I think the statement in Line 70 generalizes too much in the way it is formulated, considering that many types of wind farm models are applied in industry and research. In the same paragraph, widely used turbulence models are mentioned, but we can also think of electrical modelling, environmental impact modeling etc.

   The respective paragraph was rephrased with a clearer and more precise formulation: “Current wind farm modeling efforts in the context of industrial control and power prediction often rely on highly parameterized, physics-based steady-state models that require careful calibration (Kölle et al., 2022). These analytical wake and power models can provide reliable estimates of farm-wide energy yield, but they are not designed to capture the unsteady aerodynamic and structural dynamics needed for accurate and fast component-level fatigue load estimation."

5. Table B1 has a formatting problem (does not respect the page margins).
   Corrected.

**Citation**: https://doi.org/10.5194/wes-2026-45-RC1

## Reviewer 2

### RC2: 'Comment on wes-2026-45', Anonymous Referee #2, 21 Apr 2026

The paper introduces a new approach to model numerically loads and performance of wind turbines within a farm balancing physics-based models with surrogate models. The work is one of the scientific advancements achieved during thee EU-funded project TWAIN and the scientific approach is rigorously documented.

The paper deserves to be published in the journal of Wind Energy Science. This said, I have some optional recommendations for improvements for the authors.

Dear Reviewer, thanks for your general comment. We appreciate that you have a favorable opinion on the paper. We tackle the remaining recommendations for improvements points by points, with answers below.

1. Readability can be improved: I spent several hours (~16?) across multiple days to digest the document. I initially started reviewing digitally, and I then had to switch to a printed copy to effectively go back and forth multiple times across different sections and figures. I personally did not enjoy that a paper of 39 pages, excluding references and appendices, is only made of four sections. Also, some sections are long and dense, and the use of subsections would greatly improve readability. Introduction is made of 4 dense pages (in page 4, wouldn't a bulleted list be a better option to clearly highlight the novel contributions of this work?). Also, Section 3.4 is 6 pages long. Maybe break it up into subsections?

   We reviewed the manuscript to improve overall readability. As suggested, we added a numbered list in the Introduction to highlight all key novelties. Likewise, the main takeaways are now presented as a bullet list in the Conclusions. In addition, part of Section 3.5 has been moved to the Appendix to streamline the flow of the manuscript.

   Furthermore, we have broken the Introduction and the Conclusions section into various subsections and added a new subsection to chapter 2, to improve readability.

2. Modularity: the paper highlights the modularity of the approach. However, the argument in support of such modularity seems weak. The paper does not clearly explain what modularity means. What are the clear interfaces between steps of the framework? To this point, a more detailed Figure 1 would greatly help readers understand the steps of the framework, the flow of data, and would help follow the discussion of the paper.

   To address this point, we introduced a new subsection, "2.1 Framework design principles," in which we discuss the modularity of the framework and clarify the interfaces between its main components. We also describe scenarios in which changes to one module may require recalibration or retuning of another, thereby avoiding an overly broad claim of full modularity. The manuscript now provides explicit recommendations for such cases. In addition, we discuss the boundaries and limitations of the modularity of the framework in the Introduction.

3. Ability to overcome confidentiality barriers: the paper argues that the new approach helps overcome existing barriers to protect the confidentiality of the data. I find this claim overly optimistic. As you write, aeroelastic models are owned by turbine manufacturers and are seldomly shared with developers and operators (and academia). I don't believe manufacturers treat surrogate models capable of accurately predicting performance and loads much different. Although it's true that surrogate models offer an additional layer of protection to manufacturers against the inverse engineering of their designs compared to aeroelastic models, the claim that these surrogates will be shared with partners seems overoptimistic. Maybe a better discussion about the modularity of the new approach would also help here. Your industry co-authors can also help shed more light on these aspects.

   We agree that our original claim of overcoming confidentiality barriers was overly strong. Our intention was not to suggest that turbine manufacturers will necessarily distribute surrogate models or that confidentiality concerns disappear. Rather, our point is that the proposed framework reduces the amount of proprietary information that must be exchanged by encapsulating turbine-specific behavior within compact surrogate models, while keeping the underlying high-fidelity aeroelastic models entirely confidential. Compared with sharing full aeroelastic models, this modular architecture provides an additional level of abstraction that may lower barriers to collaboration, although any exchange of surrogate models would ultimately remain subject to commercial agreements and the manufacturer's intellectual property policies. To better reflect this distinction, we have added a dedicated section was added "2.1.3 Confidentiality barriers" to soften the claim and explicitly emphasize that the proposed approach reduces, rather than eliminates, confidentiality barriers and enables a modular integration of turbine-specific models where appropriate. We also provided a specific example to illustrate.

4. Computational efficiency: the paper argues that the new framework is more computationally efficient than traditional approaches. While this is not too hard to believe, the paper does not quantify the advantages. Also, previous approaches based on thousands of DWM simulations are well suited for embarrassing parallelism. So, in presence of large computational resources and optimized computational setups, the actual computational time can be quite small (few hours for the training of a surrogate).

We have added a section "3.1.4 Computational cost" in the Results discussing the computational cost of the various steps in using the framework, starting from the creation of the library of synthetic inflows, the aeroelastic simulations and training of surrogates, until the actual application of the framework.

For a fair comparison of computational time, we used comparable computational resources for both the expensive aeroelastic simulations and the efficient final application of the toolchain, namely an Intel i7-4790 CPU with 16 GB RAM. We acknowledge that parallelization on high-performance computing resources can substantially reduce simulation time. However, such workflows require dedicated expertise and computational infrastructure, which may not be available to all stakeholders envisioned as users of this framework. The main motivation of the proposed approach is therefore to support practical deployment in collaborative environments, where stakeholders often differ in their areas of expertise and available software tools. This was also our experience in the present project, in which the generation of the waked inflow library and the aeroelastic simulations were naturally carried out by different partners. By decoupling these tasks, the proposed modular framework allows each stakeholder to contribute within their area of expertise, while avoiding the need for every participant to establish and maintain a complete DWM simulation workflow.

5. Figures: figures often appear much after the paragraph discussing them. In part, this issue will be resolved in the final publication, but during review this aspect was quite painful. One extreme example is in page 5, where Figure 3(e) is mentioned. Figure 3 appears 5 pages later. Also, only 3e is called out, without explaining the other subfigures.

   Figure placement has been improved throughout the manuscript. The previously long-distance reference has been removed, and the concept of sector-averaged quantities is now introduced earlier in Sect. 2.2, which provides a smoother reading flow and avoids interrupting the narrative.

6. Page 7 line 182: this deliverable is hard to find, and the DOI seems to point to the webpage of the project Twain, and not to the actual deliverable. Also, the deliverable itself is not that useful if you are interested in the details of the implementation. Maybe point readers to your code repository? Lastly, the reader is not being helped by the fact that Figure 2 is a subset of Figure 1, but the inputs and outputs to the flowcharts are not 100% clear nor consistent.

   We have updated the reference with a correct link to the public deliverable. Furthermore, we have added a reference to the project code repository where details of the implementation can be found. We have updated the styling of Figure 2 to make clear that it represents a focused subset of Figure 1, enriched with additional detail.

7. Guillore 2026b is not publicly available. The paper will hopefully appear in the coming weeks, but this was another burden on readability. Also, the new sampling approach described in Page 11 lines 274-280 is a clear aspect of novelty compared to Guillore 2024. The paper does not really discuss why it was introduced compared to the old, gridded approach. Maybe this is discussed in the 2026 paper? Regardless, I would highlight and discuss this aspect of novelty compared to already published work.

   Guillore 2026b (and other Torque 2026 papers) have now been published and all references have been updated where needed in the manuscript.
   The LHS-based sampling strategy was indeed also used in Guilloré et al. (2026b), as both studies rely on the same waked inflow library. However, the present manuscript provides the most detailed description and discussion of the sampling methodology, including its motivation, implementation, and role in reducing the computational cost of the framework. Following the reviewer's comment, we have clarified and strengthened the discussion of this aspect as a contribution of the present work, particularly in the introduction where the key contributions are summarized, in the dedicated Section 2.3.3, and in the new Section 3.1.4 on computational cost.

8. Page 12 lines 316-320: I don't fully follow why shutdown and startup cases were extended to 10 minutes. This choice does not look consistent with IEC standards. Wouldn't it be more rigorous to include the fatigue damage caused by a pre-assumed number of startups and shutdowns over the life of each turbine?

   We thank the reviewer for this observation. We agree that the adopted approach differs from standard IEC load assessment practice, where start-up and shut-down load cases are typically simulated only over the duration of the transient event and combined with an assumed lifetime occurrence frequency. In the present study, however, the objective is not certification-oriented lifetime load assessment but the development of a surrogate model that can be linked directly to operational histories available at a 10-minute resolution. To achieve this, all operating modes, including start-up and shut-down, are represented by standardized 10-minute intervals containing a single transition event. The occurrence and frequency of these events are not prescribed a priori but are obtained from the operational history or control strategy being analysed, after which the corresponding load contributions are accumulated in post-processing. We have revised Sect. 2.4 to clarify this distinction and to state explicitly that the proposed approach is intended for operational-history-based load estimation rather than as a replacement for IEC fatigue assessment procedures and that by arbitrarily placing the start-up or shut-down within the 600 seconds- you invariably affect somehow the DELs of that time window.

9. I would split Section 3 into four different sections: 3) Turbine models and simulation setup 4) Surrogate setup, training, and performance 5) Single turbine loads and performance 6) Wind farm loads and performance. I think different readers would be interested in different sections, and you'd help their reading by clearly separating the different topics.

   We agree that clearly separating the different topics is important to facilitate navigation for readers with different interests. However, we prefer to keep all results within a single main Results section, as the different aspects of the proposed framework are strongly interconnected and we believe that splitting them into several independent high-level sections would fragment the presentation and make the overall workflow less cohesive. In addition, we aim to maintain a concise structure with a limited number of top-level sections, as introducing several additional high-level sections would make the paper structure more complex. The current organization already follows the structure suggested by the reviewer, with dedicated subsections covering the turbine models and simulation setup (Section 3.1), surrogate model setup and performance (Section 3.2), validation of the surrogate models against farm-level simulations (Section 3.3), comparison of turbine response (Section 3.4), and the wind-farm-level application case study (Section 3.5). We believe that this subsection-based organization provides a clear separation of topics while maintaining a coherent presentation of the complete framework.

10. Page 26 lines 613-623: again, a proposed improvement to the readability, but this paragraph ends up masking a highlight of the present work, which starts at line 624.

    We have moved this transition text to just before the start of Sect. 3.4, so that this section starts directly with the main highlight of this work.

11. Page 28: currently, the discussion feels verbose and somewhat contorted. Unless I am missing something, the study presented in Section 3.4 and in Figures 11 and 12 could have been performed in a standard aeroelastic solver without the need of the more sophisticated approaches described here. It was a bit of a disappointment to read so many pages to then read some discussion about the impact of TI on loads, which does not seem to require fancy surrogates. Related to this aspect, I am confused about the intended takeaway of Figures 11 and 12. Figure 11 is designed to compare the effect of TI. Wouldn't it have been more impactful to discuss how different turbine designs react differently to same levels of TI? Naively I would have split the TI levels across different subplots, where for each subplot one could immediately compare turbine A vs turbine

B. I thought this would help the narrative of the paper more than a less impactful story about the impact of different levels of TI. Figure 12 raises a similar doubt. I don't really follow what the intended takeaway here is. And why does start up stop at 12 m/s? I would guess that turbines need to be able to start up at any wind speed.

Thanks for this thoughtful and constructive comment. The primary objective of Section 3.4 is not to provide new aeroelastic insights into turbulence effects, but to demonstrate that the proposed framework can be applied consistently across different turbine designs and operating modes using a unified methodology. Following the reviewer's suggestion, we have revised the discussion to better emphasize this intended takeaway. In particular, we now highlight that the framework enables direct comparison of the responses of different turbine models and operating modes within a common computational workflow, rather than focusing on the expected trends associated with turbulence intensity. Although we have retained the original figure layout, we have expanded the discussion to draw more explicit comparisons between the different turbine designs.

We agree that wind turbines can in principle start up at wind speeds above rated, depending on the controller design and restart logic. The restriction to below-rated wind speeds in the present database is therefore not intended as a general operational assumption. Rather, the below-rated range was prioritized because it covers the primary applications considered in this study, particularly temporary environmental curtailment and related operational interruptions. In addition, the start-up implementation of one turbine-controller model produced unsuccessful simulations above rated wind speed due to interactions between the normal pitch-control logic and the prescribed start-up maneuver. Extending the database to above-rated start-up conditions would therefore have required substantial controller modifications beyond the scope of the present work. We have revised the manuscript in Sect. 2.4 to clarify this point and explicitly acknowledge it as a limitation of the current database.

12. Page 28 line 664: less important, but why "As expected"? In my experience many smaller ~2MW turbines are driven by storm loads (without opening the lid on the accuracy of standard engineering models for storm loads...).

We agree that the wording "As expected" was too general, as the dominant loading conditions depend on the turbine design and the considered load cases. We have therefore removed this expression and clarified in the manuscript that this observation applies to the operational conditions investigated in this study.

13. Page 34: wouldn't it be better if these formulas and paragraphs appear in the methodological sections 2.x?

    We agree with the Reviewer and have moved the equation to Sect. 2.6.

14. Figure 17: again, I am not sure what the intended takeaway here is. Also, being picky here, but stylistically I would have put the color bar on the right of each plot, and not right below the x-axis. Currently the labels of the x-axes seem to refer to the color bar.

    This plot illustrates how the DELs correlate with rotor inflow. In free-stream conditions, the DEL magnitudes are primarily governed by the rotor-averaged wind speed and turbulence intensity. Under waked conditions, however, it becomes essential to also parameterize the wake position using sector-averaged inflow quantities. We have revised and expanded the discussion around Figure 17 to make this distinction explicit and to clarify the underlying physical interpretation.

15. Figure 18 has again a lot of data and its takeaways are not exactly immediate. If the conclusion is that wind farm control through wake steering does not negatively impact loads nor clearly benefit performance for this given wind farm for either turbine design, I think that's a conclusion worth highlighting, and I don't think Figure 18 clearly brings home this message.

    Our objective in Section 3.5 is not to draw general conclusions regarding the impact of wake steering on fatigue loads or energy production, as such conclusions are highly dependent on the wind farm layout, site-specific inflow conditions, turbine characteristics, and the selected control strategy. Instead, the purpose of Section 3.5 is to demonstrate the capabilities of the proposed modular framework through a representative application case. We have revised the discussion in Section 3.5 (at its beginning and at its end) to better highlight this objective and clarify the intended takeaway of the presented case study.

16. Conclusions are again 2.5 pages of dense text. Why not splitting them to highlight the takeaways of the work? I suspect that the paper would be easier to read, be more accessible, and ultimately become more impactful.

    We have created a new bullet point list of the main conclusions of the paper, to improve readibility of this conclusion section. Furthermore, we have splitted the conclusions into subsections to improve readability.

Thank you for letting me review your work and congratulations on a successful and impactful research project.

Thank you for the constructive review.

**Citation**: https://doi.org/10.5194/wes-2026-45-RC2

# 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 response 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~~ Performance is evaluated through an exemplary wind farm configuration composed of six turbines, demon-

strating the 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

## 1.1 Background and motivation

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.

## 1.2 Key challenges

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.

- 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~~ evaluates power output and component-level ~~load prediction across~~ loads for given environmental conditions, wind farm characteristics, and control actions, covering key turbine subsystems ~~(e.g.,~~ such as blades, tower, foundations, drivetrain, and pitch~~and yaw systems ), as functions of environmental conditions, wind farmcharacteristics, and control actions~~/yaw systems – for each individual turbine in the wind farm. 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.

### 1.3 State-of-the-art

Current wind farm modeling ~~is primarily focused on power estimation in industrial wind farm control , relying on physics-based tools that are highly parameterizedand sensitive to~~ efforts in the context of industrial control and power prediction often rely on highly parameterized, physics-based steady-state models that require careful calibration (Kölle et al., 2022). ~~While these steady-state analytical models can accurately predict farm-wide~~ These analytical wake and power models can provide reliable estimates of farm-wide energy yield, but they are not designed to capture the unsteady aerodynamic and structural dynamics needed for accurate and fast ~~estimation of component-level~~ component-level fatigue load estimation. As a result, predicting turbine-specific fatigue loads remains challenging despite the availability of more detailed modeling approaches in other domains of wind-farm analysis (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~~ is unable to provide turbine-specific wake effects 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~~loads, these approaches are generally limited to isolated, non-generalizable case studies, with transferability across turbines, sites, and operating conditions not yet consistently demonstrated~~, pointing~~ . This points 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 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 (i.e., they are independent of turbine position within a farm) and control-oriented (~~capturing~~ i.e., they capture 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~~design. 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).

### 1.4 Contributions of this work

~~To this end~~Within this landscape, 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~~

1. 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. However, in some cases, a modification of one component of the model can require the retraining of another –for example, when the change alters the inflow–response relationship learned by the surrogate models. In such situations, the claim of full modularity must be interpreted more cautiously, as the independence between components is no longer strictly maintained.

   ~~Second, it allows physics-based wake models to~~

2. Physics-based wake models can be complemented with correction terms learned from operational data (Braunbehrens et al., 2023), 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~~

3. 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~~ The method introduces a standardized representation of partial and full wakes for aeroelastic simulations~~, based on two hypotheses: in the self-similar far-wake regime,~~. It relies on two assumptions: (i) 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~~ behaviour in the self-similar far-wake regime, and (ii) complex multi-wake flow fields can be represented by equivalent local inflow states. Furthermore, a smart sampling approach is introduced for the surrogate training in order to cover the full realistic operating range while keeping efficient computational time. 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~~

4. 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 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~~ steady-state flow model to a particular ~~installation~~situation. 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.

In summary, the key novelty of this work lies in the modular modeling framework that lets different stakeholders contribute improvements while keeping proprietary information protected. It combines physics-based wake models with data-driven corrections to ensure accuracy, and introduces standardized wake representations that support consistent, comparable simulations. Overall, it provides a unified and transferable workflow that strengthens accuracy, standardization, and usability across the wind-energy community.

### 1.5 Paper structure

The remainder of the paper is structured as follows. Section 2 ~~introduces~~ presents the proposed approach for modeling the wind farm power and load response. It begins ~~with the~~ by outlining the framework design principles in Sect. 2.1, followed by the wind farm flow modeling using steady-state flow models in Sect. 2.2~~, followed by the~~. The synthetic inflow generation described in Sect. 2.3 ~~, which~~ then provides the input for the ~~aero-servo-elastic~~ aero-servo-elastic simulations outlined in Sect. 2.4. These simulations form the training dataset used to develop the surrogate models presented in Sect. 2.5. 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~~ Then, in Sect. 3.3 the generalizability of the synthetic wake inflow database is evaluated 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~~, which in general may present extra- and intra-plant heterogeneity. 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~~a steady-state flow model (SFM), as described in Sect. 2.2. The ~~flow model~~ SFM 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~~ surrogate models

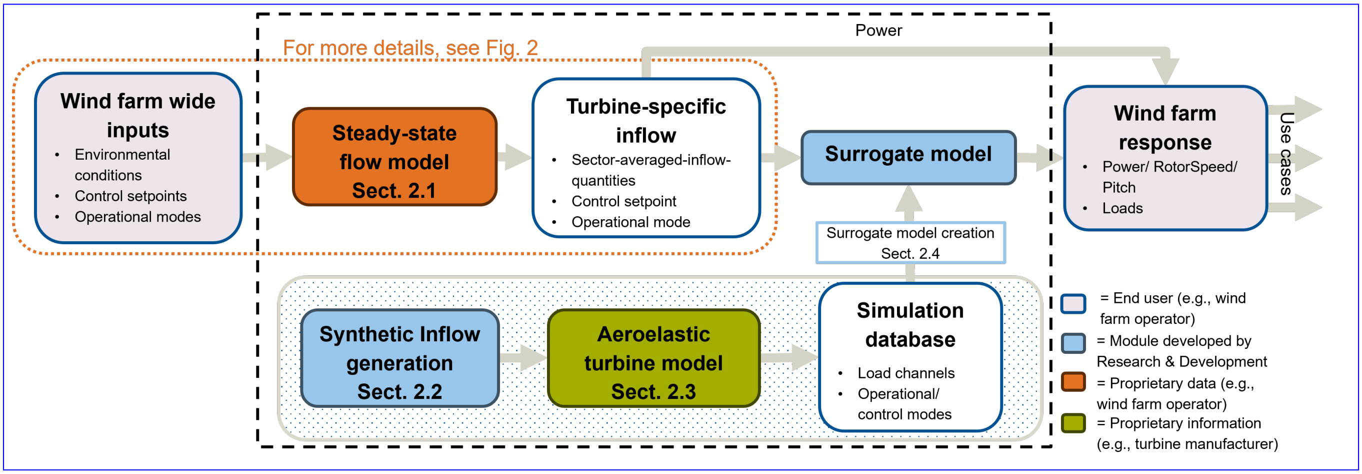


**Figure 1.** Flowchart of the multi-module framework used to model the wind-farm response. Colored elements indicate typical contributions from different stakeholder groups. The dashed black box marks the scope of the present code framework (Vad et al., 2026b). The gray-patterned box summarizes the 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.

trained on a dataset of aeroelastic simulations. The main task of the surrogate models is to map the turbine-specific inflow to the corresponding damage equivalent load (DEL) channel. Using surrogate models involves three distinct phases.

1. Synthetic inflow generation – In the first phase a dynamic flow model (DFM) produces time-varying inflow fields. These inflows are then supplied to an aero-servo-elastic model of the turbine, which computes the resulting structural loads. The load time series are subsequently converted into DELs. The procedure for generating synthetic inflow is detailed in Sect. 2.3, and the details regarding the aeroelastic simulations are presented in Sect. 2.4. While the aeroelastic simulations can be carried out with any commercial or open-source software – potentially involving proprietary data~~and thus supporting collaboration~~ , which might not be freely shared 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. 2.3, and the details regarding the aeroelastic simulations are presented in Sect. 2.4. The response surrogates are described in~~ As a result, they can be reused for turbines of similar size, enabling consistent and transferable surrogate-model development.

2. Surrogate training – In the second phase the surrogate models are synthesized, i.e. they learn the mapping between inflows and DELs. The surrogate models used in this study (Sect. 2.5~~as machinelearning models trained~~ ) are machine-learning frameworks designed to predict turbine responses ~~based on four sector-averaged~~ from a compact set of flow descriptors. Each model is trained on four sector-averaged wind speeds (SAWSs) and four ~~sector-averaged~~ sector-averaged turbulence intensities (SATIs) computed 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~~ , supplemented by the relevant control setpoints that govern turbine operation.

3. The use case – In the third phase the surrogate is fed with steady inflows generated by a SFM, which yields the loads corresponding to those inflows. This phase follows the upper branch of the flowchart and represents the intended operational pathway for applications.

When applied in the use case, the framework efficiently ~~produces~~ delivers wind-farm~~agnostic quantities —~~-agnostic outputs – 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. 2.6.

## 2.1 ~~Wind farm modeling~~Framework design principles

~~The flow model represents~~

### 2.1.1 Choice of surrogate inflow representation

The training phase of the surrogate models can be accomplished in two different ways. A first option is to train the surrogate model to represent the mapping between time averages of the DFM inflows and the time-aggregated loads represented by DELs. Since in the use phase the surrogate is fed with inflows generated by the SFM – but the surrogate learnt the mapping between DFM inflows and loads – it is crucial for the SFM inflows to accurately match the time-averaged DFM ones. When this is not the case, one can try to improve the match by tuning the SFM.

A second option is to train the surrogate to represent the mapping between SFM inflows and loads. In this second case, an accurate match between the two inflows is not strictly necessary, as long as there is a one-to-one correspondence between each element of the two sets. For example, the wake-added TI of the SFM could somehow differ from the time-averaged one of the DFM; however, if the value in one model uniquely defines the value in the other (there is a one-to-one relationship between the two), then the surrogate will yield the correct loads when fed with the SFM inflows. When one adopts this second option, a careful tuning of the SFM might not be necessary (or could be less demanding). However, this means that changing the SFM requires retraining the surrogate, which is not necessary when adopting the first option. In summary, the modular design of the framework provides flexibility, but the need to retrain surrogate models remains application-dependent and must be evaluated by the user.

In the present application, the first option is selected, i.e. the load surrogates are trained on DFM inflows and applied with SFM inflows.

### 2.1.2 Local agnosticity and limitations of surrogate models

By defining inflow quantities at the rotor level together with the control setpoints, the load surrogates become location-agnostic: the loads depend solely on the inflow conditions at the rotor disk, regardless of how the wake evolved upstream. As long as the DFM inflows used to train the surrogate encompass the relevant range of possible flow characteristics, the mapping between inflow and loads will remain valid. This implies that – within certain limits – the surrogate can be driven by an SFM that includes physical effects absent from the DFM. For example, consider the following situations:

- Blockage alters the speed distribution within the farm. Suppose the DFM does not model blockage, but the SFM does. A blockage-aware SFM may predict a slower inflow at a turbine compared to a non-blockage model. The surrogate will still return accurate loads as long as the inflow values remain within its training domain. The origin of the inflow change (in this case, blockage) is irrelevant; only the inflow at the rotor disk matters.

- Terrain-induced speed-up can increase local wind speed, turbulence intensity, or induce direction shifts. Even if the DFM does not include orography, an SFM that does can still provide valid inputs to the surrogate, provided these inflow quantities fall within the training range of the surrogate model. The same principle applies also in this case: the surrogate responds to the inflow itself, not to the physical mechanism that produced it.

There are other effects that – in this work – were not included in the DFM for simplicity, but that could readily be considered, thereby extending the applicability of the surrogate to other scenarios. These include – for example – the presence of veer, of non-conventional shears (e.g. those caused by low-level jets), of terrain-induced vertical flow components (e.g. those experienced by a turbine located close to the crest of a hill), of turbulence intensity (e.g., induced by vegetation), etc. All these and other inflow features can in principle be included in the DFM during the generation of dynamic loads. When this is done, a SFM that accounts for such flow effects (for example, because it contains a terrain model that produces vertical upflows close to a hill crest) will be able to correctly retrieve loads from the surrogate, given that all necessary information to describe these effects are given as inputs.

### 2.1.3 Confidentiality barriers

As emphasized by the multi-stakeholder approach, users may have different perspectives, levels of access to confidential data, and application use cases (e.g. research institutes, turbine manufacturers, or wind farm operators). In some contexts, a single user may possess all necessary simulation tools, aeroelastic models, and site-specific information. In many real-world cases, however, collaboration across stakeholders is essential, and confidentiality constraints limit the exchange of proprietary information. The proposed framework addresses these constraints by minimizing the amount of sensitive data that must be shared. Turbine-specific knowledge is encapsulated within surrogate models, enabling a modular workflow in which detailed aeroelastic models remain confidential. Only lightweight surrogate models are exchanged, and only when permitted by data-sharing agreements. One example is individual pitch control (IPC), whose algorithms and tuning parameters are commonly embedded within proprietary controller implementations. Conventional aeroelastic load assessments often require access to the controller model or detailed information about its logic to reproduce IPC effects. Under the proposed framework, the model owner can run aeroelastic simulations internally without disclosing controller logic or parameters. The only shared output is the turbine response, expressed through the corresponding DELs. During training, the surrogate model learns the influence of IPC implicitly, allowing wind-farm-level analyses to incorporate IPC effects without revealing underlying controller designs or intellectual property.

### 2.2 Steady-state flow model

The SFM calculates 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~~ Throughout this work, we use the term steady-state flow model to denote what can be elsewhere referred to as flow model, farm model, or engineering model. SFMs provide an effective balance between predictive accuracy and computational cost, leading to their broad adoption in both industry and research (Porté-Agel et al., 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.~~ steady-state flow model.

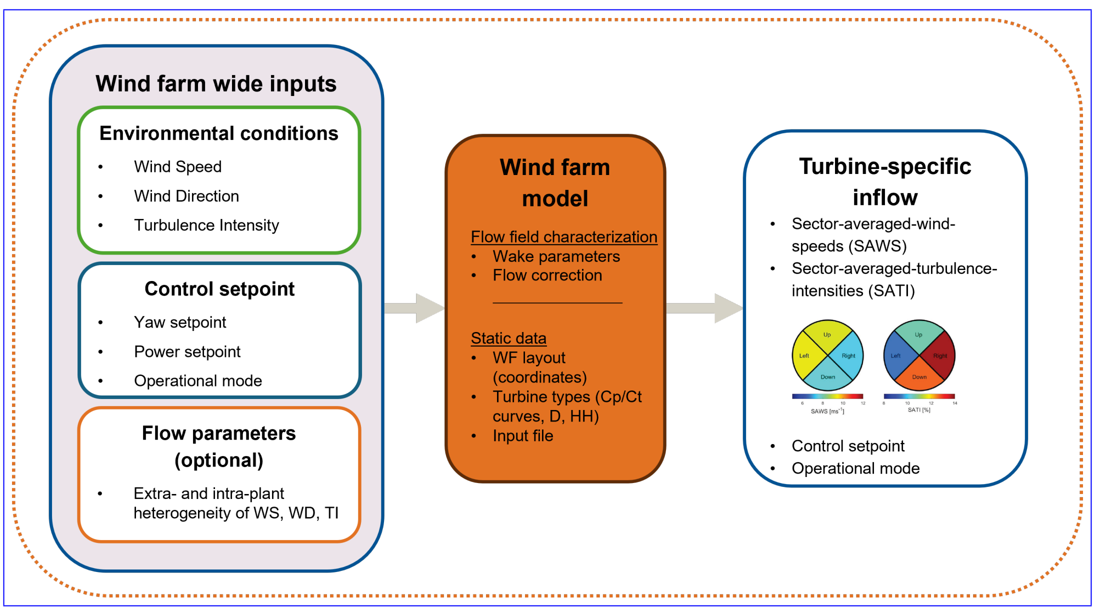


**Figure 2.** ~~Wind~~ Subset of Fig. 1 that defines the interfaces of the wind farm model in more detail. Static wind farm data is used to build up the model. Environmental conditions, control setpoints, and optional flow parameters drive the flowfield simulation and turbine-specific inflow extraction.

Before the ~~wind farm model~~ SFM 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)~~names (González-Salcedo et al., 2025; TWAIN, 2026), 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~~ SFM yields turbine-specific inflows that account for changing ~~velocity~~ speed, direction and turbulence distributions induced by upstream turbine wakes and terrain-induced flow ~~speed-ups~~effects. 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~~ SFMs 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).

For this study, the engineering flow-modeling tool `Floris` (NREL, 2025) is employed ~~for wind farm representation~~as the SFM. 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).

### 2.3 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~~designed to be adaptable, 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~~flows, whereas waked inflow profiles account for the influence of upstream turbines and their operating conditions.

#### 2.3.1 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~~ inflows to their specific aero-servo-elastic simulation requirements.

### 2.3.2 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). Accurate 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 (reflecting the R&D ~~and research institutes~~ perspective), wind turbine structural and control design (reflecting the turbine manufacturer perspective) and wind farm operation under varying environmental conditions and control optimization strategies (reflecting the 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 as the DFM 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-

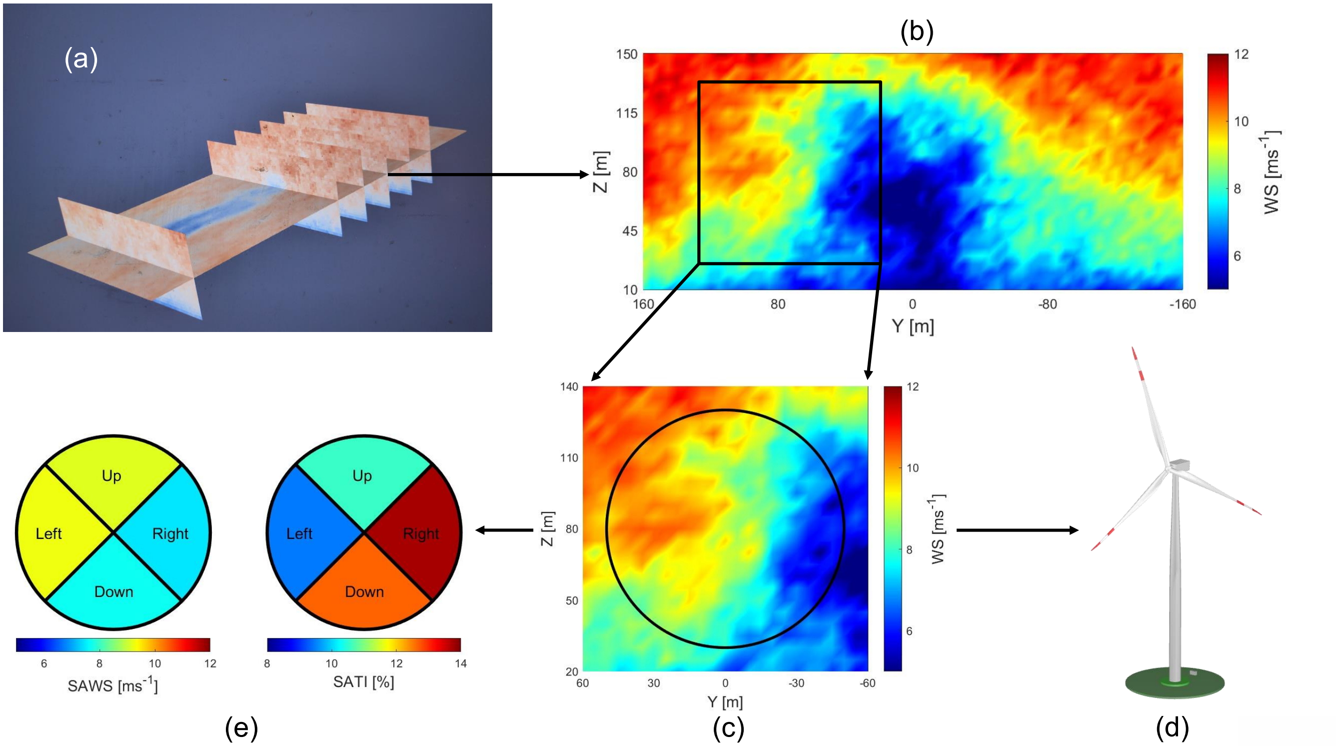


**Figure 3.** The approach to generate and apply to a wind turbine a library of wake slices. **(a)** simulations of a wake based on DWM for various ambient inflow conditions on a farm scale. **(b)** sampling of lateral slices at different streamwise distances. **(c)** extraction of rotor-scale turbulent 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.

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,~~ three underlying hypotheses, which are assessed through the validation and testing results presented in Sect. 3.2 and Sect. 3.3. First, turbine loads are assumed to depend exclusively on the local inflow experienced by the rotor. Consequently, if two turbines are subjected to the same local inflow conditions, their response can be represented through a common set of inflow descriptors, independent of the upstream flow history that generated those conditions. Second, 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~~ generate wakes of similar normalized velocity-deficit and meandering characteristics. ~~Second~~Third, 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 mapping from local inflow to fatigue loads, please refer to 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

### 2.3.3 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 wind fields spanning the entire 8-dimensional grid would require a significant number of farm-level dynamic simulations, which involve a ~~relatively~~ significantly higher computational effort. To address this, a smart sampling approach is developed to identify the relevant environmental conditions for running dynamic farm-level simulations. Compared with earlier work (Guilloré et al., 2024), where cases were manually selected to populate the load surrogate dataset, the proposed sampling approach provides a systematic and reproducible strategy for covering the relevant environmental conditions.

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~~ DFMs is computationally prohibitive. However, Ardillon et al. (2023) demonstrated that ~~steady-state engineering wake models~~ SFMs are able to reproduce the same key trends in wake characteristics~~as dynamic wake models~~. Therefore, the wake model introduced in Sect. 2.2, `Floris`~~(NREL, 2025),~~ is employed as a reliable proxy to generate the 8-dimensional distribution of sector-averaged quantities.

### 2.4 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 2.3 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. 2.5.

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.~~

All operating modes are simulated over 800 seconds using the synthetic inflow winds described in Sect. 2.3. The first 200 seconds are discarded as transients to allow the turbine to reach its nominal ~~point~~ operating condition and the wake inflow to develop fully, while the remaining 600 seconds are retained for evaluation. For each simulation case, ~~multiple~~ six 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~~

In conventional IEC fatigue load assessments, start-up and shut-down load cases are typically simulated only over the duration of the ~~maneuver itself, which is on the order of 20–40 seconds, depending on the controller design. In the~~ transient maneuver itself. The objective of the present work, however, ~~load responses for~~ is different. The surrogate model is intended to support load estimation from operational histories represented at a 10 min resolution, as commonly available from SCADA and control-system records. In such datasets, a 10 min interval may be identified as containing a start-up or shut-down event, while the exact timing of the transition and the corresponding fractions of time spent in power production and parked operation are generally unavailable. To ensure compatibility with this representation, all operating modes~~–~~, including start-up and shut-down~~ – are computed over a fixed 10 min time window, in order to ensure direct comparability with normal power-production conditions. To this end~~, are evaluated over a common 600 seconds interval.

For the transient operating modes, a standardized ~~start-up or shut-down~~ start-up or shut-down maneuver is embedded within the ~~600-second evaluation interval. In the shut-down~~ 600-second window together with prescribed periods of power-production and parked operation. Specifically, in the shut-down case, the turbine operates in ~~power-production~~ power-production mode for the first 150 seconds, after which the ~~shut-down~~ shut-down maneuver is initiated ~~,~~ and the turbine transitions ~~into idling mode~~to idling operation. Conversely, in the ~~start-up~~ start-up case, the turbine remains parked until the start-up maneuver is initiated, after which it enters ~~power-production~~ power-production mode and operates for 150 seconds. ~~This approach was adopted here as it reflects realistic operational conditions, in which individual~~ Because the maneuver can be placed arbitrarily within the 600-second interval, the resulting damage-equivalent loads inevitably depend on the chosen initiation time – for example, on how long the turbine idles before starting up. A more neutral assessment, while retaining the common 600-second duration, would require varying the maneuver initiation time and averaging the corresponding load outcomes. This was beyond the scope of the present work and is expected to have limited influence on long-term fatigue evaluations. The reported damage-equivalent loads therefore represent a full 10 min ~~intervals may include multiple operatingstates with varying duration of normal operation, while ensuring that load responses are consistently evaluatedover identical time windows across all operating modes.~~ interval containing a single transition event rather than the transient maneuver in isolation. This choice constitutes a modelling approximation and is not intended to reproduce the exact timing of individual field events. Instead, it provides a consistent mapping between operating-state histories and load estimates, allowing each 10 min interval to be assigned a mode-specific response. The common 600-second duration ensures compatibility with operational histories represented at a 10 min resolution.

~~Furthermore, for~~ The approach does not rely on a predefined lifetime number of start-up or shut-down events, as is common in certification-oriented load assessments. Instead, transitions arise directly from the operational history or the control strategy being evaluated, and the associated 10 min load contributions are accumulated afterward. This formulation is therefore not intended to replace IEC load-assessment procedures; rather, it serves as an operational-response model suited to applications where transition frequencies are strongly shaped by site conditions and control actions, such as wildlife-protection curtailment strategies (Pettas et al., 2026).

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~~ 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.

For the start-up operating mode, the database is restricted to below-rated wind speeds. This range covers the primary applications considered in the present work, including temporary environmental curtailment and other operational interruptions, for which restart events are typically expected under below-rated conditions. The extension of the start-up database to above-rated wind speeds was additionally constrained by the implementation of one turbine-controller model, for which the normal pitch-control logic interfered with the prescribed start-up maneuver. As a result, above-rated start-up events are not represented in the current database and should be regarded as a limitation of the present implementation rather than a general operational assumption.

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~~ Normal operation (NO)**: Wind turbine normal operation using baseline controller (also referred to as greedy operation in literature).

– **~~Yaw-steering~~ Yaw steering (YS)**: Static yaw-offset for wake steering flow control.

– **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~~ 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.

### 2.5 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.

#### 2.5.1 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. 2.4. 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~~ characterize the local inflow conditions ~~, represented using the SAWS and SATI (see Fig. 3 (e)).~~ of the corresponding aeroelastic simulation. These variables are the four SAWS and four SATI of the DFM-generated synthetic inflows described in Sect. 2.3. Alternatively, the surrogate models can be trained on the SAIQs extracted from the SFM for the same conditions used to generate the synthetic waked inflows, as discussed in Sect. 2.1.1. 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. 2.4. This combination of environmental and control inputs allows the surrogate models to account for both external operating conditions and turbine control actions.

#### 2.5.2 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. 2.3.1) and the waked inflow profiles (refer to Sect. 2.3.2). 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~~ outputs 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 )~~ 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

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

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).

#### 2.5.3 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 hyperparameter 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.

### 2.6 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 $j$, $N_{\mathrm{IEC}}$ 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~~ . Since damage evolution may differ substantially among turbines and load channels, the individual RUL estimates contain information that can be valuable for detailed condition assessment and optimization studies. In the present work, however, a single scalar indicator is required to compare different wind-farm configurations in Sect. 3.5. Therefore, a conservative representative value is defined as the minimum RUL across all turbines and channels, corresponding to the most critical turbine-load combination.

Furthermore, time series of 10-minute damage-equivalent-loads can be aggregated into a (single value) long-term DEL using

$$M_{\mathrm{eq,LT}} = \left(\frac{1}{N}\sum_{k=1}^{N} M_{\mathrm{eq},k}^{m}\right)^{1/m}, \tag{5}$$

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

To quantify the variability in turbine response within a wind farm, the deviations of individual turbines from the farm-wide mean can be expressed as

$$\Delta M_{\mathrm{eq},WT} = \frac{M_{\mathrm{eq,LT},WT} - \overline{M}_{\mathrm{eq,LT}}}{\overline{M}_{\mathrm{eq,LT}}}, \tag{6a}$$

$$\Delta P_{WT} = \frac{P_{WT} - \overline{P}}{\overline{P}}. \tag{6b}$$

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 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 a 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~~ within 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~~ DFM for capturing the dynamic ~~wake behavior~~ behavior of wakes.

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, which is $80\,\mathrm{m}$ above the ground. Wind speeds ranging from 3 to 26 $\mathrm{ms}^{-1}$ are considered with six different turbulent classes (A+, A, B, C, S10, S6) based on the IEC normal turbulence model guidelines ~~International Electrotechnical Commission (201~~ (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 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~~ DFM 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

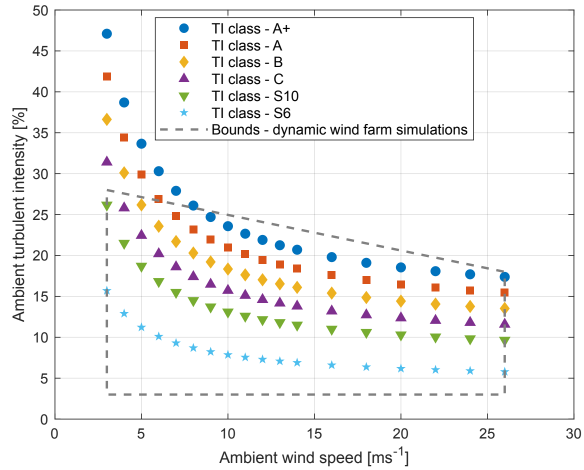


**Figure 4.** Clean inflow ambient cases based on IEC normal turbulence model and bounds for case selection for dynamic farm simulations.

fields generated using `TurbSim` (Jonkman, 2009). Inflow slices are extracted at multiple downstream and lateral positions, as described in Sect. 2.3.2. ~~To~~ The smart sampling approach outlined in Sect. 2.3.3 is used 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. 2.3.3 is applied.~~.

For each combination of ambient wind speed and turbulence intensity, SAIQs are extracted at streamwise locations of $[-0.5, 4, 5, 6, 7, \text{and } 8\,D]$ 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~~.~~, as `FAST.Farm` (Jonkman et al., 2018) does not model an upstream induction region. Rotor induction is instead resolved by the aerodynamic models of the aeroelastic solvers.

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 $\mathrm{ms}^{-1}$) and TI (between 3 and 28 %), as shown by the grey dashed line in Fig. 4. Lower turbulence intensity 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

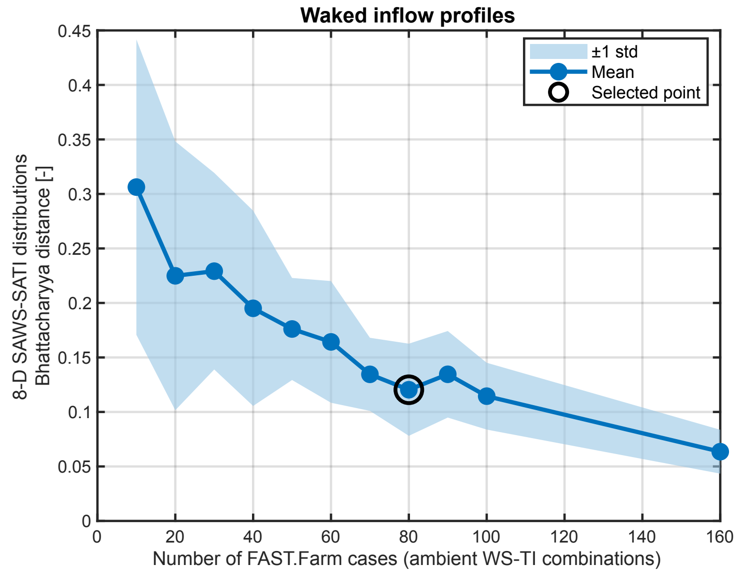


**Figure 5.** Convergence plot of the farm-level cases selection using the Bhattacharyya distance to generate waked inflow profiles.

the wide range of local inflow conditions. ~~Eventually, the~~ The turbine response surrogates performance for inflow cases not present in the training dataset are shown later in Sect. 3.2.

For waked inflow profiles, step (a) from Sect. 2.3.2 ~~(Fig. 3)~~ is performed using the DWM implementation in `FAST.Farm` (Jonkman et al., 2018). The NREL-1.79-100 turbine (rotor diameter ~~100 $m$~~ $100\,\mathrm{m}$) is used to generate the wake (NREL, 2025). The turbine is operated at ~~its full power curve~~ full power (no curtailment) and without any yaw misalignment. Consequently, the resulting library of waked ~~inflow~~ inflows 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~~ a substantial increase in the number of dynamic wake simulations required ~~for~~ to cover 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 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 $5\,\mathrm{m}$ spatial resolution) and are sampled at $10\,\mathrm{Hz}$. 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)~~ Guilloré et al. (2026a).

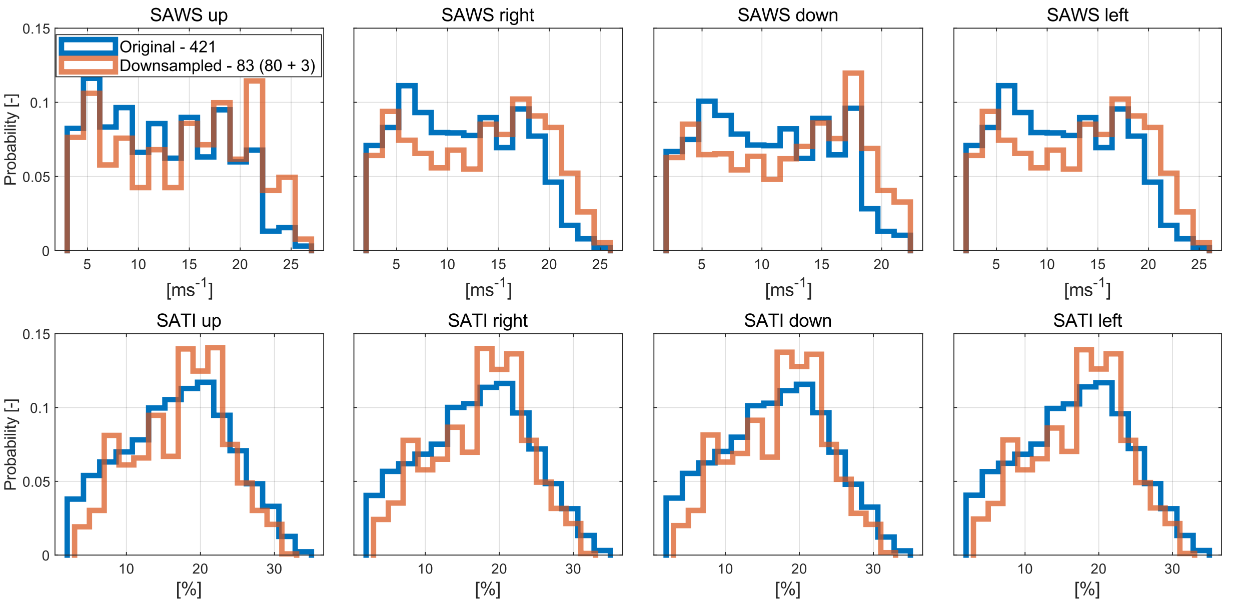


**Figure 6.** Comparison of probability distributions of sector-averaged quantities after downsampling ambient cases for waked inflow profiles.

### 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 $1.79\,\mathrm{MW}$, rotor diameter of $100\,\mathrm{m}$ and hub height of $80\,\mathrm{m}$. 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~~ commercial turbines from two different manufacturers. These turbines are modern onshore ~~turbines~~machines, with rated power around $2\,\mathrm{MW}$ 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~~OpenFAST (NREL, 2025) and the simulation code of Turbine-B is implemented in ~~HAWC2~~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~~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~~ 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~~ 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~~codes, 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.

**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 |

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)~~Guilloré et al. (2026a).

### 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~~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~~ extracted from the DFM waked slices (characterized by their corresponding 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.

As discussed in Sect. 2.1.1, the definition of the surrogate model inputs is potentially an important aspect of the framework, particularly when different flow models are used during training and deployment. In the current implementation, the surrogate model is trained on inflow conditions derived from high-fidelity DFM simulations, whereas during deployment it receives inflow conditions predicted by the lower-fidelity SFM. This could in principle introduce additional discrepancies. To assess the impact of this training-deployment mismatch, surrogate model predictions from the full framework were compared with the outputs of the aeroelastic simulations (App. C). Their close agreement indicates that retraining the surrogate models using SFM-derived inflow parameters was not necessary in this case, but cannot be excluded when substituting the SFM.

**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 |

#### 3.1.4 Computational cost

A primary motivation for developing the proposed framework is to substantially reduce the computational cost of wind-farm load assessment while keeping the workflow accessible to a broad range of stakeholders. To illustrate this, we evaluate the computational effort associated with three representative use cases: generating the synthetic inflow library (Sect. 2.3), performing the aeroelastic simulations (Sect. 2.4), and applying the toolchain itself, which is by far the most computationally

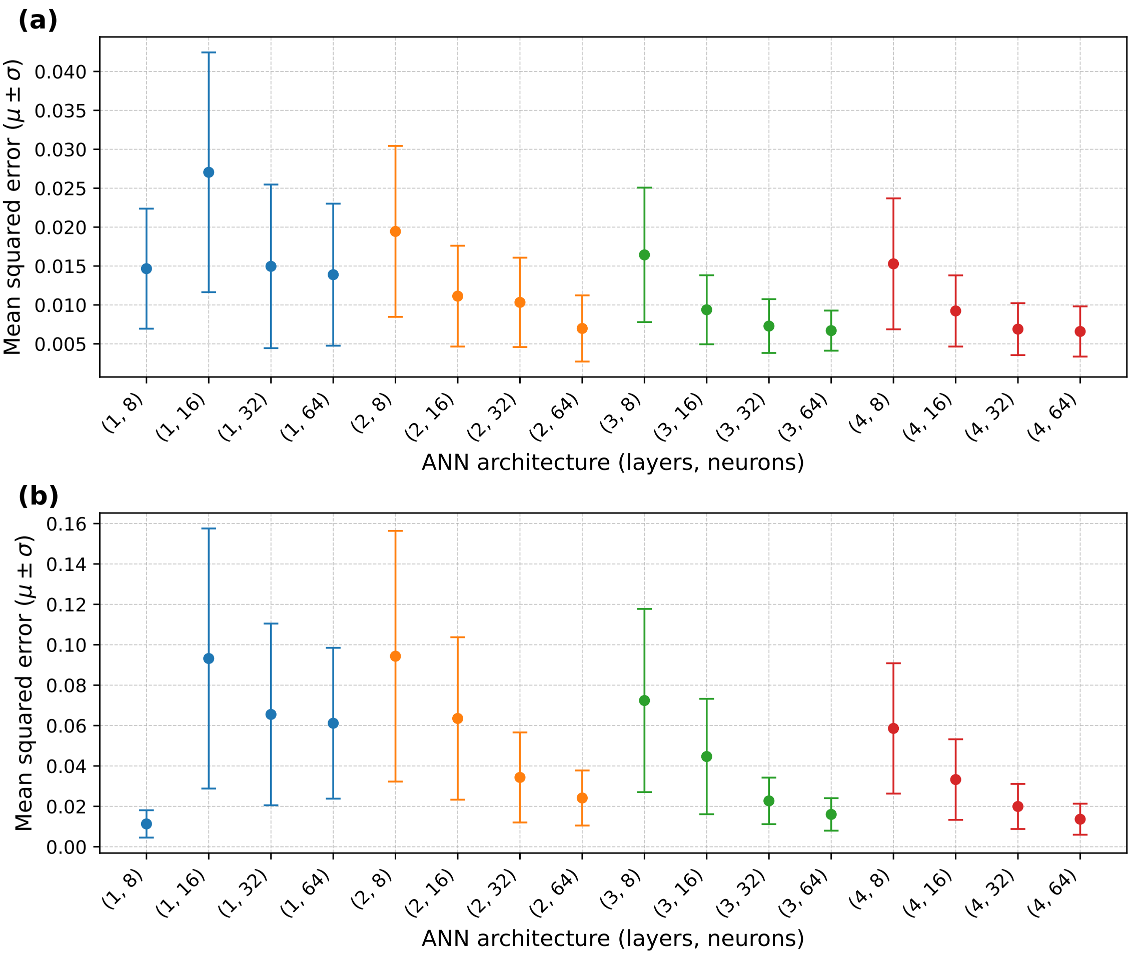


**Figure 7.** Grid-search results of the ANN architecture for ~~the~~ Turbine-A for the two operating modes: ~~(a)~~ **(a)** power production, and ~~(b)~~ **(b)** start-up. Shown are the error bars for the scaled mean-squared errors averaged across all target channels for different hyperparameter combinations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

efficient component. All execution times reported in this section were obtained using identical hardware (Intel i7-4790 CPU, 16 GB RAM) to ensure comparability.

It must be emphasized, however, that raw CPU timings inevitably lose relevance as hardware capabilities advance due to faster processors, wider vectorization, and increasingly widespread access to large-scale parallelization. Many of the computations considered here are highly repetitive and can, in principle, be parallelized almost perfectly when suitable computing infrastructure is available. Under such conditions, the wall-clock times reported below could drastically be reduced. The values presented here should therefore be interpreted as indicative orders of magnitude that contextualize the computational burden of each workflow component. They are not intended to prescribe performance expectations. While DWM-based aeroelastic simulations are feasible for selected cases and limited intervals of operation, their application to continuous long-term wind-farm assessment over periods ranging from months to decades remains computationally prohibitive. Instead, the reported timings primarily serve to illustrate the computational burden associated with such studies, thereby motivating the development of the proposed framework, which enables long-term assessment within practical computation times.

1. **Synthetic inflow generation:** If the new turbine differs substantially in size from the reference turbine, the synthetic waked inflow fields must be regenerated using a turbine model of comparable scale. This is achieved by performing single-turbine DWM simulations (implemented here in `FAST.Farm`) and extracting the local inflow at multiple downstream and lateral locations. Although this approach is computationally efficient in terms of the number of simulated turbines, it remains computationally demanding because the full time-varying three-dimensional wind field must be stored throughout the simulation and subsequently post-processed into waked turbulent inflow fields compatible with standalone turbine simulations. Consequently, the workflow is highly data-intensive. The simulation and post-processing of a single ambient condition case required approximately $9.5 \pm 0.5$ hours, with six turbulent seeds processed in parallel. Due to this substantial computational cost, the LHS-based methodology was employed to identify a reduced set of representative ambient conditions for the generation of the waked inflow library. In total, 83 ambient conditions were selected, resulting in an estimated computational cost of approximately 33 days for the complete waked inflow library generation.

2. **Aeroelastic simulations:** When a new turbine type is introduced, single-turbine aeroelastic simulations must be performed to retrain the surrogate models using the existing library of synthetic inflow fields. A key computational advantage of the proposed framework is that the library of waked inflow slices eliminates the need to regenerate DWM simulations whenever a new turbine type is considered. The inflow database can be reused provided that the new turbine is of comparable size to the turbine used to generate the synthetic inflow fields. Producing a robust training dataset for the surrogate models requires approximately 10,000 operating/inflow cases. With six turbulent seeds considered for each case, the database generation involves approximately 60,000 individual aeroelastic simulations. Using OpenFAST, the six seeds associated with a given case were processed in parallel, resulting in a wall-clock time of approximately 3 minutes per operating/inflow case. For a new turbine type, this corresponds to a total computation time of approximately 21 days. This effort can be increased by adding more inflow or control cases. It can also be substantially reduced by allocating additional computational resources and parallelization.

3. **Application of the toolchain:** Once the wind farm flow model and surrogate models have been prepared, only these computationally efficient models need to be executed. For example, a one-month simulation of the six-turbine wind farm presented in Sect. 3.5 required $17 \pm 2$ s. This corresponds to 4,320 calls to the flow model and surrogates, i.e. one for each 10-minute time step in one month. Of the total runtime, 85.4% was spent simulating the flow field with FLORIS, 5.1% extracting the flow features used as inputs to the surrogate models, and 3.6% evaluating the surrogate models for ten output channels. For comparison, a single 10-minute dynamic simulation of the same six-turbine wind farm using the DWM implementation in `FAST.Farm` under one ambient operating condition required approximately $7200 \pm 300$ s (2 h $\pm 5$ min) on a computer with comparable specifications, with six turbulent seeds simulated in parallel to ensure statistical convergence. Extrapolating this computational cost to a one-month simulation yields an estimated runtime of $31,104,000 \pm 1,296,000$ s ($360 \pm 15$ days). Relative to the proposed framework, this corresponds to a computational speed-up of approximately $1.8 \times 10^6$, or equivalently a reduction in runtime from about one year to 17 s.

### 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.

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

$$\mathrm{NRMSE} = \frac{\sqrt{\frac{1}{N}\sum_{l=1}^{N}\left(y_l - \hat{y}_l\right)^2}}{y_{\max} - y_{\min}}, \tag{7}$$

where $y_l$ ~~denotes the true (measured) value, and~~ and $\hat{y}_l$ ~~denotes the corresponding predicted value for the point~~ denote the true and predicted values associated with sample $l$, $N$ is the total number of samples in the dataset, 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~~ where values are 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~~ prediction error associated with sample $l$ is defined as

$$\epsilon_M = \frac{y_l - \hat{y}_l}{y_{\max} - y_{\min}}\cancel{,}. \tag{8}$$

~~is also included.~~ The probability density functions of these errors (expressed in %), quantified using the mean ($\mu$) and standard deviation ($\sigma$) 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~~ square error (RMSE) values are normalized by the range of the corresponding channel in the training dataset, following Eq. ~~7~~(7), 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 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 $\sigma$ 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~~partial and full waking) as well as control ~~strategies~~ modes (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.

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 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.

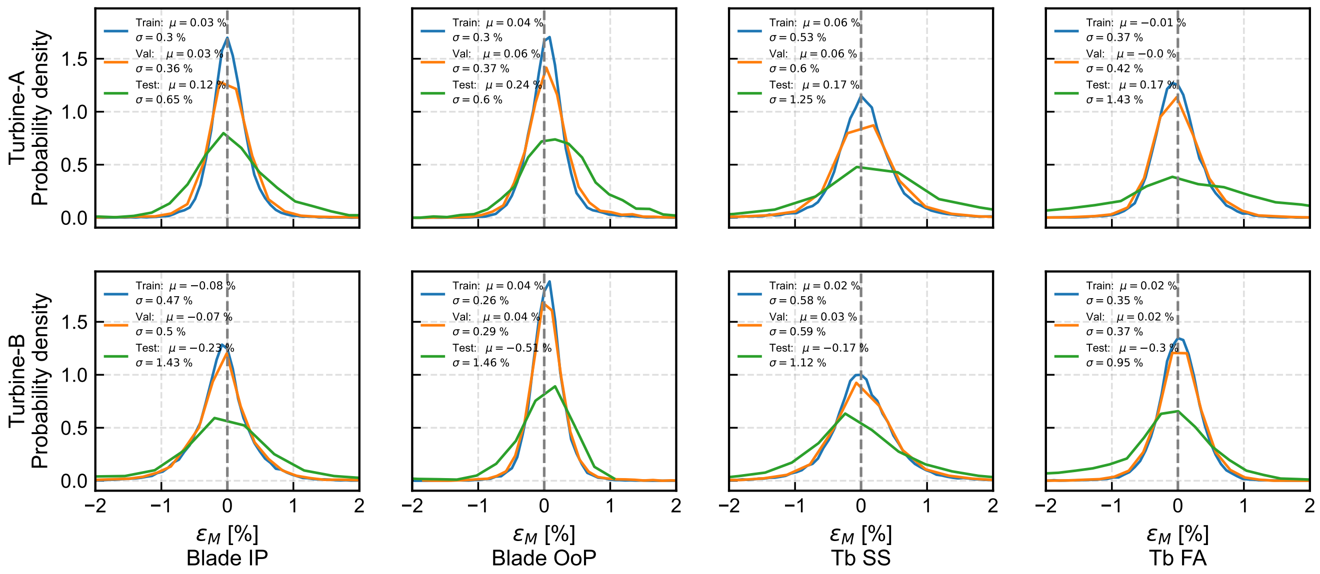


**Figure 8.** Probability density function of the error distributions (expressed in %) between the surrogate prediction and the target data from the aeroelastic simulations for the PP mode, normalized to the respective target value ranges, for training, validation and testing datasets. The rows of the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

**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 % |

**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 % |

### 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 DFM simulations. The DWM implementation in `FAST.Farm` (Jonkman et al., 2018) is employed as the DFM, as described in Sect. 3.1.1. The farm-level simulations comprised six ~~instances of the NREL-1.79-100 turbine~~NREL-1.79-100 turbines (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 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~~ SAWSs and SATIs 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. 2.3.2, are evaluated against this dataset, which captures realistic wind farm operating conditions, including different levels of wake impingement and multiple wake interactions.

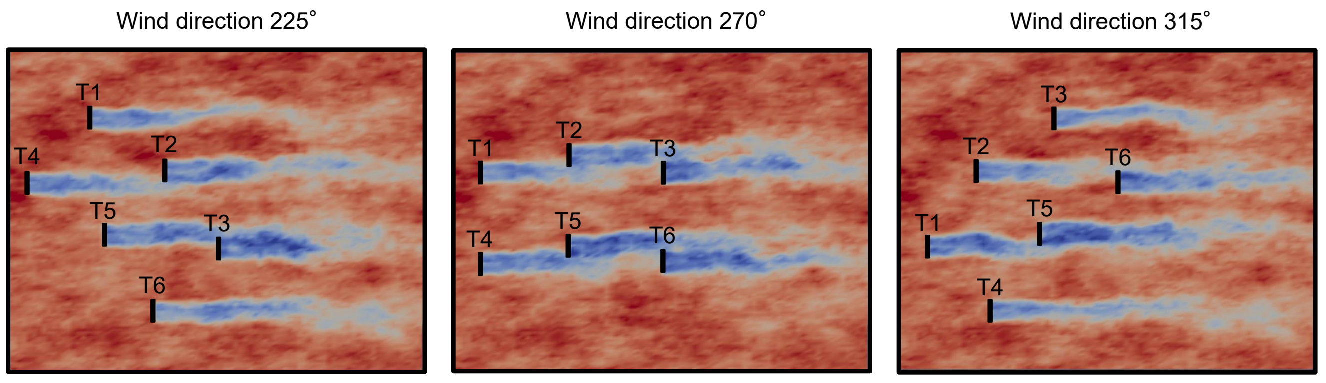


**Figure 9.** Wind farm layout (with three exemplary wind directions) used for the farm-level testing of the approach. The background ~~illustrate~~ illustrates instantaneous view of the flow field at hub height from the ~~dynamic wind farm~~ DFM simulation based on DWM.

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 $\mathrm{ms}^{-1}$, with the wind direction serving as the primary varying parameter that induces different levels of wake impingement across the farm. For relative ~~comparison~~ comparisons across 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 conditions 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~~SAWSs, SATIs, 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 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, 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.

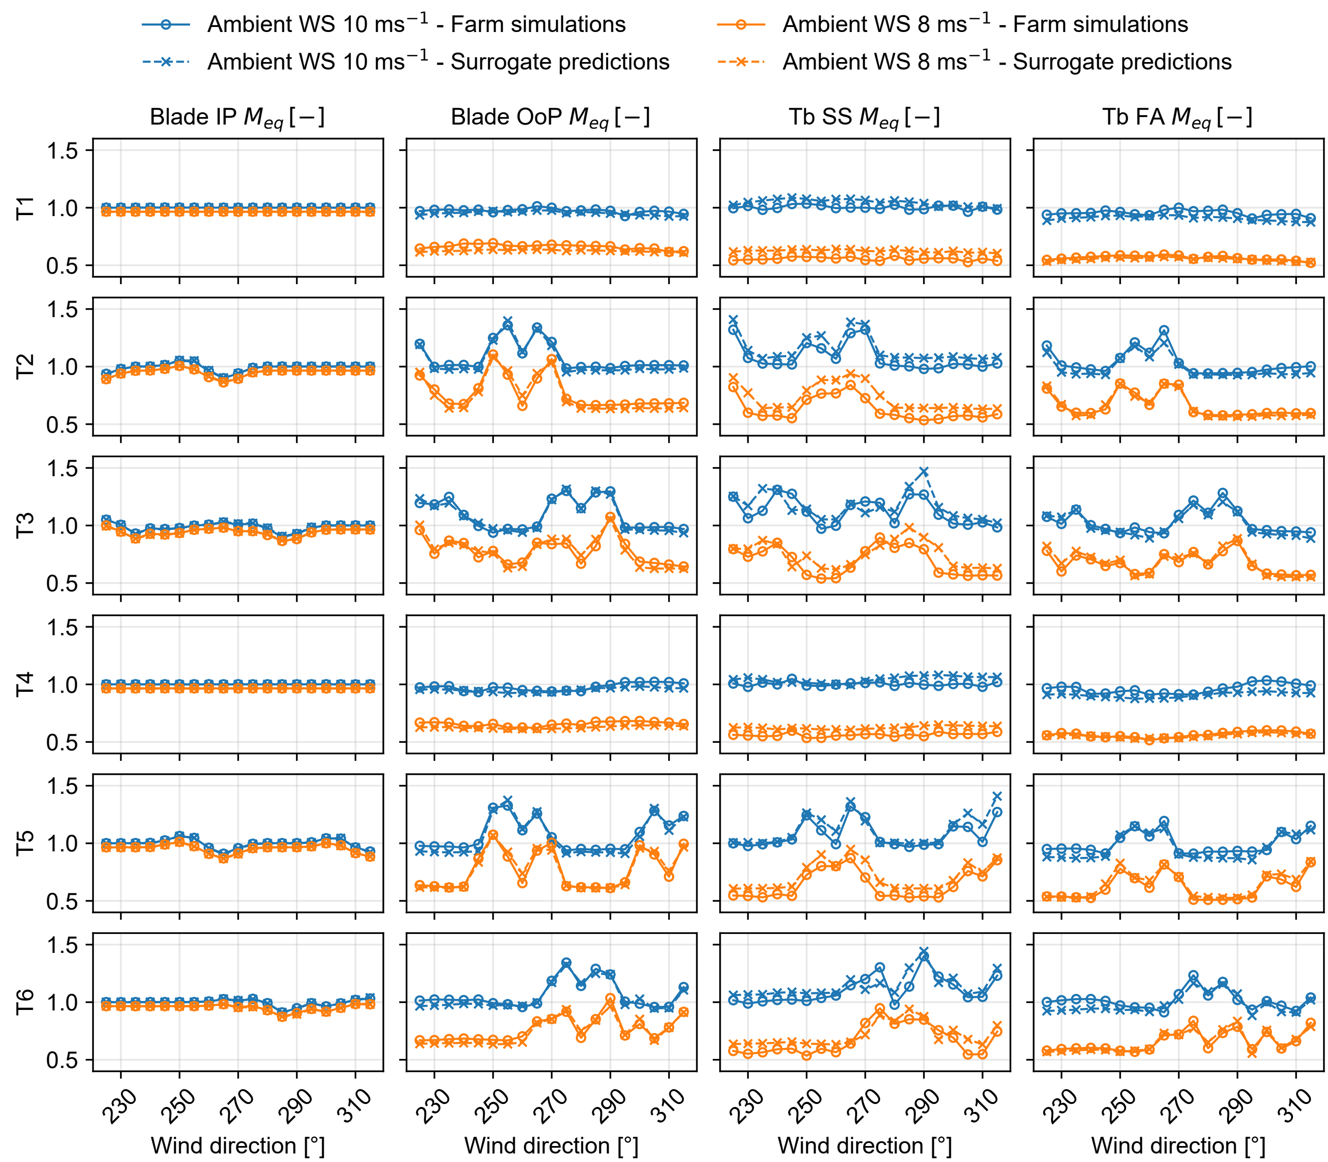


**Figure 10.** Surrogate predictions (dashed lines and crosses) ~~versus~~ and reference values (solid lines and circles) obtained from ~~DWM~~ DFM simulations for the wind-farm-level test cases (six ~~instances of the~~ NREL-1.79-100 turbines arranged as illustrated in Fig. 9), shown as ~~a function~~ functions 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.

### 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 unphysical 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 Sect. 2.3. 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.~~

### 3.4 Comparison of turbine response

A central contribution of this work lies in demonstrating that, by training surrogate models with an identical architecture and using a common dataset of aeroelastic simulations driven by the same wind inflow conditions, it ~~becomes feasible~~ is possible 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 stakeholders, simulated with different aeroelastic solvers, to be evaluated on a comparable basis under both clean and waked inflow conditions.

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. 2.2, 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~~ As outlined in Sect. 2.1.1, the surrogate models can be optionally trained on the inflow generated by the wind-farm model to compensate for systematic statistical discrepancies between the wind farm flow representation and the inflow used in the aeroelastic simulations –for example, a consistent underestimation of wake-added turbulence intensity. In the present application, this correction step was unnecessary because the chosen SFM, `Floris`, provided sufficiently accurate agreement. Appendix C presents a comparison of the performance of the framework under free-stream and various partially waked conditions based on the aeroelastic simulation outputs, demonstrating a close agreement across multiple inflow conditions and load channels.

The plots for Figs. 11, 12, and 14 ~~present results for an isolated turbine operating in freestream~~were produced using a single-turbine wind-farm configuration operating under free-stream conditions, whereas Fig. 13 ~~considers~~ is based on a two-turbine configuration of ~~identical~~ the same turbine type with ~~varying~~ streamwise and lateral spacings expressed in rotor diameters.

Figure 11 illustrates the ~~trends of five surrogate~~ surrogate model predictions for five output channels as ~~a function~~ functions of the ambient wind speed, with different ~~colors indicating the respective~~ colours representing the ambient turbulence intensity~~levels~~. The first row presents the power normalized ~~to~~ by the rated power, ~~as it provides the most intuitive basis for comparison. The subsequent four rows show the~~ followed by the Blade IP, Blade OoP, Tb SS, and Tb FA DELs. Beyond illustrating the expected influence of wind speed and turbulence intensity, the purpose of this comparison is to demonstrate that the proposed framework can be applied consistently across different turbine designs using the same inflow descriptors and surrogate model architecture. This enables a direct comparison of the load response of different turbines under identical environmental conditions.

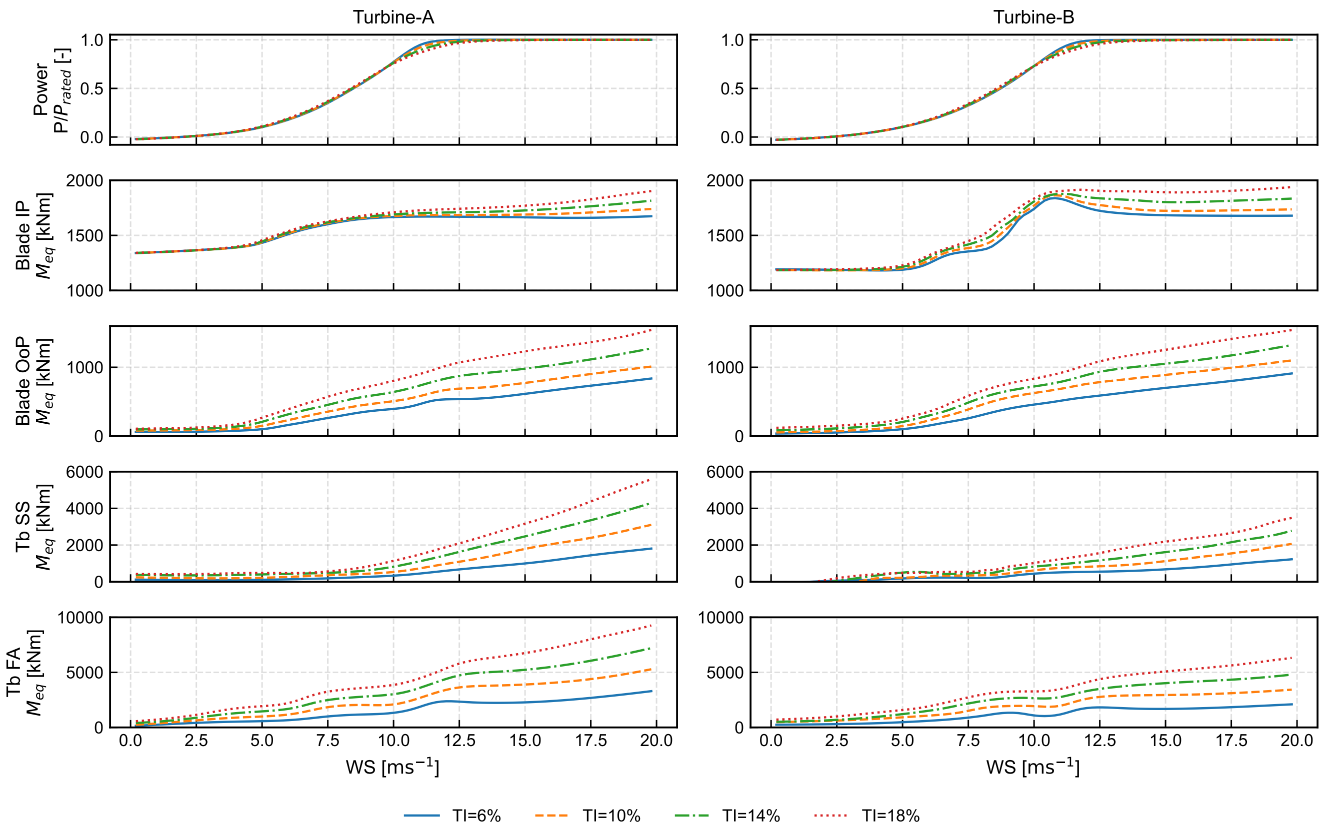


**Figure 11.** ~~Surrogate-model~~ Surrogate model outputs as functions of wind speed for varying turbulence intensities. The left column shows the response of ~~the~~ Turbine-A, the right column that of the Turbine-B.

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 $\mathrm{ms}^{-1}$. 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~~ increased pitch actuation required for power regulation at ~~elevated~~ higher 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~~ changes in the aerodynamic forces 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 both turbine models exhibit similar qualitative trends with respect to wind speed and turbulence intensity, noticeable quantitative differences are observed in their load responses. In particular, while the blade loads of the ~~Turbine-A and Turbine-B are similar in magnitude, the Turbine-A exhibits higher tower-base loads . This occurs~~ turbines remain of comparable magnitude, Turbine-A exhibits consistently higher tower-base loads despite its generally lower thrust force ~~across the investigated wind-speed range and its~~ and smaller hub height~~, indicating that additional factors – such as differences in~~ . These differences highlight the influence of turbine-specific structural and control characteristics, including mass distribution, ~~control strategies, and tower eigenfrequencies – play a significant role~~controller behaviour, and structural dynamics. More importantly, they illustrate that the proposed framework can represent these turbine-specific responses within a common computational workflow, enabling consistent comparisons across different turbine designs without modifying the overall methodology.

Figure 12 ~~shows the selected target channels as a function of wind speed, but now differentiated by operation mode~~demonstrates that the same surrogate framework can uniformly represent multiple operating modes (power production, start-up, shut-down, parked) for different turbine models, allowing direct comparison of their respective fatigue responses. 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,~~ The results show that, for the wind conditions considered here, parked conditions consistently result in the lowest loading, whereas 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 $\mathrm{ms}^{-1}$. 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, over the same duration.

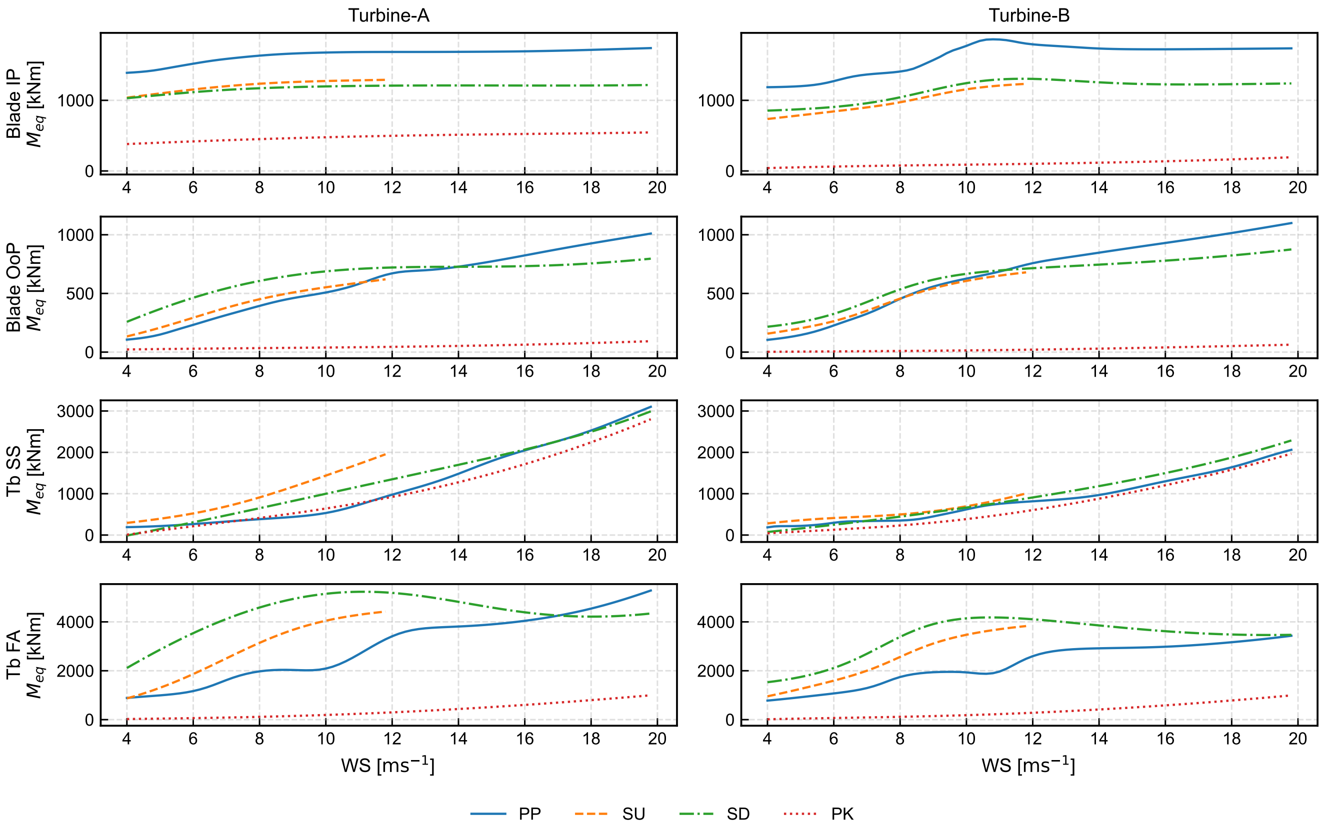


**Figure 12.** Surrogate-model outputs as a function of wind speed for various operating modes at a turbulence intensity of 10 %. SU results are presented only within the simulated training range, i.e. below rated wind speed.

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. 2.3, 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 $\Delta$Y on the $x$ ~~-axis~~ axis denotes the lateral offset between two turbines. A value of $\Delta$Y=0 corresponds to perfect alignment, whereas negative offsets indicate that the downstream turbine is shifted to the right relative to the inflow direction looking downstream, causing the wake to impinge on the left side of its rotor. The three curves in each subplot represent 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 ($\Delta$Y<0) and right-impinging ($\Delta$Y>0) wakes is governed by the direction of rotation of the rotor. When the

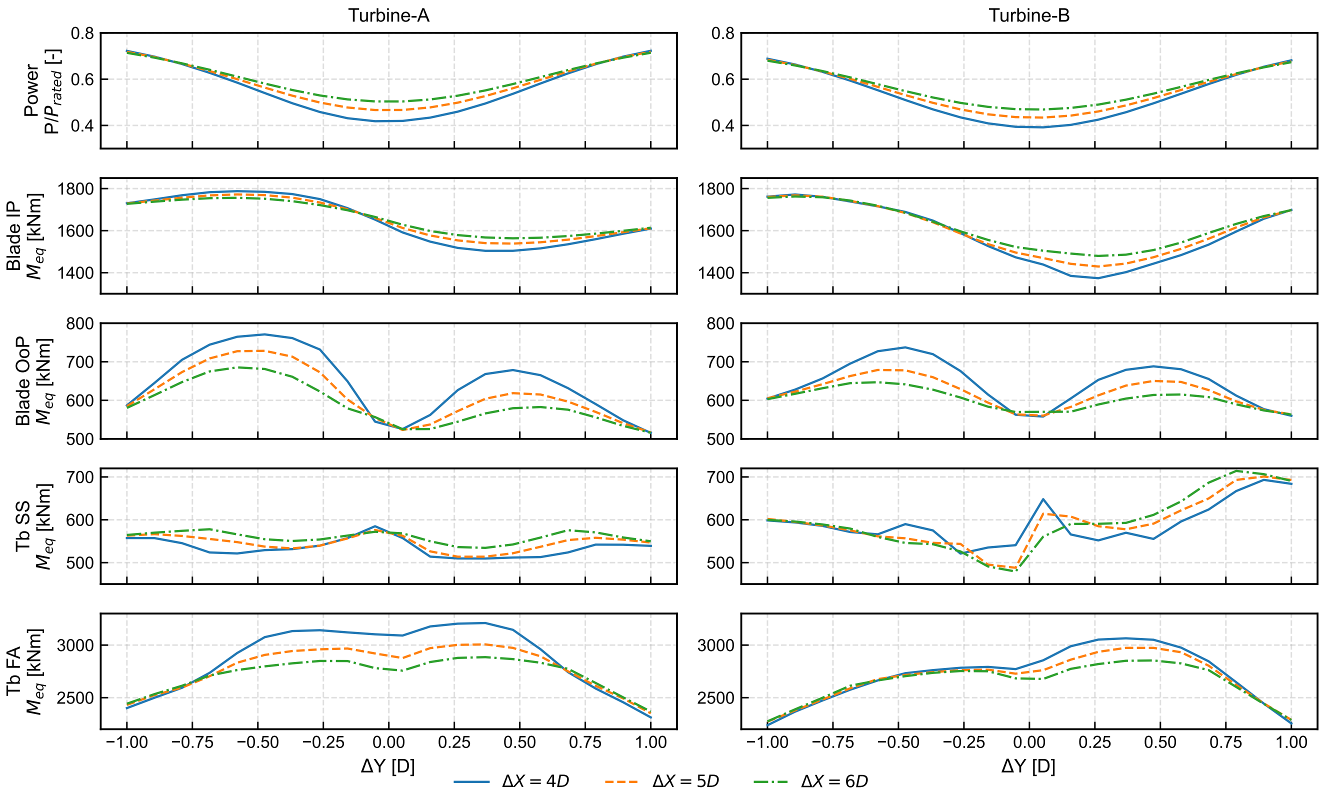


**Figure 13.** Surrogate-model predictions for partial wake overlaps at a wind speed of 10 $\mathrm{ms^{-1}}$ and a turbulence intensity of 10 %.

wake impinges on the left side of the rotor (~~lateral position Y~~$\Delta Y$<0~~relative to the wake center~~), the upward-moving blade encounters ~~reduced aerodynamic momentum~~ a reduced flow speed while working against gravity, resulting in increased cyclic load amplitudes. Conversely, a wake impinging on the right side of the rotor (~~Y~~$\Delta Y$>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 a less consistent behavior between the two turbine types. Whereas ~~the~~ Turbine-A exhibits the highest loading under full-wake conditions, ~~the~~ 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.

~~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 %.~~

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~~ observed. For 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 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.

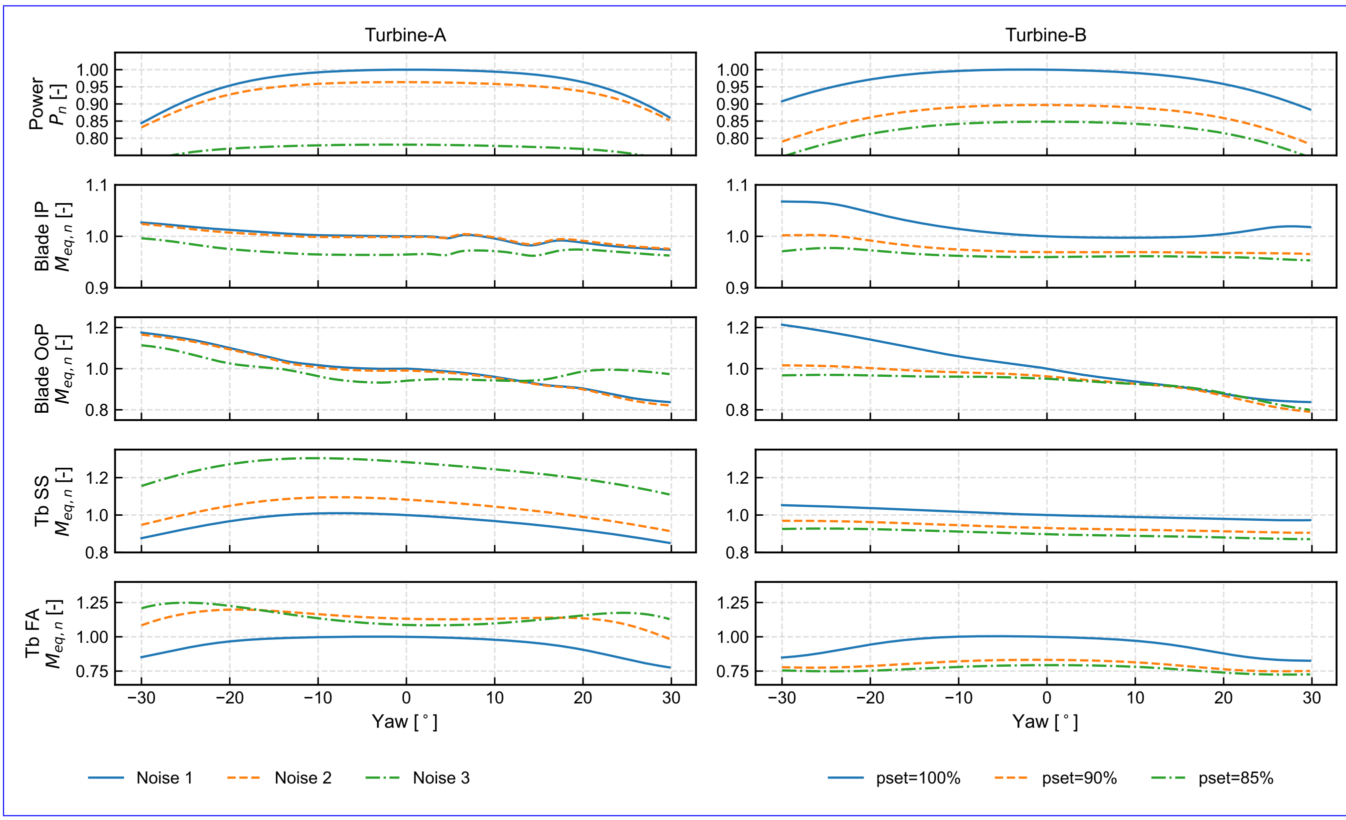


**Figure 14.** Surrogate-model outputs for various yaw angles and power/noise setpoints, normalized to the 0 degree-yaw output of the normal operating mode (Noise 1, 100 %) at a wind speed of 12 ms$^{-1}$ and a turbulence intensity of 10 %.

### 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~~ an example 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.

The objective of the following case study is not to assess the general impact of wake steering on wind-farm fatigue loads, but to demonstrate the applicability of the proposed modular framework to a representative wind-farm scenario.

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

- 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.

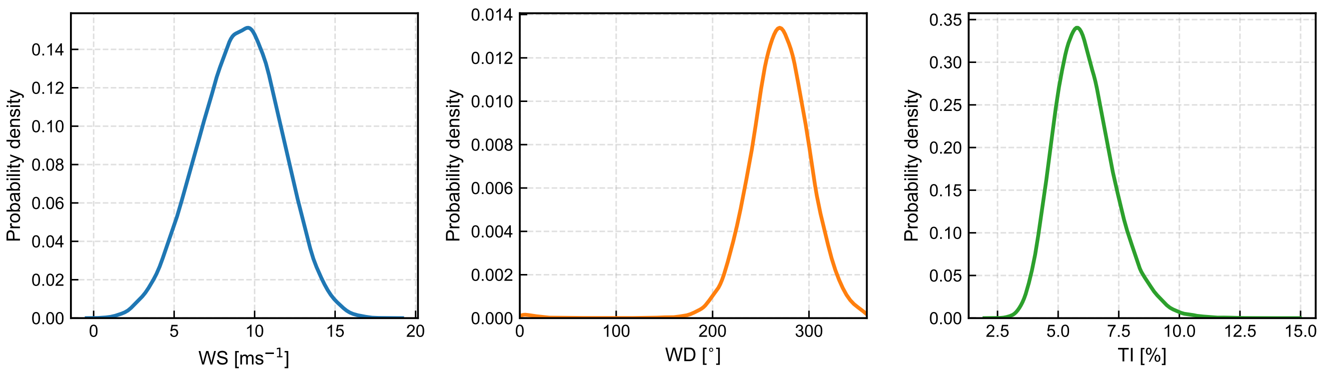


**Figure 15.** Distributions of wind speed (left), wind direction (center), and turbulence intensity (right).

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 $\mathrm{ms}^{-1}$ between 4 and 10 $\mathrm{ms}^{-1}$. The resulting look-up-table is depicted in the Appendix D in Fig. D1 for a wind speed of 8 $\mathrm{ms}^{-1}$. 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.

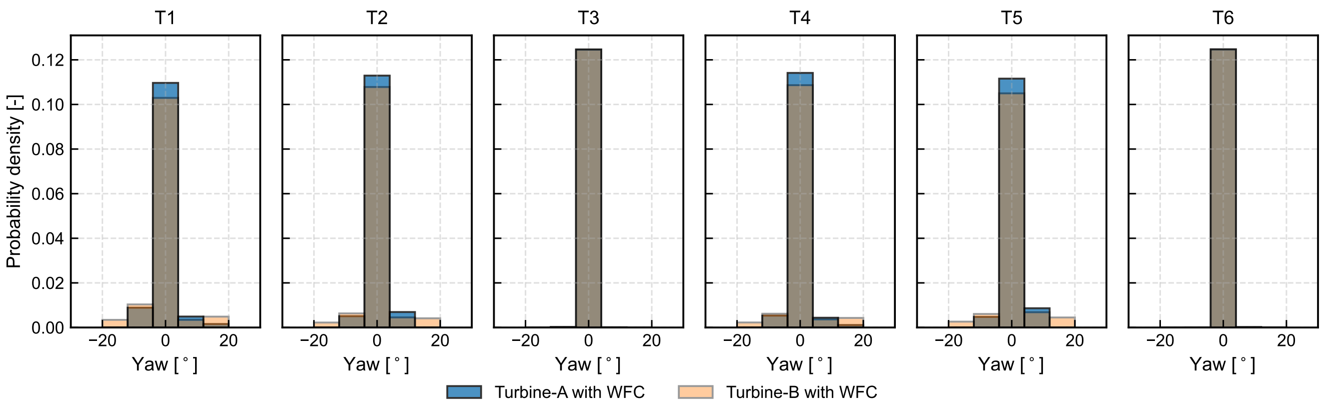


**Figure 16.** Distribution of yaw setpoints for all six turbines over the one-year analysis period for each turbine type.

By supplying the ~~defined~~ environmental conditions and yaw ~~setpoints to the~~ set-points to the proposed framework, the ~~resulting~~ response of each turbine in the wind farm can be efficiently evaluated. Figure 17 ~~depicts scatter plots of the corresponding DELs~~ illustrates the resulting DEL predictions for turbines 1 and 3 ~~as a function of the rotor-averaged wind speed~~over the first simulation month. The objective of this figure is to demonstrate the application of the complete framework to a wind-farm simulation, highlighting its ability to capture turbine-specific load responses under continuously varying operating conditions. 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~~ colour 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 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).

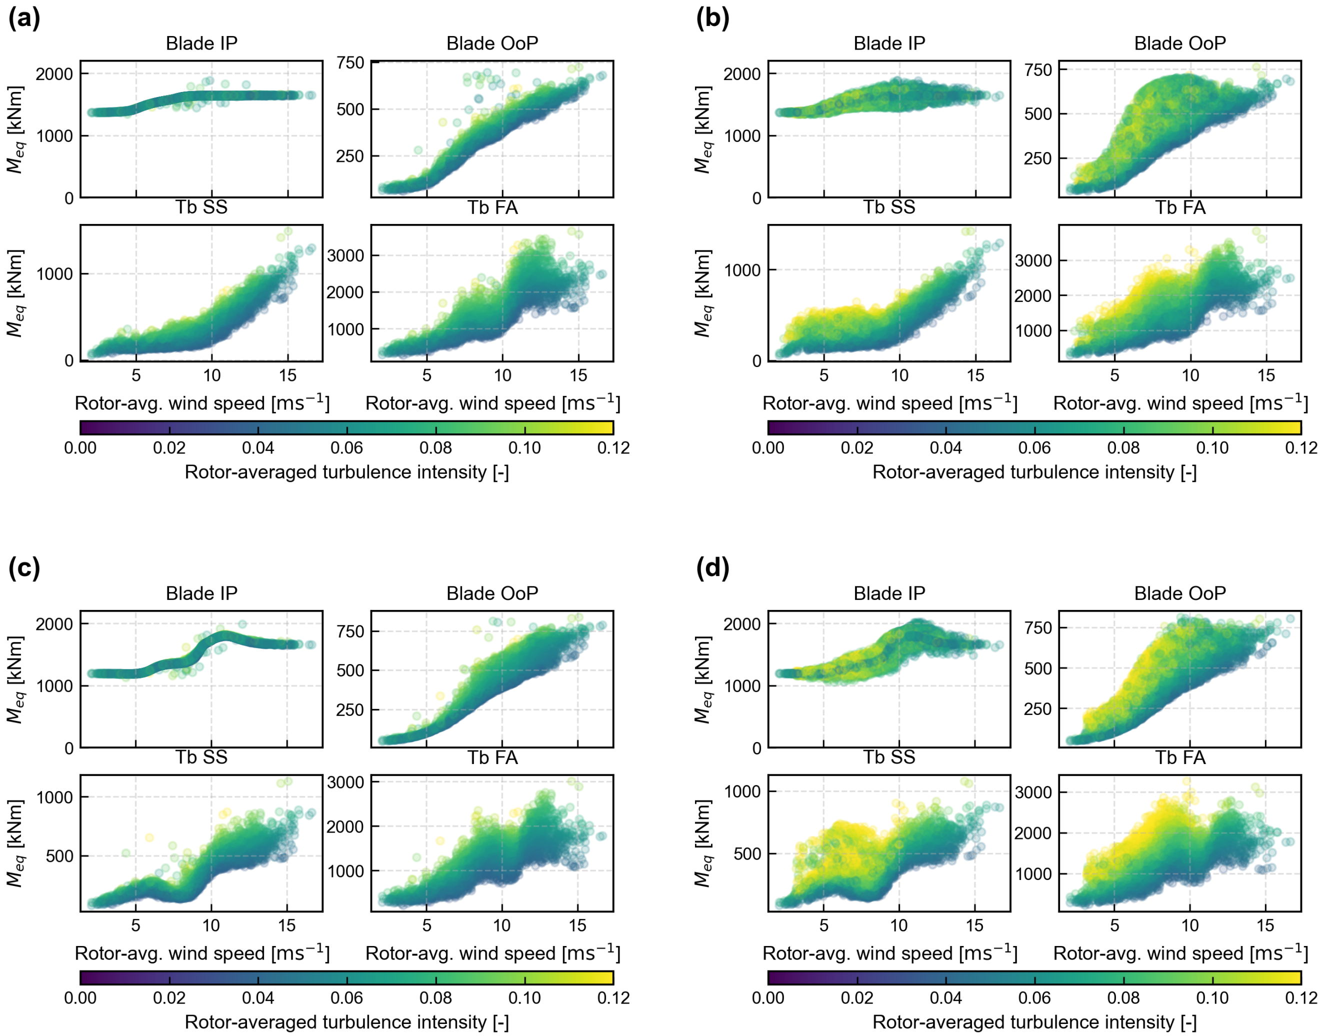


**Figure 17.** Scatter plots of the DELs for Turbines 1 and 3 (see wind-farm layout in Fig. 9) over the first simulation month without WFC. Panels show ~~(a)~~ **(a)** T1, Turbine-A; ~~(b)~~ **(b)** T3, Turbine-A; ~~(c)~~ **(c)** T1, Turbine-B; and ~~(d)~~ **(d)** T3, Turbine-B.

~~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~~

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

~~where $M_{\text{eq},k}$ denotes the damage equivalent load in time step $k$ of the one-year time series of length $N$, $M_{\text{eq,LT}}$ is the resulting long-term damage equivalent moment, and $m$ is the Wöhler exponent~~Overall, Fig. 17 demonstrates that the proposed framework is able to efficiently reproduce the variability of fatigue loads across turbines experiencing different inflow conditions within the same wind farm. Although the governing trends with wind speed and turbulence intensity remain consistent across both turbine models, the predicted DEL distributions differ due to turbine-specific aerodynamic, structural, and control characteristics. This illustrates how the proposed modular framework can be readily applied to different turbine designs while preserving their individual load responses in wind-farm-scale simulations.

Figure 18 shows the relative deviations in long-term damage equivalent loads $\Delta M_{\text{eq},WT}$ and power $\Delta P_{WT}$ 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~~

$$\Delta M_{\text{eq},WT} = \frac{M_{\text{eq,LT},WT} - \overline{M_{\text{eq,LT}}}}{\overline{M_{\text{eq,LT}}}},$$

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

The turbine-wise deviations ~~are presented~~ as defined in Eqs. (6) are evaluated for both turbine types and for the greedy ~~as well as the yaw-steering~~ and 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 side–side loading. These load channels also exhibited the strongest turbine-specific variations in ~~Sec. 3.2~~ Sect. 3.4 (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.

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 Sect. 2.6. Blade IP loading proved to be the most restrictive, and the corresponding farm-level RUL values are reported in Table 7. Overall, wake steering appears to yield a modest yet beneficial improvement in the wind farm RUL for the Blade IP loading in this specific case study. The example is intended primarily to demonstrate the broader capabilities of the framework rather than to draw firm conclusions about the suitability of wake steering for lifetime-extension strategies.

The predicted DELs are highly case-specific and turbine-specific, making general statements on the impact of wake steering on fatigue difficult. The results should therefore be interpreted as an illustration of the framework capabilities rather than as

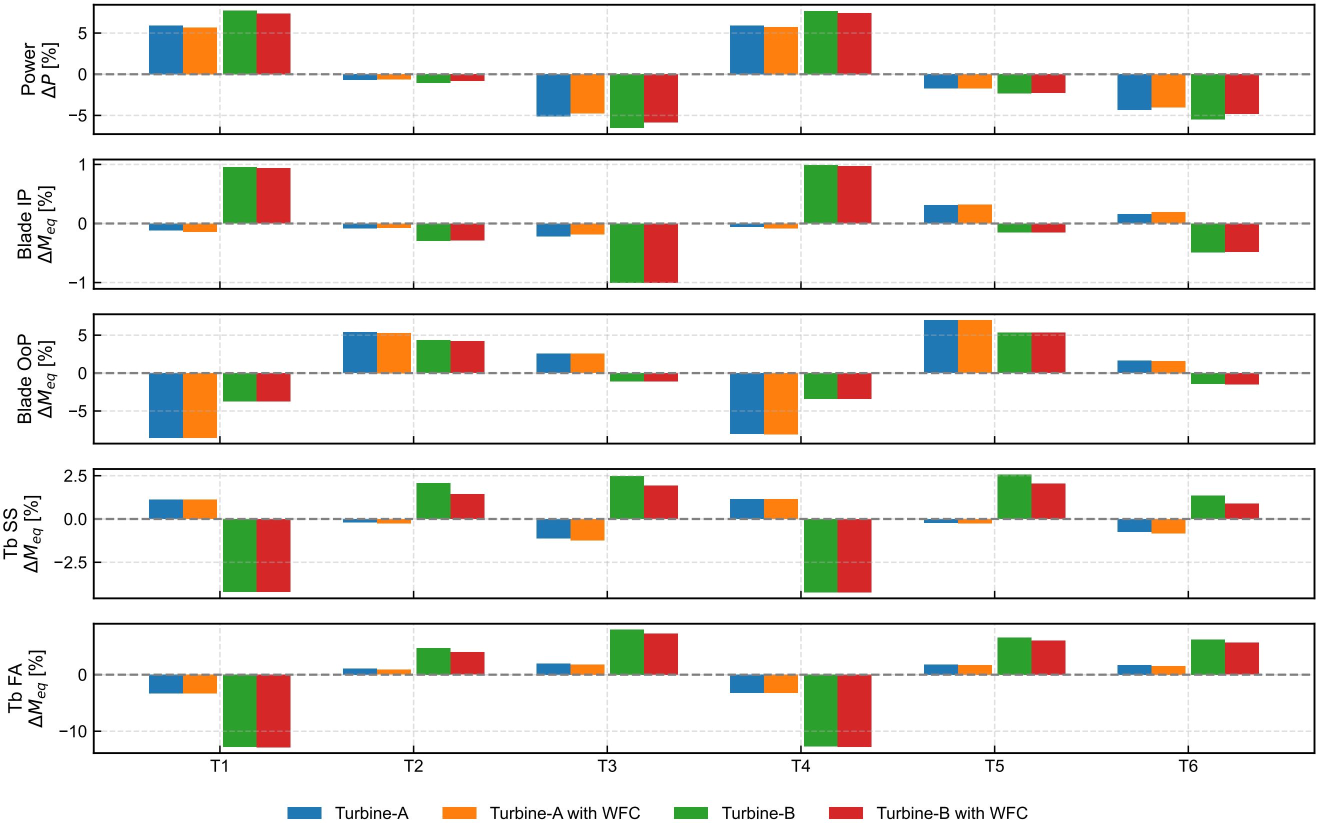


**Figure 18.** Deviation of power (top) and long-term damage equivalent loads (bottom four subplots) from the farm-wide mean across all six turbines.

**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 |

a general conclusion on the effectiveness of wake steering, which remains highly dependent on the turbine design, wind-farm layout, site conditions, and control strategy.

## 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.

### 4.1 Key findings

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~~This approach creates a coherent link between high-fidelity aeroelastic simulations and large-scale wind-farm analyses, ~~avoiding~~ supporting fair comparisons across different turbine types and reducing the need for computationally ~~expensive farm-level~~ intensive farm-level aeroelastic modeling in many ~~practical~~ real-world applications.

The applicability of the proposed toolchain has been demonstrated across multiple turbine types, including one open-source reference turbine and two ~~commercially deployed~~ commercial turbines. This confirms that the framework can be efficiently applied to turbines with different levels of data availability, structural ~~layouts, and control philosophies~~and control designs, 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~~ shut-down and start-up 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~~ among 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~~, DELs at blade root (in- and out-of-plane)~~and at the~~, and DELs at tower base (fore–aft and side–side). The performance and applicability of the framework were evaluated across multiple dimensions, including the accuracy of the individual surrogate models, the validation of the framework for wind-farm load predictions, and its application under varying inflow conditions and farm layouts. These analyses lead to the following conclusions:

– Overall, the load surrogate models exhibit very satisfactory performance across various datasets, turbine models, and control modes. In ~~power production~~ 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 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~~ load levels 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~~ an example 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~~ can be used to provide 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.

## 4.2 Limitations

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~~ diameters between 90 m and 120 m and hub ~~height~~ heights 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~~ some application-specific inflow conditions may fall outside the domain represented in the surrogate library, potentially requiring suitable 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~~ the turbine response 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~~Consequently, while the framework itself remains applicable across different wind farms, the underlying flow model may require site-~~up factors, the framework can no longer be considered wind -farm agnostic~~specific calibration when terrain effects are significant. Nevertheless, the modular architecture of the framework allows individual components to be replaced with more advanced models, mitigating current limitations as improved capabilities become available.

## 4.3 Future work

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~~ will 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~~wake mixing. 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.

## Appendix A: Nomenclature and abbreviations

| | |
|---|---|
| ADC | Actuator duty cycle |
| ANN | Artificial neural network |
| DEL | Damage equivalent load |
| DE | Derating |
| DFM | Dynamic flow model |
| DR | Down-regulation |
| DWM | Dynamic wake meandering |
| FA | Fore-aft |
| IEC | International electrotechnical commission |
| IP | In-plane |
| IPC | Individual pitch control |
| LFFC | Low-frequency fatigue cycle |
| LHS | Latin hypercube sampling |
| MISO | ~~Multiple-input–single-output~~Multiple input single output |
| NC | ~~Noise-curtailment~~Noise curtailment |
| NO | ~~Normal-operation~~Normal operation |
| NRMSE | Normalized root mean square error |
| OoP | Out-of-plane |
| PK | Parked |
| PP | ~~Power-production~~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 |
| SFM | Steady-state flow model |
| 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 |
| YS | ~~Yaw-steering~~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~~ Longitudinal wind velocity component |
| $v$ | ~~lateral~~ Lateral wind velocity component |
| $w$ | ~~vertical~~ Vertical wind velocity component |
| $y$ | Variable denoting surrogate output feature |
| $\square_{\mathrm{eq}}$ | Equivalent |
| $\square_{\mathrm{IEC}}$ | Meteorological and operational conditions as per IEC standard |
| $\square_{\mathrm{init}}$ | Initial value |
| $\square_{\mathrm{LT}}$ | Long-term accumulated value |

| | |
|---|---|
| $\square_{\text{max}}$ | Maximum value over a sample in a dataset |
| $\square_{\text{min}}$ | Minimum value over a sample in a dataset |
| $\square_{\text{ref}}$ | Reference |
| $\square_{\text{WT}}$ | Value for a wind turbine (compared to wind farm average) |
| $\Delta\square$ | Relative variation compared to a mean |
| $\hat{\square}$ | Predicted value obtained from the surrogate |
| $\overline{\square}$ | Mean value over a sample of points |
| $\epsilon$ | Normalized error for each data point |
| $\mu$ | Mean error |
| $\sigma$ | Standard deviation of error |

## 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 |
|---|---|---|---|
| Blade/hub bearing | wrot_Bl1Rad0InPlnMnt | Blade bearing ~~in plane~~ in-plane moment (fixed) | 4 |
| | wrot_Bl1Rad0OutPlnMnt | Blade bearing ~~out of plane~~ 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)~~ torsional moment | 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 | – |

## Appendix C: Comparison of outputs of wind farm response framework to aeroelastic simulations.

Figure C1 compares the surrogate outputs generated by the wind-farm response framework with the corresponding aeroelastic simulation results for clean and various waked conditions, evaluated for both turbine types. The flow fields used by the framework are produced with `Floris` (NREL, 2025), employing the default parameters of the widely used Gaussian–Curl–Hybrid, the Crespo–Hernandez turbulence model (Crespo and Hernández, 1996), a vertical wind-shear profile defined by a power-law exponent of 0.2 referenced to hub height, and a 10×10 rotor discretization. Clean conditions correspond to an isolated turbine operating in free-stream flow, whereas waked conditions are reproduced using the same two-turbine configuration applied for generating the synthetic waked inflow (cf. Sect. 2.3).

Under free-stream conditions, the surrogate outputs closely match the aeroelastic simulation results. Under waked conditions, the agreement varies across output channels, with the closest match observed for tower DELs and the weakest for blade IP loads.

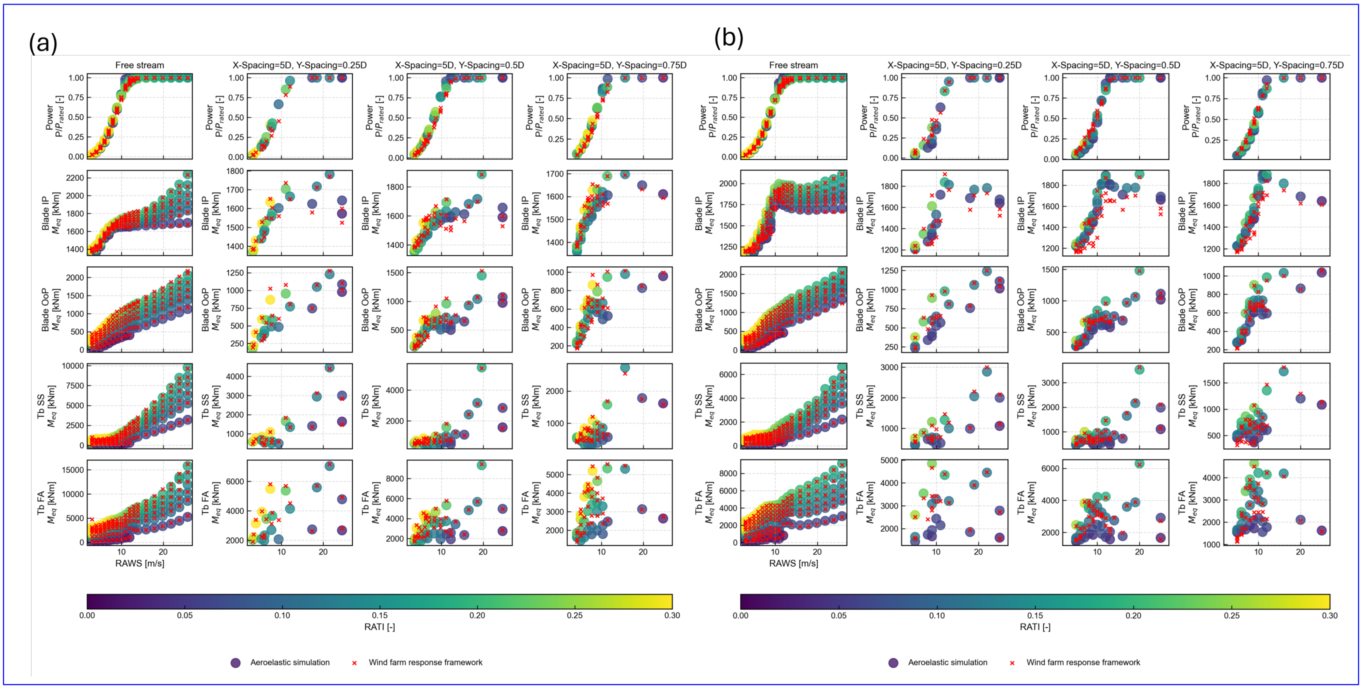


**Figure C1.** Comparison of the wind-farm response framework outputs with aeroelastic simulations for **(a)** Turbine-A and **(b)** Turbine-B under free-stream and three partially waked conditions as a function of wind speed.

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

Figure D1 presents the optimal yaw angles as a function of wind direction and turbine for a wind speed of 8 $\mathrm{ms}^{-1}$, 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.

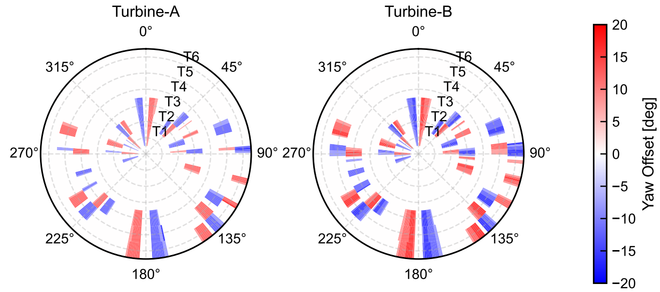


**Figure D1.** Optimal yaw angles as a function of wind direction and turbine. The color scale encodes yaw magnitude, and stacked bars represent the values for the individual turbines.

*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~~.18633503 (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~~.18634031 (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~~ steady-state flow model, and location-agnostic load surrogates, supervised the research, and ~~reviewed~~ contributed to the writing of 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 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.

## 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 D1.3: Data Standardization Methodology, Technical Report Version 0.3, TWAIN Project (Horizon Europe Grant Agreement No. 101122194), https://ec.europa.eu/research/participants/documents/downloadPublic?documentIds=080166e51cfec67e&appId=PPGMS, 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.18634031, data set, 2026a.

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, 3224, 032 106, https://doi.org/10.1088/1742-6596/3224/3/032106, 2026b.

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.

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.

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.

Pettas, V., Göçmen, T., Duc, T., and Vad, A.: On the effects of bat protection strategies on energy production and structural loads of wind farms, Wind Energy Science Discussions, 2026, 1–42, https://doi.org/10.5194/wes-2026-56, 2026.

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, 3224, 062 064, https://doi.org/10.1088/1742-6596/3224/6/062064, 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.

TWAIN: Code repository, https://gitlab.windenergy.dtu.dk/eu-twain/twain, accessed: 24 Jul. 2026, 2026.

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.

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.18633503, data set, 2026b.

Vad, A., Shah, A. H., Aktan, D. H., and Bottasso, C. L.: Identifying harmful stress cycles in wind turbine operations, Journal of Physics: Conference Series, 3224, 072 015, https://doi.org/10.1088/1742-6596/3224/7/072015, 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.