Efficient generation of location-agnostic wind turbine load surrogate models using wake slices

A Guilloré[1], A Chaudhary [1], A Anand [1], A Vad[1], A H Shah[1], V Pettas[2], T Göçmen[2] and C L Bottasso[1]

[1]Wind Energy Institute, Technical University of Munich, Garching, Germany
[2]Department of Wind & Energy Systems, Technical University of Denmark, Roskilde, Denmark

E-mail: adrien.guillore@tum.de

Abstract. Accurate yet computationally efficient models of turbine fatigue response are essential for wind farm layout design and flow-control optimization. Wake effects dominate turbine fatigue loading, and fatigue behavior varies across turbine designs and control strategies, requiring costly wind-farm-level simulations in most existing data-driven surrogate approaches. This work demonstrates that location-agnostic fatigue load surrogates can be trained using only single-turbine simulations. A large library of inflow slices is generated from dynamic simulations of a single wake and reused to drive aero-servo-elastic simulations of different turbine models of comparable scale. The methodology is demonstrated for three distinct turbines and validated against full dynamic simulations of a six-turbine wind farm under varying wake impingement, including wake-steering scenarios. The results show that the proposed approach accurately captures relative variations in turbine fatigue loads across the wind farm, both with and without wake steering, enabling efficient fatigue-aware wind farm optimization.

1. Introduction

In wind farms, wake effects influence not only the power production of downstream turbines but also their structural fatigue loading. While steady-state wake models can achieve high accuracy in predicting farm-wide energy yield, fast yet accurate estimation of component-level fatigue loads remains challenging [1]. Damage Equivalent Loads (DELs) are typically obtained from time-domain aero-servo-elastic simulations, such as those implemented in [2]. Wake interactions strongly affect the DELs of downstream turbines due to the combined effects of local velocity deficits, wake meandering, and wake-added turbulence [3, 4]. For long-term design purposes, standards for lifetime fatigue assessment often adopt conservative representations of wake effects, such as increased effective turbulence [5]. However, these approaches are insufficient when short-term impacts of wind farm control strategies on DELs must be evaluated, particularly under partial-wake conditions [4]. Alternatively, the Dynamic Wake Meandering (DWM) model provides a practical framework to incorporate wake effects into aero-servo-elastic simulations [3]. These simulation models, however, are not well suited for integration into optimization problems due to their high computational cost. Consequently, recent research has focused on developing fast, accurate, location-agnostic surrogate models for load prediction [6–8]. Such studies show that these surrogates can achieve high accuracy while being both layout-agnostic (independent of turbine position within a farm) and control-oriented (capturing the influence of control

setpoints). Nonetheless, they remain turbine-type dependent, since each turbine exhibits distinct aeroelastic behavior and controller dynamics. As a result, adopting a new turbine model requires generating an entirely new training dataset of aero-servo-elastic simulations that include dynamic wake effects. By contrast, single-turbine simulations can be performed much more efficiently.

This work addresses whether location-agnostic load surrogates can be trained using only single-turbine simulations. Building on the surrogate architecture presented in [8], we generate training datasets for three new turbine models. The proposed approach incorporates dynamic wake effects through inflow slices derived from the DWM. This methodology is described in section 2 and demonstrated via farm-level testing in section 3. A key advantage of this approach is that any changes in turbine model or controller dynamics can be rapidly reflected in the surrogates without the need for extensive dynamic wake simulations. The method is applicable to any turbine of similar size and any control setting, addressing a central concern in both academic and industrial wind energy research.

2. Approach

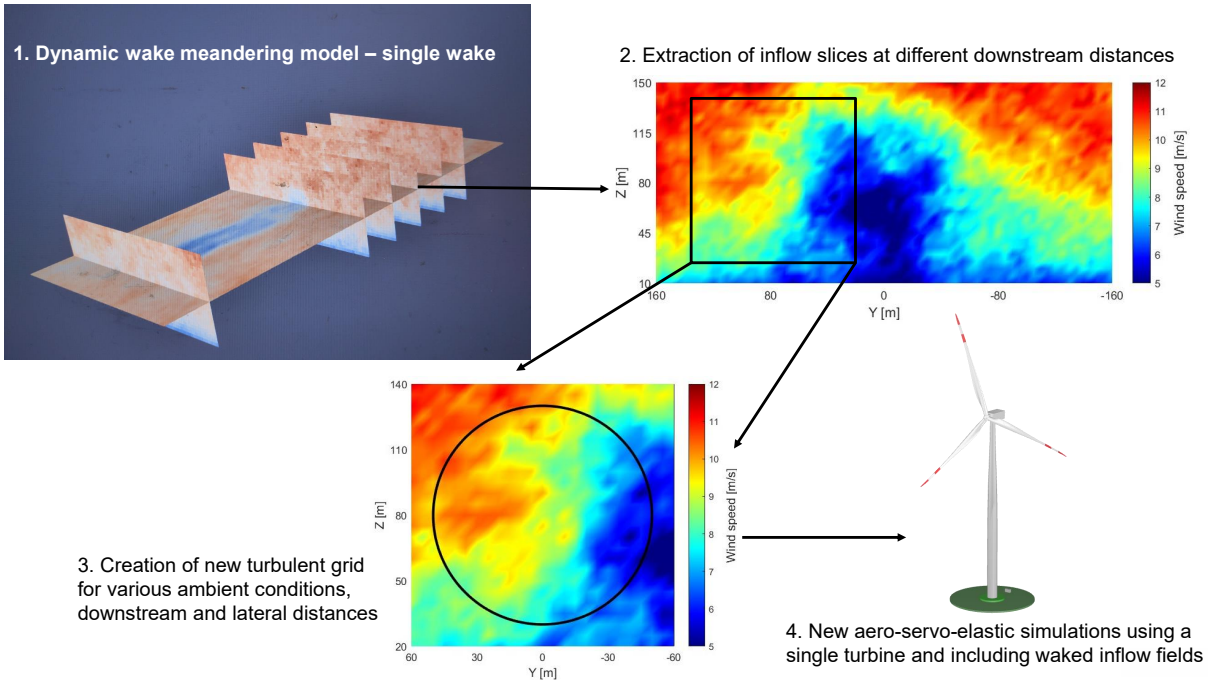

Figure 1. Overview of the proposed approach for generating a library of wake slices using DWM and applying them to single-turbine simulations.

Figure 1 illustrates the workflow for efficiently generating and using a large number of wake slices to train location-agnostic response surrogate models. First, the DWM model [3] is employed to generate a time-resolved 3D field of a single wake (step 1). Cross-stream slices are then extracted across a wide domain at the free-stream and multiple downstream locations (step 2). In step 3, reduced time-varying 2D turbulent fields are obtained at various streamwise and lateral positions, capturing different levels of wake recovery and impingement efficiently. Finally, these turbulent fields serve as inflow for aero-servo-elastic simulations of any new (similar-sized) turbine (step 4). Steps 1–3 are performed once using a single turbine type to build a comprehensive library of representative dynamic wake fields, whereas step 4 can be efficiently repeated for any (similar-sized) turbine model or control strategy.

Following [8], the response surrogates are machine-learning models trained to predict DELs based on four sector-averaged wind speeds (SAWS) and four sector-averaged turbulence intensities (SATI) across the rotor disk, together with control setpoints (yaw misalignment and power demand). Figure 2 illustrates the extraction of SAWS and SATI from the turbulent inflow slices. While figure 1 shows instantaneous views of the time-varying turbulent flow field, the top row of figure 2 presents the 12-seeds-averaged local time-averaged wind speed (left) and local turbulence intensity (right) over 10 minutes. The bottom row shows the corresponding sector-averaged quantities at the rotor disk. Figure 2 depicts a representative case with an ambient wind speed of $10\,\mathrm{m\,s^{-1}}$, ambient TI of 6%, and a slice taken at a streamwise distance of $\mathrm{X} = 6\,\mathrm{D}$ and lateral distance of $\mathrm{Y} = 0.5\,\mathrm{D}$. In this scenario, the turbine of interest is partially impacted by a wake on the right side of its rotor (viewed from upstream). As a result, the right side of the rotor experiences both lower wind speeds and higher turbulence intensity, with the local effective TI reflecting the combined effects of wake meandering and wake-added turbulence.

It was shown in [8] that these four SAWS and SATI are simple yet sufficiently accurate reduced-order quantities for capturing the spatio-temporal influence of the local inflow on structural DELs, including wake effects in a wind farm. While [8] extracted the local SAWS and SATI from farm-wide simulations, in the present work they are obtained directly from the turbulent wake slices. It is important to note that, although the wake slices are generated using a fixed wake diameter, they can be applied to turbines of different sizes. If the rotor diameter of the turbine of interest changes (black circle in figure 2), the corresponding SAWS and SATI are adjusted accordingly to the area coverage. The core principle of the present surrogate formulation is that the DELs of a turbine depend solely on the local inflow it experiences, regardless of its origin (free-stream flow heterogeneity, combinations of wakes, wake size, etc.). The validation and testing results reported in both [8] and the present study confirm the suitability of this approach to a satisfactory level of accuracy. However, there is certainly a limit to the difference in rotor diameters where this approach is usable.

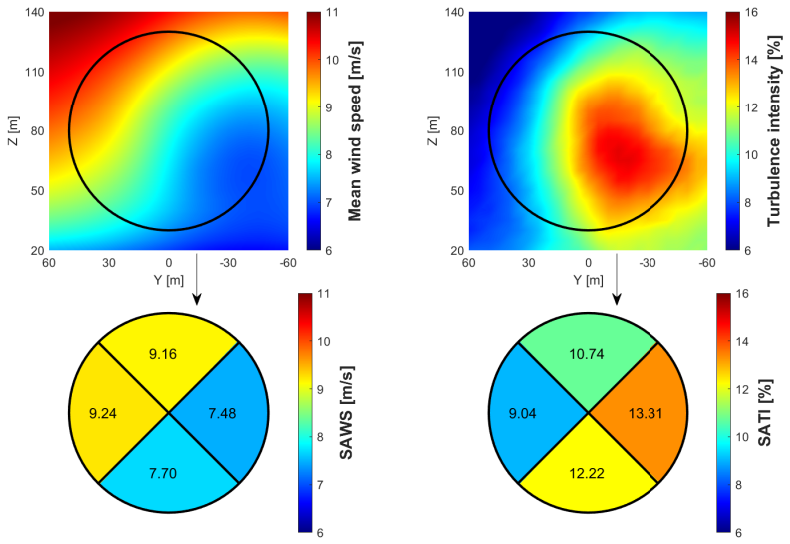

Figure 2. Inputs to the load surrogates: four sector-averaged wind speeds (SAWS, left) and four sector-averaged turbulence intensities (SATI, right). Here, the ambient wind speed is $10\,\mathrm{m\,s^{-1}}$ and TI is 6%. The streamwise distance is $\mathrm{X} = 6\,\mathrm{D}$ and the lateral distance is $\mathrm{Y} = 0.5\,\mathrm{D}$ from the wake-generating turbine. This inflow case is indicated by the triangle markers in figure 5.

In this work, step 1 (figure 1) was performed using the DWM implementation in FAST.Farm [9], which incorporates the latest wake-added turbulence and curled-wake models. The *NREL-1.79-100* turbine [10] (rotor diameter 100 m) was used to generate the wake. A total of 83 combinations of ambient wind speeds (from 3 to 25 $\mathrm{m\,s^{-1}}$) and turbulence intensities (from 3 to 28%) were selected via Latin Hypercube Sampling [11] and generated using the Kaimal model in TurbSim [12]. To generate the wake, the *NREL-1.79-100* turbine was operated at its full power curve (no curtailment) and without any yaw misalignment. Consequently, the resulting library of waked inflow does not include cases with deflected or curled wakes. The assumption was made that partial-wake cases are sufficient for the load surrogate to remain applicable in steered-wake scenarios. This simplification avoids an exponential increase in the number of dynamic wake simulations required for varying yaw misalignments. While the results in section 3.3 show a satisfactory match of the surrogates even when the farm applies wake steering, future work will expand the wake-slice library to include misaligned conditions.

Once the 3D DWM simulations of the single wake were completed, turbulent fields of 120 m by 120 m were extracted at streamwise locations of [−0.5, 4, 5, 6, 7, and 8 D] and lateral locations of [−1 : 0.25 : 1 D] relative to the wake-generating turbine (with D = 100 m referencing the *NREL-1.79-100*). This procedure yielded a total of 4 482 local wake slices. The turbulent slices consist of 25×25 points (i.e., a 5 m spatial resolution) and are sampled at 10 Hz. To ensure statistical convergence of both the inputs (SAWS and SATI) and the outputs (DELs) of the surrogates, 12 turbulent seeds were simulated for each inflow case.

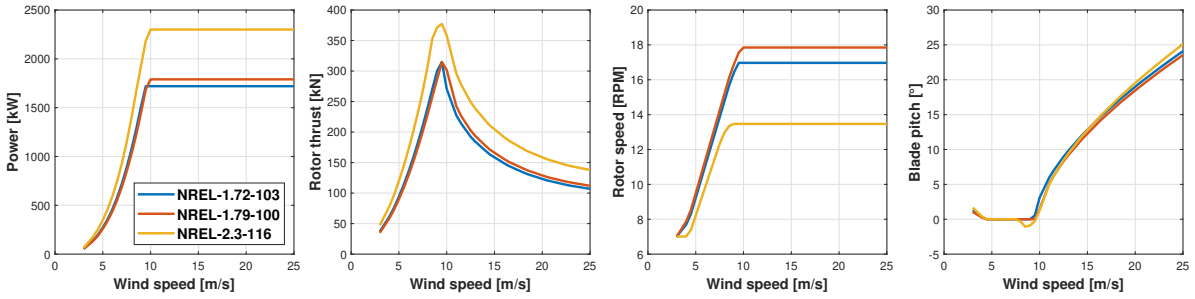

Figure 3. Regulation trajectories of the three turbine models considered, adapted from [10].

To demonstrate the generality of the approach, step 4 (figure 1) was applied using three different openly available wind turbines: the *NREL-1.72-103*, *NREL-1.79-100*, and *NREL-2.3-116* [10]. These turbines differ in size and control trajectories, as illustrated in figure 3. The hub height was fixed at 80 m for all turbines to ensure the wake remained vertically centered. The open-source ROSCO controller [13], with default control parameters provided in [10], was used. For each turbine, a total of 10 042 aero-servo-elastic simulation cases were conducted using OpenFAST [2]. This includes the 4 482 local inflow cases combined with nacelle yaw misalignments from −30° to +30° and curtailment levels from 100% to 50%. Curtailment was implemented by entering the rated region earlier, with consequent adjustments in rotor speed and blade pitch, as in [8].

For each 10-minute time series, the 1 Hz DEL is computed using conventional rainflow counting and the equivalent load formula from [14]. Four load channels are considered in this work (although the method can be applied to any available load measurements): the blade-root in-plane and out-of-plane bending moments (using a Wöhler exponent of 10 for the polymer material) and the tower-base side–side and fore–aft bending moments (using a Wöhler exponent of 4 for the steel material). After random shuffling, the datasets are split into 80% for training and 20% for validation. Load surrogate models were trained for each of the three turbines

following the methodology in [8]. In contrast to [8], Gaussian Process Regression (GPR) [15] with automated hyperparameter optimization implemented in Matlab [16] is employed instead of artificial neural networks. The surrogates use a multi-input, single-output (MISO) configuration; consequently, twelve MISO surrogate functions are trained in this study (four load channels for each of the three turbines).

To address the main research question posed in section 1, the wake-slice approach was tested using independent full wind farm simulations. FAST.Farm [9] simulations were performed for a cluster of six turbines, comprising two machines for each of the three turbine models [10]. The selected farm layout is shown in figure 4. Ambient wind speeds of 8, 10, and 12 $\mathrm{m\,s^{-1}}$ with TI of 6%, and wind speeds of 9 and 11 $\mathrm{m\,s^{-1}}$ with TI of 14%, were included, with all turbines operating at maximum power (greedy mode). Additionally, cases with wake steering were simulated for ambient wind speeds of 8 and 10 $\mathrm{m\,s^{-1}}$ at TI of 6%, using optimal yaw misalignments for each turbine derived from FLORIS v4.5 [17]. For both controlled and uncontrolled operating conditions, the wind farm layout was rotated to represent inflow wind directions ranging from $-45°$ to $+45°$ in $5°$ increments. This resulted in a total of 133 FAST.Farm simulation cases, each evaluated with 12 independent turbulent inflow realizations. Since the farm includes two instances of each turbine model, the dataset comprises 266 testing cases for each turbine-specific load surrogate. These surrogates, trained using single-turbine simulations driven by wake-derived inflow slices, were evaluated against this dataset, which captures realistic wind farm operating conditions, including multiple wake interactions, varying rotor diameters, and different levels of wake impingement.

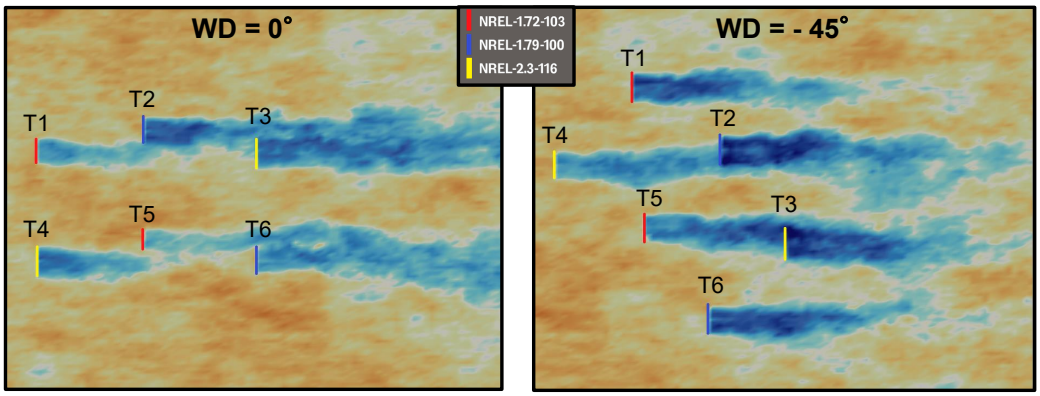

Figure 4. Wind farm layout (with two exemplary wind directions) used for the farm-level testing. The different turbine models are indicated by the colors shown in the legend.

3. Results

To demonstrate the suitability of the approach, the results are presented as follows. Section 3.1 shows some outputs of the aero-servo-elastic simulations for a variety of wake-impingement scenarios using the library of turbulent slices. Section 3.2 reports performance metrics of the load surrogates across the entire training, validation, and testing datasets. Finally, section 3.3 presents some detailed results of the load predictions in the wind-farm-level tests using six turbines.

3.1. Simulation dataset for varying wake impingement

Using the approach illustrated in figure 1, a large dataset of single-turbine simulations including dynamic wake effects was efficiently generated for the three turbine models. Figure 5 presents

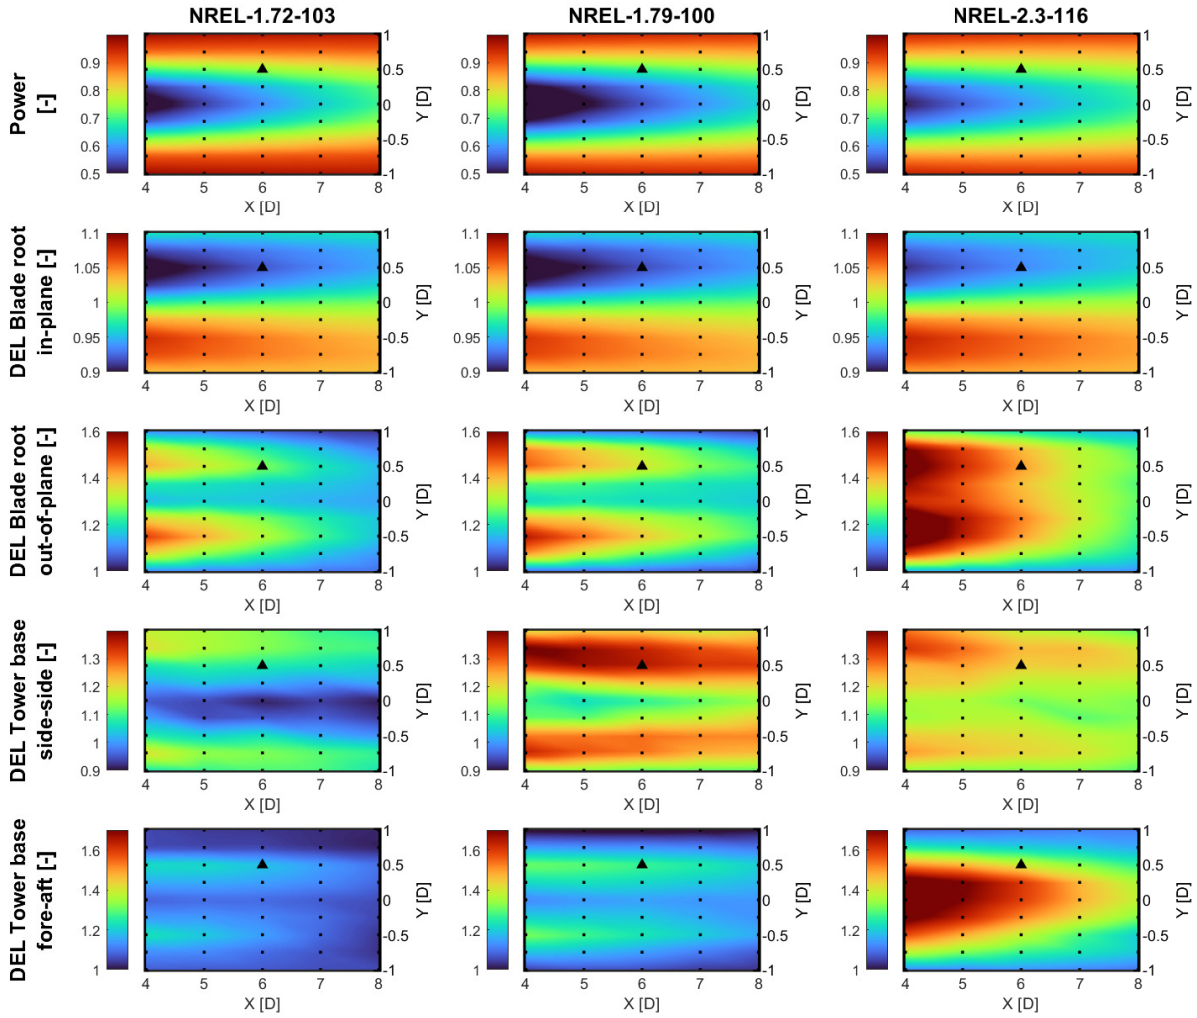

Figure 5. Contour plots of simulation output channels (by rows) for the three turbine models (by columns) as functions of streamwise (X) and lateral (Y) distance from the wake-generating turbine. Black dots indicate single-turbine simulation cases (averaged over 12 turbulent seeds), and triangle markers denote the inflow condition shown in figure 2. All quantities are normalized by the corresponding free-stream values (here at $10\,\mathrm{m\,s^{-1}}$, TI 6%) for each turbine and channel.

example outputs of these simulations, showing power, blade-root in-plane and out-of-plane DELs, and tower-base side–side and fore–aft DELs for all turbines under one ambient condition ($10\,\mathrm{m\,s^{-1}}$, TI 6%). Each subplot is constructed from 45 simulations (black dots), each averaged over 12 turbulent seeds, spanning the range of streamwise and lateral positions relative to the wake-generating turbine. Linear interpolation is used to generate the contours. The triangle markers indicate the partial-wake case illustrated in figure 2 for reference.

The power outputs (first row) behave as expected and are consistent with predictions from state-of-the-art steady-state wake models: the Gaussian-shaped velocity deficit is laterally centered within the wake and gradually recovers further downstream. The remaining four rows show DELs, which cannot be captured by steady-state wake models unless enhanced with load surrogates such as the ones presented here. Notably, each DEL channel exhibits distinct behavior that differs from power and varies across turbine models. The blade-root in-plane DEL shows

a clear asymmetry relative to the wake center: a turbine impacted by a wake on the right side of its rotor ($Y > 0$) experiences reduced in-plane DEL, whereas DEL increases when the wake hits the left side ($Y < 0$). This phenomenon arises from the combined effects of gravity and aerodynamic forces, which amplify or reduce loading cycles depending on the direction of rotation. These findings are consistent with observations reported in [18]. Further downstream, the in-plane DELs converge toward their symmetric values as the wake strength diminishes.

For the blade-root out-of-plane DELs, a pronounced double-peak pattern is observed as a function of lateral displacement. This indicates that partial-wake conditions are particularly detrimental for this load channel, as the blades repeatedly rotate in and out of the wake, leading to increased out-of-plane loading cycles. Similar double-peak features are also observed for the tower-base side–side and fore–aft DELs, although this behavior is less consistent across turbine models. Tower DELs appear to be more turbine-specific, reflecting differences in both structural design and controller behavior. Tower-base DELs are primarily driven by the turbulence intensity experienced at the rotor. In partial-wake conditions, this turbulence intensity can either increase or decrease due to a complex balance between wake meandering, which tends to increase turbulence under partial-wake conditions, and wake-added turbulence, which dominates in full-wake situations. This balance is further influenced by rotor size when a fixed wake diameter is assumed, as in the present study. Notably, the larger *NREL-2.3-116* turbine (rotor diameter 116 m) exhibits the most distinct DEL behavior. Since the wake diameter is fixed at 100 m (generated using the *NREL-1.79-100*) and this value is also used to define the spatial normalization (D in figure 5), turbines of different sizes experience different wake fractions at the same normalized lateral position. For example, at $Y = \pm 0.5\,D$, the *NREL-1.79-100* and *NREL-1.72-103* operate under clear half-wake conditions, whereas the larger *NREL-2.3-116* remains closer to a full-wake situation. While this size mismatch represents a limitation of the present approach for generalizing to turbines of arbitrary size—the *NREL-2.3-116* likely representing a practical upper bound—it also reflects a realistic operating scenario. In real wind farms, turbines of different sizes may indeed be exposed to wakes that are smaller or larger than their rotor diameters. From this perspective, the approach provides a generic framework for creating location-agnostic load surrogates by mapping the local inflow conditions to the corresponding local DELs. Additionally, differences in control trajectories, as shown in figure 3, further explain the observed variations in DEL behavior. For instance, the *NREL-2.3-116* exhibits a steeper variation of rotor thrust with wind speed, which contributes to the larger relative increase in tower-base DELs compared to the other turbines.

Overall, the DEL results shown in figure 5 visually demonstrate the need to use sector-averaged inflow quantities (SAWS and SATI) to capture the complex fatigue-loading behavior arising in full- and partial-wake conditions. They also highlight the turbine-specific nature of DEL responses, driven by differences in structural design and control strategies, thereby motivating the development of an efficient method for generating location-agnostic load surrogates across multiple turbine models.

3.2. Performance of trained load surrogates

Table 1 summarizes the normalized root-mean-square errors (NRMSEs) of the response surrogate models for all four DEL channels and the three turbine models across different datasets. For each turbine type, RMSE values are normalized by the range of the corresponding channel in the training dataset. The training dataset consists of 8 034 single-turbine simulations with wake slices used to fit the GPR models. The validation dataset comprises 2 008 independent single-turbine cases with wake slices that were not used during training. Finally, the testing dataset consists of the 266 farm-level cases illustrated in figure 4 for each turbine model.

Columns 3 and 4 show that surrogates trained on the extensive single-turbine dataset with dynamic wake effects perform well, with all NRMSEs below 0.7% of the prediction range. The

NRMSEs for the validation dataset are comparable to those obtained for the training dataset, indicating that the surrogates generalize well to unseen cases and that overfitting is avoided. Notably, the NRMSEs are consistent across all turbine models, demonstrating that the surrogate methodology of [8] is broadly applicable and that the proposed wake-slice approach delivers high predictive accuracy across diverse datasets. Column 5 of table 1 reports the NRMSEs obtained when surrogates trained exclusively on single-turbine simulations using the proposed approach are evaluated on the far more complex farm-level cases shown in figure 4. Despite the increased complexity of these scenarios—characterized by multiple wake sizes, wake superposition, and wake steering—the NRMSEs remain below 2% of the prediction range, confirming the suitability of the proposed approach for farm-level applications. Slightly higher NRMSEs are observed for the *NREL-2.3-116* turbine compared to the smaller turbines, consistent with its greater sensitivity to local turbulence and partial-wake conditions, as discussed in section 3.1.

Table 1. Normalized root-mean-square errors (NRMSE) of the trained surrogate models across different datasets, DEL channels, and turbine types. For each turbine type, RMSE values are normalized using the corresponding channel range from the training dataset.

DEL channel	Turbine	Training set: 8 034 single-turbine cases	Validation set: 2 008 single-turbine cases	Testing set: 266 farm-level cases
Blade root in-plane	*NREL-1.72-103*	0.33%	0.35%	0.72%
	NREL-1.79-100	0.35%	0.41%	0.60%
	NREL-2.3-116	0.31%	0.35%	1.06%
Blade root out-of-plane	*NREL-1.72-103*	0.39%	0.55%	1.67%
	NREL-1.79-100	0.39%	0.56%	1.39%
	NREL-2.3-116	0.50%	0.61%	1.97%
Tower base side-side	*NREL-1.72-103*	0.35%	0.40%	0.71%
	NREL-1.79-100	0.37%	0.39%	0.92%
	NREL-2.3-116	0.51%	0.60%	1.00%
Tower base fore-aft	*NREL-1.72-103*	0.30%	0.46%	1.43%
	NREL-1.79-100	0.35%	0.44%	1.34%
	NREL-2.3-116	0.43%	0.66%	1.59%

3.3. Testing against full wind farm simulations

Finally, figure 6 illustrates the performance of the different load surrogates (shown by columns) in predicting DELs for the farm-level testing cases and for each turbine instance (shown by rows) depicted in figure 4. The results correspond to an ambient wind speed of $10\,\mathrm{m\,s^{-1}}$ and a TI of 6%, with wind direction as the primary varying parameter, inducing different wake impingement conditions across the farm. Cases without wind farm control (WFC) are shown in blue, whereas cases with optimal yaw misalignments are shown in red.

Overall, the results demonstrate that surrogates trained exclusively on single-turbine simulations successfully capture the detailed variations of all four DEL channels under changing wind directions in this farm-level testing. For turbines operating solely in free-stream conditions (T1 and T4), the DEL variations arise from the prescribed yaw misalignments under nearly constant local inflow. These variations are accurately reproduced by the trained surrogates. For the remaining turbines, DEL variations result from a combination of changing local inflow conditions due to wakes and the application of yaw control. In these cases, the surrogates

adequately capture the observed DEL trends using only the local SAWS, SATI, and control set-points.

Among the four load channels, the blade root in-plane DELs are predicted with the highest accuracy, as these loads are primarily driven by gravity and relatively simple partial-wake effects. In contrast, the tower base side-side and fore-aft DELs are more challenging to predict because of their increased sensitivity to local turbulence levels arising from wake meandering, wake-added turbulence, and curled wakes under yawed conditions. A small number of cases are not well captured by the surrogates in this farm-level testing, particularly for turbine T3 at wind directions of $-40°$ and $+10°$. Further investigation of these cases revealed unusually high and horizontally sheared SATI values, without a clear corresponding increase in the tower base DELs. Future work will examine these discrepancies in more detail to determine whether they are systematic or stochastic in nature. Notably, these mismatches occur only in cases with active wind farm control, suggesting that the absence of steered and curled wake conditions in the surrogate training dataset may contribute to the observed discrepancies.

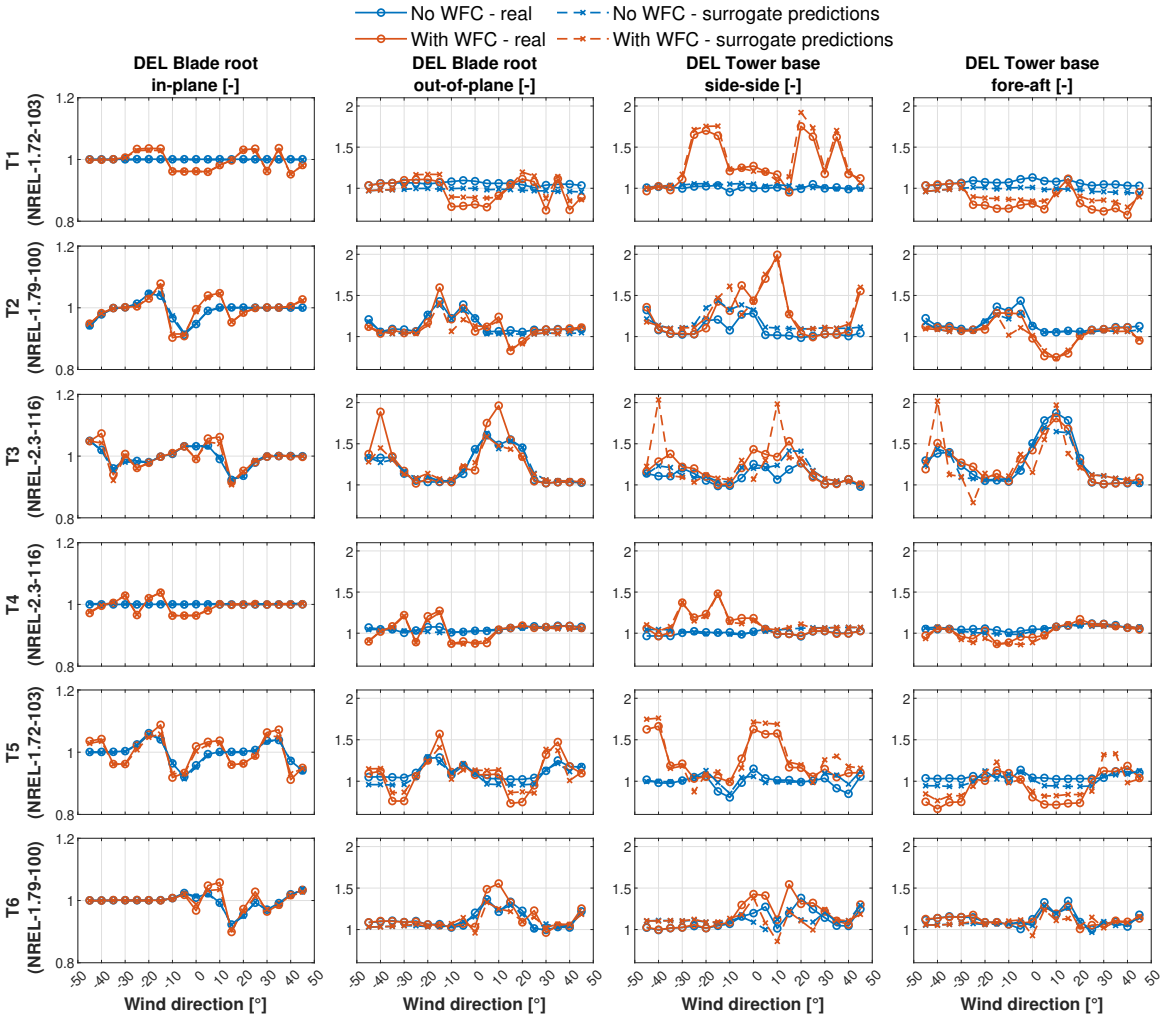

Figure 6. Surrogate predictions (dashed lines and crosses) versus reference values (solid lines and circles) for some wind farm test cases as a function of wind direction. Blue: no wind farm control (WFC); red: optimal wake steering. All quantities are normalized by the corresponding free-stream values (at $10\,\mathrm{m\,s^{-1}}$ and turbulence intensity of 6%) for each load channel and turbine.

4. Conclusions and future work

This study demonstrates that location-agnostic fatigue load surrogate models can be efficiently trained using only single-turbine simulations, provided that the training data adequately represent dynamic wake effects. This is achieved through a wake-slice methodology based on the DWM model, in which a large library of turbulent wake inflow fields is generated from a single-wake simulation and reused to drive aero-servo-elastic simulations of multiple turbine models.

Turbine-specific load surrogates were trained using sector-averaged wind speeds, sector-averaged turbulence intensities, and control setpoints as inputs. They demonstrated strong predictive performance not only on the training and validation datasets but also on independent farm-level test cases. These tests included multiple wake interactions, varying wind directions and inflow conditions, heterogeneous turbine models, and scenarios with wake steering.

Two main conclusions emerge. First, turbine fatigue loads are strongly affected by both wake interactions and turbine-specific structural and control characteristics, limiting the applicability of generic or intuitive load models. Second, combining single-turbine simulations with dynamic wake slices provides an efficient and sufficiently accurate framework for generating fatigue load surrogates for multiple, similarly sized turbines.

Future work will expand the wake-slice library to include steered wakes and apply the methodology to additional real turbine models. Ultimately, the proposed surrogates will be coupled with steady-state wake models to enable rapid wind farm layout and control optimization that accounts for both power production and fatigue loading.

Acknowledgments

This work has been supported by the TWAIN project, which receives funding from the European Union's Horizon Europe Programme under the grant agreement No. 101122194.

References

[1] Meyers J, Bottasso C, Dykes K, Fleming P, Gebraad P, Giebel G, Göçmen T and van Wingerden J W 2022 *Wind Energy Science* **6**
[2] NREL 2025 Openfast v4.0.2 URL `https://doi.org/10.5281/zenodo.14847846`
[3] Larsen G 2007 Dynamic wake meandering modeling Tech. Rep. Risø-R-1607(EN)
[4] Doubrawa P, Shaler K and Jonkman J 2023 *Wind Energy Science* **8**
[5] IEC 2001 Ts 61400-13:2001 URL `https://webstore.iec.ch/publication/2525`
[6] Shaler K, Jasa J and Barter G E 2022 *J. of Phys. Conf. Ser.* **2265**
[7] Liew J, Riva R, Friis-Møller M and Göçmen T 2024 *J. of Phys. Conf. Ser.* **2767**
[8] Guilloré A, Campagnolo F and Bottasso C L 2024 *J. of Phys. Conf. Ser.* **2767**
[9] Jonkman J, Doubrawa P, Hamilton N, Annoni J and Fleming P A 2018 *J. of Phys. Conf. Ser.* **1037**
[10] NREL 2024 GitHub repository, accessed March 12, 2025 URL `https://github.com/NREL/openfast-turbine-models/tree/main/IEA-scaled`
[11] McKay M D, Beckman R J and Conover W J 1979 *Technometrics* **21**
[12] Jonkman B J 2009 Turbsim user's guide: Version 1.50 Tech. Rep. NREL/TP-500-46198
[13] Abbas N J, Zalkind D S, Pao L and Wright A 2022 *Wind Energy Science* **7**
[14] Sutherland H J 1999 Tech. Rep. SAND99-0089 Sandia National Laboratories
[15] Rasmussen C E and Williams C K I 2005 (Cambridge, MA: MIT Press)
[16] The MathWorks Inc 2025 URL `https://www.mathworks.com/help/stats/fitrgp.html`
[17] NREL 2025 Floris v4.5 URL `https://doi.org/10.5281/zenodo.17186844`
[18] Stanley A P J, King J, Bay C and Ning A 2022 *Wind Energy Science* **7**