A control-oriented load surrogate model based on sector-averaged inflow quantities: capturing damage for unwaked, waked, wake-steering and curtailed wind turbines

A Guilloré, F Campagnolo, C L Bottasso

Wind Energy Institute, Technische Universität München, Boltzmannstr. 15, 85748 Garching bei München, Germany

E-mail: carlo.bottasso@tum.de

Abstract. This paper presents a novel approach to constructing a load surrogate model. It aims to estimate any wind turbine's damage equivalent loads (DELs) regardless of its position and under various farm control actions. The model relies solely on local inflow quantities (sector-averaged wind speeds and turbulence intensities) and local control parameters (rotor speed, pitch angle, and yaw misalignment). Despite its highly simplified representation of the complex behavior of the turbulent wind field, wake effects, and controller dynamics, these quantities prove sufficient to characterize DELs. The paper demonstrates the training of this load model within a simulation environment. Validation results using a different wind farm configuration indicate that the surrogate can accurately predict fatigue loads for both unwaked and waked turbines, encompassing scenarios of wake steering and induction control.

1. Introduction

The collective power output of a cluster of wind turbines can be improved by accounting for turbine-to-turbine interactions. Previous research has demonstrated that optimizing wind farm layout and implementing coordinated wind farm control are effective strategies for mitigating wake effects, thereby maximizing the efficiency of wind power plants. However, the effects of this collective power boosting on structural loads remain unclear. For instance, a turbine experiencing partial wake may potentially face increased fatigue. This uncertainty currently hinders the widespread adoption of these wind farm methodologies in full-scale industrial applications [16, 23]. Short-term damage equivalent loads (DELs) [22] are significantly influenced by dynamic parameters, such as local wind speed, turbulence levels, wake impingement, yaw misalignment, and local control set-points. Predicting fatigue loading at the turbine components level usually necessitates costly aero-servo-elastic simulations at the farm level. For each specific ambient condition, multiple turbulent realizations must be run to ensure convergence of the fatigue quantities. To address this computational challenge, prior research has proposed various surrogate modeling approaches [3, 5, 13, 19, 20]. However, these models are either site-specific (requiring repetition of training when the layout changes) or do not include the various complex effects of yaw misalignment and curtailment on self-induced loads and loads experienced by downstream turbines. This work introduces a novel generic surrogate approach to efficiently predict DELs of various turbine components based solely on simple local inflow quantities, in addition to the operational state of the machine. Given that reduced-order inflow models [6, 17] or wind sensing techniques [4, 21] can readily estimate these inflow quantities, the new load surrogate is generally applicable to load-aware layout and farm flow control optimizations.

The manuscript is organized as follows. Section 2 presents the proposed approach of constructing a load surrogate model based on sector-averaged inflow quantities. Section 3 details the application in a simulation environment and provides predictions of DEL trends in this dataset. Section 4 demonstrates the validation of the surrogate using a different wind farm configuration. Finally, Sect. 5 summarizes the main findings and explores future work.

2. Approach: a load surrogate model based on sector-averaged inflow quantities

Figure 1 depicts the conceptual framework of the proposed load surrogate modeling. The red background highlights the training and validation process. The blue background on the right illustrates the anticipated application of the surrogate.

2.1. Wind farm plant (panel (a))

The surrogate approach is broadly applicable to various sources of wind farm data as designated by panel (a) in Fig. 1. It may include numerical simulations, scaled wind tunnel experiments, or full-field measurements. Suitable sensors must be available to estimate all inputs (wind mapping at rotor disks and control set-points) and outputs (DELs at desired locations). Constructing a large and diverse dataset for training and validation cases is necessary.

2.2. Sector-averaged inflow quantities and control set-points (panel (b))

The primary challenge in employing a surrogate lies in identifying suitable reduced-order quantities that can be rapidly estimated at a lower computational cost while effectively capturing the variation of the desired predicted quantities with acceptable accuracy. The cyclic variations



Figure 1. Schematic overview of the proposed load surrogate model.

of structural loads fundamentally create the DELs of a wind turbine as the rotor is spinning (space dependency) and the turbulent wind field is varying (time dependency). Therefore, DELs depend not only on the rotor-averaged wind speed, but predominantly on the spatiotemporal variation of this wind speed across the rotor disk. Previous work demonstrated that relying solely on rotor-averaged wind speeds inadequately captures DELs [13]. In a wind farm, wake effects significantly amplify this spatiotemporal variation (due to local speed deficits and dynamic wake meandering). The present surrogate formulation employs four sector-averaged wind speeds (SAWSs) to capture spatial variation at the rotor disk. It uses sector-averaged turbulence intensities (SATIs) to capture temporal variation (and its coupling with spatial variation). Panel (b) in Fig. 1 illustrates the reduction from a high resolution inflow mapping into four sectors for a case of partial wake impingement.

Introducing higher-order inflow quantities, such as an increased number of sectors or refined statistics, has the potential to enhance the informational content for predicting DELs. However, obtaining these higher-order quantities for fast surrogate application may increase complexity and computational cost, and their use may lead to overfitting the training points. In this study, four sectors and their SAWS and SATI proved sufficient to deliver the desired accuracy.

Various techniques can be employed to estimate local SAWS and SATI, enhancing the broad applicability of this surrogate approach:

- In dynamic numerical simulations: utilizing virtual flow sensors at specific locations [9, 18], as demonstrated later in this work.
- In steady-state engineering flow models: estimating local wind speeds and TI, as available in [6, 17].
- In the field or wind tunnel: ground-based or nacelle-mounted lidar technology can estimate the spatiotemporal wind field [16].
- In the field or wind tunnel: wind observers were demonstrated to predict local inflow quantities based on the measurement of blade root loads [4, 21].

SAWS and SATI are, by definition, not influenced by the operational state of the machine. Nevertheless, DELs also vary with the local control set-points. This surrogate utilizes the mean turbine yaw misalignment (γ), rotor speed (Ω), and blade pitch angle (β) to capture the effects of wake-steering and curtailed turbine operation.

2.3. Architecture of the load model (panel (c))

Panel (c) in Fig- 1 illustrates the chosen architecture of the surrogate. It incorporates eleven inputs: four SAWSs, four SATIs, the mean nacelle yaw misalignment γ , rotor speed Ω , and blade pitch angle β . The output is a single selected DEL channel. A unique surrogate is trained and applicable independently of a turbine's position in a farm. Various possible mathematical ways of relating the inputs to outputs are possible. Multidimensional look-up tables proved challenging to use. While polynomial chaos expansion is a viable option, it requires an expensive training process and may not perform as well as artificial neural networks (ANNs). Extensive tests reported in [20] favored the latter option. Therefore, an ANN is selected as a surrogate fitting function, demonstrating efficient training and good predictive performance.

3. Application of the approach using aero-servo-elastic wind farm simulations *3.1. Simulation set-up*

A comprehensive dataset incorporating relevant farm effects is generated using the multiphysics tool FAST.Farm [9]. Previous validation work confirmed this tool's ability to predict DELs compared to full-field measurements [10]. The OpenFAST modules AeroDyn, ElastoDyn, and ServoDyn simulate the high-frequency dynamics of each turbine [18]. Three IEA 3.35MW



Figure 2. Control trajectories adapted from [2] illustrating the implemented curtailment option.

turbine models [2] are placed in a row with varying wind direction (by adapting the layout). Panel (a) in Fig. 1 displays an exemplary snapshot of the FAST.Farm simulations, for a case where the three turbines are each spaced 6D and the wind direction is -4° (where 0° represents fully aligned turbines). Synthetic ambient turbulent wind fields of 10 min are generated using the Kaimal model in TurbSim [8]. For each case, 24 turbulent seeds are run to ensure statistical convergence of both inputs (SAWSs, SATIs, and mean control parameters) and outputs (DELs for different load channels). Previous studies have shown that at least 20 turbulent seeds are necessary to converge all DEL channels below a 5% uncertainty threshold [12]. Each simulation runs for 600s, after a transient of 300s allowing full wake development throughout the domain. An updated version of FAST.Farm including wake-added turbulence is used [10], as it significantly affects DELs. For the incorporation of wake-added turbulence, secondary Mann turbulent boxes are generated [14] for each ambient turbulent seed according to the method detailed in [10]. All parameters of the dynamic wake model are maintained at their default tuned values [9, 10].

3.2. Curtailment implementation

The simulations use a Bladed-style DISCON turbine controller adapted from [2, 18]. The controller's dynamics and related DELs have undergone validation against a state-of-the-art open-source option [1]. Curtailment has also been implemented in the controller so that the turbine produces a desired maximum power, as illustrated in Fig. 2. This allows to include the effect of induction control on the fatigue of both the curtailed and downstream turbines. Employing a conventional PI-loop, the controller dynamically tracks the rotor speed and blade pitch angle for given wind conditions and desired maximum power.

3.3. Simulation cases for the training dataset

The surrogate must undergo training across a wide spectrum of ambient wind conditions, wake overlaps, and control strategies. To achieve this, the simulation training dataset includes ambient wind speeds of 6, 8, 10, and 12 m/s, along with ambient turbulence intensities of 6% and 12%. The wind field exhibits a vertical shear with a constant power law exponent of 0.2. Simulations account for turbine spacing at both 4D and 6D (expressed as a multiple of the rotor diameter). The wind direction varies within the range $-8^{\circ}:4:+8^{\circ}$. The most upstream turbine features yaw misalignments within the range $-30^{\circ}:10:+30^{\circ}$, while the middle turbine has yaw misalignments within the range $-20^{\circ}:10:+20^{\circ}$. Lastly, curtailment levels within the range 100:5:20 (in percentage of rated power) are applied to the front two turbines. Due to computational constraints, not all cross-cases were executed. However, the specific selection is

available upon request. In total, 2264 cases were run (each with 24 turbulent seeds). With three turbines, the resulting dataset comprises 6792 cases establishing links between local inflow quantities, control parameters, and measured DELs.

3.4. Extracting sector-averaged inflow quantities

Within the AeroDyn module [18], nine virtual sensors measuring the undisturbed wind speed are placed on each blade. Their radial positions are selected so that each sensor spans an equal disk surface, and their average radial position approximates 2/3 of the span (as per the demonstration in [4, 21]). In a post-processing step, these 27 turning sensors map the entire mean wind speed and turbulence intensity across the rotor disk, before spatially averaging them over four sectors (up, right, down, left). Panel (b) in Fig. 1 illustrates this process for a specific realization of the middle wind turbine in the farm case visualized in panel (a). In the case of a -4° wind direction, the middle turbine experiences a partial wake impingement on its left, which in turn creates the lower SAWS and higher SATI on the left compared to the right sector.

3.5. Damage equivalent loads projected along maximum damage direction

Four load channels are considered in this work (but the method is applicable to any available load measurements), using projection along the maximum damage direction. For each load time series, the 10 min DEL is computed using a conventional rainflow counting algorithm and the formula:

$$DEL = \left(\frac{\sum (n_i L_i^m)}{N_{eq}}\right)^{\frac{1}{m}},\tag{1}$$

where n_i is the number of load cycles of amplitude L_i , N_{eq} is the equivalent number of cycles (taken as per convention at 2e6 cycles over a lifetime of 20 years), and m is the Wöhler exponent of the S-N curve of the material (chosen as m = 10 for the polymer material of the blades and m = 4 for the steel material of the main shaft and tower) [11, 22].

For each of the load locations listed below, and at each time step, the two orthogonal components of a given load are projected along all possible directions, with an angular step of 10° . All corresponding DELs are computed using Eq. (1), and the maximum value is then stored as the *"projected DEL."* This projection is necessary to identify the most damaged direction on each cross section. The considered projected load channels and their components are:

- "Blade root DEL": blade root in-plane and out-of-plane bending moments.
- "Shaft DEL": shaft out-of-plane and yaw bending moments.
- "Yaw bearings DEL": tower top fore-aft and side-side bending moments.
- "Tower base DEL": tower bottom fore-aft and side-side bending moments.

3.6. Training of the Artificial Neural Network (ANN)

Once the dataset is constructed, it is randomly shuffled before being split in half to create training and validation subsets. Shuffling ensures similar distributions between both subsets. Additionally, the validation cases presented in Fig. 4 were manually inserted in the second half, meaning that they are not utilized for training. This practice ensures an unbiased validation.

The Matlab Deep Learning toolbox [15] is used for the creation and training of the ANN. The mean of squared errors is minimized using hyperbolic tangent sigmoid activation functions and Bayesian regularization. The ANN comprises 1 layer with 30 neurons.

Figure 3 presents the training and validation results of the four trained ANNs for the projected blade root, shaft, yaw bearings, and tower base DELs. Overall, minimal scatter is observed, confirming the validity of this surrogate approach. The resulting root mean squared percentage errors (RMSPE, see [7]) are relatively low compared to other surrogate approaches [3, 13, 20].



Figure 3. Scatter plots of the training and validation subsets resulting from the ANNs training for the four considered DELs. The RMSPE [7] are indicated for each load channel and subset.

The surrogate for blade root DEL performs the best, with a validation RMSPE value of only 0.48%. The surrogates for the shaft and yaw bearings DELs also perform very well with a validation RMSPE value of around 2%. The surrogate for tower base DELs shows more outliers. This component is the most sensitive to stochastic wind turbulence fluctuations. Nevertheless, the surrogate performs quite well with a validation RMSPE value of only 3.21%.

Importantly, all validation RMSPE values are not significantly higher than the training RMSPE values, indicating avoidance of overfitting and a general applicability to new cases.

3.7. Predictions of DELs trends by unknown external factors

Figure 4 presents the surrogate predictions of DELs for three selected farm trends. These cases originate from the same FAST.Farm dataset based on three turbines: WT1 (upstream), WT2 (middle), and WT3 (downstream). However, the cases presented in Fig. 4 were not part of the ANN training subset. The ambient average wind speed is 10 m/s with 6% turbulence intensity. The spacing between turbines is 6D.

In a baseline condition, the three turbines are aligned with the wind, without yaw misalignments, and at full rated power. The left column in Fig. 4 displays the effect of changing wind direction. The middle column illustrates the effect of misaligning the front turbine, steering its wake away from the downstream turbines. Lastly, the right column depicts the effect of curtailing the front turbine, reducing the speed deficit in its wake.

The surrogate model exclusively receives information on local SAWSs, SATIs, and control setpoints. Nevertheless, it predicts all the DELs with remarkable accuracy across the various farm configurations in Fig. 4. In the case of the upstream turbine WT1, SAWSs and SATIs remain constant. The surrogate effectively anticipates changes in its DELs based on the varying local control set-points (yaw, rotor speed, and blade pitch angle). Conversely, deliberate variations in control set-points are not introduced for the downstream turbines WT2 and WT3, as they are maintained in maximum power mode. The surrogate accurately forecasts changes in their DELs based on fluctuations in SAWSs and SATIs caused by wakes.



Figure 4. Surrogate predictions of DEL trends for all three turbines in the original dataset (but not included in training subset). Left column: DELs vs. wind direction. Middle column: DELs vs. yaw misalignment of the front turbine. Right column: DELs vs. curtailment of the front turbine. Each row corresponds to a load channel. Solid lines: actual DEL values. Dashed lines: surrogate predictions of DELs.

4. Application of the load model to a different wind farm configuration

To validate the desired broad applicability of the load surrogate for layout or farm control applications, its capabilities are tested using a different wind farm configuration with various farm control set-points.

4.1. Configuration of the validation wind farm

A wind farm comprising five IEA 3.35MW turbines [2] serves as a validation case. Figure 5 illustrates the layout that was selected to include various spacings between turbines and various degrees of wake impingement. WT1 is in free stream. WT2, 3, and 4 are in different partial wakes. WT5 is in a full wake. The simulations are conducted with FAST.Farm including wake-added turbulence [9, 10]. 24 turbulent seeds are run for each case. The ambient average wind speed is 10 m/s with 6% turbulence intensity. The inputs required by the surrogate model are extracted from FAST.Farm as explained in Sect. 3.4.

Twenty-one different cases of farm control set-points are run to test the ability of the surrogate to capture their effects on loads. The set-points of WT1, 2, and 3 are collectively varied. Their yaw misalignment angle varies as $-30^{\circ}:10:+30^{\circ}$. Their levels of curtailment are 100%, 80% and 60% of rated power.

4.2. Results: predictions of the load surrogate

Figure 6 compares the actual DELs with those predicted by the surrogate. The figure includes data for all five turbines (organized by columns) and four load channels (organized by rows). Overall, the surrogate demonstrates a very good ability in predicting the amplitudes and variations of all DELs for various control strategies. WT1 yields the best match, as wake effects do not influence it. The surrogate effectively captures the noteworthy effects of both yawing and curtailing WT1 on its fatigue. Curtailing the turbine mitigates the increase in blade root DELs for negative yaw misalignments. It also reduces the shaft and yaw bearings DELs for all yaw angles. However, for high misalignment amplitudes, tower base DELs in curtailed operation appear to be larger than those in fully rated operation.

The surrogate also performs well at predicting DEL variations for all the downstream turbines experiencing complex wakes and control interactions. The tower base DELs remain relatively low even for partially-waked turbines, but exhibit a significant increase in full wake situations.



Figure 5. Visualization of the layout used for validation. The yaw angle convention is indicated next to WT1. The background is a snapshot of the instantaneous velocity field at hub height using FAST.Farm [9, 10].



Figure 6. Surrogate predictions of DELs trends for the validation wind farm. Turbines are organized by columns, whereas load channels by rows. The yaw misalignment angles of WT1, 2 and 3 are collectively changed on the x-axis, while the colors blue, red and yellow respectively correspond to curtailment levels of 100%, 80% and 60% of rated power for WT1, 2 and 3 collectively. Solid lines: actual DEL values. Dashed lines: surrogate predictions of DELs.

This is evident at WT5, consistently positioned in the full wake of WT3. This is also evident at WT2 and WT3 for yaw angles that deflect the wake of WT1 fully towards them. The wakeadded turbulence is presumably driving this phenomenon. The shaft, yaw bearings, and the blade root DELs are more sensitive to partial wake impingements, exhibiting asymptric effects.

The surrogate shows a slightly reduced ability in predicting the shaft and yaw bearings DELs of WT2 and WT3. This could be attributed to the intricate coupling between the varying wake of WT1 affecting WT2 and WT3, and the fluctuations in yaw misalignments and curtailment levels for these turbines. Enhancing the surrogate's predictions in these complex scenarios might be achieved by incorporating additional training points that account for the coupling of partial wakes and local control changes.

Nevertheless, the surrogate successfully captures the variation in DELs influenced by wind farm control, even under different wake overlaps.

5. Summary and future work

This paper introduced a load surrogate model relying exclusively on local sector-averaged inflow quantities and local control set-points. The study was based on data from unsteady aero-servoelastic wind farm simulations. The surrogate reveals significantly low scatter and RMSPE across all considered DEL channels. Furthermore, validation was conducted using a different wind farm configuration. The model successfully predicts variations in DELs resulting from upstream wake steering and induction control for all turbines.

Future research should include a more detailed analysis of optimal input selection, and an investigation into the number of training cases required to achieve efficient yet accurate predictions of various DELs. The surrogate approach should be further validated using wind tunnel experiments and full field measurements. Additionally, combining the proposed surrogate with a reduced-order inflow model holds the potential to conduct load-aware layout and farm control optimizations.

Acknowledgements

This work has been supported in part by the PowerTracker project, with funding from the German Federal Ministry for Economic Affairs and Climate Action (FKZ: 03EE2036A). This work has been partially supported by the MERIDIONAL and the SUDOCO projects, which receive funding from the European Union's Horizon Europe Programme under the grant agreement No. 101084216 and No. 101122256, respectively.

References

- [1] Abbas N, Zalkind D, Pao L and Wright A 2022 Wind Energy Science 7 53-73
- [2] Bortolotti P, Canet H, Dykes K et al. IEA Wind TCP Task 37: Systems Engineering in Wind Energy WP2.1 Reference Wind Turbines. United States: N. p., 2019. Web. doi:10.2172/1529216.
- [3] Bossanyi E 2022 J. Phys.: Conf. Series 2265 042038
- [4] Bottasso C L, Cacciola S and Schreiber J 2018 Renewable Energy 116 155-68
- [5] Dimitrov N, Kelly M C et al. 2018 Wind Energy Science 3 767–90
- [6] DTU 2023, PyWake version 2.5, https://topfarm.pages.windenergy.dtu.dk/PyWake/notebooks/Overview.html
- [7] Göcken M, Özcalici M, Boru A and Dosdogru A T 2016 Expert Systems With Applications 44 320-331
- [8] Jonkman B J 2016: Turbsim user's guide: Version 2.00.00. Technical report, NREL
- Jonkman J and Shaler K 2021: Fast.farm user's guide and theory manual. Technical Report NREL/TP-5000-78485, NREL
- [10] Kretschmer M, Jonkman J, Pettas V and Cheng P W 2021 Wind Energy Science 6, 1247–62
- [11] International Electrotechnical Commission, IEC 61400-1: Wind turbines part 1: Design requirements. International Electrotechnical Commission, 2005.
- [12] Liew J and Larsen G C 2022 J. Phys.: Conf. Series 2265 032049
- [13] Liew J, Riva R and Gocmen T 2023 J. Phys.: Conf. Series 2626 012050
- [14] Mann J 1998 Probabilistic Engineering Mechanics 13(4) 269-82.
- [15] The MathWorks, Inc. (2021). Deep Learning Toolbox: User's Guide (r2018a), https://de.mathworks.com/products/deep-learning.html, last accessed July 2023.
- [16] Meyers J, Bottasso C L, Dykes K, Fleming P, Gebraad P, Giebel G, Göçmen T and van Wingerden J-W 2022 Wind Energy Science, 7, 2271–306.
- [17] NREL 2023, Floris version 3.5, https://github.com/NREL/floris
- [18] NREL 2021, OpenFAST version 3.0, https://github.com/OpenFAST/openfast
- [19] Mendez Reyes H, Kanev S, Doekemeijer B and van Wingerden J-W 2019 Wind Energy Science 4, 549-61.
- [20] Shaler K, Jasa J and Barter E 2022 J. Phys.: Conf. Series 2265 032095
- [21] Schreiber J, Bottasso C L and Bertelè M 2020 Wind Energy Science 5, 867–84.
- [22] Sutherland H J 1999: On the Fatigue Analysis of Wind Turbines. SAND99-0089. Albuquerque, NM: Sandia National Laboratories, doi:10.2172/9460.
- [23] Veers P et al. 2019 Science **366** eaau2027.