Journal of Physics: Conference Series

PAPER • **OPEN ACCESS**

Validation of new and existing methods for time-domain simulations of turbulence and loads

To cite this article: P Doubrawa *et al* 2024 *J. Phys.: Conf. Ser.* **2767** 052057

View the article online for updates and enhancements.

You may also like

- Mathematical Modeling in the Training of Future Mining Engineers
  E V Sergeeva and N A Ustselemova
- Method for determining dynamic coefficient of friction of bodies
  Yu P Borzilov, V N Yerovenko, D N Misirov et al.
- Indicators of the critical state of the ship's stability
  Yu Kochnev, E Ronnov and I Gulyaev

This content was downloaded from IP address 192.174.37.51 on 28/06/2024 at 21:32

Journal of Physics: Conference Series **2767** (2024) 052057 doi:10.1088/1742-6596/2767/5/052057

Validation of new and existing methods for time-domain simulations of turbulence and loads

P Doubrawa[1], A Rybchuk[1], J Friedrich[2], D Zalkind[1], P Bortolotti[1], S Letizia[1], R Thedin[1]

[1]National Renewable Energy Laboratory, Golden, CO United States
[2]ForWind, Institute of Physics, University of Oldenburg, Oldenburg, Germany

E-mail: paula.doubrawa@nrel.gov

Abstract. We seek to obtain a second-by-second match between the simulated and measured structural loads of a utility-scale wind turbine. To obtain the one-to-one load simulations, we start with the furthest upstream component of the modeling chain: the turbulent inflow. We consider new and existing methods to generate constrained-turbulence flow fields. The new method is based on large-eddy simulations (LES) and machine learning (ML). The existing methods include Kaimal-based TurbSim and the superstatistical wind field model. The inflow measurements used to constrain these simulations are obtained with a nacelle-mounted scanning lidar. We compare the flow fields for the different inflow simulation approaches and validate their associated load predictions against measurements collected in the Rotor Aerodynamics, Aeroelastics, and Wake (RAAW) field campaign. We find that the rotor-position control developed for this study is key in enabling the time match between measurements and simulations. When this control approach is used, the load simulation performance tracks with the inflow simulation fidelity, with LES+ML yielding errors $\leq 4\%$ for the damage-equivalent loads of flapwise bending moment, and tower fore-aft bending moments.

1 Introduction

It is extremely difficult to pinpoint the source of model errors when simulating large wind turbines embedded in a real atmosphere [1]. Some of this difficulty is related to the model chain complexity, from turbulent inflow through wind turbine response and control to wake morphology. Due to the coupling and sequencing of these models, the performance of a downstream component might contain a summation or cancellation of errors from the upstream model components, starting from the undisturbed inflow [2].

To improve the validation of aeroelastic models, it is therefore important to maximize the accuracy of the entire chain. In this work, we start by addressing the modeling of the furthest upstream component: the inflow. One way to improve inflow simulations is to rely on experimental measurements and force numerical outputs to match measured values at specific points, leaving the underlying model to fill in the gaps where observations are not available. Depending on the model and its application, this match might target different spatiotemporal scales: mesoscale-microscale coupling methods were developed for simulations of longer duration over larger domains [3], whereas the so-called constrained-turbulence approaches are used to produce flow fields with a spectral energy content that matches observations in at least one point for simulations of shorter duration in smaller domains [4, 5, 6, 7]. Regardless of the application, the underlying numerical model is responsible for filling in the space around the observed points, and the quality of the generated inflow is still bound to the limitations of the model itself.

Here, we compare the ability of different models to fill in the unobserved points around lidar measurements of the inflow of a utility-scale wind turbine. One novelty of this work is the introduction

Content from this work may be used under the terms of the Creative Commons Attribution 4.0 licence. Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.
Published under licence by IOP Publishing Ltd

of a reconstruction method based on large-eddy simulations (LES) and machine learning (ML). This novel method, which we refer to as LES+ML here, was primarily developed to serve as an observation-constrained boundary condition generator for high-fidelity, blade-resolved simulations. However, the inflows generated with the LES+ML technique can also be directly used to drive mid-fidelity aeroelastic simulations, as is done here. We compare the new inflow generation approach to a Kaimal-based reconstruction method [8] and to the superstatistical wind field model [7]. We assess the ability of each method to simulate measured loads on the wind turbine. The intercomparison of several turbulence reconstruction methods and their validation against measurements is also a unique asset of the work.

2 Data and Methodology

The data used for this study was collected during the Rotor Aerodynamics, Aeroelastics, and Wake (RAAW) field campaign conducted in 2023 in Lubbock, Texas [9]. The wind turbine being simulated is operating in isolation in flat terrain, has a rotor diameter of 127 m, and a hub height of 120 m.

2.1 Inflow Measurements

The period of study comprises 11 minutes and 44 seconds between 15:30:25 and 15:42:08 UTC (10:30:25 and 10:42:08 local time) on July 24, 2023. The time series duration is limited by the lidar measurements, which were set to cycle through four different scanning geometries every hour, leaving only < 15 minutes for the arc scans used here. Atmospheric conditions during this period are summarized in Table 1. The horizontal layout of the available inflow instruments relative to the mean nacelle yaw and hub-height wind direction is given in Fig. 1.

Table 1: Atmospheric conditions during the study period.

Quantity	Value	Unit	Height	Quantity	Value	Unit	Height
Wind speed	8.9	m s^{-1}	120 m	Rotor-layer shear	-0.45	m s^{-1}	183.5 m
Wind direction	220.7	deg	120 m	Rotor-layer veer	-5.1	deg	to 56.6 m
Turb. intensity	11.24	%	120 m	Heat flux	0.184	K m s^{-1}	2.5 m
Air density	1.03	kg m^{-3}	120 m	Inversion height	900	m	0 – 10 km

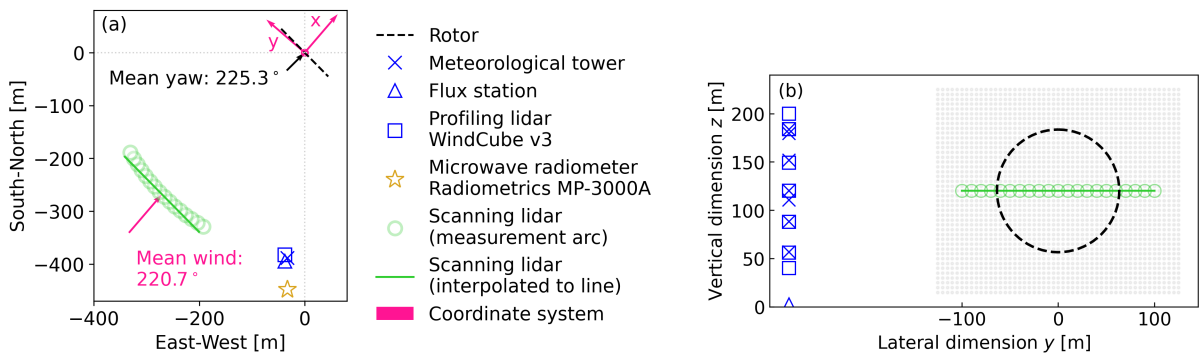

Figure 1: Layout of inflow instrumentation (a) in a horizontal plane relative to the mean wind and nacelle yaw; and (b) in a vertical plane, relative to the inflow simulation grid (dots).

Figure 1 also shows the scanning lidar (Halo Photonics Streamline XR) measurements that are used to constrain the inflow simulations: zero-elevation scans with 16 beams that span an azimuth arc of 30° with a resolution of 2°. These measurements undergo several processing steps to yield the velocities that are ultimately used as turbulence constraints (Fig. 2). First, the radial velocities are quality controlled using a dynamic filter [10]. Next, they are projected onto the center beam direction and thereby converted to longitudinal (u) velocities. The projection error standard deviation due to wind direction variability is estimated to range between $0.10 - 0.27$ m s^{-1} over this period, being proportional to the beam azimuth.

Lateral (v) and vertical (w) flow components are not considered. The temporally disjunct u measurements are then regularized in time (t) and upsampled to 1 Hz [11]. Next, the measured arc (green circles in Fig. 1) is interpolated onto a cartesian grid (green line in Fig. 1). At this point, several data gaps are still present due to blade passage in front of the lidar. These gaps are filled in using a diffusion model [12] that is trained on synthetic lidar scans derived from LES. We leverage the LES that was performed for the LES+ML inflow generation (Section 2.3.3). The diffusion model learns the relationship between a perfect scan and one that was intentionally corrupted to match the real-world measurements. As such, the algorithm is able to address blade-passage gaps, noise, and Gaussian smoothing effects with an associated uncertainty < 0.045 m s^{-1}.

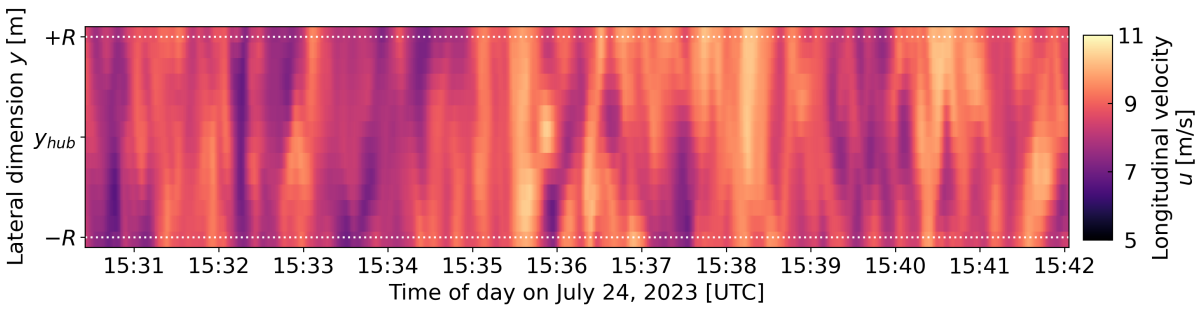

Figure 2: Lidar-derived longitudinal wind velocities $u(y, t)$ three rotor diameters upstream of the wind turbine at hub height. The dotted white lines mark the rotor span.

2.2 Wind Turbine Measurements

Operational and structural load measurements used in this study correspond to the period between 15:31:09 and 15:42:52 UTC. The offset of 44 s between the inflow and wind turbine measurements was added to account for the approximate advection time over the three rotor diameters (381 m) that separate the scanning lidar measurements and the rotor. The advection speed was taken as the laterally averaged, time-averaged lidar-derived $u = 8.66$ m s^{-1}.

The measured quantities used to constrain the wind turbine simulations include time series of blade pitch and azimuth position. While the wind turbine did yaw during the period of interest (Fig. 3), the simulated turbine does not perform any yaw maneuvers. This is justified by the fact that the inflow is constrained on the scanning lidar measurements, which are always aimed along the dominant wind direction because the lidar is mounted on the nacelle and yaws with it. The measured quantities used in the validation include near-root (at $\sim 2\%$ span) blade flapwise and edgewise bending moments, tower-base (at $\sim 7\%$ span) and tower-mid (at $\sim 46\%$ span) fore-aft and side-side bending moments. All of the wind turbine measurements used here were calibrated during the field campaign.

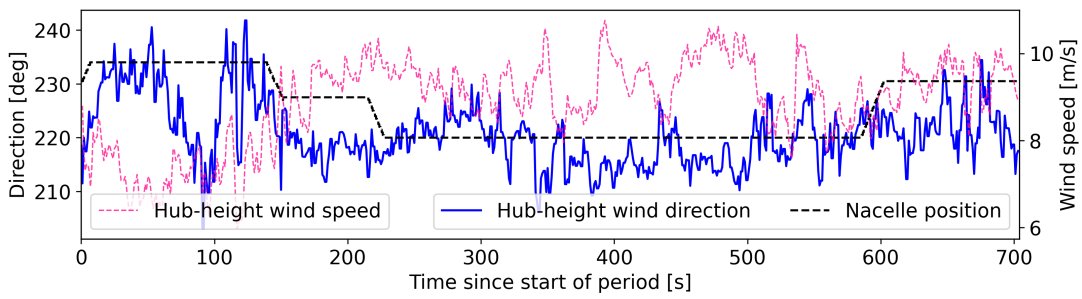

Figure 3: Time series of hub-height wind speed and direction as measured at the meteorological tower, and nacelle yaw position. Time is relative to the start of the period of interest: 15:30:25 UTC for the inflow and 15:31:09 UTC for the wind turbine.

2.3 Inflow Simulations

Time series of two-dimensional (lateral-vertical) flow planes are generated with three different methods. The extent of the planes is slightly different for each tool, but all of them follow the same grid spacing of 10 m in both directions, ensuring a point at the hub location. The time resolution is 1 s for all inflows, a hard limit imposed by the lidar measurements.

2.3.1 TurbSim Here, the turbulence constraint is enforced using the time series option for the turbulence model. The v and w velocity components are not simulated. The time-averaged wind speed profile is set to match the cup anemometer measurements collected at five heights across the rotor span at the meteorological (met.) tower. Davenport's exponential coherence model [13] is used to ensure a spatially coherent flow field. Instead of using default values, the coherence parameters are fit to *in-situ* wind speed measurements vertically separated by one rotor radius (from hub height to top tip). The fit considers a 4-hour window centered on the period of study. The best-fit values $a_K = 18.479$ and $b_K = 3.087 \times 10^{-4}$ m $^{-1}$ (see [14] for a description) represent a good fit for multiple vertical separation distances, ranging between 10 m and 127 m, which is the rotor diameter.

2.3.2 Superstatistical The superstatistical wind field model [7] was developed to reproduce the empirically observed anomalous, non-Gaussian wind field fluctuations with a superposition of Gaussian-distributed wind field statistics that are compliant with International Electrotechnical Commission (IEC) standards. The advantage of this model is that the statistical significance of extreme wind gusts measured from a met. mast or lidar, which common IEC wind field models greatly underestimate, can be precisely controlled by matching the so-called intermittency coefficient (μ) to measurements. Here, this and other model parameters (such as the correlation length, L and the Hurst coefficient, H) were estimated from an 18.5-hour time series of u as measured by an ultrasonic anemometer located on the met. tower (Fig. 1). The extent of this time series was sufficient to determine $\mu = 0.17 \pm 0.03$ from a fit of the flatness $\left\langle (\delta_\tau u)^4 \right\rangle / 3 \left\langle (\delta_\tau u)^2 \right\rangle \sim \tau^{-4\mu/9}$ of wind speed fluctuations $\delta_\tau u = u(t+\tau) - u(t)$. For a Gaussian wind field (e.g., Kaimal or Mann), the flatness is strictly 1, whereas the superstatistical wind field captures the measured increase in the occurrence of fluctuations $\delta_\tau u$ at small-scale separations τ. Furthermore, since the superstatistical wind field model is based on a scale mixture of Gaussian-distributed wind fields, it can also be constrained on real-world measurements. The u-component of the superstatistical wind field was constrained on the lidar measurement, whereas v and w components were left unconstrained with standard deviations matching the Kaimal model. As the measurements indicated no clear presence of wind shear, the latter was not incorporated in the u-profile. In contrast to integral length scales of wind field models in the IEC standard, which are typically assumed to be on the order of $10^1 - 10^2$ m, here we obtain a larger value of roughly 1.5 km.

2.3.3 Large-eddy simulation and machine learning In the LES+ML inflow generation approach, we fill in the unobserved points above, below, and off to the side of the lidar measurements using a combination of LES and machine learning. To do this, we first simulate the atmospheric boundary layer with an LES that approximately matches upwind hub-height wind speeds, turbulence intensity, rotor-disk shear, and surface heating. During this initial simulation, we sample the domain to match the layout of the real-world measurements. This sampling is performed repeatedly over the entire LES domain to maximize the amount of data that can be used to train the machine learning model. This sampling includes: a lateral line at hub height representing the lidar measurements and the two-dimensional plane that intersects the line representing the desired inflow plane. We train a diffusion model [12, 15] to generate planes from time series of the hub-height lines. More details on the model implementation can be found in a published repository [16]. The LES+ML approach generates inflows with all three correlated velocity components and air temperature, as these are available from the LES dataset. Temperature is not used for the present study. Once the ML model is trained, the real-world lidar measurements (Fig. 2) are used as input to the diffusion model, producing the lidar-constrained flow field.

2.4 Wind Turbine Simulations

The wind turbine model used here was implemented into OpenFAST version 3.5.1. The model was extensively tuned to reference data provided by the wind turbine manufacturer and subsequently validated against long-term measurements [9]. To enable a time-domain comparison, we use a rotor-position controller (RPC) that is implemented in ROSCO [17] version 2.9.0 [18]. This control feature enables the simulated blades to have the exact position and pitch angles of the blades in the field. The pitch time series are directly prescribed to OpenFAST, and the generator torque is controlled to minimize the

error in blade azimuth position. All simulations were also performed without RPC, using the ROSCO controller that was developed for this wind turbine and calibrated against extensive reference data. All simulations are performed for the 704 s of available inflow and employ the geometrically exact beam model BeamDyn for the blades, the reduced-order beam model ElastoDyn for the tower, the dynamic blade-element momentum theory with unsteady airfoil aerodynamics in AeroDyn15, and the control features from ServoDyn.

3 Results

3.1 Inflow

Time series of vertical profiles of longitudinal wind velocity are shown for the three simulation approaches in Fig. 4. Time series of lateral profiles are not shown, as they are identical for all simulation methods and follow the lidar constraint (Fig. 2). In Fig. 4, the flow fields are sampled at the lateral location of the wind turbine ($y = y_{hub} = 0$) and span slightly above and below the vertical extent of the rotor. Every method was set up to enforce the lidar-derived constraint at hub height, as marked by the white lines. The constraint led all simulation approaches to generally agree on when to speed up or slow down the flow, but the vertical extent of the turbulent structures turned out different for each tool. TurbSim produced the most vertical heterogeneity, and the superstatistical model the most vertical homogeneity. Some vertical heterogeneity is also present in the LES+ML approach in the form of vertical wind shear due to ground effects and due to the fact that it is not possible to enforce a specific shear value in the LES. Rather, the LES shear is a consequence of the lower boundary conditions: surface heat flux (0.184 K m s^{-1}, refer to Table 1) and roughness length (set to 0.1 m for this work).

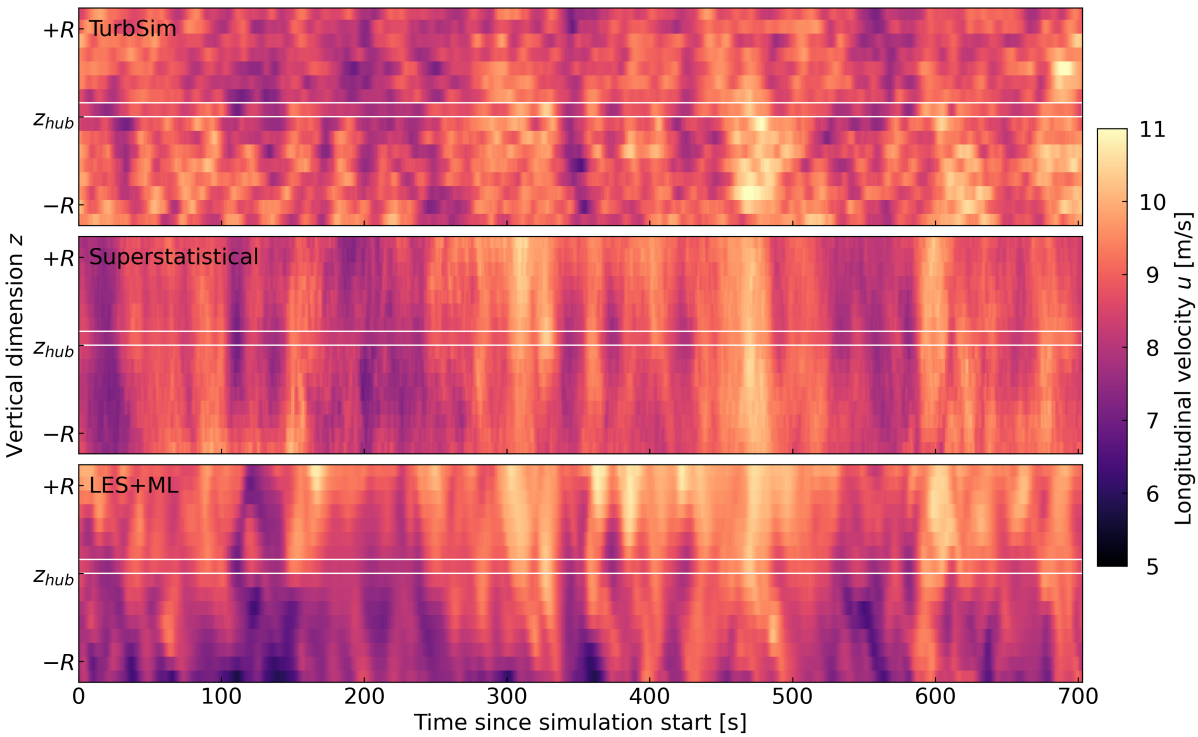

Figure 4: Time series of turbulent inflow generated with the three methods and sampled at the lateral location of the wind turbine over the vertical span of the rotor. The rectangle marks the area where the lidar constraint was enforced. The lidar data are not overlaid onto the figure.

To better visualize the differences in spatial heterogeneity, take for example the low-speed period between 200 s and 250 s in Fig. 4 and corresponding to 15:33:45 UTC – 15:34:35 UTC in Fig. 2. The lateral-vertical snapshots during this time (Fig. 5) further emphasize the differences between the inflow simulation approaches, revealing similar findings laterally to what was seen vertically: smoother flow fields for the superstatistical method, and the most spatial variability for TurbSim. The higher vertical

shear of LES+ML apparent in Figs. 4 and 5 is quantified in Fig. 6a: approximately 1.25 m s^{-1} across the rotor span. The other methods more closely match the negligible shear observed during the period. The measured shear of -0.45 m s^{-1} between rotor bottom and top is comparable to the standard error on the mean, which was estimated to be ~ 0.2 m s^{-1}. This error metric was obtained from $\sigma_u/\sqrt{N_{eff}}$ where σ_u is the standard deviation of u and $N_{eff} = T/(2T_u)$ is the effective number of samples, a function of time series duration ($T = 704$ s) and the flow timescale ($T_u \sim 16.5$ s at hub height) [19].

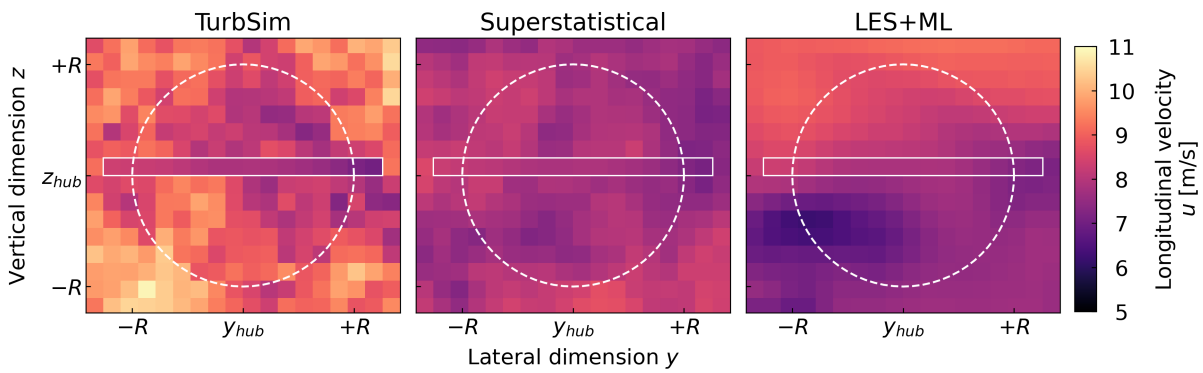

Figure 5: Snapshot of turbulent inflow generated with the three methods, shown 210 s after the simulation start (equivalent to 15:33:55 UTC in Fig. 2). The dashed line marks the rotor area and the rectangle marks the area where the lidar constraint was enforced. The lidar data are not overlaid onto the figure.

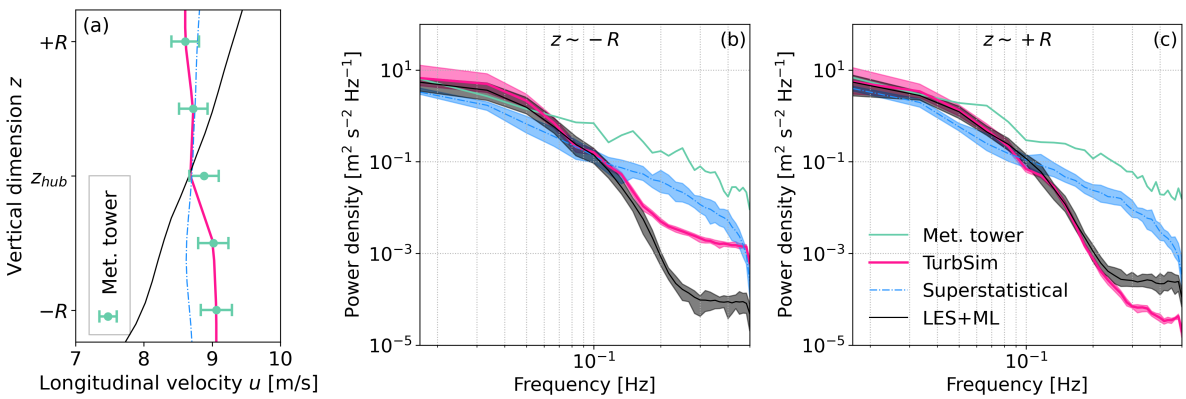

Figure 6: Wind shear (a) and power spectra (b-c) for met. tower anemometers and simulations. The shear is given by the average over time and the lateral dimension. The simulation spectra are computed for every lateral point spanning the rotor extent (15 points) considering the Welch method and 60-second segments. The shaded area shows the maxima and minima of the 15 spectra and the line marks the mean. The spectra are evaluated at the grid points closest to the rotor bottom (b) and top (c).

These spatial variability findings can be further confirmed with a spectral analysis of the time series of longitudinal velocities, shown in Figs. 6b and 6c for two heights: rotor bottom and top. The TurbSim spectra are different for the two heights, with substantially more high-frequency energy below than above. The spectral energy content of the other two methods does not vary as much across the vertical extent of the rotor disk, which is also true for the measurements. Note that the measurements in Fig. 6 are obtained at the met. tower because lidar scans were limited to hub height. The superstatistical inflow yielded the lowest spectral power content for frequencies $f < 0.1$ Hz, but produced the best match to

measurements above that frequency. The LES+ML and TurbSim show a better match to measurements at low frequencies but drop off near 0.05 Hz. Spectra at hub height are not shown, as they are identical for all methods due to the enforced constraint.

3.2 Wind turbine response

The lateral and vertical velocities are not shown but are different depending on the simulation approach. TurbSim does not simulate these components, as they are not part of the constraint. The superstatistical method produces v and w fields that are uncorrelated with u. LES+ML produces flow fields that are physically consistent for all three components, because the full velocity vector was considered when training the ML model. These differences in the wind vector affect the rotor aerodynamics, leading to rotor inflow skew angles that range $0° - 20°$ in LES+ML and $2° - 13°$ for the superstatistical inflow, while remaining constant near the shaft tilt angle in TurbSim.

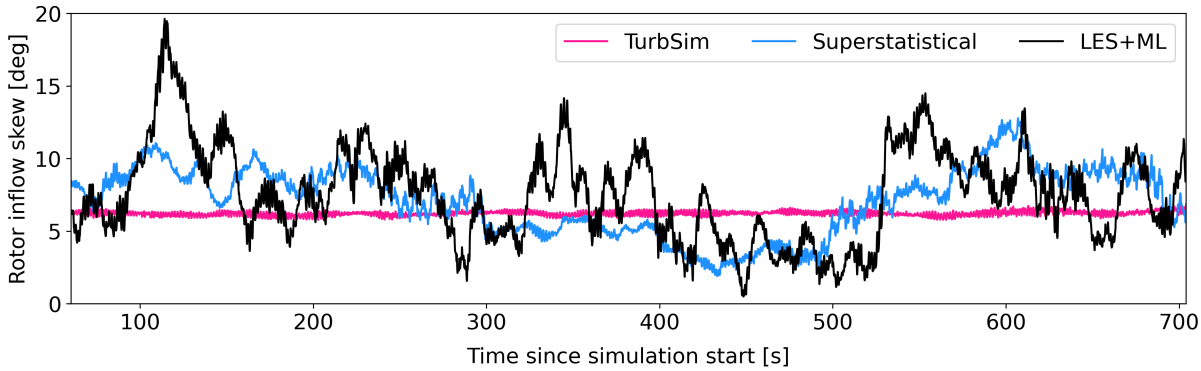

Figure 7: Time series of inflow skew angle at the rotor, a quantity computed by AeroDyn during the OpenFAST simulations.

The aerodynamic differences, in turn, affect the fore-aft structural response of the tower and blades. The remainder of this section considers statistics, spectra, and time series of structural load signals. The loads analysis discards the first 60 seconds of simulation, leaving 644 s for analysis. All results are normalized to protect the proprietary nature of the data. To obtain an equivalence between the measurements and simulations, the simulated bending moments at discrete tower (from ElastoDyn) and blade (from BeamDyn) nodes are linearly interpolated to the location of the strain gage measurements. The blade moment statistics are given as an average of the three blades. The damage-equivalent load (DEL) calculations consider a Wöhler exponent of 10 for the blades and 4 for the tower. The power spectra uses 120-second segments and the Welch method with a Hanning window.

Load statistics normalized by the measured values are given in Fig. 8. These results indicate that the means (Fig. 8a) are less sensitive to inflow differences, with the ratios for a given quantity staying approximately constant as the inflow changes. The mean errors stay below 10%, with the simulations overestimating the mean load for the tower and underestimating it for the blades. The mean was also not very sensitive to the control strategy used. When using the RPC, the mean loads were consistently higher by a few percent: $2\% - 4\%$ higher across inflow strategies for blade flap moment, and $1\% - 3\%$ for tower fore-aft moment.

The standard deviations were more sensitive to inflow and controller approach (Fig. 8b). The use of RPC considerably reduces the load standard deviations. For the superstatistical and LES+ML inflows, this reduction brings the ratios closer to unity. For TurbSim, it does the opposite, leading to a decrease in performance in terms of statistics. For the simulations with RPC, the error in load fluctuations tracks with the level of fidelity in the model, with TurbSim yielding the largest errors in the form of underestimations. The superstatistical and LES+ML results are similar, with ratios of 1.0 for blade flap and 1.1 for the tower fore-aft moments. The DEL results are similar to the standard deviation results in that RPC model performance tracks with model fidelity, with LES+ML showing a ratio of ~ 1 for all load signals considered. Hereinafter, we only consider results from the simulations with the RPC.

Load spectra (Fig. 9) indicate that the lower load variability of TurbSim is likely related to the inflow, by revealing lower values at low frequencies for this simulation approach. Note that TurbSim did not show an underestimation of spectral power in the wind speed (Figs. 6b and 6c). Therefore,

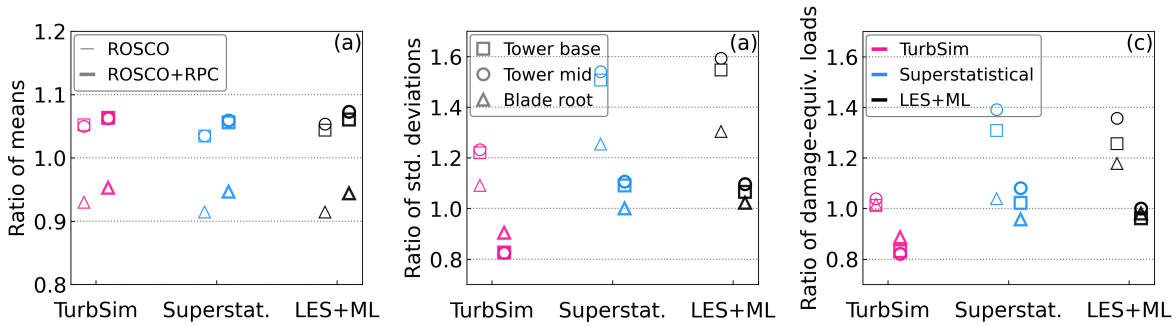

Figure 8: Ratio between simulated and measured fore-aft load statistics: mean (a), standard deviation (b) and damage-equivalent load (c). In each subplot, ratios are given for three quantities (marker shapes), three inflows (marker colors) and two controller approaches (marker thicknesses).

this load variability underestimation by TurbSim is likely related to the lack of lateral and vertical flow components in this inflow. Starting at the frequency of one blade passage (P), TurbSim matches the mesurements well, capturing the frequency and magnitude of spectral peaks better than the other inflows. The superstatistical inflow yielded the opposite results from TurbSim for the three quantities considered: a good match for frequencies < 1P, and a growing underestimation of spectral power for higher frequencies. Consistent with the statistical analysis of Fig. 8, the LES+ML results show the best performance here: a close match to measurements at low frequencies that is maintained until 5P for the blades and 3P for the tower.

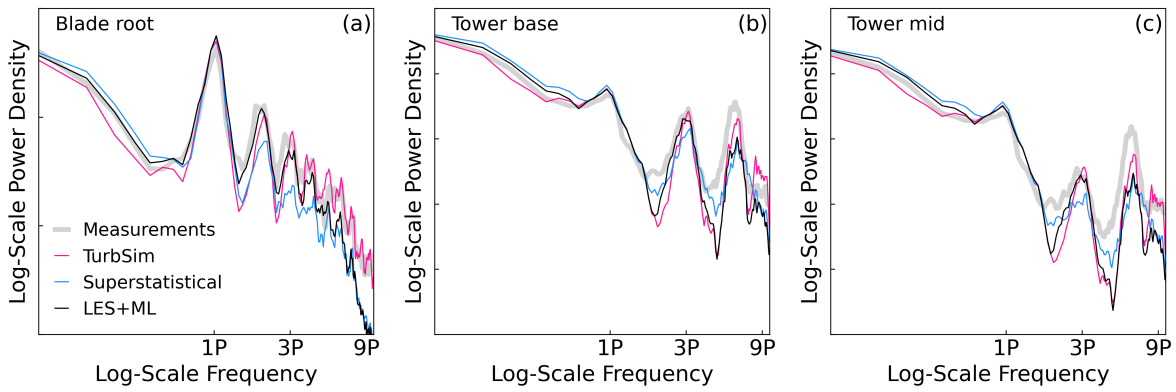

Figure 9: Power spectra of near-root blade flapwise bending moments (a), tower base (b) and tower mid (c) fore-aft bending moments. The tower (b and c) spectra are on the same axes.

4 Discussion

The main novelty of this work is the level of detail that was dedicated to reproducing in a simulation environment the second-by-second structural response of a multimegawatt wind turbine. This effort addresses the need for a comparison procedure in the time domain between calculated and measured properties under turbulent inflow, which was identified by experts working on International Energy Agency Wind Task 29 [20]. This undertaking can be separated into four main thrusts: (i) a comprehensive field campaign with a strong emphasis on inflow characterization, the development of (ii) computational methods to assimilate the measured inflow into a high-fidelity simulation environment, of (iii) a robust, validated wind turbine model, and of (iv) an RPC capability that enables precise time-space matching between the real-world and simulated blades (Fig. 10). While the work presented here required the execution of

these four efforts, the objective of the analysis was to assess the value of (ii), the development of data assimilation techniques for high-fidelity simulations. A portion of the results also quantified the added benefit of (iv). Note that the motivation behind the development of these methods goes beyond the work presented here: It is meant to enable time-resolved validation of high-fidelity wind turbine simulations. Here, we do not undertake any high-fidelity wind turbine simulations. Rather, we focus on the precursor to these simulations: the high-fidelity, turbulence-constrained inflow. We compare such inflows to those generated with existing, lower-fidelity approaches and to measurements.

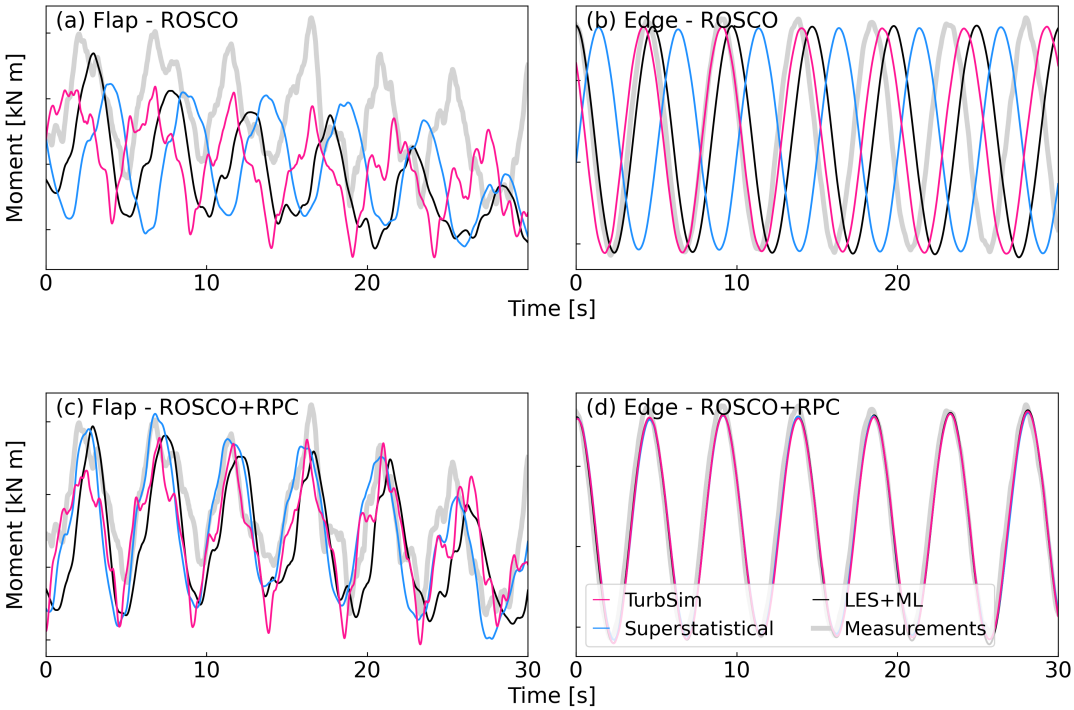

Figure 10: Zoom into time series of near-root flapwise (a,c) and edgewise (b,d) blade bending moments over 30 seconds for measurements and simulations with ROSCO (a,b) and ROSCO+RPC (c,d). Shown for only one blade.

The main takeaway from the work performed here is that using an RPC and higher fidelity inflows substantially improves the accuracy of the simulated structural loads. When an RPC is used and the inflow stays the same, the fore-aft load variability decreases substantially. This decrease represents an improvement in the simulations with higher-fidelity inflows (superstatistical and LES+ML) but increases the error magnitudes for the low-fidelity inflow (TurbSim). The greatest improvement was seen for the standard deviation of LES+ML, where RPC alone reduced the errors from 60% to less than 10%. When an RPC is used, we see that model accuracy tracks with inflow fidelity for standard deviations and DELs, with LES+ML yielding the best performance for the quantities considered here. While the load averages were also considered, they were not very sensitive to controller strategy or inflow choice.

We conclude that the techniques developed to produce high-fidelity inflows constrained on field measurements work well: They reproduce the constraint and fill in the unobserved space with physically consistent flows. However, we keep in mind that the simulated flow field can only be as good as the inflow measurements themselves. The limited spatial (10 m) and temporal (1 Hz) resolution of the lidar measurements available for this study impose a hard limit on how faithful to the real world the simulated inflow can be.

To further expand on the differences between the inflow simulation approaches and pinpoint areas for model improvement, the analysis presented here needs to be expanded. Further work will consider: (i) more cases with different atmospheric inflow conditions; (ii) multiple seeds in the inflow generation; (iii) structural load sensitivity to the spatiotemporal resolution of the inflow; (iv) other existing turbulence-

constraint inflow generation techniques; (v) more load quantities such as edgewise blade loads, side-side tower loads, and mainshaft loads for which measurements are available. In the meantime, we envision that the case presented here becomes a model validation benchmark for single-turbine turbulent inflow generation and loads simulation.

References

[1] Boorsma K, Schepers G, Aagard Madsen H *et al.* 2023 *Wind Energy Science* **8** 211–230

[2] Doubrawa P, Quon Eliot W E, Martinez-Tossas L A *et al.* 2020 *Wind Energy* **23** 2027–2055

[3] Allaerts D, Quon E and Churchfield M 2023 *Wind Energy* **26** 469–492

[4] Dimitrov N, Borraccino A, Peña A *et al.* 2019 *Wind Energy* **22** 1512–1533

[5] Rinker J M 2020 *Journal of Physics: Conference Series* **1618** 052053 ISSN 1742-6596

[6] Conti D, Pettas V, Dimitrov N *et al.* 2021 *Wind Energy Science* **6** 841–866 ISSN 2366-7443

[7] Friedrich J, Moreno D, Sinhuber M, Wächter M *et al.* 2022 *PRX Energy* **1** 023006

[8] Veers P S 1988 Three-dimensional wind simulation Tech. Rep. SAND-88-0152C Sandia National Laboratories

[9] Brown K, Bortolotti P, Branlard E *et al.* 2024 *Wind Energy Science Discussions* **2024** 1–28

[10] Beck H and Kühn M 2017 *Remote Sensing* **9** 561 ISSN 2072-4292

[11] Beck H and Kühn M 2019 *Remote Sensing* **11** 867

[12] Song Y, Sohl-Dickstein J, Kingma D P *et al.* 2021 *arXiv preprint* 2011.13456

[13] Davenport A G 1961 *Quarterly Journal of the Royal Meteorological Society* **87** 194–211 ISSN 1477-870X

[14] Jonkman B 2016 Turbsim user's guide v2 unpublished report URL `https://www.nrel.gov/wind/nwtc/assets/downloads/TurbSim/TurbSim_v2.00.pdf`

[15] Rybchuk A, Hassanaly M, Hamilton N *et al.* 2023 *Physics of Fluids* **35**

[16] Rybchuk A 2024 GitHub repo: Latent Diffusion Models for Atmospheric Large Eddy Simulation Online URL `https://github.com/rybchuk/latent-diffusion-3d-atmospheric-boundary-layer`

[17] Abbas N J, Zalkind D S, Pao L *et al.* 2022 *Wind Energy Science* **7** 53–73

[18] Zalkind D 2024 GitHub repo: ROSCO v2.9.0 Online URL `https://github.com/NREL/ROSCO/releases/tag/v2.9.0`

[19] George W K, Beuther P D and Lumley J L 1978 *Proceedings of the Dynamic Flow Conference 1978 on Dynamic Measurements in Unsteady Flows* pp 757–800

[20] Schepers J G, Boorsma K, Madsen H A *et al.* 2021 IEA Wind TCP Task 29, Phase IV: Detailed Aerodynamics of Wind Turbines Tech. rep. URL `https://zenodo.org/records/4813068`

Acknowledgments

This work was authored in part by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding provided by the U.S. Department of Energy Office of Energy Efficiency and Renewable Energy Wind Energy Technologies Office. Support for the work was also provided by GE Renewable Energy under CRADA 21-18140. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.