Response to Reviewers' comments on
"Investigating wake reproduction of a model-scale wind turbine: experimental measurements versus Large Eddy Simulation with actuator line"
Wind Energy Science, WES-2025-242

Emmanuel Gillyns, on behalf of all authors

May 27, 2026

We thank the editor and both reviewers for their time and valuable feedback regarding our manuscript, which helped improve the quality of our work. Please find below all comments addressed. For completeness, we have also attached a track-of-changes document.

Response to Reviewer 1 (R1)

R1: General comments

The authors present a LES framework (Nek5000) for a model-scale wind turbine in an atmospheric boundary-layer wind tunnel, with two main methodological components: (i) a recycling–inflow strategy augmented by a controller-based volume forcing, and (ii) an extended actuator line method (ALM) that represents blades, tower, and nacelle. The manuscript is well written, and the documentation of the experimental setup, LES validation, and uncertainty assessment is valuable for reproducibility. The authors also provide quantitative validation of wake profiles at two downstream locations. Overall, the work is technically competent and scientifically sound. However, several central claims require stronger supporting evidence before the work can be recommended for publication.
(Reviewer 1 comments citation: `https://doi.org/10.5194/wes-2025-242-RC1`)

We thank Reviewer 1 for their thorough and constructive comments, which helped improve the clarity and rigor of the manuscript. We appreciate the recognition of the manuscript's strengths (clear writing, thorough documentation, and quantitative validation) and have addressed the concerns about several central claims. These changes are reflected in the revised manuscript and detailed in our point-by-point response hereafter.

R1: Specific comment 1

The manuscript attributes improved near-wake gradients and enhanced tip-vortex representation to specific implementation choices, including local force evaluation at every grid point, avoiding radial smoothing, and disabling the optional tip-spreading correction. These decisions are physically defensible, but the current presentation does not provide a baseline comparison that isolates their incremental contribution. To substantiate the novelty and quantify the benefit, the paper should include a controlled ablation study (e.g., standard ALM versus the proposed extended ALM), reporting the impact on clearly defined wake metrics (velocity deficit, shear gradients, turbulence intensity, and spectra). Without such comparisons, the improvements remain qualitative and cannot be unambiguously attributed to the proposed modifications.

Response:

We thank the reviewer for recognizing the physical basis of our implementation. It is worth noting that avoiding radial averaging has been explored with satisfactory results in prior work [1, 2]. A baseline comparison between a standard ALM and this continuous implementation would indeed improve the completeness of this work. As described in Section 3.3, this implementation allows the use of a tip spread, which behaves similarly to a 3d-gaussian kernel, used in standard ALMs. Using a tip spread similar to the 2d-gaussian kernel used along the actuator line, this can reproduce the effect of a standard ALM in the tip region. Although it cannot be compared on the entire profile, we believe that this serves as a reasonable way to quantify the added benefit of the continuous implementation. We compared the simulation at +3 AoA with and without the tip spread. The vertical shear (derivative of the streamwise velocity with respect to the vertical position) at the tip is compared between the two and we observe that the proposed method captures a steeper gradient, which is closer to the experimental results. The vertical shear with tip spread is $15\,\mathrm{s}^{-1}$, and without tip spread it is $47\,\mathrm{s}^{-1}$. This clear increase in gradient demonstrates that reduced radial smoothing in our approach drives the improved near-wake gradients and enhanced tip-vortex representation. For completeness, the vertical shear of both runs are provided in Figure 1. However, to avoid diluting the main message of the paper, this graph was not added to the revised manuscript, but the values of this comparison are included in the text in Section 4, line 453.

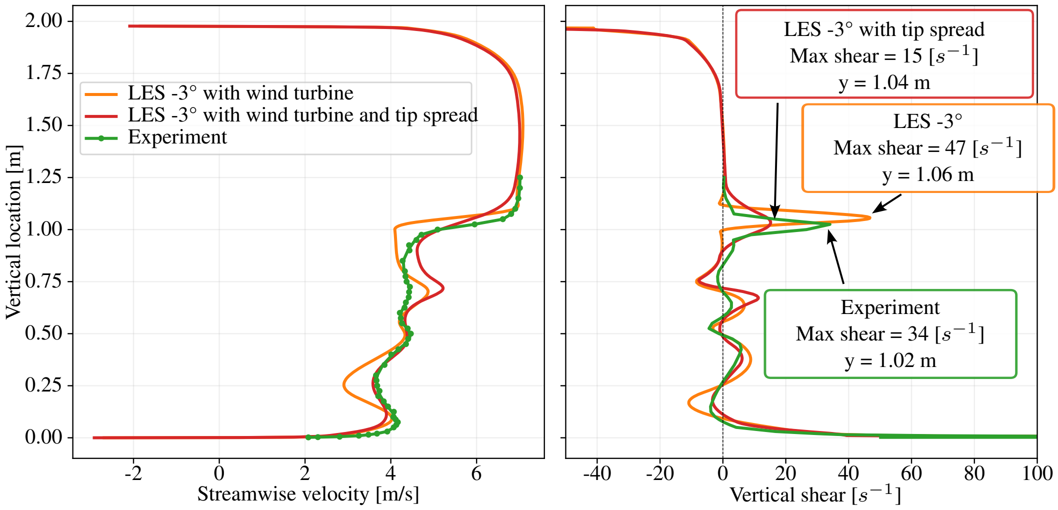

Figure 1: Shear profile with and without tip spreading.

R1: Specific comment 2

The nacelle drag coefficient is explicitly adjusted empirically until the wake appears to match observations. This compromises the independence of validation, especially for near-wake structure where nacelle effects are non-negligible. Please discuss a sensitivity study showing impact of nacelle Cd on wake metrics.

Response:

We appreciate the reviewer's observation. We agree that calibrating a model component to impose agreement with measurements can undermine the independence of validation. It is important to note, however, that the comparisons reported in the manuscript address the wake generated by the blade-swept area rather than the nacelle effects on the wake.

During preliminary testing we performed simulations with nacelle drag deactivated (effectively representing unobstructed, "floating" blades, which is what is used in standard ALMs). In that configuration the perturbation attributable to the absence of the nacelle remained confined to a narrow vertical band between 0.52 m and 0.65 m at location 2 (1.41D downstream), and the resulting difference in streamwise velocity did not exceed $0.15\ \mathrm{ms}^{-1}$. Given the limited spatial extent and small magnitude of this effect, we judged it unlikely to influence the principal blade-generated wake features and the principal conclusions of the study.

Nevertheless, to avoid the potential for a localized disturbance to propagate into the wake further downstream, we adjusted the nacelle drag coefficient (Cd) so that the modeled nacelle-induced blockage better matches the experiment for the presented configurations. This adjustment was made solely to prevent an artificial downstream influence caused by the unresolved nacelle geometry and does not alter the principal findings concerning the blade-swept wake. The adjusted value of this drag coefficient is $Cd = 2$, which appears realistic for such bluff body without a streamline casing.

The text has been revised to reflect this in Section 3.3, line 416.

R1: Specific comment 3

The controller-based forcing reduces TI MAPE below $0.75\ \mathrm{m}$ and argues that imposed perturbations have dissipated by the recycling plane. Matching TI does not guarantee matching spectra, integral length scales, anisotropy, or coherence, all of which affect wake recovery/meandering. Since downstream errors increase, we cannot determine whether residual discrepancies are dominated by inflow-structure mismatch or turbine forcing/model assumptions. Please address this point thoroughly, with clear justification and supporting evidence.

Response:

We agree with the reviewer that matching turbulence intensity (TI) alone does not ensure agreement in higher-order inflow characteristics such as spectral distribution, integral length scales, anisotropy, or coherence, all of which can influence wake recovery and meandering. This aspect has been carefully considered and is now clarified in the revised manuscript in Section 3.2, line 310.

The applied perturbation is constructed as a sum of sine waves with wavelength on the order of the integral length scale of the flow. These waves are defined with dependence on the three spatial positions (x, y, and z), and in each of them, the base wavelength is altered in a slightly randomized way. This results in a total of 30 closely spaced but distinct frequencies, which broadens the spectral content of the perturbation and avoids artificial concentration of energy at discrete modes. This design intentionally injects energy at relatively large scales, which then naturally redistributes across the spectrum through the turbulent cascade. The use of multiple sine components, rather than a single mode, avoids introducing a distinct spectral peak that would otherwise require a longer development length to blend into the background turbulence.

In addition, the phases of the sine waves evolve in time, promoting enhanced mixing and decorrelation. As a result, by the time the flow is recycled back, the perturbation exhibits a more randomized character. Spectral analysis at the recycling plane confirms that the initially imposed discrete peaks are no longer identifiable and are fully blended into the broadband turbulent spectrum.

It is also important to note that the forcing introduces only a limited amount of additional energy relative to the already well-developed turbulence in the flow. This helps ensure that the perturbation acts as a mild correction rather than a dominant structural modification of the inflow. Based on these observations, we believe that residual discrepancies downstream are unlikely to be primarily driven by persistent inflow spectral artifacts from the forcing.

To further support this point, we have added a power spectrum density of the velocity magnitude at a streamwise distance of 5.4 m (at the recycling plane), shown in Figure 2. In this figure, the power spectrum is shown both using a fast fourrier transform (FFT), and using the Welch method which is averaging windowed FFT. The imposed perturbations correspond to a wavenumber range of approximately $k = 0.34\text{–}3.4\ \mathrm{m}^{-1}$, which, using Taylor's frozen turbulence hypothesis, translates to an equivalent frequency range of approximately 2–20 Hz. The resulting spectrum shows no distinct peaks within this range, confirming that the injected energy is fully embedded within the broadband turbulent spectrum.

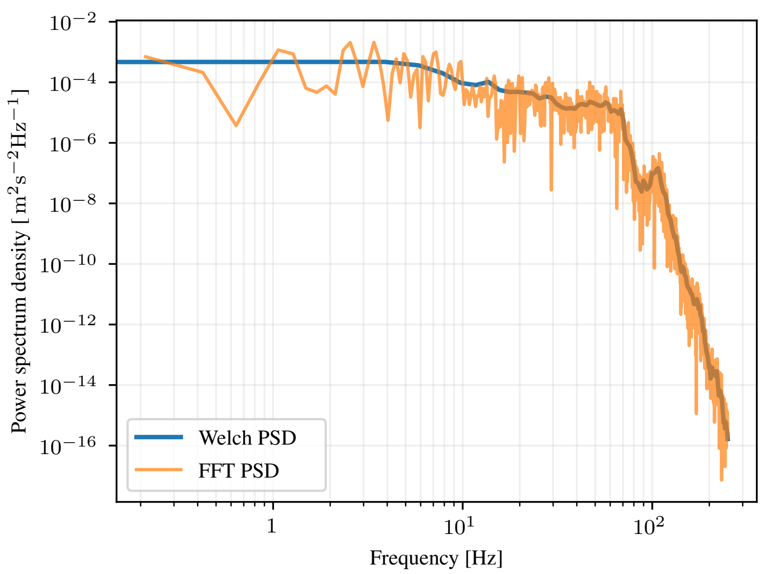

Figure 2: Power spectrum density of the wind velocity magnitude, in the center of the cross section at a streamwise distance of 5.4 m.

R1: Specific comment 4

The manuscript compares $+1°$ and $-3°$ pitch cases and notes different agreements near tip vs root. If the nacelle Cd is tuned and the root region is known to be problematic for ALM, then pitch-induced differences may be difficult to isolate. Please quantify the pitch sensitivity to demonstrate the extent to which the reported error is dominated by a region where the model assumptions are known to be invalid.

Response:

Thank you for this insightful comment. We agree that isolating pitch-induced effects is important, particularly given the known limitations of the ALM in the root region and the tuning of the nacelle drag coefficient.

In the present configuration, we have analyzed the spatial influence of the different modeling components at location 2 (1.41D downstream). We observe that the nacelle primarily affects the wake in a limited region between $0.52\,\mathrm{m}$ and $0.65\,\mathrm{m}$. Disabling completely the nacelle drag has a minor effect on the velocity deficit in this region of less than $0.15\ \mathrm{ms}^{-1}$. The blade root region, where ALM assumptions are known to be less reliable, influences the flow between $0.65\,\mathrm{m}$ and $0.75\,\mathrm{m}$ as well as between $0.42\,\mathrm{m}$ and $0.52\,\mathrm{m}$. The remainder of the swept area lies outside these regions and is governed by the actuator line formulation in a regime where it is generally considered valid.

Therefore, the differences observed in the remaining portion of the swept area can be attributed to variations in blade angle of attack induced by the pitch change. This allows us to isolate the pitch sensitivity in regions where the model assumptions are expected to hold, and confirms that the reported discrepancies are not dominated by the nacelle or root-affected zones.

This has been refined in the text in Section 3.2 and Section 6.

R1: Specific comment 5

Blade deformation mismatch should be connected to wake interpretation more directly. Since deformation affects angle of attack and thus loads, the wake validation may be partially right for the wrong reasons. Please elaborate further on this point.

Response:

We agree that the relationship between blade deformation and wake interpretation should be made more explicit. Blade deformation leads to local changes in the angle of attack along the span, which affects the aerodynamic loading distribution on the blade. Since the wake development is driven by the integrated effect of these loads, deformation can therefore influence the resulting wake characteristics, including velocity deficit and wake shape. As a result, the observed agreement in the wake profile may be partly influenced by the way deformation alters the loading distribution, and should therefore be interpreted with care when assessing the validity of the underlying aerodynamic model.

We have revised the manuscript to clarify this point and to better highlight the potential influence of blade deformation, as well as other numerical factors (such as local velocity sampling resolution, the use of two-dimensional airfoil polar data, and numerical dissipation introduced by the sub-grid scale model) on the wake comparison in Section 6.

R1: Technical comment 1

Replace “looses its coherence” with “loses its coherence.”

Response:

Thank you for pointing it out, it is now corrected in the text.

R1: Technical comment 2

Figure annotations contain spelling errors: “Settling chanmber” $\rightarrow$ “Settling chamber”; “Font turntable” $\rightarrow$ “Front turntable”; “Substract” $\rightarrow$ “Subtract.”

Response:

Thank you for pointing these out, they are now corrected in the figures.

R1: Technical comment 3

Correct “balde mean deformation” $\rightarrow$ “blade mean deformation.”

Response:

This is also now corrected in the text.

R1: Technical comment 4

Standardize unit formatting throughout (e.g., "ms-1" vs "m s-1") and ensure consistent usage of D (rotor diameter) vs "ø" notation in figure labels/captions.

Response:

These formatting discrepancies are now corrected and consistent in the text and figures.

Response to Reviewer 2 (R2)

R2: Overview

The manuscripts presents a combined experimental and numerical approach to the description of a wind turbine wake. In particular, experiments performed with ta turbulent boundary layer are used to validate an actuator line model of the same turbine. The results are promising in terms of overall agreement between the two approaches. The use of a high-order spectral method to solve the N-S equations coupled with an actuator line model is an interesting prospect. However, the manuscript lacks crucial details that significantly hamper it's long term impact. Some comments below. (Reviewer 2 comments citation: https://doi.org/10.5194/wes-2025-242-RC2)

We thank the reviewer for the overall positive assessment of our work and for recognizing the interest of combining experimental data with a high-order spectral numerical approach coupled with an actuator line model. We are pleased that the reviewer finds the results promising in terms of agreement between the experimental and numerical approaches.

We also appreciate the concern raised regarding missing methodological details that could affect the long-term impact and reproducibility of the study. In response, we have carefully revised the manuscript to address all specific comments provided below. In particular, we have added clarifications and additional details on the numerical setup, model implementation, and experimental configuration, with the goal of improving the completeness and clarity of the presentation.

We believe these revisions significantly strengthen the manuscript and improve its suitability for publication.

R2: Specific comment 1

The main concern is the scarce amount of details provided regarding the numerical setup and the corresponding lack of any independence study. This is a crucial aspect as a CFD code that uses high-order spectral decomposition rather than "traditional" finite difference is used. "standard" best practices for actuator line models require 80-120 spanwise elements per blade. The kernel size is typically refined based on how many elements are used, with finer grids allowing for smaller and more realistic projection kernels. This ammount of discretizaion may still not be enough to correctly resolve the tip vortices (see: https://doi.org/10.1016/j.compfluid.2024.106477). Based on the information provided in section 3.1 a lot less elements were used to discretize the rotor area. While not directly comparable, authors performing ALM LES on scaled models using finite difference used orders of magnitude more cells (https://iopscience.iop.org/article/10.1088/1742-6596/2767/5/052050, https://doi.org/10.5194/wes-10-1707-2025). Especially because this is, in my understanding, the first time this ALM setup is presented, the quality of the results is hard to judge without grid independence and a comparison to the established literature, even if a different method to solve the N-S equations is used.

Response:

We thank the reviewer for highlighting the need to provide further detail on the numerical setup and the resulting difficulty in fully assessing solution quality in the absence of a formal independence study, as well as for the broader comments regarding the discretization level used in the actuator line model.

In the case presented in this paper, the rotor is discretized using 9 spanwise elements per blade (ignoring the localized tip refinement). Given the seventh-order polynomial representation within each element, this corresponds to 63 degrees of freedom along each blade, which provides a more meaningful measure of resolution than the number of elements alone when comparing to classical discretization approaches, and can be directly compared to the number of cells in a finite difference code. When including the localized tip refinement, the total increases to 77 degrees of freedom.

Regarding the spreading used in the ALM, the guideline value for the kernel width applies to a configuration without tip refinement. In this case, the corresponding dimensionless parameter is given by $\sigma/l = 2.5/(9 \times 7) = 0.040$, with σ referring to the kernel width, l the blade length, 9 the number of elements (ignoring the tip refinement), and 7 the polynomial order.

In addition, the velocity sampling in the ALM is performed at each nodal point rather than at the center of the actuator line. As a result, the velocity used for force evaluation is obtained independently at every nodal point, i.e., at every degree of freedom within each element. Combined with the continuous kernel representation, this ensures that the force computation remains locally resolved at the nodal level, while also maintaining a formulation that is well-suited for parallel execution and computational efficiency.

Regarding grid independence, byproduct of the spectral element method can be used as indicator from the behavior of fluctuating quantities at element interfaces. In Nek5000, the solution is continuous at the interface between each element, but its derivative is not. Insufficient spatial resolution typically manifests as Gibbs-type oscillations, appearing as localized spikes at element boundaries. This effect has been explored in a previous publication [3]. We have therefore examined the fluctuating components of the velocity field in the wake on a plane located at $1.41, D$ downstream of the turbine, where wake development and tip-vortex structures are most sensitive to potential under-resolution effects. In this region, only minimal Gibbs-type oscillations are observed at element interfaces. The grid can therefore be considered sufficiently refined, and do not affect the resolved wake dynamics.

These clarifications have been included in the text on Sections 3.1 and 3.3.

R2: Specific comment 2

A second major issue is the lack of a comparison between experiment and simulations in total thrust or power. Similar wake characteristics can arise because of different levels of power extraction.

Response:

We thank the reviewer for this valid and important point.

Regarding the power comparison, the power measured experimentally corresponds to the electrical power on the DC bus, whereas the numerical power from the simulation represents the aerodynamic power extracted from the flow. In the experiment, there are several sources of power loss between the aerodynamic input and the measured electrical output, including mechanical losses in the bearings and bevel gear, as well as the generator operating at significantly lower voltage and rotational speed than its rated conditions, which leads to reduced efficiency. Additional losses also arise from the AC-to-DC conversion stage which operates at low voltage. For these reasons, a direct comparison between experimental electrical power and simulated aerodynamic power can only be considered meaningful in an approximate sense, for which it is matching.

Regarding the thrust, the total thrust in both the experiment and the LES simulation can be inferred from the integrated velocity deficit over the wake, as shown in Figure 19. Both cases exhibit a similar integrated deficit, suggesting that the momentum extracted from the flow (and thus the effective thrust) is comparable between the experiment and the numerical simulation. As a further improvement, a more direct measurement of thrust would be possible by instrumenting the turbine support with strain gauges. This would allow a direct comparison between measured and simulated thrust; however, this instrumentation is planned for a future experimental campaign and lies outside the scope of the present study.

This discussion has been included in Section 6.

R2: Specific comment 3

The scope of the manuscript does not appear to be crystal clear. If the focus is on validation why does the pitch angle in simulations not match the one in experiments? On the other hand if the scope is to study the influence of pitch angle, the gap that the manuscript is trying to fill should be better motivated. Moreover, experiments with a full rotor model and ABL are not easy to come across. This could be a major selling point for the study if better explained.

Response:

We thank the reviewer for this constructive comment. We appreciate the opportunity to clarify the scope of the study, as we agree that this was not sufficiently explicit in the original manuscript.

To address this, the abstract and conclusion (Section 7) have been revised to better highlight that the objective of this work is not strict validation, but rather a combined experimental–numerical investigation aimed at providing a high-resolution experimental dataset of a model-scale wind turbine under controlled inflow conditions, and assessing the capability of an improved ALM to reproduce key near-wake features and their trends.

Regarding the pitch angle, we would like to clarify that the simulations are not limited to a direct one-to-one reproduction of the experimental operating point. Instead, they also include a sensitivity analysis to investigate the influence of operating conditions on wake development and to assess whether the model correctly captures these trends. This point has been clarified in the revised manuscript to avoid ambiguity.

We also fully agree with the reviewer that experiments involving a full rotor model under controlled inflow conditions are relatively scarce. This aspect has now been more clearly emphasized in the manuscript as a key strength of the study, particularly in relation to the TWIST dataset and its relevance for future model development and assessment.

R2: Specific comment 4

No details regarding how the airfoil coefficients used in the numerical model are derived are provided beyond Figure 17. Crucial details are missing such as the Reynolds number and the Ncrit that was used in XFoil. The NCrit value influences where transition occurs, which in turn can influence the airfoils performance significantly at low Reynolds numbers. I would expect the Ncrit value to be tuned according to the TI level in the wind tunnel

Response:

In our study, the airfoil polars were generated using a Reynolds number of $Re = 5 \times 10^4$, based on the chord length, which was kept constant along the blade, as described in Section 2.2. The transition parameter was set to $Ncrit = 9$ in XFoil.

While $Ncrit$ can influence the transition location and thus the airfoil performance at low Reynolds numbers, it was not adjusted to match the turbulence intensity level of the wind tunnel. This choice was made intentionally, as ALMs applications use standard XFoil settings without case-specific tuning of $Ncrit$. Therefore, the aim of this work is to evaluate ALM performance under typical usage conditions rather than under specifically calibrated settings.

The manuscript has been revised to explicitly include these parameters in Section 3.3.

R2: Specific comment 5

Where any steps taken to control the surface finish of the blade? This can influence BL transition and 3D printing usually generates quite rough surfaces.

Response:

We thank the reviewer for raising this important point regarding surface finish and its potential influence on boundary layer transition.

Surface finish was considered during the design and manufacturing process. For the present model-scale wind turbine, the Reynolds number based on the chord is approximately 5×10^4, which lies well within the laminar flow regime for airfoils. Transition to turbulence occurs at Reynolds numbers at least one order of magnitude higher. Therefore, the boundary layer remains laminar in our case. As to the influence of the moderate surface roughness of our airfoil on the aerodynamic performance, it is expected to be very limited due to the laminar character of the boundary layer.

Regarding the manufacturing process, the blades were fabricated using vertical 3D printing, as shown in Figure 3(a), in order to avoid the use of support material. This choice prevents surface inconsistencies that typically arise from support removal. Each printed layer begins with a continuous outer contour before the internal infill is deposited, ensuring good geometric fidelity of the airfoil profile.

To further minimize surface roughness associated with the layer-by-layer process, the printer was operated at the minimum available layer thickness. This results in a relatively smooth surface finish for as-printed parts. No additional post-processing (e.g., sanding or chemical smoothing) was applied, as such treatments could inadvertently alter the airfoil geometry and compromise shape accuracy.

The text has been adapted in Section 2.2, line 196 to reflect these additional informations.

R2: Specific comment 6

No details are provided regarding the turbulence model

Response:

We thank the reviewer for this comment. We agree that additional clarification was needed regarding the sub-grid scale (SGS) formulation. The manuscript has been updated to include a more detailed description of the SGS used in Nek5000.

In particular, we now clarify that the SGS model is based on an approximate deconvolution approach using a high-pass filtering operation to estimate unresolved scales from the upper portion of the resolved spectrum. We also explicitly define the role of the two main model parameters: the relaxation (filter-weight) parameter `filterWeight`, which controls the dissipation and is typically the most influential parameter, and the cutoff parameter `filterCutoffRatio`, which defines the fraction of the resolved spectrum used in the reconstruction (default value 0.9, corresponding to the upper 10 % of the spectrum). A reference has also been added to support this formulation [4].

These clarifications have been added to the text in Section 3.

R2: Specific comment 7

How long do the simulations run? What is the timestep? Statistical convergence must be accounted for in scale-resolving simulations

Response:

The simulations were advanced using a fixed time step of $\Delta t = 2 \times 10^{-4}$ s, chosen to maintain a sufficiently low CFL number for stable and accurate time integration, while keeping computational cost reasonable. Each case includes an initial spin-up phase to allow the flow to develop from the initial condition and to remove transient effects. This spin-up period consists of advancement until all monitored integral quantities reach a statistically stationary state, followed by an additional period corresponding to five times the time constant of the exponential moving average used in the control loop for matching the turbulence intensity. This ensures that the control variable reaches at least $99.3\ \%$ convergence to its target value. After spin-up, data are collected over a physical time interval of 20 s. This sampling window was chosen based on inspection of the wake statistics to ensure that the resolved turbulent fluctuations are sufficiently converged for the quantities of interest.

We acknowledge that statistical convergence is a key requirement in scale-resolving simulations, and we have now clarified these details in the revised manuscript in Section 3.

R2: Specific comment 8

The ALM description also lacks crucial information. For instance, what strategy was used to sample the velocity in the flowfield? How many actuator points are used? What is the actual width of the kernel compared to the grid (eq. 10).

Response:

We thank the reviewer for this important comment, which allowed us to further clarify the numerical implementation of the actuator line model.

In the present formulation, the actuator line is not represented as a set of discrete pre-defined actuator points. Instead, it is treated as a continuous line distribution evaluated on the spectral element mesh. The mesh is therefore used as the discretization points of the continuous actuator line formulation, without introducing an additional discrete sampling grid along the blade. Consequently, velocity sampling and force evaluation are performed at the same mesh nodes, resulting in a collocated formulation embedded in the underlying spectral element discretization. This avoids any separate line-wise sampling procedure and ensures consistency with the element-based representation of the flow field.

Regarding the spreading used in the ALM, the guideline value for the Gaussian kernel width applies to a configuration without tip refinement. In this case, discarding the tip refinement, the blade is discretized using 9 spectral elements in the spanwise direction, each of polynomial order 7, leading to 63 DoF along the blade. The corresponding non-dimensional kernel width is then defined as $\sigma/l = 2.5/(9 \times 7) = 0.040$, where σ is the Gaussian kernel width and l is the blade length. This formulation explicitly relates the smoothing scale to the effective nodal resolution of the actuator-line representation. This ensures a smooth and numerically stable force distribution while remaining consistent with the spatial resolution of the mesh.

From a computational standpoint, this formulation is advantageous because it reuses existing mesh nodes for both evaluation and projection operations. As a result, no additional actuator-line discretization, auxiliary sampling points, or dedicated communication structures are required, which improves efficiency and parallel scalability within the spectral element framework.

These clarifications have been added to Section 3.1 of the revised manuscript to improve the description and reproducibility of the ALM implementation.

R2: Specific comment 9

Are all experiments performed with $0°$ of blade pitch? Why compare to $+1°$ and $-3°$ in numerical simulations? Moreover, there seems to be higher thrust force extraction at the rotor edges, which may suggest inaccuracies in the tip-loss modelling.

Response:

Yes, all experiments were conducted at a blade pitch angle of $0°$. The numerical cases at $+1°$ and $-3°$ were chosen as part of the parametric study around the experimental operating condition. The intent was to assess the sensitivity of the flow response to modest deviations in pitch angle on either side of the experimental setup. These two cases produce slightly different wake characteristics, with a correspondingly lower and higher velocity deficit relative to the baseline case, which is useful for comparison. However, we do not observe a systematically higher thrust force in the simulations; rather, the differences in thrust remain consistent with the change in operating condition.

Regarding the apparent higher loading near the rotor tips, the blade geometry is represented with a sharp edge in the numerical model, which is an intentional modelling choice to reflect the actual physical configuration as closely as possible within the simulation framework. Capturing this tip-region behaviour is in fact one of the objectives of the study, and we consider it a relevant feature rather than a modelling inaccuracy.

This is now explained in more details in Section 4.

R2: Specific comment 10

Figure 23: Are the forces time-averaged? Why is a certain instant considered over another? Moreover, profiles are much smoother at the tip, what is this caused by? We are also not seeing the classical decline in axial forces close to the tip. This could be due to inaccuracies in tip effect modelling, or due to the large chord near the blade tip in the model, which does not feature taper as typical blades would.

Response:

The forces shown in Figure 23 are not time-averaged. The figure represents a single instantaneous snapshot, selected arbitrarily from the simulation time series. No particular physical significance is assigned to this specific instant; it is simply used to illustrate the spatial distribution of forces at a given time.

Regarding the smoother force profiles observed near the blade tip, this behavior is primarily attributed to the local mesh refinement in the tip region, which improves numerical resolution and reduces local fluctuations in the computed loads (known as Gibbs phenomenon in a spectral method). Near the tip, we also have a sharp transition worth emphasizing which avoids radial averaging, and has been explored with satisfactory results in prior work [1, 2].

The absence of the classical decline in axial forces near the tip is linked to the modeling approach used in this study. In contrast to standard ALM, which typically introduce a degree of radial smoothing or filtering of forces, the present method does not apply such radial averaging. As a result, the characteristic tip-force redistribution commonly observed in ALM-based approaches does not appear here. This difference is intentional and reflects one of the key distinctions of the method proposed in this work.

These clarifications were added in the text in Section 5.1.

References

[1] F. Trigaux et al. "Near-tip correction functions for the actuator line method to improve the predicted lift and drag distributions". In: *Journal of Fluid Mechanics* 989 (June 2024). URL: http://dx.doi.org/10.1017/jfm.2024.461.

[2] F. Trigaux, P. Chatelain, and G. Winckelmans. "Impact of rotor size on the aeroelastic behavior of large turbines: a LES study using flexible actuator lines". In: *Journal of Physics: Conference Series* 2767.2 (June 2024), page 022062. URL: http://dx.doi.org/10.1088/1742-6596/2767/2/022062.

[3] E. Gillyns, S. Buckingham, and G. Winckelmans. "Implementation and Validation of an Algebraic Wall Model for LES in Nek5000". In: *Flow, Turbulence and Combustion* 109.4 (Oct. 2022), pages 1111–1131.

[4] P. Schlatter, S. Stolz, and L. Kleiser. "LES of transitional flows using the approximate deconvolution model". In: *International Journal of Heat and Fluid Flow* 25.3 (June 2004), pages 549–558. URL: https://doi.org/10.1016%2Fj.ijheatfluidflow.2004.02.020.