WES-2024-6: Load assessment of a wind farm considering negative and positive yaw misalignment for wake steering

Author's Response to the Reviewers

The authors thank both reviewers for their time in reviewing the manuscript and valuable comments provided. Individual replies to all points raised are given below. A marked-up version of the article highlights all the removed and added text.

Reviewer 1

In the paper "Load assessment of a wind farm considering negative and positive yaw misalignment for wake steering," the Authors present a parametric study on the effect of yaw misalignment on power production and loads for an array of turbines under different wind directions. The study is based on an engineering model calibrated with a set of large eddy simulations under different yaw misalignments and one wind direction for which the five turbines in the array are aligned.

The topic is of interest to the community and the paper presents a large and valuable dataset on yaw misalignment/wind direction configuration. Nevertheless, there are some points, which I think should be addressed before publication in WES. I list below my main points.

Major points

Reviewer Point P 1.1 — Definition of "power gain" and "baseline case." For their analysis, the Authors take as the baseline the case of zero degree wind direction (for which the wind turbines are all aligned) and evaluate power differences with respect to this case. I think it is misleading to call these differences "power gains." The zero-degree wind direction is the worst case scenario for the array considered and using it as a baseline overestimate the effect of yaw control. It appears to me that these "power gains" on the order of 30%–40% are due to a very underperforming baseline, and not an actual measure of yaw control effectiveness. I believe the appropriate baseline should be, for each wind direction, the case gamma=0 (no misalignment). I understand the Authors may want a unique reference to normalize the data (perhaps the rated power?), but the discussion should be adjusted accordingly to avoid misunderstandings.

> **Reply**: We thank the reviewer for raising an important point. We have re-worded and refrained from using power "gains", and highlighted instead power "differences" throughout the manuscript. We agree with the reviewer and can see where it can be misleading to call it power gains.
>
> Having a single baseline allows us to discuss in terms of that baseline. Further comparisons between two other non-baseline points can be inferred from their differences from the baseline. We agree with the suggestion that a baseline for each wind direction would be ideal, but that would come with more complex plots showing the results, and some substantial explanation. In the interest of clarity, we have not changed the baseline, but have added notes throughout the manuscript pointing out the baseline aspect the reviewer raised here.

Reviewer Point P 1.2 — Related to the previous point, the Authors claim that the "bottom right quadrant" contains "attractive" operational conditions. I am not sure I completely agree. For instance, estimating from figure 12, for the case at 10deg wind direction, it seems there is not much percentage difference between the highlighted point gamma=7.5 and gamma=0. I do not think it would be attractive or beneficial to incur in the potentially increased fatigue loading due to yaw misalignment for a slight increase in power production. This is a known limitation of yaw control as previous studies in the literature have proposed methods to identify 'clusters' or 'partitions' within the wind farm where yaw control might be beneficial based on some performance metric (e.g. Kanev, 2020 Renew. Energy.; Bernardoni et al. 2021 J. Renew. Sust. Energy). To a certain extent, it seems that the scenario (wind turbine array and wind conditions) considered by the Authors is overall not very amenable to yaw optimization, because even in the worst-case direction (0° deg) the power improvement for different gammas is very mild (at least judging from the contours in fig. 12).

Reply: We agree with the reviewer that it would likely not be beneficial to incur potentially increased fatigue loading for marginal increase in power. Our farm is essentially "one row" of a more realistic farm. Therefore, when the bulk of the wake has been steered away from the downstream turbines, the lack of other rows of turbines come into play in the sense that not much more change is observed (that is, the power of the nonexistent turbines are not affected). Figure 12(d) illustrates this problem well (and is in line with the reviewer's comparison). For that wind direction, a zero degree yaw would result in the wake already mostly missing the downstream turbines. The figure shows a 7.5deg yaw, which steers the wake slightly more (and thus generating slightly more power) at the expense of increased fatigue loading. To better show the differences between the 0deg and 7.5deg yaw (e.g.), we agree with the reviewer that a more complex wind farm layout is warranted. We appreciate the reviewer's comment–we have created a subsection on "Discussions" with comments on the limitation of the model and added a note on the limitations of this farm in capturing such effects.

Reviewer Point P 1.3 — Some of the differences analyzed are quite small (2-4%) and appear to be of the same order of magnitude of the difference between the FAST model and the LES (Figure 8). Are these differences statistically significant? Have the Authors considered to perform a few additional LES to validate the model results (e.g. for one or some of the conditions they highlight through the paper)? I understand the computational cost is significant, but this is an important point to validate the discussion, since based on Figure 8 the model-LES discrepancy seems to increase with yaw misalignment.

Reply:

The reviewer raises a valid point. Parameter-tuning of engineering-fidelity codes are a balance between trying to have a representative value and trying to not over-fit the data. We understand Fig. 8 at the point chosen (red mark) indicate some errors on the power when compared to the LES (and that the errors are not necessarily symmetric between positive and negative yaw). Our goal with the current analysis is to identify trends and relative differences across the cases investigated. The final differences are indeed small and carry some of the errors presented in Fig. 8, but we believe that our analyses are still valid for relative differences. We do not interpret the errors in Fig. 8 as "uncertainty bounds" but rather as a constant bias such that relative comparisons of trends in FAST.Farm are still valid, but also acknowledge that this is an assumption and not a proven conclusion. We have added text at the end of the Calibration Section to emphasize this assumption.

Expanding on the asymmetry in the errors presented in Fig. 8, such asymmetry is due to underlying physics that are not yet considered in FAST.Farm. The asymmetries in power from a positive and negative yaw case for LES comes from veer, a non-zero height-varying cross-wind component of the wind v, due to the Coriolis force. Our LES simulations contained Coriolis (for a latitude of about 40 degrees, see updated Fig 4(b)) and FAST.Farm takes into consideration a mean v within the rotor swept zone. The mean v within the rotor from the LES is close to zero due to the underlying methods used for the simulation of the atmospheric boundary layer. Typically, and as done in this study, we select a height and chose the three-component wind speed—in our case, we used the hub height (150 m) and (8.6, 0, 0) m/s. The background driving body force considered in these LES cases is such that it enforces the requested wind speed at the requested height (in a planar-averaged sense). Because of that, the planar-averaged v behaves like shown in Fig. 4(b), with a mean that is close to zero. This mean value is then used in FAST.Farm, leading to an asymmetry in the errors.

For the remainder post-calibration analysis presented, we have used synthetic turbulence as the inflow. These flow fields are generated without accounting for veer, which we understand it is a limitation of the model. To the reviewer's point, we have added discussions on the limitations of FAST.Farm in accounting for veer in the flow. We have also added clarifying details regarding the inflow generated, both LES (which includes Coriolis) and synthetic turbulence (which does not include Coriolis).

Finally, a deeper study of wake flow physics and dynamics relative to positive and negative yaw misalignment could help to more deeply explain some of the nuanced findings in this work. A higher-fidelity study would also help to confirm that the trends observed in FAST.Farm, and the conclusions we derive, remain consistent as the wake physics are better resolved and Coriolis considered. We discuss such study as a needed future work.

Reviewer Point P 1.4 — In addition, the calibration in Section 3 only reports metrics for the mean power production, what is the difference for the other load 'channels' analyzed in Sections 4 and 5? Why using the wake position for model calibration instead of the metrics used for the analysis and discussion?

Reply: The calibration process was more thorough than what we showed on the manuscript– we tried to avoid this section taking too much space when it's not the focus of the paper. The wake center is used to ensure the parameters that control the wake deflection within the dynamic wake meandering model are consistent and agree well with LES. With the wake at the location matching the LES, we are likely to obtain the correct performance of the downstream, waked turbines and thus similar power, which is the quantity we match next. The calibrated parameters shown in Tables 2 and 3 are only related to the wake. Power and loads are consequences of the wake being (or not being) correct. Other tuning, not shown on the manuscript, was related to the radial discretization of the wake model. For that, we increased the number of points and, looking at the same load channels investigated in this work, noted convergence of their standard deviation values. Figure 1 shows part of that convergence exercise taken to ensure the discretization was appropriate for our problem. Similar convergence of the standard deviation values of the load channels were conducted during the calibration phase, but ommited in the article.

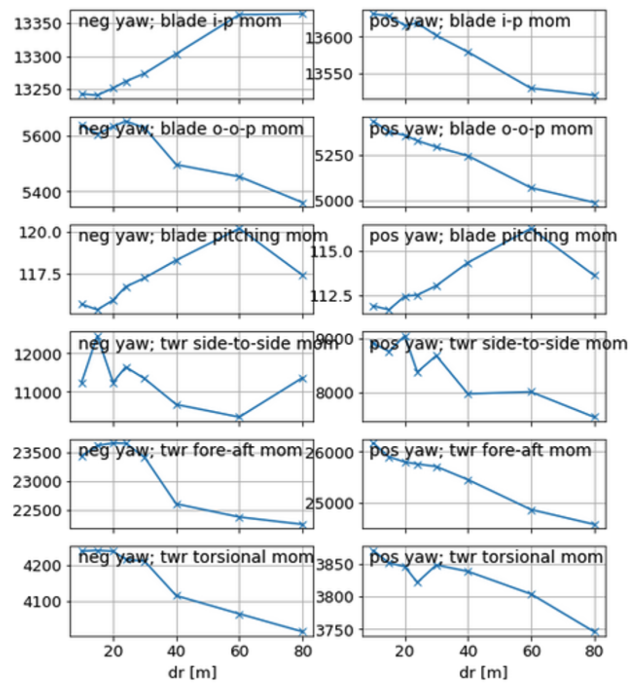

Figure 1: Example of the convergence study on the radial discretization of the wake model within FAST.Farm done during the calibration phase.

Reviewer Point P 1.5 — In the methodology description, Section 2.4, line 136, the Authors report that a "surface roughness" of 0.75m was used for the LES. I assume this refers to the aerodynamic roughness length in the classic log-law velocity profile ("z_0"). In such case, this seems an uncommonly very large value for "z_0," and I am concerned about how realistic are the inflow conditions considered in the study and the consequent wake evolution in the array.

> **Reply**: The reviewer is right, this is the aerodynamic surface roughness. We added that small clarification to the manuscript. We started this study with the goal of having an atmospheric boundary layer wind speed profile that matched a logarithmic profile with 0.2 shear exponent. With this requirement, we iterated generating different conditions using LES and obtained different flowfields. While a shear of 0.2 is closer to what a stable boundary layer would have, we aimed for neutral for simplicity. Neutral atmopsheric stability means no temperature stratification within the surface layer and lower veer than a stable boundary layer. Neutral atmospheric stability would also mean we are able to use synthetic turbulence to generate the flowfield, which was part of our goal. For this iterative study using LES, we found that the combination of 8 m/s at 150 m with an aerodynamic surface roughness value of 0.75 resulted in the shear profile we were looking. The resulting turbulence and other quantities varied within the general ballpark of a typical boundary layer (see Fig 4 for the vertical profiles).

Minor points

Reviewer Point P 1.6 — The Authors often refers to the "symmetry" of the results, or some trends as being "monotonic." As essentially all the quantities of interest are functions of two variables (yaw and wind direction), it may be ambiguous at time to follow what symmetry or monotonicity are the Authors referring to. I suggest to specify throughout the manuscript what variable is analyzed, e.g. "monotonic with respect to the wind direction."

> **Reply**: Thank you for the suggestion. We have added several clarifications throughout the manuscript following the reviewer's suggestions.

Reviewer Point P 1.7 — Figure 4, panel (e): perhaps the x-axis of the figure should be centered around 0deg?

> **Reply**: We have it centered around 270 degrees on purpose. 270 degrees is Westerly winds (meaning the wind is blowing from the west towards the east). That frame of reference is typical of the meteorology and wind energy communities, which is why we decided to keep it. A Westerly flow is consistent with the prior plot showing the farm layout (Fig. 2). In any case, we have added a clarification on the caption of Fig. 4.

Reviewer Point P 1.8 — Table 2 and 3: perhaps units are missing in Table 2? Have the Authors considered summarizing the results in one table only? I was expecting the selected value when reading table 2 and the discussion of the calibration process.

> **Reply**: You're right and we have failed to include the units in Table 2. Thanks for catching this—the manuscript has been updated to include the units. Regarding combining tables, we have attempt so, but in our opinion the resulting table was more complicated than a single one. Adding columns of "calibrated" and "default" next to the sweeps are counter-intuitive and requires explanation. We also wanted to show the calibrated values for the polar wake model, even though this model wasn't investigated further in the current study. Adding many extra columns results in a table will span more than the width of the text, or spaces will need to be reduced, decreasing the legibility of the values we are trying to convey. The tables with the sweeps and tuned parameter can, when split into two, be added near the appropriate text (in the current formatting they are three pages apart), which another benefit of keeping them separate. We appreciate the reviewer suggestion in making the manuscript look better.

Reviewer Point P 1.9 — Section 5: the Authors mention that some "quadrants" are not of interest because are not increasing power production, however they could provide useful information if interpreting the angle "gamma" as an unwanted misalignment (instead of an intentional yaw control effect).

> **Reply**: The reviewer is correct and this is an interesting point to highlight. While we kept most of our discussions focused on the net positive power increase, we have added notes about this point. Thank you.

Reviewer Point P 1.10 — Page 21, line 371: "as well with the exceptions"?

> **Reply**: Thank you for catching this mistake! We have fixed it.

Reviewer Point P 1.11 — line 396: should this be analogous the scenario (b) instead of (a)?

> **Reply**: Yes. We appreciate the attention of the reviewer in catching this and several other details that we have missed during our revisions.

Reviewer 2

The paper addresses a very important aspect of wind farm flow control through wake steering by including the effect of the steering direction on turbine power production and loads. Overall, the paper is well written, and the scientific quality is high. However a few points could be better elaborated upon to improve understanding.

Reviewer Point P 2.1 — Figure 1, It is not clear what the red label 'direction of pos alpha' signifies. While alpha is described later in the text, it should be defined here already with its first use.

> **Reply**: Thanks for the suggestion. We have made changes to the labels in this figure to be more clear.

Reviewer Point P 2.2 — Line 92, it would be beneficial to state the wall time of the LES simulation here to get some context of computational cost.

> **Reply**: The actual runtime for these simulations depend on the amount of resources used and may vary from cluster to cluster. We have now included some numbers on the rough order of magnitude of these runs to give the readers more context, but refrained from giving exact numbers. Exact number would imply our choice of hardware was made based on proper strong scaling studies for optimal performance, which was not the case in this work.

Reviewer Point P 2.3 — Line 104, why was alpha varied in increments of 2 degrees? Was a sensitivity analysis conducted regarding the minimal increment? Would the results be different if it was 1 degree, 0.5 degree?

> **Reply**: The increment of 2 degrees was chosen such that the resulting number of cases to be executed was computationally feasible. We have not performed a sensitivity analysis on the discretization of the design variables. Given the results we obtained (smooth contour plots), we did not iterate on the discretization interval as 2 degrees appeared to be enough.

Reviewer Point P 2.4 — Same comment for the yaw angle. Increments of 2.5 degrees seem arbitrary. A reader would benefit from knowing why these increments were chosen.

> **Reply**: It was somewhat arbitrary, yes. It was just chosen so that we would be able to create "heat maps" of results and those be smooth enough to draw conclusions. As mentioned above, the main driving factor for this choice was the number of cases to be executed. Given the discretization of the wind direction and reference yaw angle, a total of 483 FAST.Farm cases were executed. We have added an explanation of the discretization choice to the revised manuscript.

Reviewer Point P 2.5 — Line 113 seems grammatically incorrect, the structure could be improved.

> **Reply**: We have re-structured the sentence.

Reviewer Point P 2.6 — Figure 2, most of the legend items for the different wind directions are very difficult to tell apart from each other in this figure. It is also not clear what the high res and low res legend items are referring to, as they have not been introduced in the text yet.

> **Reply**: We have re-structured the legend box to make it more clear. We have also expanded on what high-res and low-res boxes are in the caption.

Reviewer Point P 2.7 — Line 120, while it is reasonable to select a single wind speed, the choice of 8.6 m/s again seem arbitrary. Additionally, a comment on whether choosing other below rated wind speeds could have an impact on the analysis presented.

> **Reply**: We have clarified our wind speed choice in the manuscript. The IEA Wind 15 MW reference wind turbine has a rated wind speed of 10.6 m/s. We have chosen a wind speed that is 2 m/s lower than that. The 2 m/s choice was driven by the desire of having all turbines operating in Reg. 2. We have not performed additional analysis on different wind speeds to be able to comment on whether or not they would have a different impact on the turbine, so have not added a comment about that.

Reviewer Point P 2.8 — Line 135, what was the resolution of the LES simulation? Is the resolution sufficient to capture the turbine performance (represented using actuator line model) accurately?

> **Reply**: We have added a new paragraph further describing our LES runs. We included the temporal and spatial resolution for both the precursor and the turbine simulations. We have also added other parameters such as the turbulence model used and added a note about the resolutions used and the use of the actuator line method. To the reviewer question, yes, the resolution and the number of actuator points used are on par with the recommendation from the literature for accurate computations.

Reviewer Point P 2.9 — Line 136, a surface roughness of 0.75m seems exceedingly high. Earlier on in the text, the authors cite the Massachusetts waters development as the motivation for the turbine spacing, indicating offshore application for this. The chosen surface roughness does not well represent offshore conditions, and is possibly too high even to represent onshore conditions. This could significantly affect the wind inflow conditions, the resulting wake recovery and the overall conclusions.

> **Reply**: The reviewer is correct that this value might be too high for an offshore application. The other reviewer made similar comments, and we point this reviewer to our prior answer in Point P 1.5. In a nutshell, we were aiming for a certain shear profile, and the combination of the desired wind speed and aerodynamic surface roughness yielded the conditions we were looking for under a conventionally neutral atmosphere.

Reviewer Point P 2.10 — A citation for the DEL methodology should be included

> **Reply**: We did not follow any specific methodology that is cite-able. Nevertheless, accounting for mean loads and using load roses to decompose orthogonal signals when calculating DELs is common industry practice even it is not often seen in research publications. Because of that, we

have strived to describe all the steps we have taken, so that our DEL section is self-contained and avoid having to cite articles that we did not necessarily use. We would also like to note that the original source of what is commonly referred to as "Goodman correction" nowadays is a book from 1899, which we have not cited.

Reviewer Point P 2.11 — Line 220, are both wake centerline error and power error metrics used in the minimization or just one? It's a bit unclear. Were any steps taken to avoid local minima or overfitting during the minimization? More information on how the minimization was carried out could be included.

Reply: Both are used during the calibration process. The paragraph that contains line 220 starts with the sentence "The quantities used for the comparison [...] are power and lateral wake center". We have clarified this further in the revised manuscript.

Regarding local minima, all of our comparisons were made in a semi-manual way, meaning we created plots like the ones shown in Figs. 6, 7, and 8. With that, we were able to intuitively discern trends and pick the location of minima.

Regarding overfitting, we have tried to avoid it by minimizing all three cases (positive, negative, and no yaw) at once. That discussion is presented in the second-to-last paragraph of section 3 and is illustrated in Fig. 8 by the red star.

Reviewer Point P 2.12 — Figure 11 is too small and can be increased in size to improve readability.

Reply: We have increased the font sizes and the figure size. The overall size of the figure is such that it will be within the text column in a two-column paper format used by WES.

Reviewer Point P 2.13 — Figures 16, 17, 18 for the load channels, the fore aft tower loads and combined together. Same is done for the blade flapwise and edgewise blade root bending moments. Including an additional plot showing the impact on the individual load components would be valuable to see the impact of control. The plot labels are also too small and barely legible.

Reply: The overall size of several figures (including the ones the reviewer explicitly mentions) have been increased.

Regarding the inclusion of the separate load channels, we have combined them using a load rose and presented the worst-case scenario. If we show the separate signals, we believe the results will look misleading and harder to interpret. For example, we could show the tower base with its fore-aft and side-to-side bending moments separately. With the imposed yaw sweep, the fore-aft will be highest when the turbine has no yaw misalignment. The side-to-side will keep increasing as the yaw angle deviates more and more from zero, while the fore-aft will keep decreasing. The frame of reference is the tower, and not the rotor. We know the worst direction of the bending moment will be the one aligned with the rotor, and the load rose captures that worst-performing bending moment magnitude. With a single combined signal, it is easier to compare across all the different yaw cases. We have decided to not include the separate signals, as the identification of trends are much more complicated. In earlier stages of this research, we had looked at the signals separately and decided it was best to combine them, so we certainly appreciate the reviewer's comments on presenting the separate signals.