Review for "*Yaw-based wake steering control under field uncertainties: a line of 4 industrial wind turbines investigated with LES and engineering wake models*"

The present work investigates the practical robustness of yaw-based wake steering in an open-loop implementation against real-world uncertainties, specifically static errors in the absolute wind direction and inaccuracies in nacelle yaw positioning. A four-turbine aligned layout, derived from an actual offshore wind farm, is used as the test case. Optimal yaw angles are computed using engineering wake models in FLORIS, and subsequently validated through high-fidelity LES using YALES2-BhawC. Both modeling approach are calibrated against SCADA data. The authors then provide a detailed comparative analysis across multiple uncertainty scenarios, examining impacts on power production, gains, flow fields, wake interaction and meandering etc. Overall, the wake models show good agreement with LES, while the most notable discrepancies arise under strong wake overlap conditions. Based on the results, the authors conclude that the nacelle yawing error only marginally impacts the power gains, whereas even small errors in the measured wind direction can substantially degrade performance.

General comments

Overall, the paper is clearly written and scientifically sound, and it provides a valuable contribution to the growing body of research on yaw-based wake steering. The robustness of yaw control under real-world uncertainties is indeed a topic that often remains underemphasized in the literature. The author's main conclusion – that nacelle yaw errors only marginally influence performance, whereas even small static errors in absolute wind direction can significantly degrade or even reverse expected gains – is clearly supported by the analysis and provides a practically relevant insight.

One notable strength of the work is that the paper is grounded in a real-world four-turbine offshore configuration, calibrated with SCADA data. The relatively close agreement between the engineering wake models and the LES results is surprisingly good, and there is substantial value in that comparison. In my opinion, the coupling to the real-world data may be emphasized more, especially with the title and abstract suggesting more of an uncertainty quantification in simulation environment.

At the same time, some caution is warranted when interpreting the conclusions. The scenarios selected to represent wind-direction errors effectively correspond to very adverse (or even worst-case) conditions, in which the optimized yaw settings become particularly suboptimal for the actual inflow. Combined with the deliberately aligned layout – which in practice should correspond to a relatively infrequent wind direction in a well-optimized farm layout – the resulting power losses naturally appear more dramatic. While these scenario choices are understandable for clearly illustrating sensitivity trends, it would be helpful for the authors to underscore even more explicitly that these results do not reflect a broad statistical uncertainty quantification, but rather a small targeted parameter study designed to explore a specific range of extreme cases. Clarifying this would help manage expectations, especially given the title and abstract, which may lead readers to anticipate a more systematic, full-fledged uncertainty analysis.

The discussion section is extensive and provides a detailed overview of the dynamical processes governing wake interactions and wake steering. However, a substantial part feels like reiterating well established knowledge in the wind farm community (and thus contributing less novelty on their own). In my opinion, the added value mostly lies in the connection to the real-world SCADA data and the consistency and/or discrepancies across modeling approaches. Emphasizing this integration – rather than revisiting well-established phenomena in isolation – could help sharpen the scientific contribution of the discussion.

In summary, this is a well-executed and relevant study, providing both practical insights and a solid comparative evaluation of wake models and LES under realistic conditions. With a slightly clearer framing around the extremity of the chosen scenarios and a stronger emphasis on the value of the model–LES–SCADA agreement, the paper could be further improved.

Specific comments

- In the abstract, the authors state that even "*a few-degree error in the absolute wind direction … is enough in some circumstances to turn a significant expected power benefit into slight losses*". This wording is too strong, as – in this context of this study – "*in some circumstances*" effectively refers to a deliberately constructed worst-case scenario. This should be clarified to avoid overstating the generality of the result.
- When stating the questions that the paper aims to address on p.3, Question (1) ("*How is the total variation in power output distributed across turbines in the cases of active control*") does not represent a novel research question and has been extensively investigated in prior studies. For this reason, the inclusion of the question appears redundant, or the authors must at least acknowledge that this topic has already been studied extensively (and pointing to appropriate references). A more interesting question in my opinion, which is also addressed in the paper, concerns the agreement between wake models and LES under realistic, field-calibrated settings.
- The study concludes that a few-degree error in the wind direction may be detrimental to wake steering performance. It should be made more explicit that this primarily a limitation in (pure) open-loop control. In contrast, if reliable feedback measurements are available, a closed-loop controller could in principle compensate for such biases. Although the availability of such feedback is itself a strong assumption and not guaranteed in practice, the distinction between

open-loop sensitivity and the potential robustness of closed-loop implementations should be acknowledged more clearly.

- On p.4, it is stated that "α_0 *was selected on purpose to target the maximum power increase*." However, when inspecting Figure 2 (bottom), the chosen $\alpha_0 = -2°$ does not coincide with the peak of the $\Delta P/P_R$ curve. It would be helpful if the authors could clarify this discrepancy?
- On p.4, Figure 2 (top) is referenced without explaining how the plot was generated. At this point in the manuscript, it is unclear whether this plot is based on the OUU framework described later on p. 5 or whether it was produced using a simpler or different approach. The authors should clarify how this figure was obtained.
- For the wind-direction bias of 4° and the nacelle yaw error of 2° used in Table 1, it would be helpful if the authors could comment on whether these values are representative of what is typically observed in practice. In particular, wind-direction biases of a few degrees have been reported in field studies due to calibration drift or sensor limitations, and nacelle yaw misalignment on the order of 1–3° is often cited as typical in commercial turbines. Clarifying whether the chosen values reflect common operational uncertainties – or whether they were selected for sensitivity purposes – would improve the context and interpretation of the results.
- On p.6, the manuscript states that the wake model uses "*the default parameters*", but no reference is given for where these default values are defined. It would be helpful if the authors could cite the appropriate literature.
- What are typical values for the wind direction error
- On p.7, the manuscript states: "*To unify the optimization results and thus factor in wake-model diversity, the applied yaw-command sets in all LES simulations, then possibly combined with a nacelle positioning error, were taken as the average of the sets individually given by each combination of sub-models*." To me, it was not entirely clear what this procedure entails. Does this mean that the OUU optimizer was run separately for all combinations of wake-model sub-components, and that the yaw angles applied in the LES were then computed as the average of these independently obtained sets? It would be helpful if the authors could reformulate this part for clarity.
- Regarding Table 2, it is not entirely clear how the yaw-angle values for each case were obtained. As I understand it, the yaw angles in Case 2 originate from the OUU optimizer. The 2° offsets when going from Case 2 to Case 4 and from Case 3 to Case 5 are attributable to the imposed nacelle-yaw error, which is straightforward. However, the procedure used to generate the yaw-angle set for Case 3 is unclear. My understanding is that the OUU optimization should only need to be run once, producing a baseline yaw-command in idealized conditions (i.e. Case 2), after which all other yaw-angle sets (Cases 3–5) would be derived systematically by applying the specified wind-direction and nacelle-offset errors. In that sense, it is unclear to me how the numerical values for Case 3 were obtained, and it would be helpful to clarify this. More generally, if the reasoning above does not correspond to the proposed workflow in the manuscript, then the methodology needs to be

clarified (and motivated!), as the current description leaves some ambiguity about how the yaw commands in Table 2 were generated.
Finally, in the caption of Table 2, it would improve clarity to explicitly state that the listed yaw-angle values were obtained from the OUU framework (or derived from its outputs, depending on the correct interpretation above).

- On p.9, the roughness length is set to 0.005m. Since the simulation corresponds to an actual offshore site, it would be helpful if the authors could explain how this value was selected. Typical offshore roughness lengths are often significantly lower (of order 10^{-4}–10^{-3} m), depending on sea state, so clarifying whether this value is site-specific, assumed, or chosen for numerical reasons would improve transparency.
- On p.12, the manuscript states that, after initializing the turbine response, the flow was advanced for one flow-through time at U_∞. Is one flow-through sufficient for the wakes to propagate across the entire row (i.e., for the wake of T1 to reach T4)?
- On p.13, the manuscript describes the filtering of SCADA data to obtain a set of representative samples for comparison with the simulations. However, the study assumes neutral atmospheric stability throughout, and it is not clear whether the authors assessed the stability of the selected SCADA periods to verify that they are indeed consistent with neutral conditions. Since the four-turbine layout analyzed here represents only a subsection of a larger wind farm, deviations from neutrality in the actual atmospheric conditions could have a substantial impact on both power production and wake behavior. If no atmospheric-stability filtering or assessment was performed, this should be acknowledged, and it would be appropriate to attribute part of the discrepancy or uncertainty between SCADA, wake-model, and LES results to potential stability variations in the SCADA dataset.
- As a general comment regarding the results section: as mentioned before, much of the analysis reflects findings that have already been well-established in the wind farm community. It would improve clarity if the authors more explicitly distinguish between which observations and/or conclusions represent new insights from the present study and which reflect established knowledge in existing literature (and mention those references where appropriate).
- On p.17, the authors note that although most of the power gains are concentrated in the most downstream turbines, such gains may be difficult to achieve in practice with a classical open-loop strategy because of veer and meso-scale effects. While I agree that this sensitivity is indeed a limitation of open-loop control, I also believe that it reflects the limitations of the underlying engineering wake models themselves. More advanced control-oriented models (e.g., improved engineering models or even LES-based surrogate models) could potentially mitigate part of this issue.
The authors further state that "*a positive global gain may already be achieved by considering the two first turbines only during the optimization*." This statement is unclear to me. As written, it appears to suggest that running the OUU optimization

only on the first two turbines would already yield a positive global gain, which seems contradictory if the largest gains arise downstream. I recommend clarifying the meaning of this statement and how it aligns with the preceding results.

- On p.18, the authors write: “*The results further stress the high sensitivity of the wake steering control strategy to the absolute wind direction measured in the field. Indeed, a total power increase of more than 20% can be unlocked… Nevertheless, a 4°-error in the absolute wind-direction measurement is enough to turn this benefit into slight power losses.*” As mentioned previously, this wording is too strong. In the present study, the detrimental impact of a 4° bias corresponds to a carefully selected worst-case scenario, rather than the outcome of a broad sensitivity analysis. The manuscript should clarify that this result reflects a particularly unfavorable combination of layout orientation and bias, rather than a general conclusion about the robustness of wake steering or the result of a full-fledged sensitivity analysis.
  Furthermore, it would be helpful if the authors commented on how realistic a 4° absolute wind-direction error is in practice. While small biases can indeed occur in field measurements due to calibration drift or sensor uncertainty, the typical magnitudes at commercial installations vary, and some context would help assess how representative this scenario is.
- On p.19, the discussion on the practical challenges of assessing power gains – specifically the relationship between required averaging time and the wake-meandering timescale – is very interesting and highly relevant for real-world applications. The proposed mitigation using bootstrapping is also valuable and appears to work well. In my view, this mitigation strategy could be emphasized more strongly, as it is one of the concrete, actionable outcomes of the study. It may even be worth briefly mentioning in the abstract.
  Regarding the meandering period, the manuscript notes a value of approximately 100 s. I would suggest explicitly linking this estimate to the spectral content shown in Figure 11, so the reader can directly see how the dominant timescale is identified from the presented data.
- On p.20, the authors state that “*the vertical velocity fields highlight that all wakes are on average moving upwards, in particular when they get closer to the downstream turbine.*” However, when examining Figure 12 (right), this upward motion is not obvious to me. The plot appears to show roughly symmetric regions of positive and negative vertical velocity, rather than a clearly upward-shifted wake. Could the authors clarify how the upward displacement is identified from this figure – e.g., whether this refers to a vertical shift of the time-averaged wake centerline, a statistical imbalance in the distribution of w, or some other diagnostic not immediately visible from the colormap? A brief explanation would help the reader understand this observation.
- On p. 24, part of the observed behavior is attributed to blockage effects. I am not convinced this mechanism is relevant in the present context. Blockage effects typically become significant in large wind-farm arrays, whereas in neutral atmospheric conditions and for a small configuration of only four turbines, such effects are generally minimal. Given the limited array size and the neutral ABL

assumed in this study, I would question whether blockage meaningfully contributes to the trends described here.

Technical correction

- p3: "*seeked*" → "sought".
- Table 1: include the units of e_d and e_γ.
- In several places, the manuscript uses an unnecessary pair of parentheses when the reference appears as part of the grammatical structure of a sentence. For example: "*… was addressed in (Rott et al., 2018)*" (p2), "*as designed in (Rott et al., 2018)*" (p5). When the reference functions as a grammatical object, it should not be enclosed in brackets.
- In Eq. (2), the explicit Laplacian term is unexpected, as such diffusion is normally handled within the SGS model in LES. I assume this is a typo, but if it is intentional, the authors should clarify its purpose.
- p.12: "*was only activated in last step*" → "was only activated in the last step"
- p.17: "*litterature*" → "literature"