Author's response to Anonymous Reviewer #1

Jochem De Schutter, Antonia Mühleck, Rachel Leuthold, Moritz Diehl

April 11, 2026

Dear Reviewer,

thank you for taking the time and effort in providing us with your detailed feedback and diligent comments. We feel that they have contributed significantly to the strength of the manuscript. In this document, we address these comments point by point including pointers to the changes made in the manuscript. When referring to Figure and Table numbers, we use the numbering in the updated manuscript.

Kind regards,
Jochem De Schutter on behalf of the authors

Comment 1.
I was not familiar with the concept of wake regeneration above conventional wind farms, so perhaps this could be explained a bit in more detail: How does this work? What is the advantage? .

Response:
In this concept, kite systems (without own generation capability and thus with low complexity) could be installed strategically between wind turbines in an offshore wind farm, possibly only using existing infrastructure. By flying above the wind farm, they would be able to redirect airflow from the unaffected boundary layer into the wake of wind farms to enhance their overall efficiency through wake re-energizing. The initial simulation studies referenced in the text indicate that a wind farm performance boost of several percentage points could be attainable. We have added more explanation in the text, but removed the reference from the abstract so as to not distract the reader at that point.

Comment 2.
Line 15: "The solution of the original no-wake problem deviates with 145 % compared to this reference solution." Specify what solution you are looking at; is it power, trajectory? I only find out later that it is the force. Furthermore, what is the reference solution? The one that includes the wake? Please clarify. .

Response:
Thank you for pointing out this inclarity. We have elaborated as follows: "In a numerical case study based on a dual-kite system, we examine the sensitivity of the OCP optimal cost (average tether force) and design to the key transcription parameters. Based on this sensitivity analysis, we find an efficient transcription that solves the problem at three times the computational cost of the OCP without any wake model, while still retaining accuracy (in terms of combined error of optimal cost and design) within 5 % compared to a highly resolved transcription of the wake-aware problem. The solution of the no-wake problem deviates with 145 % compared to this highly-resolved reference solution."

Comment 3.
Introduction: The writer could refer to the following papers:

- T. Haas, J. De Schutter, M. Diehl, and J. Meyers, "Large-eddy simulations of airborne wind energy farms", Wind Energy Science, 2022.
- J.-B. Crismer, T. Haas, M. Duponcheel, and G. Winckelmans, "Large eddy simulation of airborne wind energy systems flying in turbulent wind using model predictive control", Wind Energy Science Discussions, 2026, Preprint discussion started: 09 Jan 2026; under review for Wind Energy Science (WES).
- Pynaert N., Haas T., Wauters J., Crevecoeur G., Degroote J. (2025). Aero servo simulations of an airborne wind energy system using geometry-resolved computational fluid dynamics. Wind Energy Science (WES). https://doi.org/10.5194/wes-10-2663-2025.

In these studies, the wake is simulated by coupling AWEbox to a virtual wind environment (LES-RANS) that includes the wake implicitly. The trajectories that are flown are, however, based on simplified models that do not include the wake. The aforementioned studies could benefit from your work. It could also be interesting to compare the wake and induction with these studies (perhaps in future work).

Response:
Thank you for the suggestion. We agree and have integrated this research branch into the introduction narrative.

Comment 4.
Line 55: "In conclusion, the accurate performance assessment of utility-scale kite systems requires the development of optimization-friendly engineering wake models capable of capturing self-induction effects." It could be beneficial to stress that the high controllability of airborne wind energy systems enables them to avoid wake and induction effects, potentially increasing power output, provided that this information is incorporated into the controller, further highlighting the significance of this study.

Response:
Thank you for this suggestion. We have added the sentence: "Moreover, kite systems possess many degrees of freedom and their high controllability allows them to adjust their flight path so as to avoid wake effects in consecutive flight loops, thereby boosting performance, provided that this information is incorporated in the optimization model."

Comment 5.
Line 80 "To limit the scope of this paper, we consider a kite system with fixed tether length, thereby excluding pumping-style AWE systems, which rely on tether reeling." Also specify that no power generators are involved; e.g. fly-gen systems, but purely act as actuators.

Response:
We have added this information in the paragraph.

Comment 6.
Line 94: You introduce y as a parameter for the wake state; is it possible to give a physical meaning to the "wake state" to help the reader? Also sometimes it is denoted as y(t,.), something y(t,tau); is there a (mathematical) difference? Please explain.

Response:
The full wake state $y(t, \cdot)$ represents the full spatial vorticity distribution (structure, orientation, strength, ...) that has resulted from the flight trajectory up until the current physical time t. The wake state $y(t, \tau)$ then represents an infinitesimal segment of this distribution, which has an age (or convection time) $\tau \in \mathbb{R}^+$, which represents the time elapsed since that portion of the wake was generated at the boundary $\tau = 0$. We have added this explanation to the text.

Comment 7.
Line 109: A figure would be helpful for the reader to visualize the states of the system.

Response:
We have added an illustrative figure of the system including the states.

Comment 2.
Line 131: "Incorporating more realistic wind and density profiles as a function of flight altitude is considered straightforward and omitted here for simplicity." I would argue that wind shear has a big influence on the wake development, so I am not so sure this is straightforward.

Response:
Thank you for this comment. We have changed the formulation to be more precise. Within the proposed framework, one would need to just insert any

wind shear profile $u_\infty(z)$ into the model, and it would be taken into account automatically in the wake development (different convection speeds at different altitudes). That is the "straightforward" part. However, we admit that in this case further validation simulations would be necessary.

Comment 8.
Line 156: "properly distributed"; do you mean uniformly distributed?

Response:
We have clarified: For each tether, these forces are computed for the boundary nodes such that they produce resultant forces and moments equivalent to those obtained by numerically integrating the drag distribution along the tether.

Comment 9.
Line 171: "For each wing, we ignore the induction of its own trailing wake less than a time T away, since this part is taken into account by the induced drag term." Can you discuss what would be the effect of not including the induced drag instead?

Response:
Accurately resolving this region would require a significantly finer representation of the vortex distribution close to the wing, increasing model complexity and computational cost. We have made a clarification in the text.

Comment 10.
I find "u^f" not so intuitive to indicate the induced velocity, but perhaps this is the standard notation in the literature.

Response:
The "f" superscript refers to "far"-wake induction as opposed to the "near"-wake induction included in the induced drag term. Here, we follow the notation of [Trevisi et al., 2023]. We have added this relationship in the text.

Comment 11.
Figure 1: Indicate both methods using (a), (b), or (left),(right).

Response:
Thank you for the correction.

Comment 12.
Figure 3: The legends in panels (b) and (c) are difficult to read; using black instead of white would likely improve their readability. In addition, including the system trajectory in these figures would provide valuable additional context.

Response:
Thank you for this comment, we have changed the figures accordingly.

Comment 13.
Line 250: "One problem is that the optimal orientation trajectory by construction assumes zero angle-of-attack"; I would rather state that the angle of attack is not included or represented in the model, zero angle of attack would mean zero lift (or zero angle-of-attack lift).

Response:
Thank you for the suggestion, which we have implemented in the text.

Comment 14.
Figure 4: I had some difficulties understanding the figure. Clearly indicate in the figure what is part of the DUST simulation (I assume the thin black line) and the vortex loop wake (I assume the red dots, and blue arrow?).
Response:
We have added the kite trajectory to the plot, added more information to the legend and the caption so that all visual information is clearly defined.

Comment 15.
Line 323 "For the atmospheric airborne actuator case, as well as for the ship towing case, the tether pulling force is to be maximized." Please provide additional explanation and justification for why this metric is the most appropriate or meaningful choice for the atmospheric actuator case.
Response:
The rationale behind this metric is that it is used as a simple proxy for the amount of momentum that can potentially be redirected into the wind farm. We have added this explanation to the text.

Comment 16.
Figure 6: Indicating the system trajectory in the figure would improve clarity. Additionally, including an error contour plot could further enhance the informative value of the results.
Response:
We have added the full flight trajectories as well as the wing position and orientation at the time of the snapshot. While we do not disagree that an error plot would provide additional information, we have chosen to not include it in this text and we hope the reviewer agrees that, given the added orientation by the kite trajectories, sufficient information is available to compare the two results.

Comment 17.
Combine figure 6,7,8 to enhance the readability.
Response:
We have implemented your suggestion.

Comment 18.
Table 1 describes the system variable bounds (both tether and wing properties),

and Table 2 describes the kite system model parameters (both tether and wing properties). Why are they not presented in the same table? Furthermore, how are these system parameters determined?

Response:

The tables are separate because they provide different types of information (fixed parameter values vs. min-max variable bounds). While both tables address properties of both wing and tether, there is no parameter/variable that is present in both tables. So merging both tables would lead to many empty entries. However, we have moved Table 1 further backwards in the paper, closer to Table 2, since the numerical values are only relevant in the numerical example. In short, the parameters are chosen so as to represent a very high-load utility-scale example, using typical values from the AWE literature. We have provided more specific justifications in the text (beginning of Section 6.1).

Comment 19.

I wonder if it would be valuable to also compare the discretized wake with the DUST simulations.

Response:

We have added a comparison of the aerodynamic loads for the reference OCP transcription with the analytic vortex-loop solution and the DUST simulation using the exact same boundary condition, showing very good agreement.

Comment 20.

Line 418: "The transcription includes a time transformation approach to deal with the variable time grid." This appears to be an important step and would benefit from a more detailed explanation.

Response:

This is a standard technical trick since many years, not specific to the new developments. We have added a reference should the reader want to learn about the technical details.

Comment 21.

The optimal flight trajectory obtained using the no-induction model appears physically unrealistic. For a flight speed of 190 m/s and a turning radius of 50 m, the required centripetal acceleration would be approximately 73 g, which is far beyond feasible operational limits. Please clarify this point and discuss the physical realism and practical implications of the resulting trajectory.

Response:

We agree that the concrete presented trajectory is outside the range of practical relevance. Nevertheless, induction (and tether drag) effects are roughly independent of the chosen wind speed (see [Trevisi et al, 2023], where a simplified model is derived and the induction factor only depends on the gliding ratio), and similar validation and design trade-off results are therefore expected for lower design wind speeds with lower flight speeds and therefore a more reasonable flight envelope. The main modeling error here is that the mass of the wing does not scale appropriately with the very high observed wing loading, leading to an underestimation of the gravity-induced variations in the flight speed. However, in the dual-kite set-up, such variations are exptected to be low anyway. Since the goal of the example was to generate an extreme scenario for the validation rather than to examine practical design trade-offs, we decided to not repeat the simulations at lower wind speeds. We have added this explanation in the text.

Comment 22.
Line 488: "The model is also very well able to capture the induced velocity field around the wing trajectory flown, suggesting that it is well-suited for optimization purposes." I believe this sentence is a bit misformulated: The system model, including the wake model, increases the accuracy; on the other hand, the speed will determine if it is well-suited for optimization purposes. E.g., High-fidelity LES would not be well-suited for optimization purposes, but captures induced velocity very well.
Response:
We agree and have removed this part of the sentence.

Comment 23.
Line 493: Specify the original OCP (with or without wake model?)
Response:
We have changed the text accordingly.

Comment 24.
Line 495: Specify the reference solution (with or without wake model?)
Response:
We have changed the text accordingly.

Technical corrections

Comment 25.
At the end of the introduction, the conclusion is not included in the discussion of the structure of the paper.
Response:
We have adjusted the text accordingly.

Comment 26.
To be checked with the editor, but if I remember correctly, vectors should be defined in bold.
Response:
This is indeed the guideline, and we have changed the text accordingly.

Comment 27.
Section 4.4 has almost the same title as Section 4.3.
Response:

We made the titles more precise.

Comment 28.
Figure 9: Please explain N, M, N, to make the figure self-explanatory.
Response:
We have added more explanation to the caption.

Comment 29.
Figure 9: Punctuation mark forgotten.
Response:
Thank you, we have made the correction.

Comment 30.
Line 418: "four RadauIIa collocation points" → Typo?
Response:
This is not a typo, RadauIIa is a Runge-Kutta integration method.

Comment 31.
An introduction to the sections is missing in Chapter 6.
Response:
We have added an introduction.

Comment 32.
Eq. 10: is "u_t, i" defined somewhere?
Response:
This is in essence just a helper variable to build up the formulas more slowly, but to be sure we have added a definition in the text.

Author's response to Anonymous Reviewer #2

Jochem De Schutter, Antonia Mühleck, Rachel Leuthold, Moritz Diehl

April 11, 2026

Dear Reviewer,

Thank you for the time and effort you invested in providing detailed feedback and thoughtful comments. We believe these have significantly strengthened the manuscript. In the following, we address each comment point by point and indicate the corresponding revisions made in the manuscript. When referring to Figure and Table numbers, we use the numbering in the updated manuscript.

Kind regards,
Jochem De Schutter on behalf of the authors

Major Comment 1: Validation of the simplified wake model within the optimal control framework
The continuous vortex-loop wake model is convincingly validated against DUST. However, the optimal control problem employs a simplified hybrid wake model with reduced window-of influence assumptions and finite wake history. While such simplifications are understandable for computational tractability, it remains unclear to what extent the reduced model preserves the quantitative accuracy demonstrated for the full formulation. Although Figure 12 compares the discretised wake against an analytical vortex-loop solution, it is not evident how the reduced model compares to the DUST-validated formulation. Since the dual-kite optimisation results rely on the simplified wake model, it would strengthen the manuscript to clarify this validation chain, for example through a direct comparison with DUST for the reference solution or by quantifying the modelling error introduced by the OCP simplifications in the reference solution .

Response:
Thank you for this comment. In the previous version of the manuscript, the validation chain consisted of three steps:

1. Validate the analytic vortex-loop model with DUST simulations.
2. Validate a reference hybrid model transcription with a large window of influence to the analytic vortex-loop model.

3. Discretize more coarsely, choose a smaller window of influence and compare the change in optimal cost and solution compared to when using the reference transcription.

We believe your comment relates to Step 2, which we agree is not sufficiently elaborated on. This step was previously somewhat "hidden" in Fig. 12 and we agree that it would also benefit from a direct comparison of the reference transcription to the DUST simulation. We have added a comparison of the resultant aerodynamic forces prominently and with accompanying Fig. 9 in Section 6.1, showing very good agreement.

Major Comment 2: Representativeness of the case study and interpretation of wake effects
The numerical case study considers a configuration with fixed tether length and without reeling dynamics and with a very optimistic mass scaling (see https://doi.org/10.5194/wes-7-1847 2022). This appears to represent either a fly-gen system without explicitly accounting for the additional mass and drag of onboard turbines, or a ground-gen system without accounting for its reel-out operation. In both cases, the resulting flight speed and loading would be reduced in real operation, which directly affects period duration and wake strength. The analysed dual-kite configuration therefore seems close to a worst-case scenario in terms of wake sensitivity. While this does not invalidate the results, it affects how broadly they can be generalised. A more explicit discussion of the operational interpretation and limitations of the chosen case study would help readers assess the scope and transferability of the conclusions to other systems.

Response:
The case study was chosen as a dual-wing "ship-towing" or "airborne actuator" case intentionally, to indeed create a "worst-case" scenario for induction. The reason for this is was to put our model to the test. The model is able to capture the induced velocity field in great detail when the kites are flying very close to their own wake (very high effective gliding ratio). For this, the vortex-loop model is necessary. For the significantly lower effective gliding ratios that will likely be encountered in realistic airborne wind energy operation, the distance to the wake will increase. In terms of validity, this does not pose a problem: the vortex-loop model does not become less accurate at larger distances. Rather the opposite is true: the vortex-dipole model becomes more accurate. This is also clear from the newly added comparison in Fig. 9: the hybrid model, with only one layer of vortex-loops and further only vortex dipoles, gives almost a perfect fit compared to the full vortex-loop model in this extreme scenario.
In terms of computational expense, a lower effective gliding ratio mainly makes the OCP cheaper: the wake is convected faster, creating a larger distance, up to a point wherer one can use only the cheaper vortex-dipole model. Because of the large wake convection distances, one can also cut off the wake history sooner. This we point we had already made in the conclusion.
Finally, we have also added a comment on the mass-scaling, which in this case is

indeed very optimistic considering the high loading. However, for the dual-kite case, where the kites are balancing each other, the influence of the mass on the optimal flight trajectory is typically much lower compared to the single-kite case. Hence, we deem the model acceptable for the validation purposes of this paper.

Major Comment 3: Half-periodicity assumption for the dual-kite problem

The manuscript states that, because the two airfoils follow nearly identical trajectories over a full orbit, the dynamics can be integrated over half an orbit by enforcing half-periodicity (line 411), following Zanon et al. (2013). This reduction appears insufficiently justified for the present setting. For a helical trajectory, the "up-loop" and "down-loop" segments typically experience different apparent wind conditions, speeds, and load distributions due to gravity, so the orbit need not exhibit a strict half-period symmetry. This makes the choice of the two half-orbit boundary points (where the state matching is imposed) critical, and it is not clear from the text where along the orbit the half horizon starts/ends, nor under which assumptions half-periodicity is expected to hold. Related optimal-control studies (eg. https://doi.org/10.3390/en16041900) explore dual-aircraft configurations without taking only half-period and the resulting trajectories are not generally symmetric. This suggests that enforcing half-periodicity may exclude physically relevant asymmetric solutions unless the symmetry assumptions are carefully justified. I recommend that the authors (i) clearly define the half-orbit start/end points along the trajectory, (ii) explain the conditions under which the half-periodicity constraint is valid for their chosen path class and (iii) ideally demonstrate that solving the full-orbit problem does not yield materially different (asymmetric) optimal solutions for at least one representative case.

Response:

Thank you for this comment. We agree that this step has received insufficient attention in the manuscript. The main issue here is the meaning of the term "half-periodicity". This does not mean that only one half of the loop is simulated and the other half is chosen symmetric to it. Rather, we exploit the fact that we have a two-kite system in combination with a single-loop trajectory (typical case for all non-reeling systems). In this case, while one wing flies the upstroke, the other simultaneously flies the downstroke. After half a period, the roles are reversed and the first wing flies the downstroke and the second one the upstroke. Thus, by implementing this "role reversal" as a constraint instead of imposing full periodicity, we avoid simulating the same up- and downstroke twice. We have merely imposed that both wings fly the same downstroke and upstroke paths, which can be asymmetrical (this can be seen, e.g. in the force profiles). This is of course crucial, since our central claim is that we capture the non-axisymmetric nature of kite system flight trajectories. The start- and endpoints can be chosen freely by the solver, and do not influence the solution, as the flight loop is simulated completely for all possible choices. We have

elaborated on this matter in the text and have explicitly stated the new "role reversal" constraints.
In the conclusion we had already raised the point that allowing differing up- and downstrokes ("multi-loop") for the two wings might lead to a different, more performant solution, as the wings would in this case be able to adjust their flight path to "evade" the trailing wake of the other one. While this is a very interesting research path, such an investigation falls outside of the scope of this paper.

Minor comments

Comment 1.
Line 14-15: Specify what metric are you using for the accuracy and comparison with reference solution. It should be clear by only reading the abstract.

Response:
The metric is the combined error of optimal cost and design. We have added this to the abstract.

Comment 2.
Line 80: "To limit the scope of this paper, we consider a kite system with fixed tether length, thereby excluding pumping-style AWE systems, which rely on tether reeling." Also mention that no onboard power generators are included, such that they purely act as actuators

Response:
We have added this information to the text.

Comment 3.
Line 81: "However, the developed methods in this paper can be directly applied to these type of systems as well." The present method assumes a fully developed wake, whilst in pumping systems wake development becomes increasingly important, particularly if the amount of flown loops is small. Please justify further how these methods could be applied to pumping systems or remove this phrase.

Response:
The proposed unsteady model can be used for time-domain simulations and does not in itself assume a fully developed wake. Because of its unsteady nature, it is in fact very much suited for modeling pumping-style wake patterns. However, when used in a periodic OCP, we do assume a fully developed periodic wake. For a pumping system, this just means that the pumping wake pattern (including build-up during reel-out as well as the reel-in phase) is repeated indefinitely when going down in the wake's history. This does not pose a limitation to the proposed method's applicability. We have elaborated on this in the text.

Comment 4.
Figure 1. Add a label to the methods (a,b). There are some symbols that are only introduced later in the text so the reference should indicate a) or b) as well.

Response:
We have added labels to the figures and when referencing to them in the text.

Comment 5.
Line 228: "Recall that for each wing we ignore the self-induced velocity for the first half rotation interval with length T." It says recall while it is the first time a half-rotation is mentioned.

Response:
Thank you for this comment, we have corrected the text.

Comment 6.
Line 250: additional clarification of the DUST simulation setup would improve transparency. It is not entirely clear whether the wing kinematics (positions and velocities) are prescribed and the pitch angle is chosen to match a target C_L without accounting for induction effects, or whether the induced velocities computed by DUST are already incorporated in a fully coupled manner such that the local angle of attack and resulting aerodynamic loads are determined self-consistently.

Response:
The first option is the case: the wing kinematics are prescribed, with the addition that the chosen pitch angle takes into account the induction effects predicted by the proposed wake model. We have clarified in the text.

Comment 7.
Figure 4 is somewhat difficult to interpret in its current form. The readability would improve if the flown trajectories of the kites were explicitly included in the visualisation. In addition, it would be helpful to clearly distinguish in the figure or caption which results correspond to the DUST simulation and which to the proposed vortex wake model.

Response:
We have added the kite trajectory to the plot, added more information to the legend and the caption so that all visual information is clearly defined.

Comment 8.
Lines 275-280: It would be helpful for the reader to understand how to qualitatively asses over or underestimation of convection speed in Figure 4 (eg. number of tip vortices seen in DUST vs vortex method)

Response:

We have added an explanation in this paragraph. Indeed we rely on looking at the location of the tip vortices as they cross the xz-plane.

Comment 9.
Figures 6,7: it would be helpful to mark the positions where the wing crosses this plane

Response:
We have added the full flight trajectories as well as the wing position and orientation at the time of the snapshot.

Comment 10
Figure 9. The time intervals N are difficult to see.

Response:
We have updated the plot so that the intervals (now correctly labeled N_{w}) should be clearly visible.

Comment 11.
Subsection 6.2. the discussion of discretisation sensitivity is not entirely clear. In particular, the relationship between the number of collocation intervals N (global time discretisation) and the influence window parameter N_f would benefit from clarification. Please explicitly state the value of N used and clarify how it relates to the window size $2N_f + 1$.

Response:
Thank you for this comment. We had introduced some confusion by equating the number of intervals used by the finite window approach with the number of OCP intervals. There is no fundamental relationship between the two or between the parameters N and N_{f}. They can be chosen completely independently. We now explicitly differentiate between the number of intervals for the window approach N_{w} and the number of OCP intervals N. We have added the values used for both parameters in Table 3.