Response to Reviewer 1

Andrew P. J. Stanley, Owen Roberts, Jennifer King, and Christopher J. Bay

May 2021

First, we would like the express our gratitude for your review of our paper. We realize that you took time out of a busy schedule to read this manuscript and provide feedback, for which we are very grateful. We have structured this response to be clear and easy to follow. Each of your original comments will be shown in blue, immediately followed by our response in black.

1 Specific comments:

Chapter 1 (Introduction):

Directly in the second sentence, the authors state that "wind turbines ... require no external fuel, and require little to no water". With the last statement, the reviewer guesses that the authors refer to the manufacturing of wind turbines, however, this on the other hand requires external fuels. Thus, the authors should be more precise: Either all statements in this sentence should only refer to the wind turbine during operation or more details on the separate statements should be provided.

This sentence was modified, it now reads: "Wind energy provides several advantages to the sustainable energy grid of the future. Wind turbines produce minimal carbon dioxide or other air pollution, require no external fuel during operation, and require little water throughout their lifetime."

References to relevant literature are missing in the third paragraph. Of course, most references are addressed in the cited review paper, however, some relevant literature should be added for each aspect outlined in this paragraph.

Several citations were added to each of the different optimization methods in this paragraph.

References are missing in the last paragraph. It is only referred to previous papers. Please be more precise what has already been done in other research work and what is novel in this study.

Several citations were added in this paragraph referencing previous work.

Chapter 2 (Wake Model):

Which coordinate system has been used? This does not become clear when the location (x, y, z) is introduced after Equation 2.

This entire section was reworked. The coordinate definition is now included immediately following what is now Equation 1, where x, y, and z are first mentioned. The updated text now reads: "...the velocity at a desired location (x, y, z), where x, y, and z refer to the streamwise, cross-stream, and vertical coordinates, respectively ..."

Just after Equation 4 the lateral and vertical direction are introduced as corresponding to y- and z-direction. This information would be required already earlier, directly after Equation 2. When mentioning this definition directly at the beginning, the question regarding the coordinate system might no longer be so relevant.

Refer to the response from the previous comment.

How are the nine locations (for determining the average rotor wind speed) distributed with respect to the lateral and vertical? Maybe a small figure could help explaining this.

Good question and great idea, a new figure (now Figure 1) was added visualizing these points.

Chapter 5 (Optimization Algorithms):

As the parameter for the maximum number of generations is used in the algorithm code, the specified number for the maximum number of generations should be addressed in the text (Section 5.2).

The text was modified to include a description of the maximum generations, which now reads: "Convergence was assumed after the best performance was within a tolerance of 10^{-3} for 25 generations, or a maximum generation limit of 1000 was met. For the results in this paper, the maximum generation limit was never met."

Why is at the end of Section 5.4 a range of 2-18 turbines defined, whereas in Section 5.5 the spacing constraint is only defined as two times the rotor diameter?

The range 2-18 turbines refers to the number of turbines swept across in the gradient-based optimization, **not** the turbine spacing. Gradient-based optimization does not allow for discrete design variables, such as the number of turbines, so we had to try the optimization with each potential number of wind turbines to see which gave the best answer.

Why are just 2D specified as minimum spacing? The authors themselves state later in Section 6.4.1 that this is not a realistic assumption. To work with more realistic example cases, why is not directly a spacing of for example the mentioned 5D used?

We added the following sentence to the section describing the constraints to justify this decision: "This minimum spacing is on the small side, and is used to exaggerate the differences in the optimal solutions obtained with different objective function."

Chapter 6 (Results):

Within the 1D example (Section 6.1) it is interesting to see that $70/MWh and $90/MWh yield the same optimum of 18 turbines. Maybe it can be further discussed on this result in the text as well.

We thought so to, and think that is a great idea to add discussion on this topic. The following text was added to this section: "The optimal number of turbines increases from 13 to 16 as the PPA increases from $30/MWh to $50/MWh, then again to 18 as the PPA increases to $70/MWh. However, when the PPA increases to $90/MWh, the optimal number of turbines remains at 18. This is because the number of turbines is not continuous, and is only represented by integer values. For a given scenario, different PPA thresholds could be defined, above which the optimal number of turbines would increase by one. From Fig. 6, it appears that the optimal number of turbines is more sensitive at low PPA values, and becomes less sensitive as PPA increases. "

Why was the specific value of $30/MWh taken for the further example cases (mentioned at first in the second paragraph of Section 6.2)? In Section 6.6 and Figure 15 later in the paper, the authors present already a good reasoning for going for $30/MWh. This should be mentioned already (and as well) at this point.

Justification for this decision was added into the second paragraph of Section 6.2: "The PPA was assumed to be $30/MWh, which close to the COE solutions that were achieved, and is within the range of the PPAs of real wind farms from the past few years (see Fig. 16)."

Some statements in the discussions on the results are too generic. Thus, it is not really true that "the greedy and repeated sweep algorithms do not perform very well", as the repeated sweep algorithm finds the second-best solution for the COE objective (Section 6.2.2).

This paragraph was rewritten to two paragraphs, explaining the greedy and repeated sweep algorithms separately. The new paragraphs read:

"For this small wind power plant with the unidirectional wind rose, the greedy algorithm did not perform very well. If found the worst solution for both the AEP and COE objectives, and only the third best solution for the profit objective, but still underperformed by more than 8% compared to the best solution in this objective. This algorithm relies on placing turbines far apart to get the maximum benefit possible at each step of the optimization. Because the domain for this scenario was small, this made it difficult to add additional turbines without violating the spacing constraint. Doing so would require adjusting the location of multiple turbines at once to make room, which is not something this algorithm does. Although the computational expense for the greedy algorithm in this scenario was minimal, its poor performance does not justify its use.

For the objectives with higher turbine density, AEP and profit, the repeated sweep algorithm performed poorly. The answers for these objectives were either the worst or second worst solution found. However, for the COE objective this algorithm performed quite well and found the second best solution, within 0.2% of the best solution. Like the greedy algorithm, the repeated sweep algorithm has a step that relies on greedily placing turbines in the domain if they result in an improvement of the objective function. This algorithm has difficulty placing the turbines without violating spacing constraints for objectives that have many turbines in the optimal solution. For the COE objective however, the optimal number of turbines was much fewer. Thus the repeated sweep algorithm could place the turbines and move them around to a certain extent to find an excellent solution for this objective. The negligible computational expense for this algorithm could justify its use for this scenario, for an objective with optimal turbine spacings that are sufficiently larger than the minimum spacing constraints."

The sentence "For the large wind plant, some of the turbines in the layout optimized for COE are more waked than in the small wind plant" in the first paragraph of Section 6.3.1 needs some more explanations. Based on the Figures presenting the optimal layouts, as well as based on the presented numbers for the wake losses, the reader gets rather a different impression opposite to the statement in the above-mentioned sentence.

Thank you for pointing this out, this discussion was based on a previous result and was an error. We have updated the text explaining instead the optimal profit solution, which now reads:

"One main difference between the results for this large wind plant and the small wind plant is in the optimal solution for maximum profit. For the large plant, the wake losses for the optimal profit solution are significantly higher than for the small wind plant, 12% compared to 5.4% (see Figs. 10 and 8). More turbines can fit within the boundary of the large plant, pushing the plant further down the BOS cost curve shown in Fig. 3. The reduced costs from economies of scale make up for higher wake losses in the optimal solution."

Some additional discussion on the time and number of function calls is missing in Section 6.3.2, as there are also interesting changes compared to the small plant example.

Additional discussion was added to this section focusing on the differences in computational expense between the small and large plant scenarios. We also separated out the discussion of the greedy and repeated sweep algorithms to talk about each one individually.

What is the probability of occurrence taken for each wind direction bin in Section 6.4?

The probability of occurrence is shown with the windrose in what is now Figure 11. The text in the first paragraph of section 6.4.1 was rearranged to make this more apparent.

Especially with respect to the three different objectives, a case with realistic annual wind distribution would be interesting to be investigated and of most relevant meaning for real applications.

We fully agree. More wind resource distributions (of direction probability and speed) would be interesting to examine in this context. We have also thought about different/more realistic wind farm boundaries, perhaps with different discrete sections available for development, coupled optimization of turbine design and layout for different objectives, and co-located wind and solar plants. There is so much that can be done in this space with future work!

The discussion and especially last statement ("However, the optimizer didn't find this solution from the five optimizations that we ran") in the first paragraph of Section 6.4.1 are not satisfactory. There is a significant difference between the results of the two cases (48 versus 54) and furthermore there is a large discrepancy between the results from the different optimizers, presented in Table 4. Furthermore, the last sentence leaves the question on the sensitivity and trustfulness of the results, as the results cannot be repeated and there is some random chance to score one time maybe better or not or just another time.

The following text was added to the end of this paragraph addressing the difference you point out:

"Setting up an optimization run always involves a trade-off between trying to find the best solution and minimizing computational expense. One can imagine two extremes for a genetic algorithm. The first extreme has an enormous population size and very strict convergence criteria. This optimization would theoretically find a very good, maybe the best solution, but at a restrictively high computational expense. The other extreme would have a minuscule population and very lax convergence criteria. This population would converge very quickly, but would lend very little confidence that a good solution was found. For this paper, our goal was to examine overall trends, and not to find the global solution for every scenario and optimizer combination. For a one-off optimization, it may be prudent to run more than 5 optimizations with different initialization of the design space, and maybe tune the optimizer parameters and convergence criteria to the specific problem. However, for this paper our goal was to run a large quantity of optimizations across a range of scenarios, which required us to make some decisions to keep the computational expense reasonable. Even though the optimizers for the large wind farm and unidirectional wind rose scenario failed to find the best solution for the AEP objective, we can stand behind our methodology and have confidence that the general trends we have observed are accurate and valid."

Section 6.7 could go into a separate new Chapter. In general, the reviewer suggests to have another Discussion Chapter (before the final Chapter Conclusions), in which additionally the assumptions made and considered example cases should be investigated in more detail, e.g. with respect to realistic cases (see some comments before on the wind speed distribution or the spacing constraints), required further sensitivity studies, meaningfulness of the results due to the specified optimization settings (limited number of iterations, no repeatability), . . .

Section 6.7 was split out into it's own chapter, and another section was added that addresses these points.

2 Technical corrections:

Chapter 1 (Introduction): A separate paragraph at the end of Chapter 1 (Introduction), in which the structure of the paper is presented, would be useful to help navigating the reader through the paper.

This paragraph was added to the end of the Introduction: "The rest of this paper is outlined as follows: Sec. 2 presents the wake model we used in this paper, and the relevant turbine parameters, Sec. 3 presents the power models, cost models, and how they are combined to form the 3 objective functions we explored in this paper, Sec. 4 describes the different sets of design variables we used to define the locations of wind turbines, Sec. 5 explains the optimization algorithms we used in this paper, Sec. 6 presents and discusses the results from our optimizations, and Sec. 7 contains our conclusions from this work."

Some word repetitions should be avoided (e.g. addressed in the fourth paragraph of Chapter 1).

We went through the paper specifically looking out for and fixing excessive word repetitions.

Chapter 2 (Wake Model): The definition of the parameter I0 might not be relevant, as this is not used in the equations.

This entire section was reworked, now I is used and defined.

Figures and tables should be placed in such a way that text within a paragraph is not separated if this is not required (e.g. Table 1 and Figure 1 or Figure 2).

This issue will be taken care of when the paper is typeset before final publication.

The Greek symbol should be used as well in the text instead of writing phi (Section 3.1).

This change was incorporated.

Throughout the paper it should be ensured that the tenses are used consistently.

We went through the paper to to fix inconsistent tenses.

Chapter 6, first sentence: the word "of" is missing between results and our.

This change was incorporated.

Figure 5: Please complete the legend in the right plot, such as "PPA in $/MWh".

This change was incorporated.

The reference to "previous section" in the second sentence of Section 6.6 is wrong. Please use the reference to the number of the specific section (here to 6.4) to be clear.

Except for one occurrence that was talking about previous sections in general, all mentions of "previous section" throughout the text were replaced with the specific section number they are referring to.