

Response to Reviewers

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

Reviewer 1

First, we would like to 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. It 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 I_0 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.

Reviewer 2

First, we would like to express our sincere thanks for reviewing our paper. We know it is a time consuming process, and we are extremely grateful for your thorough and thoughtful review. We have structured this response to be clear and easy to follow. Each of your original comments will be shown in red, immediately followed by our response in black. Keep in mind that the line numbers associated with each comment correspond to the original draft. After our changes, there will be some offset between the original line numbers and those in the updated manuscript.

3 Structure:

The paper has a potentially high scientific impact and while showing high quality in its results. However, although it is well structured in general, the results on the performance of the different algorithms would become more fluent and dynamic if the sections for each case study (6.2.2, 6.3.2, 6.4.2) were concentrated only in section 6.5, a section by the way already summarizing the performance of the algorithms on all cases.

Thank you! And thank you for your thoughts on the paper organization. The purpose of Sections 6.2.2, 6.3.2, and 6.4.2 is to discuss the differences in the algorithm performance **within each specific problem setup**. On the other hand, Section 6.5 discusses the overall performance of each algorithm across the entire range of problems tested in this paper. Even though there is some repetition of information, we feel that it is helpful to discuss our results from these different perspectives, and keep the discussions of the algorithm performances within each section they are referring to. We have not made any changes to the text based on this comment.

4 Specific comments:

Line 22: crossed out “Because of economies of scale”

This was reworded to say: “Because of economies of scale, utility-scale wind turbines are deployed in groups.”

Line 62: you cannot talk about results in the introduction. You have to talk in general terms from what literature says or elaborate a general rationale.

We feel that providing a small amount of specific results from our paper is helpful for the reader’s experience, and gives a preview of one of the important takeaways of our paper. After consideration, we have decided to keep text from the original manuscript.

Line 67: See Li, Wenwen, Ender Özcan, and Robert John. “Multi-objective evolutionary algorithms and hyper-heuristics for wind farm layout optimisation.” *Renewable Energy* 105 (2017): 473-482. There, they compare the impact of 2 different COE definitions on the number of turbines (Fig. 2).

and

Line 68: Again, argument with some rationale or literature.

The mentioned citation was added, and the text was reworded to say:

“Li et al. began to explore this sensitivity with multi-objective optimization of wind farm layout and turbine number, considering AEP and COE. As part of their paper, these authors studied how different formulations of the COE definition affect the final solutions (Li et al., 2017). Balasubramanian et al. also mention the

importance of appropriately defining the objective for wind farm layout optimization (Balasubramanian et al., 2020). In our paper, we include an empirically based cost model, and explore three different objectives in our single objective optimization formulation to further understand the sensitivity of wind farm layout and turbine number to the objective.”

Line 69: Please include some state-of-the-art dissertation

This was reworded to say: “In past research on wind farm layout optimization, there has been a wide variety of algorithms and problem formulations used, with little consensus on which strategies are the best (Shakoor et al., 2016; Baker et al., 2019; Hou et al., 2019; Balasubramanian et al., 2020).”

Line 89: what does delta actually represent? and which values are you using for it?

and

Line 91: 1) how is that dependency?? 2) which value do you use for I0 and how does it change for further downstream turbines??

and

Line 96: you have u_{inf} , but how do you calculate u_R and u_0 ? Please clarify/unify/simplify expressions.

and

Line 99: how do you define k ? / which values do you assign to k ?

Section 2 of the paper was reworked, with several clarifications and corrections to the original submission. See the revised section of the manuscript.

Line 107: In general for this section: You have to explain when a certain CF can be more convenient for some cases and which one for others.

We understand where you are coming from, thanks for this comment. However, the purpose of our paper is not to present one objective function or another as “correct” for different situations, but to demonstrate that the solution is sensitive to the objective function, indicating that correctly defining the objective of a wind farm layout optimization is important. It is up to those who are designing the wind farm to define an objective function that prioritizes what is most important to them. We have not made any changes to the manuscript from this comment.

Line 155: why?

This sentence was reworded to say: “Because a primary interest of most businesses is to make money, this objective would likely be of more interest to wind power plant developers, as opposed to AEP or COE previously discussed.”

Line 159: I don't see how would you be able to optimize with a non constant PPA.

This was clarified. The new manuscript says: “For this paper, we assumed that the PPA was a constant, as opposed to using time of day pricing, seasonal or yearly PPA adjustments, or including PPA incentives or penalties for power quality. In a given optimization the PPA is defined as a constant, although we will vary this constant to study its effect during different optimizations.”

Line 178: Then why some works consider a full continuous search space?

We're not exactly sure what you mean with this question, but we think you are asking "if computational expense increases with the number of design variables, how can some optimization methods treat the domain as continuous which would essentially be an infinite number of design variables?" The statement that you are referring to with this question is talking about the number of **design variables**, and not about whether they are formulated as continuous or discrete. In the grid formulation, the design space is split into a grid of Boolean design variables, which either contain or do not contain a turbine, meaning that the number of design variables increase with grid refinement. If the location of every turbine were to be defined with it's own continuous x and y coordinate design variable, then the number of variables would remain fixed with the number of turbines and would not suffer from this increase in computational expense with increasing design variables.

Line 190: In this case (BG), will each grid point hold a turbine? or they might or might not hold it? Please clarify

The following was added to clarify this point: "For the gridded domain design variables, the grid defined potential turbine locations which were assigned a Boolean value during the optimization to determine if they had a turbine. In the boundary-grid method, the design variables directly determine the location of every turbine in the farm, meaning there is always a turbine placed at the points defined by the boundary-grid parameterization."

Line 198: could you show the role of (cx,cy) in the problem encoding? These coordinates look strongly redundant, as the rotation of the parallelogram can be defined just with theta.

The entire grid moves uniformly with the center (x, y). The following was added to clarify this point: "The center of the grid is shown as the point (cx, cy), which determines the translation of all points in the grid"

Line 224: please specify the mechanism that adds/removes turbines in this algorithm

Good question. There is no mechanism that is used to specifically add or remove turbines. In a discrete genetic algorithm, the design variables are encoded in binary strings. When the strings for each design variable are appended, the resultant string is called an "individual." The entire "population" is the set of individuals, each representing a different combination of design variables and corresponding to a certain objective function value. Each set of design variable can correspond to a different number of turbines. The genetic algorithm randomly initializes the population, then proceeds to combine individuals in different ways and randomly switch individual bits ("mutate"). Combinations and mutations that improve the objective function are kept, while the others are removed. This proceeds until convergence is met. For the gridded parameterization, the 0's and 1's defining each individual directly correspond to the number of turbines (the number of 1's is the number of turbines in the farm). However, for the boundary-grid parameterization the 0's and 1's map to each of the 11 design variables, which then determine the number of turbines. We have not added any further description in the text, as this information is considered general knowledge about genetic algorithms.

Line 227: what is your criteria when setting the parameters of the genetic algorithm? Justify

We added the following statement to the paper: "We chose the tuning parameters for our algorithm with a combination of trial and error and best practice recommendations."

Algorithm 2: please specify the mechanism that adds/removes turbines

Refer to the previous explanation about the genetic algorithm.

Line 241: Could you at least mention some reference or argument on which you base your method? Otherwise it looks like this algorithm arises from nowhere. If it is the result of a trial and error process, please specify.

The following text was added to give a small background into the creation of this algorithm:

“The creation of this optimizer was inspired by attempting to apply gradient-based optimization principles to discrete design variables. As described below, the algorithm works by comparing adjacent points, and switching the values if it would improve the objective function, which could be imagined as the discrete version of a gradient.”

Line 270: sounds weird

This sentence was reworded to say: “For one case discussed in Sec. 6.2 we also used gradient-based optimization in order to compare the results.”

Line 279: what does this stand for? is this some kind of sequential quadratic programming?

This was reworded to say: “For the results in this paper, we used the open-source SLSQP (Sequential Least Squares Programming) optimizer available in SciPy.”

Line 287: of

This change was incorporated.

Line 298: very informal style

“This figure is particularly interesting when you notice that ...” was changed to “One key takeaway from this figure is that ...”

Line 339: This sounds strange placed here. Perhaps move to conclusions.

and

Line 344: I am completely lost here. You might want to say that smaller projects might be only feasible with a high PPA price..

and

Line 345: smaller plant sizes have higher construction costs??? COE might be higher, but not the overall project or construction costs... please modify

This paragraph was reworked. The new paragraph reads:

“Historically, capacity expansion models have assumed a constant power density that does not vary with the PPA. Not only does Fig. 6 demonstrate the differences between a minimum COE objective and a maximum

profit objective, but it also shows that the cost modeling assumptions can greatly affect the optimal number of turbines in a given land-constrained wind plant. Aggressive carbon reduction scenarios or other renewable energy goals typically result in high PPAs for renewables, which would lead to a higher optimal number of turbines and higher capacity densities for land constrained sites. This has important implications for capacity expansion models and could play a role in the future deployment of wind, as capacity density may often be much higher than is currently assumed.”

Table 2: do not use / unless you are dividing. it is confusing and an erroneous notation.

Tables 2—4 were modified to have these values in their own column

Line 388: By now we only know that this happens using a unidirectional wind rose. Move down or remove

We qualified the statement to say there is a potential of this behavior, rather than a guarantee. The updated sentence says: “If we were to repeat the optimization for a much larger domain we could potentially see results similar to Fig. 5, where having too many turbines could actually be detrimental for AEP.”

Line 392: yes it exists. See Figure 7

This incorrect claim was removed. The new sentence reads: “The turbines are arranged such that waking is minimal.”

Line 393: So far only in this case. Please emphasize.

“We can conclude that . . .” was changed to “For this case, we can conclude that . . .”

Line 400: A good question here would be which PPA value would be equivalent to apply a COE minimization and an AEP maximization. Could you derive that equivalence?

Good question! The short answer is no, we can't. A COE minimization and an AEP maximization would entail a multi-objective optimization, and would result in a set of solutions (a Pareto front) instead of a single point. It would be possible to find a PPA that results in the same solution for some of the individual points on the Pareto front, but that is beyond the scope of this paper.

Line 422: what are exactly function calls? you mean the evaluation via the wake model? Then please specify. On what depends the number of calls?

The following text was modified in the first mention of functions calls in the paper, which now reads “. . . shows total number of calls to the wind farm evaluation, or function calls, required to run . . .”

The number of function calls is an indication of the efficiency of the optimizer, and how many feasible solutions are tested by the optimizer. Runs that quickly converge on a solution and don't need to keep iterating will have relatively fewer function calls, and algorithms and problem formulations that tend to generate infeasible combinations of design variables will have relatively fewer function calls because the objective function is not fully called for a solution that violates the constraints.

Line 435: When a wind energy project of thousands millions is about to be deployed, I don't think 1 week or even longer optimization times are not worth.

This thought is understandable, and is actually pretty common when discussing the computational expense of wind farm optimization. For a single run that is guaranteed to find a solution (which is not at all guaranteed), a time scale of a couple of weeks or maybe even months is not particularly problematic. However, there are other considerations that are important, including setting up the optimization problem, which requires several iterations and small adjustments, wanting to try several optimizations with different parameters such as different turbine models or cost model assumptions, higher fidelity models for which the computation time is much higher per function call, and research (such as this paper) where the turn around time is lower and more configurations are generally tested. For these reasons we can see that computational expense in wind farm optimization, including wind farm layout optimization, is important to consider.

Line 443: what???

This was reworded to say: “Because this algorithm uses continuous design variables and allows full access to the wind plant domain . . .”

Line 446: which ones?

This was clarified, the new text reads: “Even though the greedy, genetic, and repeated-sweep algorithms did not provide as much freedom as the gradient-based optimization . . .”

Line 448: Within gradient-free algorithms, If an algorithm is efficient, the different launches will tend to provide very similar results each other. The extreme contrary case would be a complete random search (which might even be a good solution in some problems). Hence, it would be interesting to see the dispersion in the final results.

This is a great point, this would be interesting to consider in future work.

Line 455: Highlighted but no note written

We are not sure what you meant with this highlight, but we clarified the statement a big further. The new text says: “For the gridded domain, the plant is divided into a 20-by-20 square grid which, because of the increased size of the plant area, maintains the same spacing between grid points as in the small wind plant example.”

Table 4: Is this 44 hours? If so, please mention these kind of computation times in the text.

A new column was added to each of the results tables with the time expressed in hours, in addition to the seconds column.

Line 540: text very redundant. Summarize in 6.5

Refer back to our response to the first comment regarding this organization.

Line 545: too informal

This was replaced with “does not”

Line 579: sounds strange

“be absolutely clear” was replaced with “reiterate”

Line 603: 2

This change was incorporated.

Line 617: Why (on average) PPA decreases over time?

The following sentence was added at the end of Section 6.6 to answer this question: “This decrease in PPA prices in recent years has been the combined result of higher capacity factors, declining installation costs and operating costs, and low interest rates (Wiser and Bolinger, 2019).”

Line 640: You cannot compare to them, because your wake model is assumed to work only on flat terrain or offshore. In complex terrain, orography and high roughness lengths produce high turbulence levels, which make wakes to dissipate more rapidly, this allowing bigger power densities. Please emphasize.

The intent of this paragraph is simply to inform the reader that land-constrained sites do exist and are fairly common, not to compare our results to the specific farms mentioned. We removed part of this paragraph to remedy this.

Line 642: you have any idea why?

and

Line 644: This is not clear. Please reword.

We decided to remove this paragraph.

Line 653: Suggestion: “The number of turbines in each optimal solution showed highly sensitive to the objective function used”.

This sentence was reworded to say: “The number of turbines in each optimal solution was highly sensitive to the objective function that was used.”

Line 667: you mean the 20-by-20 grid? because in line 455 you say it has the same spacing as the 10-by-10 grid

This was reworded to say “However, for a 20-by-20 grid . . .”