Response to Sebastian Sanchez Perez-Moreno

Andrew P. J. Stanley, Christopher Bay, and Paul Fleming

May 2023

Before we begin our response, we want to sincerely express our gratitude. Sebastian, as an expert in this field and busy researcher we know that your time is valuable. We truly are grateful that you have taken your time to review this paper and provide excellent feedback so that this paper can reach the potential that we think it is capable.

We have incorporated your feedback into our revised paper. Below is a summary of your points in blue, and our response and changes in the paper immediately following shown in black

1 General comments

This paper provides a novel approach to co-optimise the wind turbine layout and yaw angles for wake steering. This method greatly accelerates the co-design process by avoiding nested optimisation loops that are computationally very expensive. Together with the understanding of the dependency of the optimal yaw angle with respect to the relative position of the nearest waked wind turbine, these are the biggest contributions from this research. The outcome of this paper is directly applicable in engineering processes in our industry.

It is enourmously appreciated the transparency and repeatability that comes with the code being publicly available. Also want to congratulate the authors on the high quality of the figures, which communicate the outcomes more clearly.

Thank you, we are very excited about this work and how similar approaches could be used for other aspects of codesign.

I would like to read follow up research ideas stemming from this first step, to reduce the uncertainty and risks associated with the assumptions made here (power density, wake model). With more research, industry would deploy this method with more confidence, as the layouts optimised with co-design perform worse than without it, when not operated with wake steering.

We fully agree. There are several opportunities for continued research. We see the primary opportunities as 1) improving the geometric yaw relationship including better quasi-optimal yaw angles, and exploring the validity for different modeling parameters, 2) improving the objective function to capture more complex operation (a wind farm will neither operate fully with wake steering nor fully without wake steering), and 3) expanding a similar approach to other aspects of codesign where we could develop an implicit relationship between layout and some other variable. We will improve this discussion about future opportunities in the revised paper.

Additionally, could the authors mention something with respect to the turbine fatigue loads expected from implementing these yaw angles?

Yes, thank you for this comment. This is lacking in the preprint paper, and we will add some discussion in the revision on the impact of yaw offset on fatigue loads.

2 Specific comments

Line 2: it says "...currently optimized separately... more and more wind plants implement wake steering as their primary form of operation" (gives impression that wake steering is already implemented in wind plants in operation, is this true?)

Great point. We will reword this in the revised manuscript.

Line 54: "...per wind condition". What is meant here, wind speed and direction? Can you be more specific?

This will be changed to "per wind speed and wind direction combination" in the revised manuscript.

Figure 1: Don't the number of wind speed bins count? Or is a single wind speed per direction used to make this figure? This could be stated explicitly. Over 150 hours for a "small" wind farm is already too long for a single wind speed.

Good point. We will clarify this in the revision. For this figure, a single speed bin per wind direction was used (which yes is already far too long).

Line 75: In my opinion, it's somewhat exaggerated the statement that this new approach alone can accelerate the deployment of wind energy and reach goals.

We agree, this is probably overstated. We will reword in the revised manuscript.

Line 85: All downstream turbines are waked (if by small amounts) according to Gaussian models, so it would good to be explicit say what model is used here (looks like a top hat Jensen profile). And can you mention what wake decay factor corresponds roughly to the wake radius formula? I am curious about how "optimistic" this wake radius is and how the conclusions can change depending on this wake expansion factor.

This is a good point, and one that we will clarify in a revised version. You are correct that all downstream turbines would be waked if using a Gaussian wake model. Therefore, for this paper we used two separate wake calculations.

1. Calculate which downstream turbines are waked. You can imagine several ways to do this. If you wanted to use a Gaussian model, you could assume the wake width is everywhere the wake deficit is greater than X%. For this paper, you are correct that we used a Jensen top hat model to determine if downstream turbines were waked. The important thing for the purposes of our geometric yaw relationship is that this is a Boolean value; a turbine is either waked or not.

2. Perform a full wake calculation for the wind farm energy production. This model can be different than the one used to determine "waked-ness," and is in this paper.

Hopefully that is clear, certainly reach out if you have further questions. Regardless, in the revised paper we will add discussion on this topic, including how the results and model may change based on different assumptions.

Fig 2c: are there orange or blue coloured points underneath the black dots? Are there clearly many more black dots than coloured? How did you determine the 1D threshold?

These are good questions. We will revisit the figure generation to make sure all of the points are appropriately represented and many points aren't hiding behind others. If we understand the last question correctly, the 1D threshold was determined manually by inspecting the figure 2c. This is clearly an opportunity for improvement in future research.

Line 121: why do you re-optimise the yaw angle after a layout has been optimised with co-design? What are the quantitative improvements before this step?

This is a good question and is similar to a question from the other reviewer. We will copy part of our response here:

"In short, yes the final continuous yaw optimization is necessary. The geometric yaw relationship underperforms the continuous yaw optimization fairly significantly in the current form. This information in itself is quite interesting however. The geometric yaw relationship does not need to perform very closely to a fully continuous optimization! It just needs to perform well enough that when used in a codesign framework, wake steering is sufficiently accounted for during the layout optimization and a better solution is achieved. In the revision, we will add discussion on this matter, and also clarify that an improved geometric yaw relationship may enable us to find even better solutions, and remove the requirement of a final continuous yaw optimization step."

Do the final re-optimised yaw values correlate nicely with the deterministic yaw angles found during the co-optimisation?

Good question, this relates to the previous question as well. In short, no not really. They are different enough that before performing the continuous yaw optimization, the codesign layout underperforms the sequential one (but, that is not really a 1 to 1 comparison). This is an area for improvement in further development of this geometric yaw relationship. In the revised paper, we will discuss and clarify this last continuous yaw optimization step to clarify all of these elements.

Line 115: doesn't SLSQP require multiple initial conditions to get closer the global optimum? What were the initial conditions (layout) in 3.1?

Good questions. As discussed in the paper, for 3.2 we used 50 different starting layouts. For 3.1 we assumed the design space was simple enough that we only used one starting set of design variables (all turbines spaced 5 rotor diameters apart). In the revised version we will repeat this optimization for multiple starting layouts for example 3.1, and add the appropriate discussion on the topic.

What is the turbine nameplate capacity in example 3.2? What is the plant power density? 2 km x 2 km seems small even for a 3.5 MW WTG rated capacity (14 MW/km2). Is the AEP increase of co-design as big as 0.8% for sites with smaller but more realistic power density (e.g. 5 MW/km2)?

Excellent questions, these are important to clarify. In the revised paper we will add explanation of the wind plant characteristics, as well as discussion on how the 0.8% figure may apply to other sites.

Does wake steering optimisation for existing plants still have any value after finding this geometric relationship? What's the trade-off between "slow" optimisation and finding the optimal angles deterministically?

These are good questions. There is value to have a deterministic relationship between turbine locations and optimal yaw angles. Because turbines are sometime down for routine or unscheduled maintenance, to is infeasible to have a look up table for every possible combination of turbines that may be operational at a given time. One possible solution to this issue is to improve this geometric yaw relationship (or something similar) to better approach the continuous yaw optimization solution. We will add discussion of this in the revised manuscript.

3 Technical corrections

Figure 1: shouldn't it say "with 24 wind direction bins" instead of "with the 24 wind direction bins"?

This will be changed to "using 24 wind direction bin" in the revised manuscript.

Line 115: typo - should say "optimizer" where it says "optimzer"

This will be corrected in the revised manuscript.