Authors' response to Review 2 of "Sensitivity analysis of wake steering optimisation for wind farm power maximisation"

F. Gori, S. Laizet, A. Wynn

May, 2023

Main comment

The authors would like to thank the reviewer for the careful reading of the paper and the in-depth review. The authors' responses (in blue) to all reviewer's comments (in black) are included in this document. All changes are marked in red in the revised manuscript.

General:

1. I'm not sure the inter-model comparisons make a good contribution to the paper. The different models give different results, but often it should be noted that this is the intention, in other words, the reason model development continued was to improve the match to higher-fidelity models of cases similar to the ones explored in this paper. The differences are therefore somewhat the point of the model, and so relate a bit awkwardly to the paper's framing of sensitivity.

   Our motivation for using inter-model comparisons has been clarified with an addition to the second paragraph of Section 2.

   For completeness, we have two reasons for studying inter-model behaviour. First, if a model is used in wake steering optimisation, it is important to ask how complex the model should be to produce broadly the correct decision variables: lower fidelity models may still perform well in an optimisation context. Second, even if two models give similar optimal variables, it is important to know which model gives good solutions more consistently (e.g. under different initialisations).

   The examples given in our submission help clarify these points. First, while the Gaussian and Multizone models can broadly give the same optimal yaw angles, the Gaussian model is much more robust (see Figure 8-b,c). Second, while the GCH model provides more physically accurate optimal yaw angles compared to the Gaussian model (due to secondary steering effects), GCH optimisation results are more variable (Figure 8-c,d), and hence less robust.

   We do not believe that, to-date, the question finding the "correct" model fidelity for wake steering optimisation has been fully resolved. Our intention is that the presented inter-model comparison makes a step towards answering it.

2. I thought the comparisons between optimization strategies were useful and interesting though. However, the description of the turbo algorithm was a little too brief for me to fully understand its approach (I'm also not familiar with several of the terms).

The description of the TuRBO algorithm in Section 3.1 has been extended to provide further insight into Bayesian optimisation and the specifics of TuRBO. Moreover, an illustrative 1D minimisation problem on a toy function using a Bayesian optimisation has been added in Appendix-C.

3. I think in general the paper would benefit if some aspects were condensed and de-emphasized:

   (a) The Jensen model does a poor job at wake steer modeling, I don't think it was really designed for this so the time spent describing the model or looking at results of this model I think could be given over to better use

   We agree that our study shows that the Jensen model performs poorly in wake steering optimisation. However, the Jensen model remains widely used by the wind energy community for wake steering applications [Kheirabadi and Nagamune, 2019, Houck, 2021, Andersson et al., 2021]. For this reason, we believe it is valuable to demonstrate its high sensitivity for at least the simple and medium-complexity cases considered in this study (the 2×1 and 5×5 farms).

   For brevity, however, we have removed the Jensen model from the results for the Horns Rev example.

   (b) Then more generally, the difference of the results between models is also less interesting I think because the models were anyway designed to produce different results, so this does not need to be proved in my opinion.

   Please refer to answer to Comment 1, above.

4. It would be interesting to know more about the Turbo method, not just the general theoretical description of the method, but how specifically it is implemented in this case and why it is out-performing SLSQP

   Section 3.1 has been expanded to give a more detailed description of the TuRBO algorithm. In addition, a simple 1D example has been added to Appendix C to help visualise the iterative behaviour of Bayesian optimisation algorithms (of which TuRBO is an example).

   These added details help to explain the main differences between the two considered algorithms: at each iteration, TuRBO's global approach can potentially sample from anywhere in the design space, while SLSQP typically takes a small step computed using local gradients. This observation underpins the discussion of many of the presented results (e.g. the sensitivity to initialisation demonstrated in Sect 4.1).

5. Is this result specific to Turbo vs SLSQP, or would it be expected to generalize to similarly structured optimizers?

   Our results (see, e.g., Figures 4 and 6) demonstrate that the farm power improvement objective function typically has multiple local maxima and, for the Multizone and GCH models, has either discontinuities or discontinuous gradients. This behaviour implies that significant differences should be expected between the performance of any typical global optimiser (e.g., TuRBO) and the performance of any typical gradient-based optimiser.

6. A diagram of the TuRBO method would be welcome.

An illustration of Bayesian optimisation applied to a simple 1D minimisation problem has been added in Appendix-C. For further examples, please also see [Shahriari et al., 2016, Figure 1] and [Eriksson et al., 2019, Figure 1].

7. Would other optimizers compare interestingly? Genetic annealing? Serial-Refine?

   Consistent with the answer to Comment 6, above, we expect an improvement in performance for global strategies (e.g. genetic annealing) over local searches, which include gradient-based (e.g., SLSQP) or gradient-free (e.g., Serial-Refine) approaches, due to the multi-modal and discontinuous/non-smooth nature of the objective function.

8. In terms of optimization, layout optimization (or coupled layout/control design) is a harder problem for optimizers, since there are many more variables. Would turbo be interesting for those doing research in layout design?

   TuRBO would be useful in this context due to its ability to reduce optimisation sensitivity compared to unconstrained gradient-based optimisation (e.g. the recent use of Bayesian optimisation in farm layout optimisation [Bempedelis and Magri, 2023]). However, a known drawback of Bayesian approaches is the computational cost involved with tuning the Gaussian processes in high-dimensional optimisation problems which require many objective function evaluations. Although TuRBO mitigates this issue by using trust regions, it may also be promising to take inspiration from our findings and add physically-inspired constraints to enable efficient gradient-based algorithms for layout optimisation problems.

Specific:

9. Section 4: Does it make sense to include the Jensen model in these investigations, such as Figure 2? For instance you state:

   "Model sensitivity is caused by the flatness of the Jensen model's downstream power curve apparent in Figure 2-b, that is, $\frac{d}{d\gamma_1}P_{T_2}(\gamma_1)$ is significantly smaller for Jensen than for other models. Flatness arises due to the uniform, or "top-hat", profile of the Jensen distribution (see Figure 1), which results in a lack of sensitivity 210 of streamwise velocity deficit to moderate yaw perturbations."

   The inability of the flat Jensen distribution to model the impact of wake steering (or fit to LES data) was a motivation for the development of the multi-zone model (Gebraad, P. M. O., Teeuwisse, F. W., van Wingerden, J. W., Fleming, P. A., Ruben, S. D., Marden, J. R., and Pao, L. Y. (2016) Wind plant power optimization through yaw control using a parametric model for wake effects—a CFD simulation study. Wind Energ., 19: 95– 114 doi: 10.1002/we.1822.).

   Please refer to the answer to Comment 3, above.

10. Figure 4: This is an interesting result. But I might recommend you also check the slightly misaligned cases (for example 1 deg off in either direction) in addition to the aligned, these problems can be special to circumstances where wind turbine rows are perfectly aligned to the inflow, which is actually not the dominant case in a practical setting.

    The results and conclusions of Section 4.2 are also valid for slightly misaligned cases. Figure 1 of this document replicates the conditions of Figure 8 in the paper but extends the results for slightly misaligned conditions. The objective function statistics are found to be consistent with those presented in our manuscript. Similar initialisation sensitivities, as well as model and optimiser dependencies, are observed for at least 2° of misalignment from fully aligned conditions. As the revised manuscript is already fairly long, we have decided not to add this figure in the revised manuscript and to only have a small comment for slightly misaligned cases.

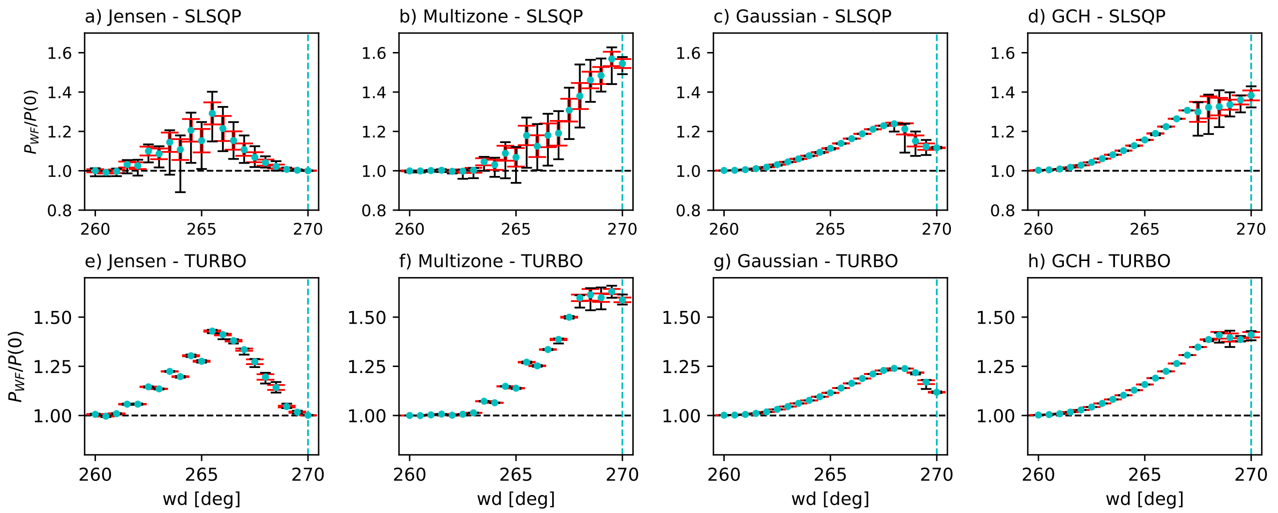

Figure 1: Comparison of the 5 × 5 objective function statistics in wind direction range [260, 270] between SLSQP (**a-d**) and TuRBO (**e-h**) optimisation algorithms for (**a,e**) Jensen, (**b,f**) Multizone, (**c,g**) Gaussian, and (**d,h**) GCH wake models. Mean values are illustrated with error bars for one standard deviation (red) and minimum and maximum values (black). Vertical dashed lines indicate fully aligned conditions, while horizontal dashed lines correspond to no power improvement.

11. Figure 7: Interesting find, is this still an issue in the latest versions of FLORIS?

    The results are independent of FLORIS, as long as the exact same wake, deflection and superposition models are used.

12. Section 4.3: The C2 constraint is a good idea, but is it complicated to carry out in all cases of wind directions? Can it be included in the optimization which turbines are upstream of which in every wind direction?

    The C2 constraint can be easily adapted to all cases of wind directions. For example, a simple permutation of the labelling of the turbines can be performed to obtain columns of turbines aligned with the downstream direction. The C2 constraint could then be applied in an equivalent manner to the one described in our study, using the new labelling.

    A small clarification has been added to the end of Section 4 to explain this point.

References

A. C. Kheirabadi and R. Nagamune. A quantitative review of wind farm control with the objective of wind farm power maximization. *Journal of Wind Engineering and Industrial Aerodynamics*, 192: 45–73, September 2019. ISSN 01676105.

D. R. Houck. Review of wake management techniques for wind turbines. *Wind Energy*, 25(2):195–220, 2021. ISSN 1099-1824.

L. E. Andersson, O. Anaya-Lara, J. O. Tande, K. O. Merz, and L. Imsland. Wind farm control - Part I: A review on control system concepts and structures. *IET Renewable Power Generation*, 2021. ISSN 17521424.

B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas. Taking the Human Out of the Loop: A Review of Bayesian Optimization. *Proceedings of the IEEE*, 104(1):148–175, January 2016. ISSN 1558-2256.

D. Eriksson, M. Pearce, J. Gardner, R. D. Turner, and M. Poloczek. Scalable Global Optimization via Local Bayesian Optimization. In *Advances in Neural Information Processing Systems*, pages 1–12, December 2019.

N. Bempedelis and L. Magri. Bayesian optimization of the layout of wind farms with a high-fidelity surrogate model. 2023. doi: 10.48550/arXiv.2302.01071.