Review WES-2026-62

"Gaussian process surrogate modeling for efficient controller tuning and fatigue load prediction of the helix wake-mixing method"

by Daan van der Hoek, Tim Dammann, and Jan-Willem van Wingerden

May 2026

Main comment

I found the paper interesting and enjoyable to read. It presents a valuable advancement of the helix approach for wind farm control, particularly through the use of GP-based surrogate modelling and the combined consideration of power extraction and structural loading. I especially appreciated the section in which energy yield and loading are combined using Gaussian Processes, as this provides a useful perspective on practical wind farm control. Overall, I believe the manuscript is suitable for publication in Wind Energy Science after satisfactory minor revisions. My comments are provided below, divided into major and minor points.

Major comments:

- From reading the paper, it is not particularly clear to me why wake mixing is a promising approach. It is quite complex to simulate, requiring higher-fidelity simulations, it involves a dynamic actuation, which introduces additional challenges for simplified modelling compared with other control approaches, and it systematically increases loading. I am not saying it is not promising, just that it is not particularly clear from reading the paper.
- Section 2.3: I find this section very specific and not very explanatory for the common reader. I believe it could be helpful to introduce at the beginning of the section a higher-level explanation of Gaussian Process regression, mentioning in a clearer and simpler way the aims, the steps (initial dataset generation, training, then predictions), and the applications.
- Section 2.4: I find this section particularly difficult to follow. I suggest first adding a short summary for the reader to understand the general concept. Then, I suggest describing the various steps in order following the workflow shown in Figure 4 in a simple and systematic manner. As an example, I find the description of the same method significantly clearer in the Conclusions.
- I think there is a substantial difference between employing GPs to represent a dimensional space as best as possible for interpolation (which appears to be the aim for the loading part of the paper) compared to using GPs in the context of optimisation, where the aim is to finding the

optimum of an objective function (aim for the controller tuning). In the latter approach, accurate predictions across the entire design space are not required, but it is more about balancing exploitation and exploration. Please clarify this important point in the manuscript. In lines 400 and 401, "First, the power output of a two-turbine array was modeled as a function of the actuation frequency and amplitude." mentions modelling of the function, so what is the aim? Modelling, optimisation, or both? Please clarify.

- Lines 261-262: "Based on the inferred power ratio, a new frequency is selected and subsequently applied in a new LES run, after which the process is repeated.". Could you please further explain and justify the criteria used to select new frequencies to simulate? As mentioned in the point above, I am confused if the aim of this power production section is modelling or optimisation. From my understanding, determining the location of the optimum control inputs, namely (i) fixed amplitude with variable frequency and (ii) combined frequency and amplitude variations, to maximise power production of the overall farm is at least one of the aims. My doubt is why an optimisation algorithm is not employed, such as gradient-based optimisers, genetic algorithms, or Bayesian optimisation. The latter in particular aims at balancing exploitation and exploration through GPs, effectively locating new data points to simulate through an acquisition function by exploiting the GP-predicted maximum and also exploring areas with high uncertainty. I believe manually locating the new locations is quite inefficient and ineffective. With the two main components of model-based optimisation being the model and the optimiser, I believe the latter is missing in this paper. Please clearly state the employed optimisation procedure and the limitations of the proposed methodology, which of course could be further expanded in future work.

- I believe the paper lacks information regarding the applicability of the proposed methodology to wind farm control. First, the computational time required to generate the datasets and to perform the Gaussian Process hyperparameter tuning is not reported. As control applications move towards larger wind farms and, ideally, closed-loop scenarios, understanding the computational effort associated with dataset generation, surrogate-model training, updating, and querying is important for assessing the practical appeal of such approaches, both in academia and in industry. Second, I suggest providing more detail on the size of the datasets used, for example the number of training points, as is already done for the loading scenario. Finally, Gaussian Process training scales cubically with the number of training samples, which can make their application challenging for larger datasets. Since the paper discusses the possible inclusion of additional dimensions or operating conditions, such as those mentioned in line 305, I believe it would be valuable to acknowledge the scaling behaviour of GPs and the associated challenges and limitations that may arise. I am not necessarily asking the authors to provide all of this information in full detail. Rather, I suggest that the authors address these points in the manuscript, at least by acknowledging their importance and discussing their implications for the practical applicability of the proposed methodology.

- In general, I found quite a few inconsistencies in terminology, symbols which are not introduced, and the nomenclature not is not always clear. Please address these throughout the manuscript.

- Several figure captions are not self-contained, as the reader is required to go back to the main text for clarifications. Please address this.

- I believe the use of subfigure labels, together with explicit references to them in the text, would significantly improve the readability of the paper. It can be quite challenging to follow the discussion of the results and trying to understand which subfigure it refers to. I found this to be particularly relevant for Figure 6, 7, 9, and 11.
- Are the dataset and code associated with this paper publicly available? I think this would benefit the wind energy community and improve reproducibility.

Minor comments:

- Please include a sentence in the abstract summarising the results from the loading assessment. Currently, only the method capabilities are mentioned.
- Lines 19-20: "Wind farms... wakes". Please slightly elaborate, or clarify turbulent wakes arise from aerodynamic interactions between turbines, for non-expert readers.
- Line 21: I would replace "or" with "and/or" as the two objectives are not necessarily mutually exclusive.
- Lines 56-57: "load surrogate models offer a more efficient solution". Please clarify what it is more efficient than. LES coupled with aeroelastic simulations? Also, consider mentioning the distinction between efficiency in prediction and the requirement for a dataset generation and a surrogate model training.
- Lines 58-59: Did Dimitrov et al. create a dataset for these site-specific inputs and then evaluate different surrogate models? In line 61, "this dataset" is mentioned but never introduced.
- Line 64: Is the LUT constructed from data generated using aeroelastic simulations? The description is generally vague and could be confusing for some readers.
- Line 81: Please replace "yield" with "energy yield" for clarity.
- Line 89: I believe it can be helpful to clarify early in the paper what variations in conditions the array will experience. Are the conditions atmospheric, operational, both?
- In Figure 1 caption, please mention the temperature profile shown in the figure.
- Line 145: Please introduce the symbol for the frequency of the pitch signals when first mentioned.
- Lines 145 and 148: To avoid confusion, please keep a consistent terminology. In this case, rotational frequency of the turbine is referred to as rotor frequency just a few lines after. Please check the terminology throughout the manuscript.
- Line 179: Load channels are not introduced here or mentioned before. Please add a short explanation.
- Line 188: How is "no wake overlap" defined considering the variations in operating and atmospheric conditions? I would imagine that already between the helix and baseline cases, wake expansion would be different. Please clarify.

- Line 185, 187, and 188: In the lines mentioned, please also express the grid and other measurements normalised by the rotor diameter for clarity.
- Line 191 and 192: "The OpenFAST simulations for the loads database also include the turbine using different control strategies to assess the impact on loads.". Weren't the simulations with control actions just specified above? I might have misunderstood, please clarify as this is not clear to me.
- Line 260: Please define the power ratio as it has not been previously introduced or explained. Also, please be consistent with the power ratio definition in the captions of Figure 6 and 7.
- In Figure 6 caption, please mention what the different LES samples correspond to in terms of variations, range, and interval if applicable.
- Please be consistent when referring to the power ratio (for example, see the difference in y-axis between Figure 6 and Figure 7).
- Figure 7: labels (a), (b), and (c) are mentioned in the caption but are not shown in the actual figure. (I find this plot very nice!)
- Line 274: I believe "hyperparameter" is incorrect here as the inclusion of the pitch amplitude corresponds to an additional GP input dimension (or input variable). Please correct this accordingly.
- Lines 279-281: In "The optimum indicates that at some point, the power loss of the actuated turbine (T1) can no longer be compensated by the power increase of the downstream turbine (T2).", shouldn't it be the opposite? The power gain from T2 exceeds (at its maximum) the power loss experienced by T1.
- In Figure 7, the uncertainty of the GP optimal prediction is quite high (ranging from approximately 5% to 10% for the overall power ratio). This is quite a large difference in prediction when quantifying the capabilities of a wind farm control strategy. As mentioned in one of the major comments, the use of an optimiser could reduce this uncertainty effectively. Please clearly state the limitation of the optimum prediction in the manuscript.
- Figure 12: Please define the relative power increase in the caption, rather than only in the main text.
- Line 367: please add more details regarding the simulation environment used to produce the "generator power" dataset. In line 395, it is specified that it is different from AME-Wind, so please explain here when first introduced.
- Wake steering is included in the tuning data, but I believe it is not presented or discussed in the results. What is the aim of running these additional simulations? I might have missed it, in which case I apologise. Please clarify