REVISION TO MANUSCRIPT DRAFT

Wind Energy Science Discussion

Lidar-assisted model predictive control of wind turbine fatigue via online rainflow-counting considering stress history

The authors would like to thank the two reviewers for their time and for the useful feedback. All inputs that they provided have contributed to the improvement of the paper.

A list of point-by-point replies to the reviewers' comments is reported in the following.

Reviewer RC2

1. **[Reviewer]** *The paper provides a comprehensive study of lidar-based Model Predictive Control using online fatigue damage estimation and is deemed of high scientific value. Lidar-based wind speed estimation, MPC of wind turbines and fatigue are introduced in a pedagogical manner. Then, the novel lidar- and fatigue damage-based algorithm PORFC is presented in detail, including a thorough description of the tuning procedure. A comprehensive cross-comparison of control algorithms is then performed showing the potential of the approach.*
   [Authors] Thank you very much for this positive assessment.

2. **[Reviewer]** *A request for revision lies in how the overall uncertainty in the fatigue cost model is considered. Although it is understood that using fatigue damage directly brings us one step closer to monetary costs compared with indirect fatigue damage metrics, the concept of "profit" directly subtracting fatigue cost from revenue as main quantitative key performance indicator is inappropriate. The uncertainty lying in this last step impairs the excellent scientific value of the paper. In general, the paper should be kept neutral and objective, not trying to subjectively sell the method. More into detail, reducing the cost to CAPEX of one single component is overly simplifying the problem. Fatigue costs involve holistic system-level analysis over the entire lifetime of the turbine. For instance, it is expected that O&M costs on drivetrain and pitch actuation system components would change the conclusions dramatically (the term "worth investigating" used in conclusion is a mild statement). As the overall gain in profit is not so evident in terms of order of magnitude, suggesting a reduction in power production in order to increase profit by decreasing fatigue should be taken with caution. However, this can be fixed without major revision through targeted changes in writing style keeping objectiveness in mind and emphasizing on revenue and fatigue costs separately, rather than on profit. The latter may still be kept for illustration, but the rough underlying assumptions should be clearly mentioned. The promising potential of the method should rather be promoted through control algorithms that do not (or not much) reduce revenue (in other words, revenue and fatigue costs should be weighted unevenly, ideally with varying weight to encompass multiple scenarios and handle the aforementioned uncertainty).*
   [Authors] We agree with the proposed changes, and reduced the emphasis on this simplified definition of profit. Consequently, the discussions of the fatigue metrics in Sections 3.1 and 3.2 , of the optimization problem in Section 3.5, of the performance indicators in Section 5.1.6, of the tuning in Sections 5.2.2 and 5.2.6, and of the results in Sections 5.3.1-5.3.4, as well as the conclusion in Section 6.1 have been adjusted accordingly.

3. **[Reviewer]** *A relation between the Moving Horizon Estimator and the much more standard Kalman filter would be welcome, both being optimal state estimators*

[Authors] Thank you very much for this suggestion. We are happy to add a relation to Kalman filters. As described by [1], the widely used Extended and Unscented Kalman filters (EKF, UKF) have significant disadvantages compared to Moving Horizon Estimators (MHE). First, Kalman filters are minimum-variance state estimators for linear dynamic systems with Gaussian noise; although assumptions on linearity and Gaussian noise behaviour can be relaxed, MHEs are formulated as more general non-linear optimization problems over a time horizon, which represent a natural complement to the similarly general non-linear optimization-based formulations behind MPC. Second, the inclusion of constraints in state estimation problems can be important to prevent non-physical results. The inclusion of state constraints is possible in Kalman filters, but not straightforward, and non-linear constraints lead to loss of optimality of the filter and may generate different results, depending on the formulation [4]. In contrast, state constraints can be explicitly and readily set in an MHE [1]. Although state constraints are not employed in the estimator used in the present work, this feature may become relevant in future research.

Third, Kalman filters are one-step recursive estimation methods, and thus have to start operation with only one measurement time sample. In case of large initial state errors, this can lead to inaccurate estimation and possibly to the divergence of the filter [1]. MHEs are less vulnerable to this danger since, right from the start, they take into account an estimation horizon comprising numerous measurement samples.

Consequently, it can be argued that indeed both MHE and KF are optimal, but on different amounts of input data.

An advantage of the recursive nature of Kalman filters is their significantly lower computational effort compared to MHEs. However, in the wind turbine context, the computational effort of MHE is low enough if suboptimal solution methods like the "Real-time iteration" are utilized [3].

This reasoning has been added to Section 4.

4. **[Reviewer]** *Even though OpenFAST has a much higher fidelity than most control-oriented models, it should not be presented as high-fidelity model. It is an engineering model, commonly defined as mid-fidelity, high-fidelity being reserved for CFD.*
   [Authors] Thank you very much for this suggestion. We changed the wording to "mid-fidelity".

5. **[Reviewer]** *Paragraph 5.2.2/Filtering: the IIR filter used behind the filtfilt function has not been specified (a Butterworth filter of order 2 would be typical for this purpose). The moving average filter may perform better than the particular IIR filter used here, but it would be surprising if this can be generalized, so the comparison should probably not be presented here. This said, zero-phase IIR filters are indeed inappropriate for online use on very short records, and the horizon length is an easier tuning parameter to work with than a cutoff frequency in this case, so the moving average filter is a good choice.*
   [Authors] Thank you very much for this suggestion. Within filtfilt we used a first-order Butterworth filter, as also done in [2]. We included this information and the following remark in Section 5.2.2: "It should be noted that this superiority of the moving average filter has only been observed for the present lidar and wind turbine configuration, and cannot be generalized. Thus, for another configuration, the comparison should be repeated."

6. **[Reviewer]** *Paragraph 5.1.6: missing right parenthesis after "based on a realistic piecewise S-N-curve"*
   [Authors] Thank you for pointing this out. We have corrected this typo.

7. **[Reviewer]** *Paragraph 5.2.2: should it be "if the MPC internal and the plant model are NOT matching"?*
[Authors] We believe that this actually is not a typo. Our observation has been that if the MPC is controlling a plant model consisting of the same differential equations like the MPC-internal model (i.e. plant model and MPC-internal model are matching), oscillations can be perfectly damped and fatigue can be reduced to unrealistically low levels.

Reviewer RC3

1. **[Reviewer]** *This is a well written paper about the author's contributions to the field of model predictive control for wind turbines. Extensive detail is given to the problem formulation and setup, where fatigue calculations, model predictive control, and lidar-based estimation are all described very thoroughly. This PORFC algorithm which uses the lidar measurements and employs online rainflow counting is shown to produce favorable results in a comprehensive study comparing the proposed controllers against the standard NREL 5MW control algorithm. Generally, I think this is a well written, novel, and thorough presentation of the completed work.*
[Authors] Thank you very much for your positive judgement.

2. **[Reviewer]** *I am curious as to why a moving horizon estimator was used instead of a Kalman filter. At least some commentary on this would be helpful, especially given the estimator offsets seen in Figure 9.*
[Authors] We decided to use a Moving Horizon Estimator since it would allow to straightforwardly introduce constraints, e.g. on the states. In the present work, this capability has not been exploited. However, we would like to have this freedom for our future research. Next, we wanted to avoid inaccurate or divergent estimation output, as it can occur with Kalman filters particularly at the beginning of operation (see explanation for Reviewer 2). This reasoning has been added to Section 4.
Finally, we believe that the estimator offset rather stems from the inaccurate system model, since this offset also pertains in the MPC prediction, which is based on the same system model. Instead of changing to another estimator type, this may provide further motivation for online model parameter estimation and adaptation. This interpretation of the estimation offset has been added to Section 5.2.4.

3. **[Reviewer]** *I agree with the other reviewer's comments regarding using "profit" as a primary metric of interest. In my opinion, I don't think a major revision is necessary, rather a more explicit acknowledgement of the blanket nature of a metric like this might be more useful. I think it is fine to propose an objective function that subtracts some estimated fatigue cost from revenue, but I think it is important to discuss the limitations of such a metric. In presenting your results and conclusions, it is important to emphasize that the novelty of this work lies in how you are including online fatigue estimations in MPC, and that you are simply using the proposed cost function to demonstrate it. The novelty is not in the cost function in and of itself, and thus, the function could be changed accordingly. You talk a little bit in the last section about how the cost function could be modified for "different business cases", but this is only suggested as a possibility late in the paper.*
[Authors] We fully agree that this profit cost function is very simplified and limited, and that this should be mentioned at early stages in the paper. Consequently, the discussions of the fatigue metrics in Sections 3.1 and 3.2 , of the optimization problem in Section 3.5, of the performance

indicators in Section 5.1.6, of the tuning in Sections 5.2.2 and 5.2.6, and of the results in Sections 5.3.1-5.3.4, as well as the conclusion in Section 6.1 have been adjusted accordingly.

4. **[Reviewer]** *I am not particularly keen on comparisons of controllers that have some inherent optimality/objective (like MPC) to highly generalized (and old) reference controllers. The results of this paper suggest a ~20% profit increase over the NREL5MW controller, but the NREL 5MW controller has not been tuned or optimized with any explicit consideration of fatigue, let alone "profitability".*
In an ideal world, it would be nice to see the novel MPC algorithm compared to a reference controller that has been explicitly re-tuned to address the same, or similar, objectives. In an even more ideal world, the reference controller would be more "modern" and have features such as thrust limiting, as this would provide a more realistic basis of what sort of advantages the MPC controller provides over more modern standards (my, obviously biased, opinion is that NREL's ROSCO controller is a good candidate for this sort of study, though a re-tuned in-house or modified NREL 5MW controller could be sufficient).
That said, I understand that it may be unrealistic to re-tune and re-optimize a reference controller for each project. I don't think this is necessary for the work presented in this article, but I do think the rudimentary nature of the NREL 5MW controller deserves some mention in this paper. And of course, I hope you keep these sentiments in mind in your work going forward.
[Authors] We fully agree that the sole comparison to the NREL5MW controller in terms of our definition of profit would not prove the effectiveness of our controller. The NREL5MW controller has been utilized since it has been a comparison object in many other publications. Therefore, we hope to facilitate cross-comparison between our and other publications. The actual comparison object within our study clearly should be the TTVP MPC, which also has been utilized in many other studies. This reasoning is added to section 5.1.6. The comparison to a modern reference controller indeed would be very interesting. Therefore, this is added to the outlook in Section 6.2.

We have taken the opportunity to make several small editorial changes to the text, in order to improve readability. A revised version of the manuscript is attached to the present reply, with additions highlighted in blue and deletions in red.

The authors

References:

[1] Rawlings, Mayne, Diehl: Model Predictive Control - Theory, computation, and design, Nob Hill Publishing, 2017.

[2] Schlipf: Lidar-assisted control concepts for wind turbines, Dissertation, 2016.

[3] Gros, S., Vukov, M., and Diehl, M.: A real-time MHE and NMPC scheme for wind turbine control, in: 2013 IEEE 52nd Annual Conferenceon Decision and Control (CDC), pp. 1007–1012, 2013.

[4] Simon, D.: Kalman filtering with state constraints: a survey of linear and nonlinear algorithms, IET Control Theory and Applications, 4,1303–1318, https://doi.org/10.1049/iet-cta.2009.0032, 2010.