

Dear Referee #1,

First and foremost, thank you for reviewing our paper “*Uniform Blade Pitch Misalignment in Wind Turbines: a learning-based detection and classification approach*” (Preprint wes-2025-153).

We are glad that your feedback was positive.

We have carefully revised the original manuscript to accommodate all your suggestions, taking this opportunity to make minor improvements to the text. In the amended version of the paper, the changes marked in orange refer to your comments. The point-by-point reply to your comments is reported here below.

GENERAL COMMENTS

The paper proposed a data-driven method for detecting collective turbine pitch offsets. The topic is relevant to wind energy, and the methods are well described. Some minor changes to the manuscript are advised in order to improve its impact and readability.

[Reviewer] The literature review is a bit too concise, and the authors should spend some time finding more references for the topic (at least 10). The topic of wind turbine proactive monitoring is quite vast, and it feels that many important works were not mentioned.

[Answer]

We thank the reviewer for this comment. We agree that the original literature review was too concise and did not sufficiently reflect the breadth of existing research on wind turbine proactive monitoring. The manuscript has now been revised to substantially expand the related works section. In particular, more than ten additional references have been included, covering both physics-based and data-driven approaches for condition monitoring and fault detection in wind turbines, with a specific focus on pitch system monitoring and SCADA-based methods. The revised literature review now provides a broader and more balanced overview of the state of the art and better positions the proposed contribution within the context of existing work.

[Action] Section Related Works has been modified accordingly

[Reviewer] The titles of sections 2 and 3 are similar, which may disorient the reader. “Preliminary” and “Exploratory” are also repeated. The paper would benefit from a restructuring into a more common layout including only Introduction, Methods, Results, and Discussion.

[Answer] We thank the Reviewer for this insightful comment. We agree that the similarity between the titles of Sections 2 and 3, as well as the repeated use of the terms “Preliminary” and “Exploratory”, may lead to confusion.

[Action] We have restructured the manuscript following a more standard and widely adopted layout. The paper is now organized into Introduction, Methods, Results, and Discussion, with the content of the former Sections 2 and 3 redistributed accordingly. Section titles have been revised to clearly reflect their scope and avoid redundancy. We believe that this reorganization significantly improves the clarity and readability of the manuscript.

[Reviewer] A point that was skimmed over is the definition of the inflow. In particular, the

turbulent class is just briefly mentioned, but it would be important to understand wind conditions and terrain type the inflow turbulence represents. The accuracy of each of the 4 models should be reported in a more consistent way..

[Answer]

We thank the reviewer for highlighting this important aspect. We agree that the description of the inflow conditions was too concise in the previous version of the manuscript. The turbulent inflow was generated according to the IEC 61400-1 (Ed.3) Normal Turbulence Model (NTM), considering Turbulence Class C. According to the IEC standard, Class C corresponds to a reference turbulence intensity $I_{ref} = 0.12$, representing relatively low ambient turbulence conditions. The turbulence intensity varies as a function of the mean wind speed following the IEC prescription. Such conditions are typically associated with low-roughness or offshore sites. To ensure statistical robustness, multiple turbulence realizations were generated by varying the turbulence seed across simulations.

[Action]The manuscript has been revised accordingly in Section 2 to provide a clearer and more complete description of the inflow characteristics. The accuracy of each of the 4 models has been reported in a more consistent way in a concise table.

[Reviewer] All sections should have a table reporting recall, precision, F1 score, and support, and possibly also a confusion matrix (particularly when using the binary detection, like in Region II).

[Answer]

We thank the Reviewer for this helpful suggestion. Tables reporting recall, precision, F1 score, and support were already included throughout the manuscript to quantitatively assess the classification performance. Following the Reviewer’s recommendation, we have now added the confusion matrix for the binary detection case in Region II

[Action] We have now added the confusion matrix for the binary detection case in Region II

[Reviewer] Also, the existing tables and figures are missing important information in their captions, so it is hard to pinpoint the data they refer to.

[Answer]

The captions of all tables and figures have been revised and expanded to include clearer and more detailed descriptions of the corresponding datasets, experimental conditions, and analysis context. This ensures that each table and figure is self-contained and that the reader can easily identify the specific data and results being presented without ambiguity.

[Action] The captions of all tables and figures have been revised.

[Reviewer] To conclude, it would be interesting to discuss the real-world challenges that implementing this method on a real turbine would inevitably bring about. Think, for instance, about how you would train a model for full-scale application (using SCADA data or still simulations?), how different turbulence conditions may affect accuracy, if industry cares more about false negatives (i.e., missed alarms) or false positives (i.e., unnecessary inspections), etc..

[Answer]

We thank the Reviewer for this insightful comment. In real-world deployment, several practical aspects should be considered. First, model training would likely rely on a hybrid strategy combining high-fidelity simulations (digital twin) and SCADA data, since uniformly distributed pitch misalignment events are rare and rarely labeled in field data. Domain shift between simulated and real turbulence, controller tuning, and sensor noise may affect generalization and would require site-specific calibration. Second, turbulence intensity and wake interactions may increase variability in the collective pitch demand, thus requiring adaptive baselines and potentially longer time windows, as already observed in the NTM scenarios. From an industrial perspective, missed detections (false negatives) are generally more critical, as prolonged operation under uniform misalignment may lead to long-term fatigue damage and efficiency losses. However, excessive false positives should also be avoided to prevent unnecessary inspections and loss of operator trust. Therefore, in practice, threshold tuning would likely prioritize high recall while maintaining an acceptable false alarm rate.

[Action]

SPECIFIC COMMENTS

[Reviewer] L91: Should “Normal Weather Prediction” be “Numerical Weather Prediction”. This is what NWP stands for, generally. Please add a reference for this as well.

[Answer] We thank the Reviewer for highlighting this point. In the original manuscript, we had indeed referred to Numerical Weather Prediction, and we have now revised and replaced the corresponding phrase accordingly.

[Action] Numerical Weather Prediction has been substituted in L91

[Reviewer] L106: Please define $1\times\text{Rev}$ and make its formatting consistent.

[Answer] We thank the reviewer for this comment. The term $1\times\text{Rev}$ is now explicitly defined as the once-per-revolution harmonic associated with the rotor rotational speed, and its notation has been made consistent throughout the manuscript.

[Action] $1\times\text{Rev}$ has been defined and formatted in the text

[Reviewer] L121: It is not clear why the pitch regulator would still maximize power in the case of a pitch misalignment. It sounds like the pitch offset, as defined, is indeed the difference between the pitch command and the actual pitch, so it should lead to suboptimal performance in Region II.

[Answer] We thank the Reviewer for raising this important point. In our definition, the pitch misalignment is modeled as a uniform offset in the blade pitch angle, rather than as a tracking error between the commanded and measured pitch. In other words, the controller still operates according to its nominal control law and computes the collective pitch demand as usual, but the aerodynamic equilibrium is altered due to the imposed offset. In Region III, the collective pitch controller actively regulates the turbine to its rated power. Therefore, even in the presence of a uniform offset, the controller compensates by adjusting the demanded collective pitch, resulting in comparable power production but different pitch-demand statistics. In Region II, the pitch controller is not active, since torque control regulates the turbine operation while the blades are kept at their fine pitch position. In this case, a uniform offset may indeed shift the operating point away from the optimal $C_p-\lambda$ condition, potentially leading to suboptimal aerodynamic performance. However, due to the relatively flat shape of the C_p curve around its maximum, small uniform offsets ($\pm 0.5^\circ$ – $\pm 2^\circ$) do not produce large observable deviations in power, which motivates the use of alternative indicators (mechanical moments) in this region.

[Action]

[Reviewer] L130: “fro” instead of “from”. L174: There is a “1” in the text. L176: Missing parenthesis for citation.

[Answer] All the indicated typos and formatting issues have been corrected.

[Action] The spelling error at L130 has been fixed, the extraneous character at L174 has been removed, and the missing parenthesis in the citation at L176 has been added.

[Reviewer] L186: It is recommended to turn the definition of the metrics into a table or formulas that use the more common definitions of False/True positives and False/True Negative (see e.g., <https://www.osti.gov/biblio/1769877>)

[Answer] We thank the reviewer for the suggestion. The definition of the performance metrics has been revised and reformulated using the standard notation based on True Positives, True Negatives, False Positives, and False Negatives.

[Action] The manuscript has been revised to explicitly report the definitions of the evaluation metrics in terms of true and false positives and negatives, following standard conventions.

[Reviewer] Fig. 4: These plots are a bit unclear: • Are gray dots all the cases? • Where are the blue dots on the left panel? • It is unclear when anomaly and anomaly predictions superimpose. It would be better to use a cross and an empty circle as markers or play with transparency, for example.

[Answer] The gray dots indeed represent all the analyzed cases. This has now been explicitly clarified in the figure caption. We agree that the superposition of anomaly labels and anomaly predictions could reduce readability.

[Action] The figures have been updated by using distinct marker styles and improved visual differentiation to clearly distinguish true anomalies from predicted ones.

[Reviewer] L215: “model” is a typo.

[Answer] Corrected.

[Action] The typographical error at L215 has been fixed.

[Reviewer] L261: The concept that a 10-minute window leads to better classification performance was already introduced at L256, consider removing it

[Answer] Corrected.

[Action] The redundant statement at L261 has been removed to avoid repetition, as the concept was already introduced at L256.

We look forward to your kind reply, and in the meanwhile, we send our warmest regards.

Sincerely yours,

Sabrina Milani, on behalf of all Authors.

Dear Referee #2,

First and foremost, thank you for reviewing our paper “*Uniform Blade Pitch Misalignment in Wind Turbines: a learning-based detection and classification approach*” (Preprint wes-2025-153).

We are glad that your feedback was positive.

We have carefully revised the original manuscript to accommodate all your suggestions, taking this opportunity to make minor improvements to the text. In the amended version of the paper, the changes marked in orange refer to your comments. The point-by-point reply to your comments is reported here below.

GENERAL COMMENTS

I think a more detailed review of the literature, further consideration of the physics (i.e. to help describe what the ML algorithm is likely identifying), and consideration of how this might be applied in reality and the potential limitations in doing so would be beneficial. The paper is otherwise well written. Thanks

[Reviewer] Line 117 do not in italics. Line 122, worthwhile instead of worthless.

[Answer]

We thank the Reviewer for the suggestion. We agree that the term should not be entirely in italics. However, worthless was intentionally used to stress that, in this specific case, the power curves do not convey meaningful information.

[Action] We changed the text accordingly.

[Reviewer] Figure 2. What’s the sampling rate?

[Answer] The sampling rate used to generate Figure 2 is 5 Hz.

[Action] We have now clarified this point in the text of the manuscript.

[Reviewer] Figure3. Can you show the pitch curve too? i.e. pitch vs wind speed? This should highlight the physics of the pitch offset above rated. Can you also state what the anomaly shown is, e.g. 2 degree positive offset?

[Answer]

We thank the reviewer for the suggestion. We agree that showing a pitch–wind speed characteristic helps to better highlight the physical effect of a uniform pitch offset in Region III (above rated), where power regulation is achieved through collective pitch control. In the revised manuscript, we added an additional plot reporting the mean collective demanded pitch versus wind speed for healthy and uniformly misaligned cases. We clarify that the pitch signal shown corresponds to the controller collective pitch demand (used by the turbine to compensate the altered aerodynamics), since the direct “faulty pitch” is not an independent measured signal in the considered setup. We also explicitly state in the figure caption and text the magnitude of the considered anomaly (e.g., a uniform -1.5deg pitch offset applied to all blades).

[Action] Figure 3 has been modified accordingly.

[Reviewer] Above rated, for the non-turbulent case, I would expect it to be trivial to compare the set point of the pitch misalignment between healthy and anomalous cases, and so determine the offset, since the controller will adapt to the offset accordingly and will simply shift the curve up/down by the offset.

The issue with a real turbine is that you do not know what the healthy state is, and even if different offsets are applied, it wouldn't be clear which one is "best" above rated power, because the power output will be the same.

Below rated, then we'd look for the offset that maximises generation. My concern would be that in reality noise would mask small changes in torque/power output, so might be difficult to detect. In addition, I don't see the advantage of looking at bending moments over just using power directly. Bending moments are likely to have much higher uncertainty in the real world than the power measurements.

Have you also considered looking at using SCADA data, which is typically stored at 10-min resolution?

Would you be able to offer some commentary on these points?

[Answer]

We agree that, in an ideal non-turbulent scenario with direct access to pitch measurements and a known healthy reference, the pitch offset above rated wind speed could be trivially identified, as the controller would adapt by shifting the pitch curve accordingly. In this case, the offset could indeed be estimated by comparing the set-points between healthy and anomalous conditions. However, this situation does not reflect realistic operating conditions, where neither the healthy reference nor the absolute pitch offset is known a priori. Moreover, as correctly stated by the Reviewer, above rated wind speed the controller enforces power regulation, so different pitch offsets may lead to indistinguishable power outputs. This intrinsic ambiguity above rated power is precisely one of the motivations of this work. As discussed in the manuscript, uniform pitch misalignment has a symmetric effect and may not manifest through power deviations, even in the presence of a fault. For this reason, power alone cannot be used to identify which offset is "best" or even to detect the anomaly. Below rated wind speed, an optimal pitch offset could in principle be inferred by maximizing power production. However, as the Reviewer correctly points out, in real turbines small changes in torque or power may be masked by measurement noise and environmental variability. This limitation is explicitly acknowledged in the paper and motivates the use of longer time windows and complementary indicators, as shown in the NTM analysis. While power measurements are typically more accurate than load measurements in operational turbines, power is strongly regulated above rated wind speed and therefore becomes weakly informative for uniform pitch misalignment. In contrast, bending moments are directly linked to aerodynamic loading and reflect the redistribution of forces induced by pitch offsets, even when power remains unchanged. We acknowledge that bending moments may be affected by higher uncertainty in real-world measurements; however, they provide sensitivity to fault signatures that are fundamentally invisible in power signals. This trade-off is now more clearly discussed in the manuscript. Regarding the use of SCADA data, we agree that SCADA measurements at 10-minute resolution are highly relevant for practical deployment. However, experimental SCADA data are not available in the context of this study. The results presented in the manuscript are based on simulation data, including aggregated time windows of up to 10 minutes, which are intended to resemble the temporal resolution typically associated with

SCADA systems. The purpose of this analysis is to assess the robustness and detectability of uniform pitch misalignment under realistic temporal aggregation, rather than to provide a full experimental validation. The application and validation of the proposed approach using real SCADA data is therefore considered a natural and important direction for future work.

[Action] -

[Reviewer] Check colors in the figures

[Answer]

We have revised all the figures in the manuscript.

[Action]

We have revised all the figures in the manuscript.

[Reviewer] Figure 8. A real turbine will be operating day-in day-out, so having a longer time window is not a limitation, especially for something like pitch misalignments, that are generally static (i.e. don't vary with time).

[Answer]

We thank the Reviewer for the comment and we agree with this observation. Indeed, for real turbines operating continuously, the use of longer time windows is not a practical limitation. On the contrary, for quasi-static faults such as uniform pitch misalignment, longer observation windows are naturally available and beneficial.

As shown in Figure 8, increasing the time window improves the robustness and accuracy of the detection and severity assessment, particularly under turbulent conditions. This result therefore supports the practical applicability of the proposed approach rather than representing a limitation. We have clarified this point in the manuscript to avoid ambiguity.

[Action] -

[Reviewer] I think potentially you need to do a search of wind turbine control literature. Perhaps searching for C_p - λ or performance coefficient optimisation. I think this has been historically approached as a control optimisation problem (i.e. what is the best fine pitch angle to maximise production), rather than offset/error identification problem.

Extremum Seeking Control (ESC) seems to come up a few times for this: e.g. "Maximizing Wind Turbine Energy Capture using Multivariable Extremum Seeking Control."

There's also an example of doing this in reality, which is pretty cool: "Experimental Evaluation of Extremum Seeking Based Region-2 Controller for CART3 Wind Turbine."

[Answer] We thank the reviewer for pointing out the connection with the wind turbine control and C_p - λ optimisation literature. We agree that performance coefficient optimisation and fine pitch tuning have been historically addressed as control optimisation problems, rather than as offset or fault identification tasks. Extremum Seeking Control and related approaches have indeed been proposed to maximise energy capture by continuously adapting the pitch angle, and have shown promising results both in simulation and experimental studies. The manuscript has

been revised to acknowledge this body of work and clarify the distinction with respect to the proposed approach.

While control-based optimisation methods aim at continuously adjusting the pitch to maximise aerodynamic efficiency, the objective of this work is to detect and identify collective pitch offsets as abnormal conditions within a proactive monitoring framework. In this sense, the proposed method is complementary to control optimisation strategies, as it focuses on fault awareness and diagnosis rather than on online control adaptation.

[Action] The related works section has been modified accordingly.

We look forward to your kind reply, and in the meanwhile, we send our warmest regards.

Sincerely yours,

Sabrina Milani, on behalf of all Authors.