

Response to Reviewer Comments

Shubham Baisthakur and Breifni Fitzgerald

June 10, 2025

We sincerely thank the reviewer for their thorough review of our manuscript titled "*Multi-task Learning Long Short-term Memory Model to Emulate Wind Turbine Blade Dynamics*" and for providing valuable feedback. The reviewer's comments have greatly contributed to improving the overall quality of this paper.

This document presents our responses to the reviewers' comments. Each comment is shown in *italic text*, followed by our response in regular text. Since Reviewer 2 provided comments directly on the manuscript rather than in a formal report, we have listed those queries here with corresponding page and line numbers to facilitate easier tracking. Additionally, a separate PDF file is attached to address Reviewer 2's detailed comments on the manuscript

Reviewer 1

***Q 1.1** The article presents a method of developing LSTM models for predicting blade response using multi-stage modelling and multi-task learning with dimensionality reduction techniques. This is a very well written article which presents a novel and relevant contribution in data driven methods for modelling blade dynamics. The authors have provided a satisfactory response to previous reviewer comments. I recommend the article to be published as it is, and have no further comments.*

Reply: We thank the reviewer for their positive evaluation and recommendation. We are pleased that the novelty and clarity of the proposed framework were appreciated. The reviewers' comments throughout the process have greatly helped improve the quality of the manuscript and are sincerely acknowledged.

Reviewer 2

***Q 2.1** Page 2, line 26: What about the accuracy?*

Reply: We appreciate the reviewer's attention to accuracy, which is indeed critical in any surrogate modelling approach. While the introduction is intended to broadly set the context and highlight the computational motivation, we have thoroughly addressed accuracy throughout the paper. Specifically, the model's accuracy is benchmarked against a validated numerical model in Section 7, with results showing consistent performance across all key response variables. To avoid ambiguity, we have now added a clarifying sentence at the end of the introductory paragraph.

***Q 2.2** Page 3, line 66: If you are using a sequential ML. If the model has a systematic approach, this statement is not valid.*

Reply: The statement in question is intended as a practical observation about the computational cost of training separate surrogate models for each DOF. Immediately following this statement, the manuscript proposes multi-task learning as an efficient solution that overcomes this challenge by predicting multiple outputs within a single model. Therefore, within the full context, the original statement is valid and serves to motivate our approach. We believe this logical flow is clear in the current manuscript.

Q 2.3 Page 3, line 88: So, is it on monoplie? Jacket? floater? not clear

Reply: Thank you for the comment. To clarify, the numerical model of the IEA-15MW wind turbine in this study is based on the monopile foundation. The manuscript is updated to explicitly state the foundation type for clarity.

Q 2.4 Page 4, line 93: How many nodes do you have on the blade?

Reply: Thank you for pointing this out. Although the blades are modelled using the modal summation method, which is not directly a node-based method, the aerodynamic and inertial forces are computed at 50 discrete nodes along the blade span, consistent with the original discretisation used by NREL in the IEA-15MW reference turbine definition. This detail is reflected in the updated manuscript.

Q 2.5 Page 5, line 127: a) At which location along the blade? b) can you elaborate on the phase change for the in-plane displacement at 7m/s?

Reply: a) Figures 1 and 2 have been updated to explicitly indicate that they present the displacement response at the blade tip. b) We thank the reviewer for pointing this out. A slight phase offset is indeed observed at 7m/s between the in-plane displacement responses of the numerical model and OpenFast. This deviation is localised to this specific operating condition and remains within an acceptable range across the overall response duration. Such minor phase differences may arise due to subtle differences in the numerical implementation of the structural damping model, integration schemes, or aerodynamic force evaluations between the two solvers. While our model and OpenFast follow the same underlying physical formulations, differences in solver architecture can introduce small discrepancies in time-domain comparisons. Importantly, this phase shift does not significantly affect the amplitude or frequency content of the response, nor the fidelity of the dataset used for training the surrogate model.

Q 2.6 Page 8, Figure 4: Why not fore-aft? Side to side is directly coming from the rotor torque, while fore-aft is coming from the thrust. I think both are necessary. Also TSS is not defined before.

Reply: While both directions are dynamically relevant, the side-to-side tower response is typically used for FFT-based validation in literature, as it is lightly damped and thus exhibits clearer spectral peaks corresponding to the structural modes. In contrast, the fore-aft response being more heavily damped due to aerodynamic damping tends to exhibit a flatter FFT spectrum, making it less informative for identifying natural frequencies. For this reason, we presented the side-to-side FFT to verify that the natural frequencies of the numerical model are captured accurately. Also, the TSS terminology is removed from the figure.

Q 2.7 Page 9, Section 3.2: Your answer to my comment: This section needs clarification. What is happening to time here? TurbSim output is time series, so $X(i,j)$ are time series! It is not clear what are you trying to do here is very helpful to understand this work better. I suggest to add it to the text.

Reply: Thanks for your suggestion. The manuscript is updated to reflect these details.

Q 2.8 Page 10, line 212: Be consistent. In figure 7 you are using y and z. Here, which one is vertical, and which one is horizontal. I assume y is vertical and x is horizontal. If that is true, why you have higher wind speed at the right hand side of figure 5. If that is not true, I think you need to rotate figure 5. I commented this to the previous review. Also, if x (or z) is the vertical axis, the reason for higher wind speed is shear too, not only turbulence.

Reply: Thank you for highlighting this point again. We have revised Figure 7 to ensure consistency with the axis nomenclature used throughout the manuscript. As correctly noted, in our convention, the x-axis represents the horizontal direction and the y-axis denotes the vertical direction.

Regarding the higher wind speeds observed on the right-hand side of Figure 5, we agree that both shear and turbulence contribute to this pattern. While vertical wind shear typically manifests in a time-averaged sense the snapshot shown in Figure 5 captures the wind field at a single time instance, where turbulent

fluctuations can momentarily dominate and distort the expected shear profile. Therefore, the observed spatial variation in wind speed reflects a combined effect of transient turbulence and wind shear.

Q 2.9 Page 11, line 233: *You answer my comment about high frequency component before. However, those high frequencies cause fatigue and by removing them your fatigue prediction is not valid anymore. I think this need to be mentioned.*

Reply: Thank you for your comment. To clarify, the high-frequency components are removed only from the wind field representation, not from the structural response. The DCT-based compression is used as a dimensionality reduction step for the input (wind field), with the goal of retaining dominant spatial patterns that generalise well across varying wind conditions. The full-resolution response data, which inherently contains higher-frequency content due to turbulence-induced effects, is retained and used during training. As evidenced in the results, the LSTM model is able to capture the dynamics in the blade displacement response, suggesting that the retained wind features are still informative for fatigue-relevant behaviour. The author intend to further investigate the impact of such models on fatigue analysis in a separate study.

Q 2.10 Page 15, line 332: *How did you decide on these bounds? If you Why the lower bound for shear is zero?*

Reply: The wind speed range of 3–25 m/s is adopted from the operational range of the reference 15 MW wind turbine, ensuring that all simulations remain within its cut-in and cut-out wind speed limits. The bounds for turbulence intensity (TI) and the shear exponent are based on the approach outlined in:

- Dimitrov, N., Kelly, M. C., Vignaroli, A., & Berg, J. (2018). From wind to loads: wind turbine site-specific load estimation with surrogate models trained on high-fidelity load databases. *Wind Energy Science*, 3(2), 767-790.

Regarding the lower bound for the shear exponent, we chose 0 as the minimum rather than allowing negative shear values. The negative shear were excluded in this study to reduce complexity and limit the parameter space. Negative shear introduces atypical inflow conditions that can lead to unusual blade loading patterns, which, while interesting, were outside the scope of the present surrogate modelling framework focused on representative and commonly occurring operating scenarios.

This explanation is added to the manuscript for further clarity.

Q 2.11 Page 16, line 349: *How did you decide on your mean wind speed? Why 16? How many simulations did you end up with?*

Reply: Thank you for your comment. As detailed in the manuscript, the range of mean wind speeds was chosen based on the operational wind speed range provided in the turbine properties. Within this range, 50 mean wind speed values were sampled using a Sobol sequence method. The selection of 16 seeds per mean wind speed point each representing a unique combination of turbulence intensity, wind shear exponent, and random seed was made by considering IEC recommendations alongside data augmentation needs, as explained in Section 5.1. Consequently, the dataset comprises $50 \text{ (wind speeds)} \times 16 \text{ (realisations)} = 800$ simulations.

We have updated the manuscript to clarify these details for the reader's benefit.

Q 2.12 Page 19, line 438: *This is an interesting finding, and worth discussion. Why this is happening?*

Reply: Thank you for your comment. However, this observation was discussed in detail during the previous round of reviews. For clarity, we briefly restate the rationale here:

PCA and DCT both reduce dimensionality of the wind field but differ in how they capture patterns. The blade pitch angle is highly sensitive to localised spatial variations, such as shear and turbulence, which are better captured by DCT's spatial frequency decomposition. In contrast, rotor speed reflects a temporally integrated effect of the inflow, which aligns better with the global variance captured by PCA. This explains the observed performance distinction. We believe this explanation is sufficient for the scope of this study, but acknowledge that a more granular comparison may be pursued in future work.

The manuscript is updated for further clarity.

Q 2.13 Page 20, line 464: *I commented this before. What are the hyperparameters and how did you end up with these bounds? You answered my comment, but I don't see it reflected in the text.*

Reply: We have now updated the manuscript to include this explanation.

Q 2.14 Page 20, line 468: *What is rotor average wind speed?*

Reply: Thank you for your comment. Rotor averaged wind speed refers to the mean wind speed over the entire rotor swept area at a given time instant. This term is commonly used in wind energy research, which is why it was not explicitly defined in the manuscript. However, to avoid any ambiguity, we have included a brief definition in the revised manuscript for clarity.

Q 2.15 Page 21, Figure 10: *I had two comments on this part before, which are not answered. The first one about training/testing/validation separation, you answered, but it is not in the text. Which I think is necessary. The second one about RPM above rated and the fact that it is constant, I don't see any discussion or comment on that. I copy paste my comment here: So, this is not surprising. The results you are showing here is the rotor speed. This is mainly depends on average wind speed, which goes in directly into the model. The fluctuation in the wind specially the higher PCA mode wouldn't effect the rotor speed that much. Also, the results that you are showing here they are all above rated, which means the rotor speed is constant. I think it is necessary to show below rated results too.*

Reply: Thank you for highlighting the missing information regarding the data splitting approach, this has now been added to the updated manuscript.

The authors acknowledge the reviewer's observation, and the fact that the results are not surprising further highlights the interpretability of the model and the underlying physical behaviour of the system. These results also demonstrate the ability of the chosen feature selection method to uncover relevant features.

With regard to the comment on presenting results for below-rated wind conditions, we would like to reiterate the explanation already provided in Lines 448–451 of the manuscript for clarity: As previously mentioned, wind turbine dynamics differ significantly between above-rated and below-rated operational regimes. Consequently, models developed for each regime may require distinct architectures and input features to accurately capture the underlying physics. Presenting a second set of results for below-rated conditions including separate feature selection and hyperparameter tuning was deemed to introduce excessive detail without proportionate benefit to the core methodology. Nevertheless, the presented modelling framework remains applicable to both above- and below-rated regions.

Q 2.16 Page 21, Equation 21: *This Z is DCT Z, right?*

Reply: Thank you for pointing this out. This typo is corrected in the updated manuscript.

Q 2.17 Page 23, Figure 13: *What is on y axis?*

Reply: The y-axis represents the magnitude of DCT component mentioned in the subfigure title. This figure is updated in the revised manuscript.

Q 2.18 Page 26, line 535: *What do you mean by constant feature here? They are not changing with time?*

Reply: Thank you for your comment. By "constant features," we refer to features that remain permanently included in the input feature set during the recursive feature addition process. This clarification has now been added to the manuscript.

Q 2.19 Page 26, line 545: *Here you talk about DOF, which is necessary. However, when you present the data in terms of displacement in E and F directions, reader needs to know at which location on the blade.*

Reply: Thank you for this observation. To clarify the relationship between the DOF and the corresponding deformation response, we have highlighted the relevant nomenclature in Section 6. Additionally, Equation (20) defines the mapping between the DOFs and the displacement response.

Q 2.20 Page 28, Figure 18: *You space is limited and your plots are small. Instead of showing 600sec of results, which not that much can be observed, show 100sec. I can not really see what is happening in the edgewise direction.*

Reply: We thank the reviewer for the suggestion. Presenting the response over a 600-second duration follows a standard convention in wind turbine dynamics studies, as it demonstrates the model’s robustness and ability to capture variations across the full turbulent inflow period. While a shorter time window may offer improved visual clarity, it risks overlooking long-term trends or transient behaviours that can emerge only over extended periods. That said, we acknowledge the need for better visual interpretation and have now split the original figure into two separate plots (Figures 18 and 19, Figures 20 and 21) for flapwise and edgewise responses. This revision ensures that the edgewise behaviour, which was previously harder to interpret, is now more clearly visible at the same time the validity of model results for the entire simulation duration is demonstrated.

Q 2.21 Page 29, line 560: *I asked this question before, I didn’t get an answer. How your numerical model performs againsts openfast model in unsteady wind? I think this question rises, as you try to answer to solve two problems:*

- a Try to build a cheap model that replaces expensive openfast.*
- b Try to build a surrogate model on a set of data using ML methods.*

However, I think you aim for this paper is showing how this LSTM architecture can work. If that is correct, I think you need to emphasis on this more at the motivation part, and be more clear that the numerical method is nothing more than a dataset generator, and the reason you are using it is to bypass OpenFAST expensive simulations. Also, it is important to mention that the fact that this LSTM is trained and validated on a numerical model, and not OpenFAST is a limitation. The question that remains there, and can be answered in future work is how this works when you do this with OpenFAST model.

Reply: We thank the reviewer for the comment. As discussed during the previous review stage, validation of the numerical model against OpenFast under steady wind conditions was deemed adequate for the objectives of this study. Furthermore, the numerical model is based on the same governing principles as OpenFast, including Kane’s dynamics for deriving equations of motion, the modal summation method for modelling flexible elements, blade element momentum theory for computing aerodynamic loads, and Morison’s equation for hydrodynamic forces—ensuring strong physical fidelity. The accuracy and application of the numerical model across multiple domains have already been established in prior peer-reviewed studies, now cited in the manuscript:

1. Sarkar, S., & Fitzgerald, B. (2021). Use of Kane’s method for multi-body dynamic modelling and control of spar-type floating offshore wind turbines. *Energies*, 14(20), 6635.
2. Sarkar, S., & Fitzgerald, B. (2022). Fluid inerter for optimal vibration control of floating offshore wind turbine towers. *Engineering Structures*, 266, 114558.
3. Sarkar, S., Chen, L., Fitzgerald, B., & Basu, B. (2020). Multi-resolution wavelet pitch controller for spar-type floating offshore wind turbines including wave-current interactions. *Journal of Sound and Vibration*, 470, 115170.

The focus of the manuscript is on the multi-stage modelling framework and the performance evaluation of the LSTM-based surrogate model through a multi-task learning approach, as emphasised throughout the manuscript and reflected in the title. This methodology remains valid irrespective of the source of the data, and the model can be directly applied to datasets generated using OpenFast without modifying the learning framework. Hence, we do not consider the use of the proposed numerical model as a limitation. The codes developed for the surrogate model are openly shared with the manuscript, and interested readers are welcome to retrain the model using data from OpenFast or other high-fidelity aeroelastic solvers, if desired.

Lastly, we believe the reviewer may have misunderstood the objective of this study. At no point do we suggest that the numerical model serves as a computationally cheaper alternative to OpenFast.

Q 2.22 Page 29, Section 7.3: *The picture that you paint here, without talking about training and RFA computational cost is not complete.*

Reply: The primary focus of this section is on the inference-time performance of the surrogate model, which is the relevant metric in applications requiring repeated simulations. The training and feature selection processes are one-time, offline operations that are not repeated for each simulation and thus do not contribute to the operational cost of the surrogate model. We believe this separation of training and inference cost is both appropriate and widely accepted in the field.

Q 2.23 Page 31, Figure 20: *There is a value to extract mean and std from you testing, for both actual and prediction and show them on a scatter plot with a regression line. That will show how good is your fit from statistical point of view. If you think about it, we generate all of these displacement/loads time series in wind turbine engineering, but at the end we look at them from statistical point of view (mean, std, min, max), so it is important that those at least mean and std match as well.*

Reply: We thank the reviewer for the helpful suggestion. While we agree that statistical comparisons can offer additional insights, the focus of this study is on capturing the dynamic behaviour. Time-domain validation through direct comparison of time series responses and NRMSE metrics sufficiently demonstrates the model's accuracy for the current scope. These provide a direct evaluation of the surrogate model's ability to reproduce the temporal characteristics of the blade response. We acknowledge the value of comparing statistical summaries and will consider this in future work aimed at applications such as fatigue assessment.

Q 2.24 Page 32, line 585: *I saw deformation, not loads, please revise.*

Reply: Thanks for pointing this out. This is rectified in the updated manuscript.

Q 2.25 Page 32, Section 8: *Maybe "Summary and Conclusion"? Also, please add future work and limitation.*

Reply: The manuscript is updated to include these suggestions.