Responses to Reviewers' Comments for Manuscript 10.5194/wes-2025-261

Graph Neural Operator for wind farm wake flow

Addressed Comments for Publication to

WES Wind Energy Science

by

Jens Peter Schøler, Frederik Peder Weilmann Rasmussen, Julian Quick,

and Pierre-Elouan Réthoré

Document Overview

This document contains the authors' responses to the reviewers' second-round comments for manuscript wes-2025-261, "Graph Neural Operator for wind farm wake flow." The document is structured as follows:

- **Section 1:** Responses to Reviewer 1's comments
- **Section 2:** Responses to Reviewer 2's comments

We thank the editor and both reviewers for their continued engagement with our manuscript. All comments have been addressed in the revised version.

Authors' Response to Reviewer 1

Comment 1

In the revised abstract, the primary motivation for developing the data-driven surrogate model is described as the need to perform numerous simulations to account for wake effects from multiple neighbouring farms. It is not clear why such a large number of simulations is required in this context. Clarify this point in the introduction.

Response: The following sentence has been added to the introduction explaining where this is applicable:

"For now, the GNO performs best in the far wake, which is relevant during WFLO when neighbours are present and have not finalised their layouts, as the uncertainty associated with the unknown layouts must be accounted for, requiring numerous simulations to evaluate multiple possible neighbour configurations."

Comment 2

Page 25 appears to be extremely large in terms of tile size (MB). Consider compressing the figures on this page to make the manuscript more manageable.

Response: Fixed. This was caused by invisible layers in the plotting routine for the flow map figure (Figure 11). The issue has been resolved and the file size is now comparable to the other pages.

Comment 3

In Table 2, error metrics have been added for a "Naïve free stream baseline". The reported metrics are surprisingly bad (e.g. MAE of 11.87 m/s). In reality, wake-induced wind speed reductions are typically much lower. How can the MAE be so large for a naïve baseline which assumes a wind speed reduction of 0 m/s?

Response: An error was identified where padding values were being included in the metric computation. This has been corrected, and the naive baseline now shows much more reasonable numbers, e.g. the MAE is 0.218 m/s.

Comment 4

Line 487: As noted in the introduction, other works have deployed GNNs for modelling multiple parameters at turbine locations. Briefly discuss how your model compares with these GNN approaches. If a quantitative comparison is not feasible, a qualitative comparison may be valuable.

Response: The following text has been added, clarifying why a quantitative comparison is not feasible and instead comparing some of the methodological differences as suggested:

"While several GNN-based models have been proposed for predicting quantities at turbine locations, a direct quantitative comparison is not feasible due to fundamental differences in farm configurations, prediction targets, training data sources, and evaluation metrics across studies. This is further complicated by the fact that detailed turbine-level error analysis has not been a primary focus in these works. For instance, Park and Park (2019) and Li et al. (2024) predict farm-level power, Ødegaard Bentsen et al. (2022) predict individual turbine power using a GAT, and Duthé et al. (2024) and de Santos et al. (2024) predict both loads and power. Furthermore, these models are specifically designed and optimized for turbine-level quantities, whereas predicting at turbine locations is not the primary objective of the GNO. Consequently, error characteristics should be viewed in the context of a model whose main purpose is to predict a spatially continuous flow field. Despite this, the GNO captures overall trends in turbine velocities well, suggesting that the learned flow representation encodes physically meaningful turbine interactions."

Comment 5

In the conclusion: "The GNO is less accurate than PyWake, as expected, since a machine learning surrogate cannot fully replicate the fidelity of its source model." Reformulate or remove this sentence. Since PyWake was taken as ground truth, the accuracy of the two methods has not actually been compared. Only with a higher-fidelity reference (e.g. CFD or field measurements) could such a comparison be made. It is not excluded that the GNO could outperform PyWake, similar to data-driven methods trained on measurement data that predict wake losses more accurately than low-fidelity physics-based models.

Response: We agree with the reviewer's reasoning. The sentence has been rephrased to: "The GNO does not fully reproduce the PyWake data, which is unsurprising given that a surrogate will inevitably introduce some approximation error relative to its source model."

Concluding Response. Thank you again for your careful review. We have addressed all comments and hope the revisions are to your satisfaction.

Authors' Response to Reviewer 2

Comment 1

Eq. 8b: The cosine cutoff distance d_c is introduced but never defined numerically. Is this a hyperparameter, or is it derived from the data? Please clarify.

Response: The numerical range has been added along with an explanation of the choice: "with cut-off distance $d_c \approx 695\,D$, chosen as the diagonal of the largest flow map to ensure that no edge is fully ignored. A more aggressive cut-off could be adopted in the future, but would require a dedicated study to determine an appropriate value."

Comment 2

Line 296: Alignment issues in the formatting.

Response: Fixed by removing a too strict formatting requirement in LaTeX.

Comment 3

Eq. 10: The edge features a are indexed with a superscript (m), suggesting they update at each message-passing step. However, as I understand it, the latent edge features remain static after encoding and throughout message passing. If so, the superscript (m) on a_{ji} in Eq. 10a should be removed to avoid implying that edges are updated at each step.

Response: You are correct. The superscript (m) has been removed; it was added by accident.

Comment 4

Above Eq. 13: Typically 'min-max normalization' involves subtracting the min from the values. Consider renaming your scaling method (e.g. range normalization).

Response: Yes, that is a better name. We have updated the terminology to "range normalization" accordingly.

Comment 5

Figure 10: The consistent AEP overestimation by the GNO relative to PyWake for both layouts could benefit from a brief discussion (one or two sentences), as this bias may be relevant for practical applications.

Response: This has been addressed with the following paragraph:

"In Fig. 10 an Annual Energy Production (AEP) measure has been provided for each layout. It was found that the GNO consistently overestimates AEP by approximately 3.4% and 3.9% for the regular and irregular layouts, respectively. This is attributed to the tendency of the GNO to slightly underpredict wake deficits, leading to higher predicted turbine wind speeds and, consequently, inflated power estimates. Notably, the relative difference in AEP between the two layouts is well captured (the GNO predicts a 1.2% increase from the regular to the irregular layout, compared to 0.8% for PyWake), suggesting the model can distinguish between layouts despite the absolute bias. Reducing this systematic offset, for instance, through improved training strategies or architectural refinements, should be considered for applications that require accurate absolute AEP estimates."

Concluding Response. Thank you for your constructive comments throughout the review process. All points have been addressed in the revised manuscript and we hope it now meets your expectations.