Response to Reviewer 1

Thank you for your time reviewing our manuscript. Your constructive comments have substantially improved the quality of our work. In the following, please find our point-to-point responses to your comments.

General comments: The authors present PhyWakeNet, a physics-integrated machine learning framework for dynamic wind turbine wake modeling under aerodynamic force oscillations. The model decomposes the instantaneous velocity field into time-averaged, meandering, and small-scale turbulent components. The time-averaged wake is governed by mass and momentum conservation with an entrainment-based closure; wake meandering uses conditional GAN-reconstructed SPOD modes with a data-driven dynamical system; small-scale turbulence is generated via a CNN. The model is trained and validated using LES data of a single NREL 5 MW turbine under transverse force oscillations at various Strouhal numbers (St_F). It successfully captures frequency-dependent wake recovery, meandering amplitude, and turbulence statistics.

Response: Thank you for your kind consideration of our work.

Comment 1: Lack of Multi-Turbine Validation and Wake Superposition. Most analytical wake models fail precisely in wake interaction and superposition within wind farms — a critical practical challenge. Despite the stated motivation of "wake management in wind farms," all results are for a single turbine. No simulation or discussion addresses how PhyWakeNet handles partial wake overlap, merged wakes, or cumulative turbulence in arrays.

Why was no multi-turbine case investigated? This omission severely limits the claimed applicability. The authors should either: Include at least one 2–3 turbine inline or staggered case (with wake superposition), or explicitly justify the single-turbine focus and discuss planned extensions to farm-scale modeling. Without this, the wind farm relevance remains speculative.

Response 1: We fully agree that the development and validation of a wake superposition model for multi-turbine scenarios are crucial. To address this concern, in the revised manuscript (Appendix C), we added a test case with two turbines in an inline configuration, explained the model setup for this case, and discussed future development of the model for cases with more turbines and partial wakes. (lines 541-558, page 33)

Comment 2: Inadequate Literature Review on ML-Assisted Wake Modeling. The core innovation is ML integration (CGAN + CNN + physics) for dynamic wake prediction — yet the introduction lacks any review of prior ML-based wake or turbulence modeling. Relevant works are not cited. A dedicated paragraph is needed comparing: a) Data-driven vs. physics-constrained approaches; b) SPOD + GAN vs. POD-RBF, LSTM, or PINN methods; c) Quantitative performance (e.g., error in deficit, TKE, meandering). Without this context, the novelty and improvement over existing ML wake models are unclear.

Response 2: Thanks for the suggestion. A dedicated paragraph has been added to the Introduction section to enhance the literature review (lines 53-72, page2-3). The novel contributions of the present work have been clarified (lines 73-80, page 3).

Comment 3: Figure Clarity and Completeness Issues. Figure 13: The difference between red and gray lines is not explained in the caption or text.

Response 3: The differences have been explained (lines 430-431, page 26).

Comment 4: Figure 13(c,d): Horizontal axis labels are illegible (overlapping or cut off).

Response 4: Corrected.

Comment 5: Figure 9 and Figure 12: Captions state "five rows (a–e, f–j)" but subplots d, e, i, j are missing in the figures.

Response 5: Corrected.

Comment 6: Figure 11: No legend — unclear which line corresponds to LES, PhyWakeNet, or submodels (only mentioned in the caption). These errors undermine result interpretation and must be corrected.

Response 5: Corrected.

Comment 6: Critical ML Methodology Relegated to Appendix. In data-driven modeling, dataset generation, model architecture, training strategy, and validation protocol are core contributions. Currently: a) LES setup, SPOD extraction, CGAN/CNN architectures, loss functions, training data split, and validation metrics are buried in appendices. b) The main text jumps from equations to results with minimal explanation of how the ML models were built or validated.

Move key ML details to the main body, including: a) Table of LES cases (St_F, turbulence intensity, length scale); b) CGAN and CNN architectures (layers, inputs, conditioning); c) Training/validation split, loss functions, and convergence; c) Number of SPOD modes (N) and sensitivity hyperparameter tuning may remain in appendix, but model design and data pipeline must be in the main paper.

Response 6: The suggested ML details have been moved to the main body, with the technical specifics (e.g., hyperparameter sensitivity) left in the appendices.

Response 7: Minor but Important Typos and Inconsistencies. Figure 1 caption: "GCAN" → should be CGAN.

Response 7: Corrected.