We thank the Associate Editor for the careful assessment of the revised manuscript and for the clear guidance provided in the decision letter. We are pleased that the essential conditions from the previous revision have been addressed. The manuscript has been revised further to address the remaining points concerning interpretation and presentation. Our point-by-point responses are provided below.

Response to the Associate Editor

Comment 1

Please distinguish between the two prediction tasks when framing the baseline result. For electrical power, the quantities are already close to sufficient. For fatigue the shortfall is of a different order, with blade-root edgewise and the main-bearing bending channels degrading by factors of five and six. Please make this distinction explicit rather than reporting a single degradation result covering both tasks, and note that the pattern is consistent with the expectation in Liew et al. (2023b). This strengthens the paper, since it resolves an earlier concern that the results ran against the accepted picture.

Response

We thank the Associate Editor for this clarification. We have revised the manuscript to distinguish explicitly between the electrical-power and fatigue-load baseline results rather than describing the baseline effect uniformly across both prediction tasks.

For electrical power prediction, the revised manuscript now emphasises that the two spatially aggregated baseline quantities already retain most of the predictive information. Under the SimpleMLP architecture, the baseline achieves an RMSE of 51.184 kW, corresponding to approximately 2.225% of the turbine's rated power, with an R^2 of 0.9966. The distributed full-grid and PCA representations further reduce the absolute prediction error, but the baseline is now described as already providing a relatively accurate representation for electrical power prediction.

For fatigue-load prediction, the revised manuscript now highlights the substantially stronger and load-channel-dependent effect of removing distributed rotor-plane information. Blade-root edgewise bending (MxBr) degrades by a factor of approximately 4.6 under the baseline representation, while the two main-bearing bending channels degrade by factors of approximately 6.4 and 6.5. In contrast, tower-base fore-aft bending (MyTB) changes only marginally.

The abstract, electrical-power results, fatigue-load results, discussion, and conclusions have been revised to reflect this distinction consistently.

We have also related the observed fatigue-load pattern to the corresponding findings of Liew et al. In the revised bibliography, the scientific article *Efficient Mann turbulence generation for offshore wind farms with applications in fatigue load surrogate modelling* is labelled Liew et al. (2023a), while the associated DTU dataset

is labelled Liew et al. (2023b). The scientific interpretation is therefore cited to Liew et al. (2023a) in the revised manuscript.

Comment 2

Please develop the load-channel dependence currently noted only in passing at line 400. Tower-base fore-aft improves marginally under the distributed input, while blade-root edgewise and the main-bearing bending channels improve by large amounts. This is physically coherent: thrust-driven channels are largely captured by a rotor-averaged quantity, whereas asymmetry-driven channels are not. A short paragraph developing this would give the fatigue results a physical interpretation rather than a set of error ratios.

Response

We agree and have expanded the fatigue-load results to provide a physical interpretation of the observed channel dependence.

The revised manuscript now explicitly contrasts the different responses of the load channels. For tower-base fore--aft bending (MyTB), RMSE decreases only marginally from 117.07 kNm for the baseline representation to 110.58 kNm for the distributed full-grid representation. By comparison, blade-root edgewise bending decreases from 196.61 to 42.67 kNm, while the two main-bearing bending components decrease from approximately 199 kNm to approximately 31 kNm.

A new explanatory paragraph relates these differences to the underlying loading mechanisms. Load channels dominated primarily by rotor-integrated thrust can be represented to a comparatively large extent by rotor-averaged inflow quantities, whereas channels that are more sensitive to asymmetric rotor loading depend more strongly on the distributed structure of the inflow across the rotor plane. The revised discussion therefore interprets the baseline result as strongly load-channel dependent rather than as a uniform degradation across all twelve fatigue-load targets.

This interpretation is also related to the corresponding expectation reported by Liew et al. (2023a).

Comment 3

Tables 2 and 3 and Figures 7 and 8 report errors without units. Please add them, and report electrical power error additionally as a percentage of rated power so the magnitudes are interpretable.

Response

We have revised Tables 2 and 3 and Figures 7 and 8 to report the physical units of the prediction errors explicitly.

For electrical power prediction, RMSE and MAE are now reported in kW in Table 2 and Figure 7. Table 2 additionally reports both RMSE and MAE as percentages of the turbine's rated power of 2.3 MW. For example, the SimpleMLP baseline RMSE of 51.184 kW corresponds to approximately 2.225% of rated power, while the PCA-128 RMSE of 5.534 kW corresponds to approximately 0.241% of rated power.

For fatigue-load prediction, all RMSE values in Table 3 are now explicitly identified as being reported in kNm, consistent with the physical units of the damage-equivalent load targets. Figure 8 has likewise been updated so that the RMSE axis is labelled in kNm.

The corresponding figure captions and results text have also been revised to use the same units consistently.

Comment 4

Please move Section 6.1 ahead of the sensitivity analysis, so the PCA configuration is described before the results that depend on it.

Response

We have reorganised the PCA material so that the former Section 6.1 has been removed and its content integrated directly into Section 6, now titled *Principal Component Analysis and Dimensionality Selection*.

The revised section first explains how PCA is applied within each cross-validation fold, including the use of training-partition standardisation and fitting the PCA basis exclusively on the training data. It then presents the dimensionality sensitivity analysis used to determine the number of retained components. The analysis identifies 128 components as the configuration used for the subsequent comparative evaluation.

This structure places the PCA methodology and the dimensionality-selection procedure together in a single section, so that both how PCA is applied and how the final 128-component representation is chosen are established before the electrical-power and fatigue-load results that depend on it.

We thank the Associate Editor again for the constructive comments. The revised manuscript now distinguishes the electrical-power and fatigue-load baseline results explicitly, develops the physical interpretation of the load-channel dependence, reports prediction errors with appropriate physical units and rated-power normalisation, and presents the PCA configuration before the associated sensitivity analysis. We believe that these revisions address all remaining points raised in the minor-revision decision letter.