

We thank both reviewers for their constructive comments. In response, the manuscript has been substantially revised and reorganised. The study has been redesigned to focus on the influence of flow-field representation rather than handcrafted descriptors, resulting in a revised title, methodology, experiments, and discussion. All reviewer comments have been addressed point-by-point below.

Major revisions include

- redesigned the study around three input representations (flattened, PCA, spatial);
- unified the surrogate architectures for both prediction tasks;
- introduced CNN and ResMLP baselines;
- added a PCA sensitivity analysis and reconstruction study;
- expanded the methodology, dataset and training protocol descriptions;
- reorganised the Results and Discussion;
- revised the title to reflect the new scope of the study.

Reviewer 1

Comment 1

Different surrogate models were used for fatigue and power prediction, making it impossible to determine whether the observed differences originated from the representation or the surrogate architecture.

Response

Thank you for this observation. We completely redesigned the experimental framework. The revised manuscript no longer compares different surrogate models for different prediction tasks. Instead, electrical power and fatigue-load prediction are both now evaluated using the same set of surrogate architectures (SimpleMLP, ResMLP, and CNN) under identical GroupKFold cross-validation. This isolates the influence of the input representation from that of the learning architecture.

Comment 2

CNN should be included for the raw spatial representation.

Response

The revised manuscript introduces a dedicated CNN operating directly on the two-dimensional rotor-plane inflow representation. This provides a topology-preserving baseline and enables a direct comparison between flattened, PCA-based and spatial representations.

Comment 3

Validate descriptors across different regressors.

Response

Rather than extending the original descriptor-based methodology to additional regressors, the study was redesigned to compare multiple neural network architectures operating on three fundamentally different input representations. This directly addresses the reviewer's concern by evaluating how representation and architecture interact under a common experimental framework.

Comment 4

Compare against PCA or Autoencoders.

Response

A PCA representation has been introduced as a third input representation and evaluated throughout the manuscript. In addition, a dedicated sensitivity analysis was added to determine an appropriate PCA dimensionality, together with qualitative reconstruction figures illustrating the information retained by increasing numbers of principal components.

Comment 5

Discuss robustness to real-world conditions and measurement noise.

Response

The manuscript now explicitly discusses the limitations associated with using the HAWC2Farm simulation dataset and notes that validation under measured wind-farm conditions remains future work.

Comment 6

Section numbering error.

Response

Corrected.

Reviewer 2

Comment 1

Describe how DELs and power were calculated.

Response

The dataset section now contains a clearer description of the publicly available HAWC2Farm dataset and explicitly references the original publication for details regarding the aeroelastic simulations and target generation.

Comment 2

Additional explanations and figures are required.

Response

The manuscript has been substantially revised and reorganised. The previous descriptor-based methodology has been replaced by a systematic comparison of three input representations (flattened, PCA, and spatial). New figures illustrating the rotor-plane inflow and PCA reconstruction have also been added.

Comment 3

Use the same surrogate for both tasks.

Response

Both prediction tasks now use the same surrogate architectures, allowing fair comparisons between electrical power and fatigue-load prediction.

Comment 4

Section numbering error.

Response

Corrected.

Comment 5

Justify architecture choices.

Response

The revised manuscript includes descriptions and motivation for the SimpleMLP, ResMLP and CNN architectures together with a common training configuration

Comment 6

Reference Figure 1.

Response

Corrected.

Comment 7

Clarify training/validation/test split.

Response

The manuscript now explains the GroupKFold protocol, how simulations are partitioned, and that preprocessing is performed exclusively using the training partition.

Comment 8

The results were not intuitive and might be caused by different surrogate models.

Response

The experimental design has been redesigned so that identical surrogate architectures are used for both electrical power and fatigue-load prediction. The revised results demonstrate that the preferred representation depends on the prediction objective rather than on differences between learning algorithms.

Comment 9

The conclusion is repetitive.

Response

The conclusion has been substantially condensed and now focuses on the principal findings and their implications.