https://doi.org/10.5194/wes-2026-26
Preprint. Discussion started: 4 February 2026

WIND ENERGY SCIENCE DISCUSSIONS
Open Access
eawe european academy of wind energy

On the Sufficiency of Low-Order Spatial Descriptors for Fatigue Load and Power Prediction from Aggregated Rotor-Plane Flow Fields in Offshore Wind Farms

Pradhyumnaa Ganapathi Subramanian[1]

[1]Bauhaus-Universität Weimar, Geschwister-Scholl-Straße 8/15, 99423 Weimar, Thuringia, Germany

Correspondence: Pradhyumnaa Ganapathi Subramanian (pradhyumnaa.ganapathi.subramanian@uni-weimar.de)

Abstract. High-fidelity wind farm simulations are widely used to generate data for surrogate modelling of offshore wind turbine fatigue loads and electrical power output, often relying on high-dimensional spatial representations of rotor-plane wind fields. However, the effective dimensionality required for accurate prediction using temporally aggregated wind-field data remains unclear. This study investigates whether compact, physically motivated spatial descriptors can provide competitive predictive performance relative to full spatial wind-field representations when only temporally aggregated inputs are available. Mean wind speed and standard deviation fields aggregated over 10-minute intervals on a 30×30 rotor-plane grid are considered. A deterministic set of 31 spatial descriptors summarizing wake effects, shear, asymmetry, and spatial variability is constructed and used to train surrogate models based on a residual multilayer perceptron. Model performance is evaluated using group-based five-fold cross-validation. The results show that the compact descriptor set achieves predictive accuracy comparable to full spatial inputs for fatigue load components, indicating that low-order spatial representations are sufficient for capturing load-relevant flow characteristics in the aggregated data regime considered. In contrast, for electrical power prediction, models based on compact descriptors exhibit reduced accuracy relative to those using full or hybrid spatial representations, suggesting that energy-related quantities retain sensitivity to higher-dimensional spatial structure even after temporal aggregation. These findings highlight a task-dependent trade-off between input dimensionality and predictive performance and clarify when compact spatial representations are appropriate for surrogate modelling of wind turbine response.

1 Introduction

High-fidelity aeroelastic wind farm simulations are increasingly used to generate data sets for surrogate modelling of offshore wind turbine fatigue loads and electrical power production. Such data sets typically combine temporally aggregated wind-field information at the turbine rotor plane with corresponding fatigue-relevant load quantities and power output, enabling efficient evaluation of wind farm designs and operational strategies within optimisation frameworks (Liew et al., 2023b).

In these surrogate modelling pipelines, wind fields are commonly represented as spatial grids of flow statistics, such as mean velocity and turbulence intensity, evaluated over the rotor plane. This representation arises naturally from mid-fidelity wind farm simulation tools such as HAWC2Farm and FAST.Farm, which resolve the spatial structure of wakes and inflow conditions at the turbine level (Larsen and Hansen, 2013; Jonkman et al., 2018). While retaining high spatial resolution is

https://doi.org/10.5194/wes-2026-26
Preprint. Discussion started: 4 February 2026

WIND ENERGY SCIENCE DISCUSSIONS
Open Access
eawe european academy of wind energy

therefore straightforward in simulation-based workflows, the extent to which this resolution contributes meaningful predictive information for downstream surrogate tasks such as fatigue load and power estimation once the data are temporally aggregated remains an open question.

In offshore wind applications, wind-field data are typically aggregated over time windows on the order of minutes, consistent with operational data products, power assessment practices, and fatigue load evaluation standards (Wang et al., 2026; International Electrotechnical Commission, 2019). Such aggregation suppresses fine-scale temporal and spatial variability, suggesting that performance-relevant information in the aggregated fields may be well captured by low-order flow characteristics associated with wake effects, shear, and large-scale asymmetries rather than detailed spatial structure (Vermeer et al., 2003). This aggregation regime also reflects practical offshore conditions, where dense, frame-resolved measurements of the rotor-plane inflow are rarely available.

When frame-resolved wind-field data are available, dimensionality reduction techniques such as proper orthogonal decomposition have been shown to be effective for fatigue load prediction (Liew et al., 2023b). However, frame-resolved data are often unavailable in offshore operational or design contexts. As a result, surrogate models based on temporally aggregated rotor-plane statistics typically retain full spatial representations for both fatigue load and power prediction, and the effective dimensionality of these aggregated inputs has not been systematically examined (Guilloré et al., 2024). Here, effective dimensionality refers to the minimum number of input degrees of freedom required to retain predictive information after temporal aggregation.

The objective of this study is to investigate the effective dimensionality of aggregated rotor-plane wind-field data for offshore wind turbine fatigue load and electrical power prediction under realistic data-availability constraints. Using the publicly available Lillgrund wind farm simulation data set (Liew et al., 2023a), we compare surrogate models trained on full aggregated spatial wind-field grids with models using a compact set of physically motivated spatial descriptors. By evaluating predictive performance across multiple fatigue load components and power output under identical modelling and validation protocols, this study assesses whether high-resolution spatial inputs provide measurable benefit beyond low-order spatial descriptors in the aggregated data regime considered.

2 Dataset Description

This study uses the publicly available simulated fatigue load dataset of the Lillgrund offshore wind farm provided by the Technical University of Denmark (DTU) (Liew et al., 2023a). The dataset is generated using high-fidelity aeroelastic wind farm simulations performed with HAWC2Farm and comprises synthetic operating data for a 48-turbine offshore wind farm with a rated turbine capacity of 2.3 MW.

A total of 1000 independent wind farm simulations were performed, of which only the publicly released training portion is used in this work, corresponding to 800 simulations. The remaining simulations are embargoed and not accessible.

Each simulation spans one hour of physical time and is driven by randomly sampled inflow conditions, including wind direction, ambient wind speed (4–20 m/s), vertical shear exponent (0.1–0.3), and turbulence intensity (2.5–10%). During each

https://doi.org/10.5194/wes-2026-26
Preprint. Discussion started: 4 February 2026

Table 1. Summary of predicted fatigue load components. All quantities are 10-minute damage-equivalent loads (DELs).

Abbreviation	Component	Load direction / type	Unit
MxBr	Blade root	Edgewise bending DEL	kNm
MyBr	Blade root	Flapwise bending DEL	kNm
MzBr	Blade root	Torsional DEL	kNm
MxTB	Tower base	Side–side bending DEL	kNm
MyTB	Tower base	Fore–aft bending DEL	kNm
MzTB	Tower base	Torsional DEL	kNm
MxTT	Tower top	Fore–aft bending DEL	kNm
MyTT	Tower top	Side–side bending DEL	kNm
MzTT	Tower top	Torsional DEL	kNm
MxMB	Main bearing	Bending DEL	kNm
MyMB	Main bearing	Bending DEL	kNm
MzMB	Main bearing	Torsional DEL	kNm

simulation, the wind field immediately in front of each turbine is recorded on a uniform 30×30 Cartesian grid covering an area of $1.1D \times 1.1D$ around the rotor plane, where the rotor diameter is $D = 92.6$ m.

The one-hour simulations are divided into six non-overlapping 10-minute intervals. For each interval, the recorded wind fields are post-processed into temporally aggregated rotor-plane wind-field statistics. Each turbine–interval pair constitutes a single observation in the dataset, resulting in $48 \times 6 = 288$ observations per simulation and a total of 230,400 observations across the 800 simulations used in this study.

For each observation, the input consists of two flattened 30×30 fields: the spatial mean and spatial standard deviation of the rotor-plane wind field over the corresponding 10-minute interval. These are stored as two vectors of length 900, yielding a fixed 1800-dimensional input representation. The corresponding outputs include electrical power and a set of 12 fatigue-relevant damage-equivalent load (DEL) quantities for the same turbine and time interval.

In this study, surrogate models are trained to predict both structural fatigue loads and electrical power output from the aggregated rotor-plane wind-field representation. The 1800-dimensional spatial representation is used directly as a baseline input and is subsequently transformed into lower-dimensional representations, as described in Section 3.

2.1 Predicted Fatigue Load Components

The surrogate model predicts twelve fatigue-relevant load components corresponding to key turbine substructures. All target quantities are provided as 10-minute damage-equivalent load (DEL) values from high-fidelity aeroelastic simulations of the Lillgrund wind farm. Table 1 summarizes the predicted load components, their physical interpretation, and units.

https://doi.org/10.5194/wes-2026-26
Preprint. Discussion started: 4 February 2026

WIND ENERGY SCIENCE DISCUSSIONS

3 Input Representations

This section describes how the temporally aggregated rotor-plane wind-field data introduced in Section 2 are represented and transformed into compact spatial descriptors.

The motivation for considering alternative input representations is to assess whether the predictive information contained in temporally aggregated rotor-plane wind-field statistics can be retained using a low-dimensional, physically motivated summary, rather than the full spatial grid. To this end, a set of deterministic spatial descriptors is constructed a priori from aggregated wind-field statistics derived from the spatial mean wind speed and turbulence intensity fields.

The descriptors are designed to summarize distinct aspects of the aggregated inflow, without introducing data-driven feature selection or representation learning. All descriptors are computed directly from the aggregated wind-field statistics, enabling a controlled comparison between full spatial inputs and low-order descriptor-based representations in the surrogate modelling tasks considered.

3.1 Wind-Field Representation and Spatial Descriptors

For each observation, the inflow conditions at the turbine are represented by spatial wind-field statistics evaluated over the rotor plane. The available inputs consist of the spatially resolved mean wind speed field

$$\mu(x,y) \tag{1}$$

and the corresponding standard deviation field

$$\sigma(x,y), \tag{2}$$

both defined on a uniform Cartesian grid.

From these quantities, the turbulence intensity field is defined pointwise as

$$\mathrm{TI}(x,y) = \frac{\sigma(x,y)}{|\mu(x,y)| + \varepsilon}, \tag{3}$$

where ε is a small constant introduced for numerical stability.

Rather than using the full spatial grids directly, a compact set of deterministic spatial descriptors is constructed from $\mu(x,y)$ and $\mathrm{TI}(x,y)$. These descriptors are designed to summarize large-scale flow characteristics relevant to fatigue loading, such as wake structure, shear, asymmetry, and spatial variability of the aggregated fields, while substantially reducing the dimensionality of the input representation.

https://doi.org/10.5194/wes-2026-26
Preprint. Discussion started: 4 February 2026

3.2 Global Statistical Descriptors

Basic global statistics are computed separately for the mean wind speed field $\mu(x,y)$ and the turbulence intensity field $\mathrm{TI}(x,y)$. For a generic field $f(x,y)$, the following quantities are included:

$$\overline{f} = \frac{1}{N}\sum_{i=1}^{N} f_i, \tag{4}$$

$$\sigma_f = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(f_i - \overline{f}\right)^2}, \tag{5}$$

$$f_{\min} = \min_i f_i, \tag{6}$$

$$f_{\max} = \max_i f_i, \tag{7}$$

where the index i runs over all grid points on the rotor plane.

These descriptors capture the overall magnitude, variability, and extrema of the aggregated wind field.

3.3 Shear and Gradient-Based Descriptors

Large-scale shear in the mean wind field is characterized using discrete spatial gradients computed via Sobel operators. Discrete approximations of the horizontal and vertical spatial gradients of $\mu(x,y)$ are denoted as

$$\frac{\partial \mu}{\partial x}, \qquad \frac{\partial \mu}{\partial y}. \tag{8}$$

The following gradient-based descriptors are included:

$$\overline{\frac{\partial \mu}{\partial x}}, \qquad \overline{\frac{\partial \mu}{\partial y}}, \qquad \overline{|\nabla \mu|}, \tag{9}$$

where

$$|\nabla \mu| = \sqrt{\left(\frac{\partial \mu}{\partial x}\right)^2 + \left(\frac{\partial \mu}{\partial y}\right)^2}. \tag{10}$$

An analogous gradient magnitude is computed for the turbulence intensity field $\mathrm{TI}(x,y)$ and spatially averaged to obtain a turbulence shear descriptor.

3.4 Wake Geometry and Asymmetry Descriptors

Wake-related characteristics are extracted directly from the mean wind speed field. The wake centre is defined as the grid location of minimum velocity:

$$(x_w, y_w) = \arg\min_{x,y} \mu(x,y), \tag{11}$$

with the coordinates normalized by the grid dimension.

https://doi.org/10.5194/wes-2026-26
Preprint. Discussion started: 4 February 2026

eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

Wake depth is defined as

$$\Delta\mu = \mu_{\min} - \overline{\mu}, \tag{12}$$

representing the velocity deficit relative to the spatial mean.

Wake width is approximated by the fractional rotor-plane area where the wind speed falls below a deficit-based threshold:

$$W = \frac{1}{N}\sum_{i=1}^{N} \mathrm{I}[\mu_i < \overline{\mu} + \alpha(\mu_{\min} - \overline{\mu})], \tag{13}$$

where $\mathrm{I}[\cdot]$ denotes the indicator function and $\alpha = 0.1$.

Large-scale asymmetry is quantified using left–right and top–bottom mean differences of the wind speed field, computed by splitting the rotor-plane grid at its geometric centre.

3.5 Structure Function Descriptors

To characterize spatial variability at short length scales, second-order structure functions are computed for both $\mu(x,y)$ and $\mathrm{TI}(x,y)$. For a given spatial lag r, the structure function is defined as

$$S_r(f) = \frac{1}{N_r}\sum_{j=1}^{N_r} (f_{j+r} - f_j)^2, \tag{14}$$

where N_r denotes the number of valid grid point pairs separated by r grid spacings.

Structure functions are evaluated for $r = 1$ and $r = 2$ grid spacings for both fields.

3.6 Radial Descriptors

Using polar coordinates (r,θ) centered at the rotor hub, radially weighted means are computed as

$$\overline{fr} = \frac{1}{N}\sum_{i=1}^{N} f_i r_i, \tag{15}$$

for both $\mu(x,y)$ and $\mathrm{TI}(x,y)$.

In addition, a normalized turbulence intensity radial decay descriptor is defined as

$$D_{\mathrm{TI}} = \frac{\overline{\mathrm{TI}\,r}}{\overline{\mathrm{TI}}}, \tag{16}$$

where $\overline{\mathrm{TI}}$ denotes the spatial mean of the turbulence intensity field.

3.7 Azimuthal Harmonic Descriptors

Large-scale angular asymmetries are captured using low-order azimuthal harmonics. For a normalized field

$$\tilde{f}(r,\theta) = \frac{f(r,\theta)}{\overline{f}}, \tag{17}$$

https://doi.org/10.5194/wes-2026-26
Preprint. Discussion started: 4 February 2026

the magnitude of the mth azimuthal mode is defined as

$$A_m = \frac{1}{\sum_i r_i} \sqrt{\left(\sum_i \tilde{f}_i r_i \cos(m\theta_i)\right)^2 + \left(\sum_i \tilde{f}_i r_i \sin(m\theta_i)\right)^2}, \qquad m = 1,2,3. \tag{18}$$

Azimuthal modes are computed for both the mean wind speed and turbulence intensity fields.

3.8 Descriptor Summary

In total, the original 1800-dimensional input is mapped to a vector of 31 deterministic spatial descriptors. No data-driven dimensionality reduction or learning-based compression is applied; all descriptors are computed directly from the temporally aggregated wind-field statistics.

4 Surrogate Model and Evaluation Protocol

Two distinct surrogate modelling approaches are employed in this study, corresponding to the different response variables considered. Structural fatigue loads are predicted using a neural-network-based surrogate, while electrical power output is predicted using a tree-based ensemble model. In both cases, identical input representations and validation protocols are used to enable a controlled comparison of different wind-field representations.

4.0.1 Fatigue Load Surrogate

Fatigue load prediction is performed using a feed-forward neural network with residual connections, hereafter referred to as a residual multilayer perceptron (ResMLP). The network maps a fixed-length input feature vector to a vector of fatigue-relevant load quantities corresponding to multiple load components and is trained in a multi-output regression setting.

The model begins with a linear projection that maps the input feature space to a fixed hidden dimension. This is followed by a stack of residual blocks operating in the hidden feature space. Each residual block consists of layer normalization, a two-layer feed-forward transformation with gated nonlinear activation, and residual skip connections. Specifically, the hidden representation is first normalized, then passed through a linear expansion followed by a SwiGLU activation, and finally projected back to the original hidden dimension. Dropout and stochastic depth regularization are applied within each block to improve generalization and stabilize training for deep networks.

Residual connections are scaled by a learnable parameter initialized to a small value, which ensures stable optimization in the early stages of training while allowing the network to adaptively increase residual contributions as training progresses. Stochastic depth is applied with a linearly increasing drop probability across layers, providing implicit model ensemble behavior during training.

After the sequence of residual blocks, a final layer normalization is applied, followed by a linear output layer that produces simultaneous predictions for all fatigue load components. No output-specific heads are used; all targets are learned jointly within a shared latent representation.

https://doi.org/10.5194/wes-2026-26
Preprint. Discussion started: 4 February 2026

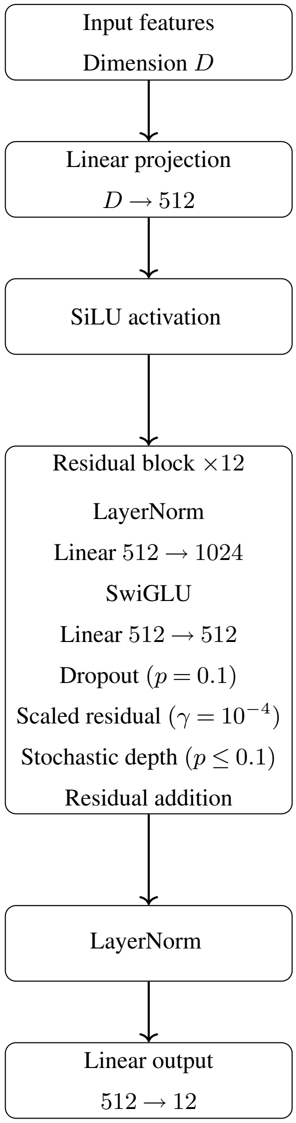

Figure 1. Residual multilayer perceptron (ResMLP) architecture used for fatigue load prediction. All layer dimensions, activation functions, and regularization mechanisms correspond exactly to the implemented model.

4.0.2 Electrical Power Surrogate

Electrical power prediction is performed using a gradient-boosted decision tree model implemented with LightGBM. The model is trained in a univariate regression setting to map the same wind-field input representations used for the fatigue surrogate
180 to electrical power output.

https://doi.org/10.5194/wes-2026-26
Preprint. Discussion started: 4 February 2026

eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

All LightGBM models are trained using identical hyperparameters across experiments, and no model-specific feature engineering beyond the considered input representations is applied. This ensures that observed performance differences for power prediction can be attributed solely to the choice of input representation rather than differences in model structure or optimization.

Table 2. LightGBM hyperparameters used for electrical power surrogate modelling.

Hyperparameter	**Value**
Boosting type	Gradient boosting decision trees (GBDT)
Objective function	Regression (squared error)
Evaluation metric	Root mean squared error (RMSE)
Number of leaves	64
Learning rate	0.05
Feature fraction	0.9
Bagging fraction	0.8
Bagging frequency	1
Verbosity	1
Boosting rounds	3000

The number of boosting rounds reported in Table 2 corresponds to the maximum number of iterations used during the group-based cross-validation experiments. Increasing the number of boosting rounds to 10,000 can substantially improve electrical power prediction accuracy under single train–test splits.

However, applying such extended training schedules within a five-fold group-based cross-validation framework leads to a prohibitive increase in computational cost. For this reason, electrical power surrogates are evaluated using a fixed budget of 3,000 boosting rounds per fold. The reported results therefore do not represent the maximum achievable performance for power prediction, but rather reflect a controlled and computationally tractable comparison of different wind-field input representations under identical modelling and training conditions.

4.1 Cross-Validation Strategy

Model performance for both surrogate models is evaluated using group-based K-fold cross-validation to prevent information leakage between training and validation sets. All samples originating from the same wind farm simulation realization are assigned a common group identifier, and samples from a given group are restricted to a single fold.

The full dataset is denoted as

$$\mathcal{D} = \{(x_i, y_i, g_i)\}_{i=1}^{N}, \tag{19}$$

https://doi.org/10.5194/wes-2026-26
Preprint. Discussion started: 4 February 2026

where x_i represents the input feature vector, y_i the corresponding response variable (either fatigue loads or electrical power), and g_i the simulation group identifier associated with sample i.

The dataset is partitioned into $K = 5$ disjoint folds such that no group identifier appears in more than one fold. For each fold k, the training and validation sets are defined as

$$\mathcal{D}^{(k)}_{\text{train}} = \mathcal{D} \setminus \mathcal{D}_k, \tag{20}$$

$$\mathcal{D}^{(k)}_{\text{val}} = \mathcal{D}_k, \tag{21}$$

where $\mathcal{D}_k$ denotes the subset of samples assigned to fold k.

For each fold, the corresponding surrogate model is trained using $\mathcal{D}^{(k)}_{\text{train}}$ and evaluated on the validation set $\mathcal{D}^{(k)}_{\text{val}}$. Final performance metrics are obtained by averaging results across all folds.

Input features are standardized independently within each training fold. For a generic feature z, normalization is performed according to

$$z' = \frac{z - \mu_{\text{train}}}{\sigma_{\text{train}}}, \tag{22}$$

where μ_{train} and σ_{train} denote the mean and standard deviation computed from the training subset of the corresponding fold. These statistics are subsequently applied to the validation data to ensure that no information from the validation set is used during training. For the fatigue load surrogate, target variables are normalized analogously on a per-fold basis.

4.2 Evaluation Metrics

Model performance is assessed using multiple complementary error metrics computed separately for each response variable. For fatigue load prediction, metrics are evaluated independently for each damage-equivalent load (DEL) component, while for electrical power prediction they are evaluated for the scalar power output. Together, these metrics quantify absolute error magnitude, relative error, and the proportion of variance explained by the surrogate model.

The root mean squared error is defined as

$$\text{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(\hat{y}_i - y_i)^2}, \tag{23}$$

where y_i and $\hat{y}_i$ denote the true and predicted values, respectively. RMSE penalizes larger prediction errors more strongly than smaller ones and therefore emphasizes occasional large deviations, which are particularly relevant for fatigue load estimation, where extreme responses can dominate accumulated damage.

The mean absolute error is given by

$$\text{MAE} = \frac{1}{N}\sum_{i=1}^{N}|\hat{y}_i - y_i|. \tag{24}$$

MAE provides an intuitive measure of the average absolute prediction error and is less sensitive to outliers than RMSE, offering a complementary view of typical prediction accuracy.

https://doi.org/10.5194/wes-2026-26
Preprint. Discussion started: 4 February 2026

The coefficient of determination is computed as

$$R^2 = 1 - \frac{\sum_{i=1}^{N} (\hat{y}_i - y_i)^2}{\sum_{i=1}^{N} (y_i - \overline{y})^2}, \tag{25}$$

where $\overline{y}$ denotes the sample mean of the true values. The R^2 score measures the fraction of variance in the target quantity that is explained by the model and facilitates comparison across different load components or response variables with varying magnitudes.

In addition, the root mean squared percentage error is used to assess relative error magnitude:

$$\mathrm{RMSPE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(\frac{\hat{y}_i - y_i}{\max(|y_i|, \varepsilon)}\right)^2}, \tag{26}$$

where ε is a small constant introduced to avoid division by zero. RMSPE normalizes the prediction error by the magnitude of the true value and thus provides a scale-independent measure that is particularly useful for comparing performance across fatigue load components and electrical power, which differ substantially in absolute range.

For each cross-validation fold, all evaluation metrics are computed separately on the corresponding validation set. Final reported values are obtained by averaging each metric across the $K = 5$ folds. This procedure ensures that performance estimates reflect generalization across independent wind farm simulation realizations and are not biased by any single data partition.

5 Results

5.1 Fatigue Load Prediction Performance

Table 3 reports the coefficient of determination for all twelve fatigue load components using three different input representations. Across all components, the absolute difference in R^2 between the full spatial representation (Raw, 1800 features) and the compact descriptor-based representation (Spatial, 31 features) remains below 0.003. No systematic reduction in explained variance is observed when replacing the full rotor-plane wind-field grids with the low-dimensional spatial descriptors.

Table 4 summarizes the relative prediction errors in terms of RMSPE. For all load components, the difference in RMSPE between the Raw (1800) and Spatial (31) representations remains below approximately one percentage point, with most components exhibiting differences below 0.3 percentage points. These results indicate that replacing high-dimensional aggregated rotor-plane wind-field inputs with a compact set of physically motivated spatial descriptors does not lead to a practically significant degradation in relative fatigue load prediction accuracy.

5.2 Energy Prediction Performance

Table 5 summarizes the electrical power prediction performance for the three considered input representations. In contrast to the fatigue load results, a clear dependence on input dimensionality is observed for power prediction. Models trained exclusively on the compact set of 31 spatial descriptors exhibit substantially higher RMSPE compared to models using the full aggregated

https://doi.org/10.5194/wes-2026-26
Preprint. Discussion started: 4 February 2026

eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

Table 3. Coefficient of determination (R^2) for fatigue load prediction using different input representations. Values are averaged over five-fold group-based cross-validation.

Load component	Raw (1800)	Spatial + Raw (1831)	Spatial (31)
MxBr	0.9904	0.9912	0.9912
MyBr	0.9819	0.9826	0.9807
MzBr	0.9545	0.9568	0.9530
MxTB	0.9386	0.9400	0.9393
MyTB	0.9373	0.9386	0.9374
MzTB	0.9795	0.9790	0.9787
MxTT	0.9836	0.9833	0.9825
MyTT	0.9711	0.9709	0.9704
MzTT	0.9794	0.9791	0.9787
MxMB	0.9814	0.9816	0.9813
MyMB	0.9814	0.9824	0.9819
MzMB	0.9825	0.9829	0.9818

Table 4. Root mean squared percentage error (RMSPE) for fatigue load prediction using different input representations. Values are averaged over five-fold group-based cross-validation.

Load component	Raw (1800)	Spatial + Raw (1831)	Spatial (31)
MxBr	5.657 %	5.560 %	5.769 %
MyBr	1.528 %	1.481 %	1.573 %
MzBr	2.210 %	2.149 %	2.228 %
MxTB	9.801 %	9.477 %	9.423 %
MyTB	14.965 %	14.666 %	14.508 %
MzTB	7.089 %	7.296 %	7.243 %
MxTT	6.411 %	6.602 %	6.719 %
MyTT	7.159 %	7.398 %	7.260 %
MzTT	7.974 %	7.785 %	7.827 %
MxMB	3.122 %	3.106 %	3.100 %
MyMB	3.138 %	3.066 %	3.078 %
MzMB	16.949 %	16.522 %	17.547 %

https://doi.org/10.5194/wes-2026-26
Preprint. Discussion started: 4 February 2026

Table 5. Electrical power prediction performance using different input representations. Values are reported as mean ± standard deviation over five-fold group-based cross-validation.

Input representation	RMSE (kW)	MAE (kW)	R^2	RMSPE (%)
Raw (1800)	11.824 ± 1.312	5.592 ± 0.404	0.999818 ± 0.000040	9.592 ± 4.942
Spatial + Raw (1831)	$\mathbf{9.949 \pm 1.190}$	$\mathbf{4.585 \pm 0.286}$	$\mathbf{0.999871 \pm 0.000029}$	$\mathbf{8.296 \pm 4.564}$
Spatial (31)	14.244 ± 1.811	6.873 ± 0.696	0.999735 ± 0.000070	15.978 ± 12.069

spatial wind-field representation. While all models achieve very high coefficients of determination, the relative error metrics reveal a pronounced degradation in predictive accuracy when fine-grained spatial information is discarded.

Combining the low-order spatial descriptors with the full spatial input yields the best overall performance across all reported metrics, indicating that the engineered descriptors provide complementary information but are not sufficient on their own for accurate power prediction. These results suggest that, even in temporally aggregated rotor-plane wind-field data, electrical power remains sensitive to higher-dimensional spatial structure, in contrast to fatigue load components for which low-order flow characteristics are largely sufficient.

6 Conclusions

This study investigated the effective dimensionality of temporally aggregated rotor-plane wind-field data for surrogate modelling of offshore wind turbine fatigue loads and electrical power. The results demonstrate that a compact set of physically motivated spatial descriptors can achieve predictive performance for fatigue load estimation that is comparable to models trained on full aggregated 30×30 spatial wind-field grids. Across all twelve evaluated fatigue load components, differences in RMSPE between models using the full spatial representation and those using only 31 engineered descriptors remain small and do not exhibit a systematic degradation in accuracy. This indicates that, within the examined data regime, much of the fatigue-relevant information is captured by a low-dimensional representation.

In contrast, electrical power prediction exhibits a stronger dependence on high-dimensional spatial input representations. Models trained exclusively on the compact set of spatial descriptors show a clear degradation in relative prediction accuracy compared to models using the full aggregated spatial fields. While all considered input representations achieve very high coefficients of determination, relative error metrics reveal that discarding fine-grained spatial information leads to substantially higher RMSPE for power prediction. Combining the engineered spatial descriptors with the full spatial input yields the best overall performance, indicating that the descriptors provide complementary information but are not sufficient on their own for accurate electrical power estimation.

These differing outcomes highlight a target-dependent effective dimensionality of temporally aggregated rotor-plane wind-field data. Fatigue load prediction appears to be largely influenced by large-scale flow characteristics, such as wake deficits,

https://doi.org/10.5194/wes-2026-26
Preprint. Discussion started: 4 February 2026

shear, asymmetry, and overall turbulence intensity levels, which are represented by the proposed low-order descriptors. Electrical power prediction, by contrast, remains sensitive to higher-dimensional spatial structure even after temporal aggregation, consistent with the nonlinear dependence of power on local velocity distributions across the rotor plane.

Previous surrogate modelling studies have shown that when frame-resolved wind-field data are available, data-driven dimensionality reduction techniques such as proper orthogonal decomposition can be effective for fatigue load prediction. For example, Liew et al. (2023b) reported fatigue load prediction errors expressed in terms of RMSPE using high-frequency rotor-plane wind-field snapshots combined with modal decomposition. Although a direct numerical comparison is not appropriate due to differences in temporal resolution, data availability, and modelling assumptions, the present results suggest that, in the absence of frame-resolved wind-field data, physically informed low-order spatial descriptors provide an effective alternative for fatigue load prediction, but not for electrical power.

Beyond demonstrating sufficiency of the 31-descriptor representation, exploratory analyses indicate that a smaller subset of dominant descriptors can retain a large fraction of the fatigue load prediction performance. However, these reductions do not materially alter the central conclusion that fatigue-relevant information in temporally aggregated rotor-plane fields is already well captured by low-order spatial structure. Accordingly, the present study emphasizes sufficiency rather than minimality of the descriptor set.

From a practical perspective, the demonstrated reduction in input dimensionality to a compact set of deterministic descriptors offers clear advantages for fatigue-oriented surrogate modelling pipelines. Lower-dimensional inputs improve interpretability, reduce computational cost during training and inference, and simplify data storage and transfer. At the same time, the power prediction results indicate that care must be taken when applying low-order representations to quantities that remain sensitive to fine-scale spatial variability. Importantly, these results characterize the information content required for fatigue load prediction after temporal aggregation, independent of how the aggregated inflow information is obtained or represented upstream.

The conclusions of this study are specific to the data regime considered, namely temporally aggregated rotor-plane wind-field statistics derived from high-fidelity offshore wind farm simulations. No claim is made regarding the applicability of the proposed descriptors to frame-resolved wind fields or to scenarios in which fine-scale turbulence structures dominate fatigue or power variability. Future work may investigate the sensitivity of these findings to alternative aggregation windows, descriptor formulations, and wind farm configurations, as well as their integration with higher-frequency wind-field data when such information is available.

Data availability. The data used in this study are publicly available from the Technical University of Denmark (DTU) data repository at https://doi.org/10.11583/DTU.21399645.

Author contributions. The author was solely responsible for the conceptualization, methodology, software development, analysis, and writing of this manuscript.

https://doi.org/10.5194/wes-2026-26
Preprint. Discussion started: 4 February 2026

Competing interests. The author declares no competing interests.

https://doi.org/10.5194/wes-2026-26
Preprint. Discussion started: 4 February 2026

References

Guilloré, A., Campagnolo, F., and Bottasso, C. L.: A control-oriented load surrogate model based on sector-averaged inflow quantities: capturing damage for unwaked, waked, wake-steering and curtailed wind turbines, Journal of Physics: Conference Series, 2767, 032 019, https://doi.org/10.1088/1742-6596/2767/3/032019, 2024.

International Electrotechnical Commission: IEC 61400-1: Wind Energy Generation Systems—Part 1: Design Requirements, 2019.

Jonkman, J., Doubrawa, P., Hamilton, N., Annoni, J., and Fleming, P.: Validation of FAST.Farm Against Large-Eddy Simulations, Journal of Physics: Conference Series, 1037, 062 005, https://doi.org/10.1088/1742-6596/1037/6/062005, 2018.

Larsen, T. J. and Hansen, A. M.: How 2 HAWC2, the User's Manual, Risø National Laboratory, Technical University of Denmark, 2013.

Liew, J., Riva, R., and Göçmen, T.: Simulated Fatigue Load Data Set of the Lillgrund Wind Farm, https://doi.org/10.11583/DTU.21399645, accessed: 18 December 2025, 2023a.

Liew, J., Riva, R., and Göçmen, T.: Efficient Mann turbulence generation for offshore wind farms with applications in fatigue load surrogate modelling, Journal of Physics: Conference Series, 2626, 012 050, https://doi.org/10.1088/1742-6596/2626/1/012050, 2023b.

Vermeer, L., Sørensen, J., and Crespo, A.: Wind turbine wake aerodynamics, Progress in Aerospace Sciences, 39, 467–510, https://doi.org/https://doi.org/10.1016/S0376-0421(03)00078-2, 2003.

Wang, S., Vidal, Y., and Pozo, F.: Recent advances in wind turbine condition monitoring using SCADA data: A state-of-the-art review, Reliability Engineering and System Safety, 267, None, https://doi.org/10.1016/j.ress.2025.111838, 2026.