Wind turbine drivetrain fault detection using physics-informed multivariate deep learning

Faras JAMIL[1,2], Cédric PEETERS[2] Timothy VERSTRAETEN[1,2], Jan HELSEN[2]

[1]Acoustics Vibration Research Group, Belgium
[2]Artificial Intelligence Lab Brussels, Belgium
faras.jamil@vub.be

Abstract

Vibration analysis is a prevalent technique in the predictive maintenance of wind turbines. It is an effective method for early fault detection and enables the creation of cost-effective maintenance strategies. Commonly used vibration analysis methods in the literature rely on signal processing techniques such as time and frequency domain approaches. However, the signal processing techniques require manual interpretation by domain experts. It is important to note that different indicators exhibit sensitivity to specific faults. Manual analysis of indicators can be avoided by fusing them to derive high-level wind turbine health status. It enables the learning of complex non-linear relationships among the indicators. This research focuses on a multivariate deep learning model, i.e., autoencoder, which fuses different signal processing indicators to provide a single high-level health status. The proposed model is a normal behaviour model that learns the indicator's normal behaviour and labels faults if it observes deviation from the normal behaviour. The proposed fusion method of indicators is robust compared to individual indicator models as it learns complex non-linear relationships among indicators. The proposed method is tested for fleet-level fault detection both with and without fine-tuning for a specific wind turbine. Furthermore, it decreases the time required for wind farm health prognosis analysis and computation. Various autoencoder architectures have been compared, including simple feedforward neural networks, convolutional neural networks, and recurrent neural networks. The proposed method is demonstrated using real-life, high-frequency condition monitoring data from offshore wind turbines over several years, including wind turbines observed faults. The method's effectiveness and performance were demonstrated through analysis of planetary stage, generator, and high-speed stage failure cases.

1 Introduction

The increasing interest in wind energy comes with the challenge of significant operating and maintenance (O&M) costs, which contribute approximately 30% of the overall energy costs. To address this challenge, understanding wind turbines' health status becomes crucial. Early fault detection helps to devise efficient control and maintenance strategies for the entire wind farm. Condition monitoring (CM) techniques enable the development of predictive maintenance strategies to plan cluster maintenance. Group maintenance in a wind farm reduces the O&M cost significantly. Furthermore, the anticipation of wind turbine health status avoids catastrophic failure, which increases overall wind farm production by reducing downtime [1].

Vibration analysis has emerged as the primary technique in condition monitoring (CM) [1, 2]. Various signal processing features are derived from raw vibration signals. Analyzing these features enables experts to detect evolving faults at an early stage. Reliable fault detection techniques require comprehensive feature analysis, as different features exhibit sensitivity to various types of faults. Time and frequency domain methods are commonly used signal processing techniques in condition monitoring. Time domain features involve statistical parameters such as root mean square, kurtosis, peak-to-peak, Moors kurtosis, peak energy index, and crest factor [3, 4]. On the other hand, frequency domain features rely on cyclostationarity methods [5], specifically envelope analysis [6] and spectral correlation [7].

Manually analyzing a vast number of features becomes an impractical task for experts, mainly when dealing with multiple machines across the fleet having many components. Moreover, non-stationary operating conditions within complex machines, such as wind turbines, further complicate the CM task. An artificial intelligence-based CM method capable of providing a comprehensive high-level health status overview has the potential to significantly streamline the task of analyzing the health of the entire wind farm.

The research community has considerable interest in the integration of deep learning methods for fault detection. However, a key challenge lies in the limited availability of fault data. Transfer learning presents a viable solution by enabling the training of machine models using data from similar machines [8]. Normal behaviour models (NBMs) hold substantial promise for fault detection applications, as they solely require healthy data for training, which is readily accessible. The NBM learns the normal behaviour of machines from healthy data and identifies potential faults when deviations are detected [9, 10]. Vibration signal processing features [9], along with derived coherence maps [11], enable the detection of faults in the rotary components of wind turbines. On the other hand, temperature signal-based NBMs are trained to identify faults in wind turbine generators [12, 13].

This research introduces an approach utilizing a physics-informed multivariate deep learning NBM that combines time and frequency domain vibration signal processing features. The proposed method provides a high-level health status assessment, eliminating the need for analyzing individual features. The proposed method is validated on wind farm data collected over multiple years and detects faults that have been confirmed through manual inspection by engineers.

2 Multivariate deep learning

The proposed method integrates time and frequency domain signal processing indicators with machine learning models to effectively identify mechanical failures. These indicators are computed from raw vibration signals and then utilized as inputs for multivariate deep learning models, such as deep autoencoders. The method establishes a normal behaviour model by learning the machine's healthy behaviour. The model can effectively predict and detect faults by observing any deviations from this learned healthy behaviour.

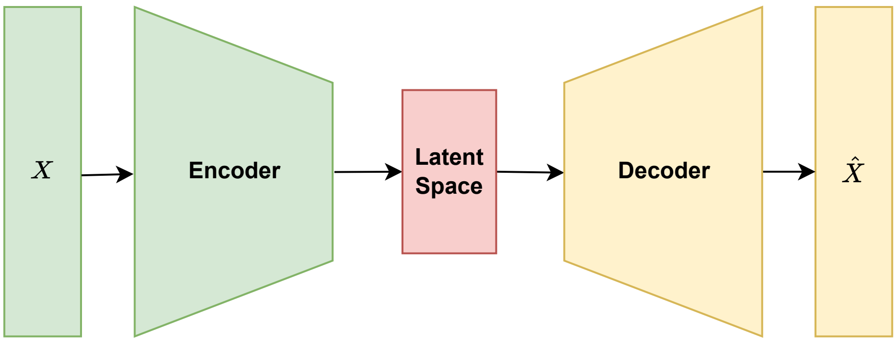

Figure 1: A schematic diagram deep autoencoder has three parts encoder, latent space, and decoder. The input X is passed at the input layer, and the autoencoder tries to reconstruct the output $\hat{X}$ resemblant as the input by minimising the RE

The autoencoder is a type of deep neural network that aims to reduce the dimensionality of input data by compressing it into a lower-dimensional representation and then reconstructing it to resemble the original input data. The common architecture of an autoencoder, depicted in Figure 1, consists of three main components: an encoder, a latent space, and a decoder. The encoder component gradually encodes the input data, transforming it into a representation in the latent space. This latent space serves as a condensed and lower-dimensional representation of the input data. The decoder component then takes this encoded data from the latent space and reconstructs it back to its original form, resembling the initial input data. Throughout this process, the autoencoder aims to minimize the reconstruction error (RE) between the actual input and the decoded output.

A preprocessing step is performed to eliminate interference removing from other components and environmental factors. This step involves separating the raw signal into deterministic and stochastic parts. Time domain and frequency domain indicators are computed on preprocessed signals. Time domain indicators are statistical indicators, while frequency domain indicators are primarily the characteristic frequencies observed in spectral and envelope domains. These indicators provide valuable insights into the underlying patterns and properties of the vibration signals, facilitating further analysis and interpretation.

The task of fault detection in wind turbines is particularly challenging due to the continuously changing operating conditions. The normal behaviour of wind turbines is constantly changing, therefore a robust fault detection method should be capable of adapting to these changing operating conditions. To address this challenge, the proposed method incorporates active power and rotation speed measurements, which introduce the operating condition information to the model. By integrating this relevant data into the multivariate autoencoder model, the method enhances its ability to accurately detect faults in wind turbines.

The proposed method consists of three essential steps:

- Normal behaviour physics-informed multivariate autoencoder (PIMA) training using healthy data.
- Computation of thresholds using healthy data.
- High-level health status prediction on new data.

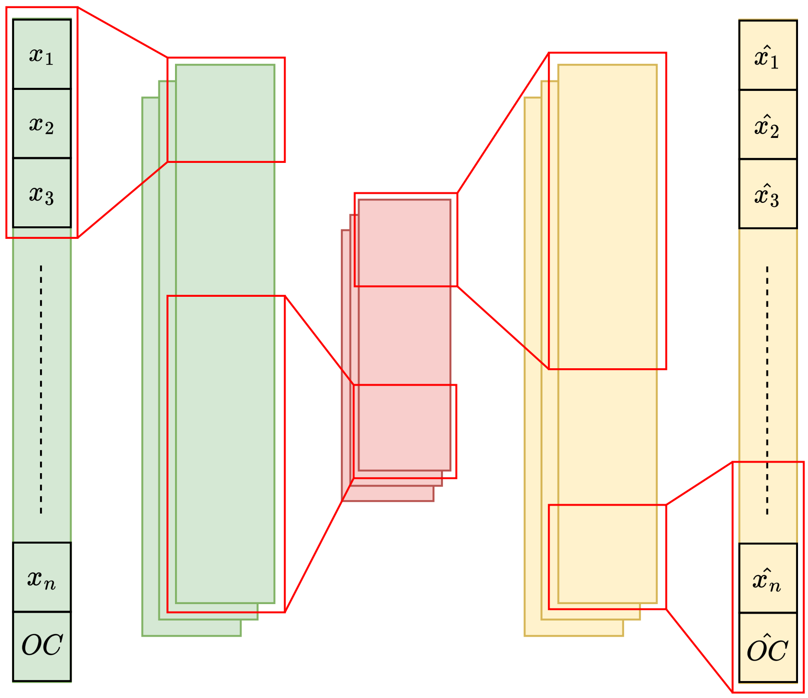

Figure 2: A schematic diagram of the multivariate convolutional neural network autoencoder architecture. The input layer takes signal processing indicators and operating conditions (OC) X as input. The encoder encodes the input to a latent space, and the decoder reconstructs the input $\hat{X}$ from the latent space.

2.1 Normal behaviour PIMA training

The proposed method utilizes a multivariate deep autoencoder, it requires the normalisation of indicators to ensure they are scaled into the same range. To achieve this, a Minmax scaler is used to normalise the indicators within the same range. By applying the Minmax scaler, all indicator values are scaled between 0 and 1, as demonstrated in Equation 1 and Equation 2.

$$X_{std} = \frac{x - min(x)}{max(x) - min(x)} \tag{1}$$

$$X_{norm} = X_{std} * (Max - Min) + Min \tag{2}$$

Where Min, Max is the normalisation range.

The normal behaviour multivariate deep autoencoder is trained on operating conditions and indicators computed during the healthy period of the wind turbine. The duration of this healthy period is determined by engineers by observing the wind turbine's condition monitoring indicators. The healthy training dataset D is a time series that encompasses M observations, denoted as t_i with $i = 1, ..., M$.

$$D = \{t_i | i = 1, ..., M\} \tag{3}$$

Each timestamp t_i within the time series contains the recorded operating conditions and the corresponding computed indicators.

$$X_i = \{x_1, x_2, x_3, ...x_n, oc\} \tag{4}$$

The normal behaviour multivariate deep autoencoder architecture is depicted in Figure 2. The specific configuration of the network architecture depends on the number of indicators and operating conditions involved. For this purpose, a one-dimensional convolutional network is employed, which comprises an encoder, a latent space, and a decoder. The encoder, denoted as function g, transforms the input X_i into a lower-dimensional latent representation Z_i. This transformation is achieved through the following equation 5:

$$Z_i = g(X_i) \tag{5}$$

The decoder f then reconstructs the original input $\hat{X}$ from the latent representation Z_i using the following equation 6:

$$\hat{X}_i = f(Z_i) \tag{6}$$

Where the $\hat{X}_i$ is decoded representation of the original input X_i:

$$\hat{X}_i = \{\hat{x}_1, \hat{x}_2, \hat{x}_3, ...\hat{x}_n, \hat{oc}\} \tag{7}$$

The mean squared error (MSE) is used as a metric for RE between the original input X_i and the decoded output $\hat{X}_i$. During the training process, the autoencoder learns to minimize MSE, aiming to reduce the discrepancy between the original input and the reconstructed output.

$$RE = MSE = \frac{1}{N}\sum_{i=1}^{N}(X_i - \hat{X}_i)^2 \tag{8}$$

Where N represents the number of indicators and operating conditions in the input. X_i represents the original input, while $\hat{X}_i$ denotes the reconstructed output. The MSE is the difference between the original input and the reconstructed output. A lower value of the RE indicates a more precise reconstruction, demonstrating that the autoencoder has effectively minimized the discrepancy between the original input and the reconstructed output. To prevent overfitting, a regularisation term is incorporated into the MSE loss function.

$$MSE = \frac{1}{N}\sum_{i=1}^{N}(X_i - \hat{X}_i)^2 + \lambda L_2(W) \tag{9}$$

Where W is neural network parameters, L_2 is the regularisation term, and λ is the regularisation parameter that controls the effect of the regularisation term. L_2 regularisation is a sum of squares of the neural network weight parameters, which adds a term MSE loss function that penalises the sum of squares of the weight parameters in the neural network.

2.2 Threshold computation

The threshold is computed using the same healthy dataset D, which consists of M observations and is used to train the normal behaviour PIMA. This threshold represents a specific deviation from the normal behaviour of wind turbines. Different threshold levels can be computed based on the requirements of the application. To validate the proposed method, two threshold levels are considered: two and three standard deviations away from the mean reconstruction error (RE) calculated on the healthy dataset D. These threshold levels serve to categorize faults into different states, such as normal, warning, and alarm states. The mean and standard deviation of the healthy dataset D are computed as depicted in Equation 10 and Equation 11, respectively.

$$\mu = \frac{1}{M}\sum_{i=1}^{M} RE_i \tag{10}$$

$$\sigma = \sqrt{\frac{1}{M}\sum_{i=1}^{M}(RE_i - \mu)} \tag{11}$$

Two levels of thresholds are calculated using the computed mean and standard deviation. The warning threshold is determined as two standard deviations away from the mean RE as in Equation 12,

$$Threshold_1 = \mu + 2 * \sigma \tag{12}$$

while the alarm threshold is set at three standard deviations away from the RE mean as shown in Equation 13.

$$Threshold_2 = \mu + 3 * \sigma \tag{13}$$

2.3 High-level health status prediction

The normal behaviour PIMA is trained to reconstruct the typical functioning of wind turbines with a focus on minimizing the RE. When the RE of the model exceeds the predefined threshold, it signifies anomalous behaviour, indicating the potential presence of a fault. The new data point is assigned a label of normal, warning, or alarm based on the two thresholds. If the RE falls below the threshold level $Threshold_1$, it is labelled as normal. If the RE is between $Threshold_1$ and $Threshold_2$, it is assigned the warning label. Lastly, if the RE exceeds $Threshold_2$, it is categorized as an alarm. The threshold-based approach enables the detection of deviations from normal behaviour and facilitates the identification of potential faults in wind turbines. To mitigate false alarms during the healthy state of the wind turbine, a sliding window technique is employed. This technique involves counting the number of alarms within a window and labelling a data point as healthy if the number of alarms falls below a certain threshold, even if the RE exceeds $Threshold_1$ or $Threshold_2$.

3 Experiments

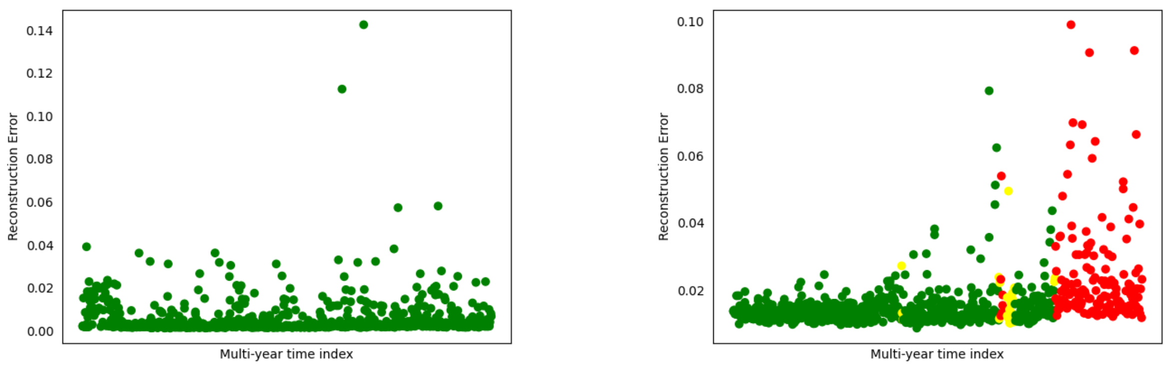

(a) Time domain statistical indicators

(b) Time and frequency domain indicators

Figure 3: The plots represent the fault trend of the planetary stage channel over multiple years. The x-axis denotes the multi-year time index, while the y-axis represents the RE. Each data point is colour-coded based on the corresponding alarm level. Green indicates healthy behaviour, yellow signifies a warning, and red represents an alarm state.

The proposed method is validated using a dataset obtained from an offshore wind farm consisting of more than 50 wind turbines. The data is collected for ten seconds every two or three days over the span of multiple years. Multiple channels are installed on the wind turbine drive train to capture the individual components' behaviour. A separate model is trained on each channel to provide a comprehensive overview of the high-level health status based on multiple signal-processing computed indicators. The normal behaviour PIMA is trained using approximately one year of healthy data. This trained model is then utilized to predict the high-level health status over multiple years, enabling the detection of faults at various stages. To ensure the accuracy and reliability of the detected faults, engineers conduct manual borescope inspections of the components to confirm the presence of the faults. This validation process involves expert evaluation and provides additional verification of the detected faults identified by the proposed method.

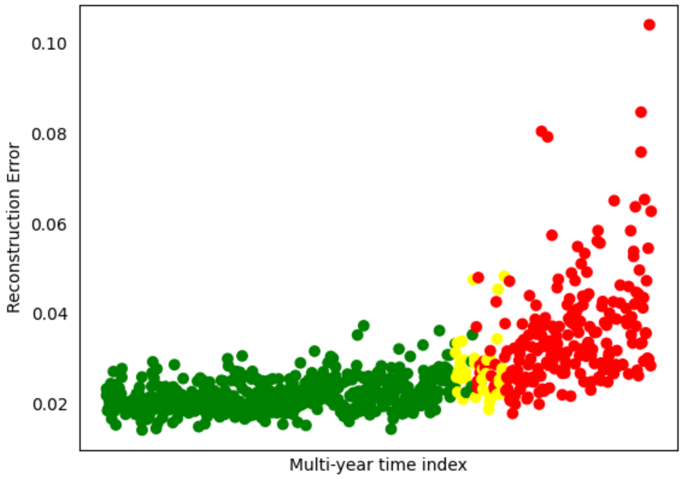

Figure 4: The fault trend of the generator channel is depicted over multiple years. The x-axis represents the multi-year time index, while the y-axis corresponds to the RE. The colours green, yellow, and red respectively indicate the healthy, warning, and alarm states.

To further emphasize the significance of utilizing multiple indicators, the following results are presented, demonstrating instances where faults are not detected solely based on time-domain statistical indicators. However, by incorporating frequency domain indicators in conjunction with statistical indicators, the proposed method successfully detected the faults. By leveraging a diverse range of indicators, the proposed method enhances its ability to capture various fault patterns and improve the overall accuracy of fault detection because different types of indicators exhibit sensitivity towards distinct types of faults. The results of the proposed method are demonstrated in Figure 3, Figure 4, and Figure 5. These figures illustrate the high-level health status labels determined by the predicted RE of the normal behaviour PIMA. The assigned colours correspond to different levels of alarms determine by threshold levels defined in Equation 12 and Equation 13. During the healthy state of the wind turbine, the model predicted RE is below $Threshold_1$, which is labelled as green and represents normal behaviour. In the presence of faults, the RE gradually deviates from the normal behaviour and exceeds $Threshold_1$, indicating a warning state and labelled yellow. As the fault severity increases, the RE surpasses the second threshold, indicating an alarm and labelled red. The proposed method demonstrates its capability to detect faults at early stages, as the RE increases with an increase in fault severity. These detected faults are further confirmed through the observation of cracks on the alarmed components during manual borescope inspection.

Figure 3 showcases the fault in the planetary stage channel. Figure 3a utilizes time domain statistical indicators, while Figure 3b combines both statistical and frequency domain indicators to compute the high-level health status. The statistical indicators alone are not sufficient as input to the normal behaviour PIMA for detecting specific faults in the planetary stage channel. However, by combining them with frequency domain indicators, the method becomes capable of detecting different types of faults. The generator channel fault case depicted in Figure 4 demonstrates a gradual increase in alarms over time. The fault trend consistently exceeds both $Threshold_1$ and $Threshold_2$, indicating a warning and alarm state, respectively. The fault detected by the proposed method was further confirmed through manual borescope inspection. Figure 5 illustrates a healthy case in the high-speed stage channel where statistical and frequency domain indicators are utilized as input to the proposed model. The plot demonstrates a consistent trend in the data, with the RE consistently staying below $Threshold_1$, which indicates the high-speed shaft healthy behaviour.

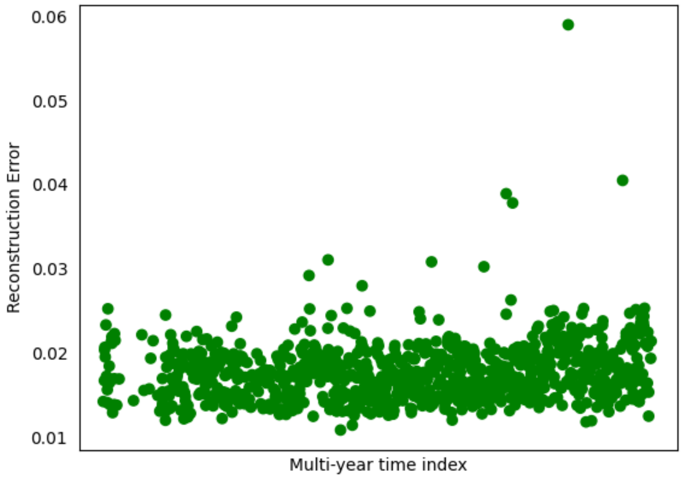

Figure 5: The plot illustrates the healthy trend of the high-speed stage channel over multiple years. The x-axis represents the multi-year time index, while the y-axis corresponds to the RE. The RE is less than $Threshold_1$ indicating the healthy behaviour of the high-speed shaft.

4 Conclusion

The proposed normal behaviour physics-informed deep learning method demonstrated its effectiveness in detecting faults in wind turbine drivetrain components. By leveraging multiple indicators, including time domain statistical and frequency indicators, the method provided a comprehensive high-level health status overview of wind turbine drivetrain components. Validation of the method on a planetary stage channel fault highlighted the limitations of using solely statistical indicators as model input, while the integration of frequency domain indicators successfully detected the fault. Through the utilization of multiple indicators, the proposed method offers valuable insights and a comprehensive health status overview for wind farm operators.

Acknowledgments

The authors would like to acknowledge FWO (Fonds Wetenschappelijk Onderzoek) for their support through the SB grants of Faras Jamil (#1S63123N), post-doctoral grant of Cédric Peeters (#1282221N), and SBO project Robustify (S006119N). Furthermore, this research was supported by funding from the Flemish Government under the “Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen” programme and from the Blauwe Cluster ICON Supersized 4.0 project. The authors are also grateful to the VSC Supercomputing Flanders centre for the support in the context of the VSC Cloud program.

References

[1] Jan Helsen, Cédric Peeters, Peter Doro, Eveline Ververs, and Pieter Jan Jordaens. Wind farm operation and maintenance optimization using big data. In *2017 IEEE Third International Conference on Big Data Computing Service and Applications (BigDataService)*, pages 179–184, 2017.

[2] Cédric Peeters, Quentin Leclére, Jérôme Antoni, Peter Lindahl, John Donnal, Steven Leeb, and Jan Helsen. Review and comparison of tacholess instantaneous speed estimation methods on experimental vibration data. *Mechanical Systems and Signal Processing*, 129:407–436, 2019.

[3] Cédric Peeters, Patrick Guillaume, and Jan Helsen. A comparison of cepstral editing methods as signal pre-processing techniques for vibration-based bearing fault detection. *Mechanical Systems and Signal Processing*, 91:354–381, 2017.

[4] Cédric Peeters, Nicoletta Gioia, and Jan Helsen. Stochastic simulation assessment of an automated vibration-based condition monitoring framework for wind turbine gearbox faults. *Journal of Physics: Conference Series*, 1037(3):032044, jun 2018.

[5] Jérôme Antoni. Cyclostationarity by examples. *Mechanical Systems and Signal Processing*, 23(4):987–1036, 2009.

[6] D. HO and R.B. RANDALL. Optimisation of bearing diagnostic techniques using simulated and actual bearing fault signals. *Mechanical Systems and Signal Processing*, 14(5):763–788, 2000.

[7] A.C. McCormick and A.K. Nandi. Cyclostationarity in rotating machine vibrations. *Mechanical Systems and Signal Processing*, 12(2):225–242, 1998.

[8] Faras Jamil, Timothy Verstraeten, Ann Nowé, Cédric Peeters, and Jan Helsen. A deep boosted transfer learning method for wind turbine gearbox fault detection. *Renewable Energy*, 197:331–341, 2022.

[9] Cédric Peeters, Timothy Verstraeten, Ann Nowé, and Jan Helsen. Wind Turbine Planetary Gear Fault Identification Using Statistical Condition Indicators and Machine Learning. volume Volume 10: Ocean Renewable Energy of *International Conference on Offshore Mechanics and Arctic Engineering*, 06 2019.

[10] Jan Helsen, Cédric Peeters, Timothy Verstraeten, J Verbeke, N Gioia, and A Nowé. Fleet-wide condition monitoring combining vibration signal processing and machine learning rolled out in a cloud-computing environment. In *International Conference on Noise and Vibration Engineering (ISMA)*, pages 17–19, 2018.

[11] Fabian Perez-Sanjines, Cédric Peeters, Timothy Verstraeten, Jérôme Antoni, Ann Nowé, and Jan Helsen. Fleet-based early fault detection of wind turbine gearboxes using physics-informed deep learning based on cyclic spectral coherence. *Mechanical Systems and Signal Processing*, 185:109760, 2023.

[12] Jan Helsen, Christof Devriendt, Wout Weijtjens, and Patrick Guillaume. Condition monitoring by means of scada analysis. In *Proceedings of European wind energy association international conference Paris*, 2015.

[13] X Chesterman, T Verstraeten, PJ Daems, F Perez Sanjines, A Nowé, and J Helsen. The detection of generator bearing failures on wind turbines using machine learning based anomaly detection. *Journal of Physics: Conference Series*, 2265(3):032066, may 2022.