

Article

Vibration Fault Detection in Wind Turbines Based on Normal Behaviour Models without Feature Engineering

Stefan Jonas ^{1,*}, Dimitrios Anagnostos ², Bernhard Brodbeck ² and Angela Meyer ¹ 

¹ School of Engineering and Computer Science, Bern University of Applied Sciences, Quellgasse 21, 2501 Biel, Switzerland

² Winji AG, Badenerstrasse 808, 8048 Zurich, Switzerland

* Correspondence: stefan.jonas@bfh.ch

Abstract: Most wind turbines are remotely monitored 24/7 to allow for early detection of operation problems and developing damage. We present a new fault detection approach for vibration-monitored drivetrains that does not require any feature engineering. Our method relies on a simple model architecture to enable a straightforward implementation in practice. We propose to apply convolutional autoencoders for identifying and extracting the most relevant features from a broad continuous range of the spectrum in an automated manner, saving time and effort. We focus on the range of [0, 1000] Hz for demonstration purposes. A spectral model of the normal vibration response is learnt for the monitored component from past measurements. We demonstrate that the trained model can successfully distinguish damaged from healthy components and detect a damaged generator bearing and damaged gearbox parts from their vibration responses. Using measurements from commercial wind turbines and a test rig, we show that vibration-based fault detection in wind turbine drivetrains can be performed without the usual upfront definition of spectral features. Another advantage of the presented method is that a broad continuous range of the spectrum can be monitored instead of the usual focus on monitoring individual frequencies and harmonics. Future research should investigate the proposed method on more comprehensive datasets and fault types.

Keywords: condition monitoring; wind turbines; fault detection; vibrations; autoencoders; convolutional autoencoders; neural networks; renewable energy



Citation: Jonas, S.; Anagnostos, D.; Brodbeck, B.; Meyer, A. Vibration Fault Detection in Wind Turbines Based on Normal Behaviour Models without Feature Engineering. *Energies* **2023**, *16*, 1760. <https://doi.org/10.3390/en16041760>

Academic Editor: Francesc Pozo

Received: 19 January 2023

Revised: 2 February 2023

Accepted: 6 February 2023

Published: 10 February 2023



Copyright: © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Wind energy is becoming an essential pillar of the energy mix, helping to decarbonize the energy system and to improve energy independence and security of supply in many countries [1,2]. Most commercial wind farms today are remotely monitored around the clock to keep the operation and maintenance costs low. The continuous monitoring allows for early detection of potential operation problems and facilitates proactive condition-based maintenance. Gearboxes and generators are of particular interest in the monitoring of wind turbines (WT), because they are especially costly to replace, and associated replacement work tends to entail long downtimes [3–6]. Thus, more and more wind turbine drivetrains are monitored with accelerometers. This makes it possible to derive and track the vibration spectra of critical components such as gearbox parts and generator bearings.

Previous studies have proposed various methods and features for vibration-based fault detection in the drivetrains of wind turbines. The proposed frequency-domain methods include the monitoring and analysis of particular frequencies, their harmonics, sidebands and the signal envelope. Recent studies in this area applied to wind turbine gearboxes include [7], in which conventional filtering, Hilbert transform and cepstrum methods are applied for detecting broken teeth faults of gears. To extract features for weak faults however, complex wavelet transform providing multi-scale enveloping spectrograms from which bearing failure can be detected is introduced. The work in [8] proposed a method

based on the Vold–Kalman filter and higher order energy separation. In [9], a method based on sparse representation theory is outlined, with the stated advantages over conventional methods being in alleviating background noise, handling the coupling effect and a more intuitive fault diagnosis process. Similarly, Ref. [10] introduced a structured sparsity method, highlighting and demonstrating the same advantages. A novel modulation model for the diagnosis of compound faults is shown in [11]. The presented approach enables adaptive decomposition of the signal, with the authors presenting superior performance compared to classic empirical mode decomposition methods. In the time domain, the existing methods focus on analysing the amplitudes of the vibration time series and tracking statistical properties of the vibration response distribution. Recent contributions include [12], presenting a fault diagnostic method based on measurement indicators such as root mean square analysis (RMS), skewness, kurtosis or the crest factor. Experimental results on a fixed-speed planetary gearbox test rig indicated that promising results in discriminating fault conditions could be obtained using RMS analysis when compared with other indicators. However, the author notes the indicator's sensitivity to variable operational conditions. In [13], mean vibration signatures were extracted from recordings of a group of healthy identical WTs. The resulting measurements were used as reference, and by detecting deviations, faults can be identified, as presented in their case study. A different approach is outlined in [14], where these traditional vibration-based features such as RMS or peak-to-peak were replaced with linear regression parameters of the feature-generator power relationship, allowing for load-independent evaluations and higher interpretability. Methods based on extracting features from multiple domains have also been proposed: Most recently, in [15], a fault diagnosis method based on fused features is presented, of which an optimal subset is chosen by a statistical feature selection. In [16], features from the time, frequency, and time–frequency domains are extracted and combined as well. A random forest is used to reduce the dimensionality of the feature set and to select an optimal subset for the fault diagnosis model. We refer to [17,18] for comprehensive reviews of the state-of-the-art methods of vibration-based fault detection in wind turbine drivetrains.

The existing methods require the upfront development and extraction of component-specific features from the accelerometer measurements by condition-monitoring engineers. The upfront definition of spectral features for the monitoring constitutes a major time investment before commissioning and operating the turbine. For instance, in-depth information about the respective gearbox design and composition are needed to this end. The characteristic frequencies of the monitored components need to be determined if the corresponding spectral lines are supposed to be tracked. Collecting this spectral information for every monitored component in every monitored wind turbine and model constitutes a major effort. The extracted spectral features tend to be turbine- and component-specific. Therefore, they can generally not be reused for new wind turbine types or when components have been updated. Moreover, feature engineering can result in more and in less effective features for fault detection. For example, in a case study by [19], spectral line and cepstrum analyses resulted in higher fault detection accuracies than time synchronous averaging and spectral kurtosis methods. Thus, depending on the chosen feature engineering approach, one may end up with more or less satisfactory fault detection accuracies.

Recently, first studies have proposed the application of fault detection methods that do not require any feature engineering. The work in [20] demonstrated the application of autoencoders for detecting blade damage in wind turbines based on the blade stress and strain signals obtained from strain gauge sensors. In a similar application, Ref. [21] made use of a convolutional autoencoder for detecting blade damages. To this end, they employed data from the wind turbine's supervisory control and data acquisition (SCADA) system to identify changes in the dynamics of the blade system. In [22], an autoencoder is applied for the transfer learning of fault diagnosis tasks on SCADA and failure status datasets. They investigated different fault types and focused on the transfer of fault diagnostics models to target turbines with few available SCADA data. When vibration sensors are unavailable,

SCADA-based modelling of the turbine's normal behaviour can enable the detection of operation faults in drivetrain components (e.g., [23–28]).

The goal of our study is to develop and demonstrate a fault detection method for vibration-monitored drivetrain components in wind turbines that does not require any feature engineering but is able to learn the characteristic spectral features of the monitored components without any human assistance from a continuous range of the spectrum, such as the full half spectrum. The method should enable simple model architectures to facilitate its adoption by practitioners. At the same time, it should neither be restricted to monitoring specific frequencies, nor should it require a comprehensive set of gearbox- and fault-type-specific observations. The latter would be needed for fault detection methods based on supervised machine learning, but in practice, such datasets are usually not available or not accessible to wind turbine operators.

To achieve this goal, the present study proposes and demonstrates the application of convolutional autoencoders for the feature learning, extraction and early detection of faults in vibration-monitored wind turbine gearboxes and generators. This study is the first to propose spectral normal-behaviour models constructed without feature engineering for the purpose of vibration-based fault detection in wind turbine drivetrains, to the best of our knowledge.

In this study, we also propose and investigate fault detection based on non-convolutional autoencoders. Moreover, we compare the performance of reconstruction-error-based gearbox and generator health indices to health indices derived by one-class classification with an isolation forest model.

This paper is organized as follows. Section 2 describes our new fault detection methods. Section 3 introduces the gearbox and generator datasets and models used for demonstrating the methods. The results of our study are discussed in Section 4, and our conclusions follow in Section 5.

2. Fault Detection Algorithms

Autoencoders are artificial neural networks that are trained to compress and then reconstruct their input (e.g., [29–31]). Autoencoders are autoassociative, in that they are trained to duplicate their input as their output and to minimize the resulting reconstruction error. In the training process, the reconstruction error is minimized by adapting the autoencoder's network weights to the training data in an iterative manner. Autoencoders consist of encoding and decoding hidden layers. The encoding layers learn a lower dimensional representation of the training data, which is subsequently decompressed by the decoding layers to reconstruct an output highly similar to the input, with the same dimensionality as the input. The dimensionality reduction of the input data has a similar compression effect as a principal component analysis. The proposed convolutional autoencoders reduce the dimension of the spectral feature space and replace human feature engineering. Convolutional autoencoders are convolutional neural networks with an encoder–decoder architecture. They are particularly suited for image data compression and decompression because of the feature learning and image reconstruction capabilities of the convolutional layers. Autoencoders can be used to detect anomalous input data instances based on the large reconstruction errors they make when trying to reconstruct those instances from their compressed representations. Therefore, convolutional autoencoders have been successful in image and video-based anomaly detection tasks (e.g., [32,33]).

In this study, we propose to apply convolutional autoencoders to learn and extract spectral features from the accelerometer measurements and to use the autoencoder reconstruction errors as an indicator (health index) of unusual operation behaviour and potential faults of the monitored components. Alternatively, the extracted features can be input to a one-class classification algorithm, such as an isolation forest model, to construct a health index of the monitored component, as demonstrated in our case study. We present two fault detection methods below that rely on two steps: (1) automated data-driven feature learning and extraction by autoencoders to model the normal vibration responses of the

monitored component, and (2) subsequent fault detection based on deviations from the spectral normal-behaviour model learnt in step (1). The two fault detection methods may be used individually or in combination. We demonstrate and discuss them in a case study in Sections 3 and 4.

2.1. Fault Detection Based on Autoencoder Reconstruction Losses

As shown in Figure 1, an autoencoder is trained with spectrogram segments constructed from the accelerometer measurements. The autoencoder's reconstruction loss is used as an anomaly score (health index) in this case. If the autoencoder has been trained only on vibration spectra from components unaffected by faults, then the reconstruction error of spectra from fault-affected components is expected to be higher than that of unaffected components. Thus, the higher the reconstruction loss of an input segment, the more likely its underlying spectrogram is to indicate a fault in the monitored component. In this approach, the autoencoder should be trained on data from healthy components, so it learns to reconstruct spectra that reflect the normal operation behaviour of the monitored component. Thus, the autoencoder learns a normal-behaviour model of the vibration responses of the component. The underlying idea of this unsupervised fault detection approach is similar to that of the normal-behaviour models in SCADA-based condition monitoring [23–28]. Training, validation, and test sets are created from the spectrograms constructed from the vibration responses of each monitored component. The mean absolute reconstruction error is calculated for each segment by comparing its autoencoder reconstruction to the original input. To identify fault-affected spectrograms, we define a loss threshold T . When this threshold is exceeded, the component associated with the spectrogram is considered to be fault-affected.

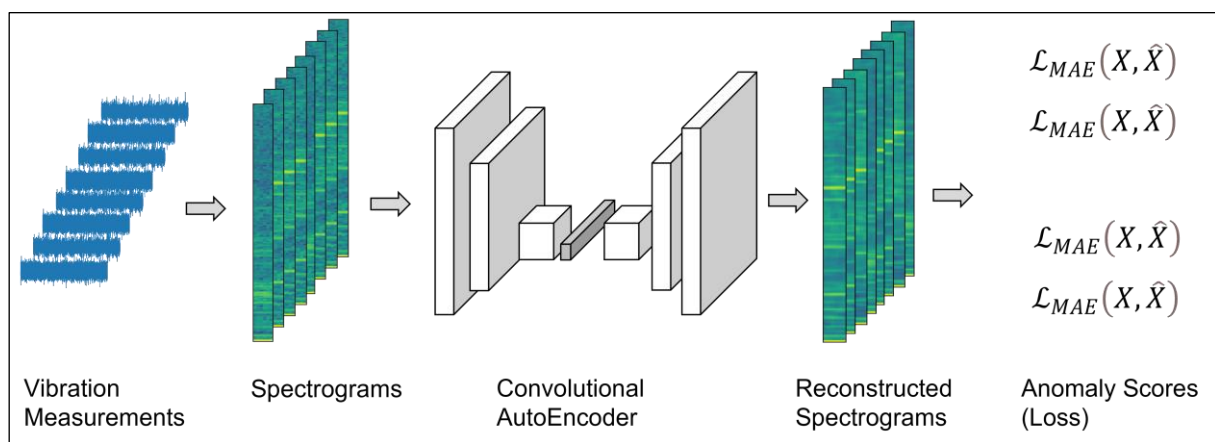


Figure 1. An illustration of the fault detection approach based on autoencoder reconstruction errors. First, the vibration measurements are converted to spectrogram segments. The segments are given as input into a convolutional autoencoder, which returns reconstructed spectrograms as output. The mean absolute reconstruction errors are calculated with respect to the original inputs. These errors are used as health indices to distinguish healthy and fault-affected components.

Formally, an autoencoder consists of encoder network ϕ_e and decoder network ϕ_d , parameterized by Θ_e and Θ_d respectively, derived by minimizing the reconstruction loss during training. The health index h assigned to a spectrogram segment X_k recorded with accelerometer s is obtained by calculating the mean absolute error \mathcal{L}_{MAE,AE_s} between the original input X_k and its reconstruction \hat{X}_k by an autoencoder AE_s that was trained on only healthy segments from accelerometer s . The classification as fault-affected is based on whether the health index h exceeds a predefined threshold T . This threshold will ensure

that normal variability in the spectrogram will not be flagged as fault conditions, whereas unusual spectral changes will be detected.

$$\hat{X}_{k,s} = \phi_{d,s}(\phi_{e,s}(X_{k,s}; \Theta_{e,s}); \Theta_{d,s})$$

$$h(X_{k,s}) = \mathcal{L}_{MAE,AE_s}(X_{k,s}, \hat{X}_{k,s})$$

$$anomaly(X_{k,s}) = \text{true, if } h(X_{k,s}) > T$$

2.2. Fault Detection Based on Isolation Forests

The second fault detection approach presented in this study combines the encoding part of a convolutional autoencoder with one-class classification by an isolation forest model, as shown in Figure 2. As in the first approach, a convolutional autoencoder is trained with the goal of reconstructing spectrogram segments. In contrast to the first approach, however, the autoencoder will now only be used to extract a compressed representation of the most important features from the input spectrogram. In a second step, the extracted features serve as input to an isolation forest model [34,35]. The features are taken from the output of the last encoding layer of the autoencoder architecture. In this study, we call this layer the bottleneck layer.




Figure 2. An illustration of the fault detection approach based on isolation forests. First, the vibration measurements are converted to spectrogram segments. The segments are given as input to the encoding part of a convolutional autoencoder, the last layer outputs of which are an encoded representation of the segment features (shaded layer). Then, an isolation forest model computes an anomaly score for each segment. The score is used as a health index to distinguish healthy and fault-affected components.

The isolation forest model detects an anomalous spectrogram based on how easily the spectrogram can be isolated from the remaining spectrograms in the decision trees constituting the forest. In this way, the isolation forest model distinguishes anomalous spectrograms from spectrograms featuring normal operation behaviour. A separate isolation forest model is created for each monitored component and accelerometer. The inputs to the forest are the feature vectors extracted from the bottleneck layer of the previously trained autoencoder. In a final step, the model evaluates all segments from the training, validation and test sets and outputs an anomaly score. We need to define a threshold to discriminate spectrograms that reflect normal operation behaviour from spectrograms of fault-affected components. This threshold is set to 0.5 in accordance with [34].

More formally, for a spectrogram segment X_k recorded by accelerometer s , the trained isolation forest model IF for s outputs the health index h based on the feature vector $z(X_k)$ output from the encoder network $\phi_{e,s}$ of autoencoder AE_s :

$$\begin{aligned} z(X_{k,s}) &= \phi_{e,s}(X_{k,s}; \Theta_{e,s}) \\ h(X_{k,s}) &= IF_s(z(X_{k,s})) \\ anomaly(X_{k,s}) &= \text{true, if } h(X_{k,s}) > 0.5 \end{aligned}$$

3. Case Study

The proposed fault detection methods were demonstrated in four wind turbine drivetrains. First, we tested the approach in two commercially operated wind turbines, which we call wind turbine 1 (WT1) and wind turbine 2 (WT2) in this study. We knew from the operators that WT1 exhibited no known damage nor functional impairments, whereas WT2 may have suffered damage in a drivetrain component. Our goal was to test the proposed fault detection methods by detecting if any of the monitored components of WT2 are affected by faults based on their vibration responses, and to identify the fault-affected components, if any.

In addition to the commercial wind turbines, we also tested the performance of our fault detection methods with test rig measurements from the National Renewable Energy Laboratory (NREL).

3.1. Commercially Operated Wind Turbines

Accelerometer measurements from the drivetrains of two commercially operated multi-MW wind turbines were analysed. The two onshore turbines were identical in design and composition, with a nominal power of 2 MW. Table 1 provides further specifications of the turbines.


Table 1. Technical specifications of the two commercial wind turbines.

Quantity	Value
Rotor diameter	114 m
Nominal power	2.1 MW
Type	variable-speed horizontal-axis pitch-controlled
Deployment	onshore
Gearbox	three stage

In each turbine (WT1, WT2), eight accelerometers (S1–S8) monitored the gearbox, generator, main shaft bearing and the main frame, as detailed in Table 2. Figure 3 illustrates the monitored components and the locations of the accelerometers. The measurements were taken at a sampling frequency of 25.6 kHz in time slices of 2 s each, over a period of three months. All time slice measurements were taken at the same generator load of 1345 ± 3 rpm. The constant load enabled a direct comparison of the spectra across the full measurement period of three months and between the two turbines. In total, the dataset contained 562 time slice recordings from WT1 and 672 recordings from WT2. The time slice recordings were converted into spectrograms. According to information provided to us by the operator, WT1 exhibited no known damage or impairments, whereas WT2 may have damage in any component. Therefore, our goal was to demonstrate our presented fault detection methods by evaluating whether any monitored components of WT2 are affected by faults, based on the fault detection model evaluations of spectrograms from each component of WT2.

Table 2. The monitored components and measured acceleration directions in the two wind turbines.

Accelerometer ID	Monitored Component and Direction
S1	Main frame, axial
S2	Gearbox, rotor side, axial
S3	Main frame, radial
S4	Gearbox, generator side, radial
S5	Main shaft bearing, radial
S6	Gearbox, generator side, axial
S7	Generator, coupling side, axial
S8	Generator, coupling side, radial

**Figure 3.** Sketch of the drivetrain of the two wind turbines. The accelerometers S1–S8 are indicated in red.

3.2. Test Rig Measurements

In addition to the commercial wind turbines, we tested our fault detection methods with accelerometer measurements from the drivetrains of two wind turbines operated on a NREL test rig. The turbines were two 750 kW machines identical in construction. The gearboxes provided a transmission of 1:81.5 and were composed of a low-speed planetary stage and two parallel moderate and high-speed stages. One of the gearboxes showed no damage or functional impairments. The other gearbox had suffered an oil loss and resulting damages in multiple components. The vibration measurements were taken from multiple components of the gearboxes as part of the Gearbox Condition Monitoring Round Robin study [36]. The measurements were conducted for ten minutes and under constant load at both gearboxes, with 22 rpm of the low-speed shaft and 1800 rpm of the high-speed shaft. The accelerometers measured at a 40 kHz sampling frequency. A comprehensive description of the gearbox, the sensing system and test environment are provided by [36,37].

The vibration measurements analysed in this study were obtained from components which exhibited different damage levels in the gearbox that had suffered the oil loss, such as scuffing damage. Their counterparts in the healthy gearbox have been examined for damage as part of the round robin study [36], and no damage was found in them. The accelerometers monitored the vibrations in the following locations in the damaged and in the healthy gearbox:

1. In a bottom-facing position at the ring gear,
2. At the bearing of the high-speed shaft,
3. At the bearing of the low-speed shaft.

All accelerometers measured the accelerations in the respective radial direction.

4. Results and Discussion

In the following sections, we will first present and discuss the results of our proposed methods for our case study with two commercially operated wind turbines (Section 3.1). To convert the accelerometer measurements into spectrograms, we applied a short-time Fourier transform (STFT; [38]) with window sizes of 250 ms and an overlap of 100 ms. Fault detection was performed based on spectrogram segments of 1 s duration and considering frequencies up to 1000 Hz. These parameters were chosen so as to enable sufficient temporal evolution and frequency resolution of the signal and at the same time allow for large enough training, validation and test sets. Our results are robust against modifications of these parameters. Prior to the model training, we applied a log-transformation to each spectrogram and a min–max normalization, such that all data points fall in the range of [0, 1]. Figure 4 shows examples of the resulting spectrograms for the components monitored with the accelerometers S1–S8 in both wind turbines.




Figure 4. Sample spectrograms of WT1 and WT2, converted from the measurements of components monitored by the accelerometers S1–S8. The figure shows one sample from each wind turbine and each accelerometer. All spectrogram segments run over 1 s of time (the vertical axis) and show normalized logarithmic amplitudes for frequencies up to 1000 Hz (horizontal axis).

Dataset splits. For our case study, we created a separate model for each accelerometer. We split the recordings of WT1 into a 70% training set, a 15% validation set and a 15% test set for each accelerometer. A sliding window was applied to extract multiple 1 s segments from each 2 s recording. The resulting segments of the training set were used to train the autoencoders and the isolation forests, while the validation set was used for the model selection and early stopping. All data from WT2 served as an additional separate test set, so we could obtain fault notifications for each component.

Model selection. We performed a preliminary model search using the Hyperband hyperparameter optimization algorithm [39] in order to find an optimal convolutional autoencoder architecture. We evaluated architectures consisting of only convolutional and pooling layers which resulted in a bottleneck size of 128 units (the flattened size of the last encoder layer). This dimension requirement was set beforehand by us. Choosing a bottleneck size significantly smaller than the input size forces the network to learn a compressed representation of the inputs. This compressed representation serves as input to the isolation forest. In terms of configurations, we evaluated models with up to five encoding layers with either four, eight, or sixteen feature maps, with kernel sizes (height, width) of either (1, 3), (1, 5), (3, 3), (3, 5), or (5, 5) and with varying learning rates between 3×10^{-2} and 1×10^{-4} . The decoder architecture was always symmetrical to the encoder. The network weights were optimized using the adaptive moment estimation (Adam) optimization algorithm [40], minimizing the mean absolute error (MAE) between

the reconstruction output and the original input. Further, we applied an early stopping mechanism to stop training when the validation loss had not improved within 15 epochs.

This hyperparameter optimization was performed only on the training and validation set of accelerometer S1. The Hyperband search algorithm resulted in the best-performing convolutional architecture (“conv-AE”) outlined in Table 3. Additionally, we compared the reconstruction performance to a minimal dense architecture (“dense-AE”) described in Table 3, i.e., to the smallest possible model configuration in terms of parameters with at least one hidden fully connected layer and the same bottleneck size.

Table 3. Specifications of the models used in our work. The abbreviations are “FC Layer”: fully-connected layer (number of units, activation function); “Conv. Layer”: convolutional layer (number of feature maps, kernel size in height \times width, activation function); “MP”: max-pooling layer (height \times width); “elu”: exponential linear unit.

Model	Architecture
Convolutional autoencoder (conv-AE)	<ul style="list-style-type: none"> • Input: $6 \times 251 \times 1$ (height \times width \times channels), zero-padded to $8 \times 256 \times 1$ • Encoder: Block 1: Conv. Layer (16f, 1×3, elu), MP (1, 2) Block 2: Conv. Layer (16f, 3×3, elu), MP (2, 2) Block 3: Conv. Layer (16f, 1×3, elu), MP (2, 2) Block 4: Conv. Layer (16f, 1×5, elu), MP (1, 2) Block 5: Conv. Layer (8f, 1×5, elu), MP (1, 2) Bottleneck Layer: Flatten ($2H \times 8W \times 8f$ to 128 units) • Decoder: Reshape: (128 units to $2H \times 8W \times 8f$) Block 1: Upsampling (1, 2), Transposed Conv. Layer (8f, 1×5, elu) Block 2: Upsampling (1, 2), Transposed Conv. Layer (16f, 1×5, elu) Block 3: Upsampling (2, 2), Transposed Conv. Layer (16f, 1×3, elu) Block 4: Upsampling (2, 2), Transposed Conv. Layer (16f, 3×3, elu) Block 5: Upsampling (1, 2), Transposed Conv. Layer (16f, 1×3, elu) • Reconstruction: Cropping Layer: ($8H \times 251W \times 16f$ to $6H \times 251W \times 16f$) Conv Layer: (1f, 5×5, sigmoid) <p>Number of parameters: 10,385</p>
Dense autoencoder (dense-AE)	<ul style="list-style-type: none"> • Input: $6 \times 251 \times 1$ (height \times width \times channels) flattened to 1506 units • Encoder: Hidden Layer: FC Layer (128 units, elu) Bottleneck Layer: FC Layer (128 units, elu) • Decoder: Hidden Layer: FC Layer (128 units, elu) • Reconstruction Layer: FC Layer (1506 units, sigmoid) Output reshaped back to $6 \times 251 \times 1$ <p>Number of parameters: 420,194</p>
Isolation Forest	<p>Number of trees: 100 Contamination: 0.0001</p>

Both the convolutional and the dense model architectures evaluated by us are capable of reconstructing visually similar spectrogram segments after finished training, as shown by the autoencoder reconstructions in Figure 5. While the training loss was lower using the dense autoencoder model (3.55×10^{-2}) compared to the convolutional autoencoder (3.63×10^{-2}), the validation loss was higher (3.82×10^{-2}) compared to the convolutional autoencoder model (3.70×10^{-2}), indicating a poorer reconstruction performance of the dense autoencoder model on the same unseen dataset. We attribute the divergence of the training and validation losses observed in the case of the dense autoencoder to its overfitting on the training set, due to the dense autoencoder’s large number of parameters. Our case study showed that the convolution-based architecture can achieve a better reconstruction performance on unseen data with a comparably small number of parameters, specifically with only 2.5% of the dense autoencoder model’s number of parameters. At the same time, the convolutional autoencoder maintained a good correspondence between training and validation losses. This suggests that the convolutional autoencoder network learns more generalizable features and is less prone to overfitting in this application. In addition, the dense autoencoder has a large number of parameters, so its training can quickly become computationally expensive when considering multiple models or larger datasets with even more components and wind turbines. Based on these results, we determined that the

convolutional architecture is better suited for our task. Consequently, we proceeded by using this configuration for all further experiments.

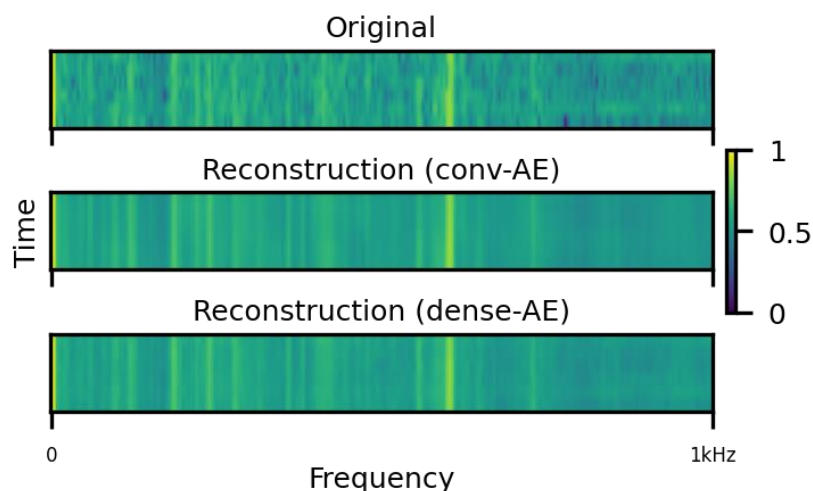


Figure 5. An example of the reconstruction capabilities of the two considered autoencoder models. An original spectrogram segment (**top**) is given as input to an autoencoder, which outputs a reconstruction. A convolutional and a dense autoencoder were adapted to reconstruct a spectrogram derived from accelerometer S1 in this illustration. The reconstruction by the dense autoencoder (**bottom**) suggests that this model also tries to fit more of the noise in the spectrogram, which is consistent with the overfitting observed with the dense autoencoder.

Fault detection based on autoencoder reconstruction losses. We trained a separate convolutional autoencoder network with the conv-AE configuration for each accelerometer, using 70% of the measurements of the monitored healthy component from WT1 as the training set and 15% as the validation set. The networks were trained with the same procedure as outlined in the model selection experiment. During the training, all eight autoencoders achieved similar performances in terms of their validation losses (3.48×10^{-2} – 4.68×10^{-2}). We evaluated the reconstruction errors obtained from the reconstructions of the segments in the training, validation, and test sets of WT1 and the test data from WT2. A threshold was determined based on the training errors, and if it was exceeded, we considered the segment as anomalous.

Our autoencoder-based fault detection method found increased reconstruction errors for the spectrograms derived from accelerometers S7 and S8, as shown in Figures 6 and 7. These sensors monitor the generator component coupling to the gearbox, as shown in Figure 3. The health index in Figure 7 displays the evolution of the health of each monitored component and the degree of anomaly of the component’s vibrational response. Based on the sensors from WT1, we estimated that an appropriate threshold value for bounding the reconstruction errors of normal-vibration-response spectrograms is around 0.6. To estimate this threshold more accurately, a larger number of fault instances would be needed. Based on the health index, we defined a custom rule for notifying the wind turbine operators when a certain number of anomalies were detected within a given timeframe. Specifically, a fault alarm was generated in our case study if the threshold was exceeded three times in a row, as shown by the shaded areas in Figure 7. No alarm was triggered for the components associated with S1–S6. Our findings indicate unusual and persistent spectral changes likely to result from fault-affected vibration responses of the monitored generator. We confirmed this finding by investigating the logs of the affected WT2. The logs specified that WT2 had suffered incipient generator damage without further detailing the type of damage. Thus, we could confirm that our proposed fault detection method is sufficiently sensitive to detect incipient generator damage from accelerometer measurements in commercial wind turbines. Note that we arrived at our diagnosis without any feature engineering but, rather,

by letting the autoencoder itself learn what spectrograms look like in the normal operation of healthy components.

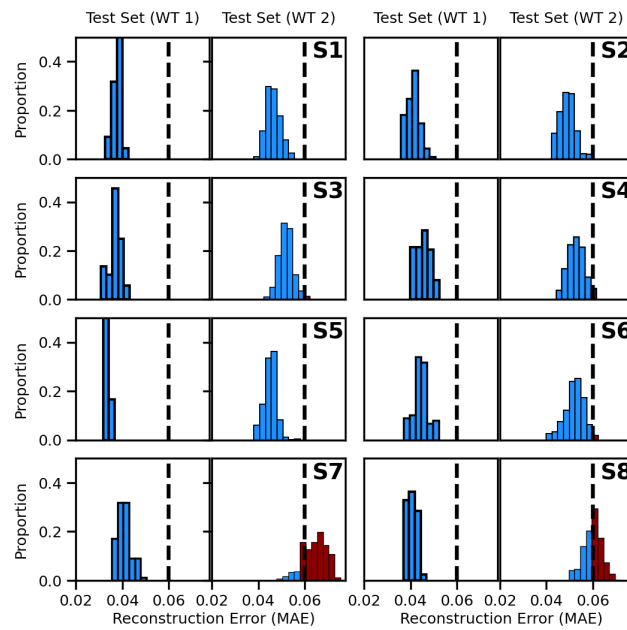


Figure 6. Fault detection based on spectrogram reconstruction. Reconstruction errors for accelerometers S1–S8 after reconstruction with a convolutional autoencoder. The spectrograms from the test set of WT1 are specified as healthy based on their comparatively low reconstruction errors, which is consistent with the logs of WT1. In contrast, the spectrograms of accelerometers S7 and S8 of WT2 indicate damage conditions in the generator. This finding was confirmed by the logs of WT2.




Figure 7. Fault detection based on spectrogram reconstruction. The reconstruction errors are considered health indices of the monitored drivetrain components of WT2. The red dashed line represents our threshold, set at 0.06, the maximal reconstruction error before a segment is classified as anomalous. Red shaded areas represent periods when the alarm notification criterion is active, following three consecutive crossings of the threshold.

Fault detection based on isolation forests. We applied the previously trained autoencoders for isolation-forest-based fault detection by using only their encoder parts. Each encoder part outputs a compressed feature representation in a vector of size 128 (the bottleneck layer output) which was then input into the isolation forests. For each accelerometer, we trained an isolation forest (IF) using the extracted feature vectors of spectrogram segments from the training set.

Each isolation forest was evaluated on unseen spectrograms from the test set of healthy WT1 and with spectrograms from WT2, as shown in Figure 8. The isolation forest models also provided an anomaly score (health index) for each spectrogram segment. Across all eight evaluated models, the isolation forests have consistently and correctly classified the WT1 test sets as healthy, just like the reconstruction-error-based method, as shown in Figures 6 and 8. When evaluating the spectrograms from WT2, the isolation forest models assigned elevated anomaly scores to the generator (accelerometers S7, S8 in Figures 8 and 9). Figure 9 displays the health index for all monitored components. As shown in Figures 8 and 9, fault detection based on isolation forests resulted in larger anomaly scores for the spectrograms derived from the accelerometers S7 and S8. This indicates a fault in the generator, which is consistent with the results of our reconstruction-error-based fault detection method.

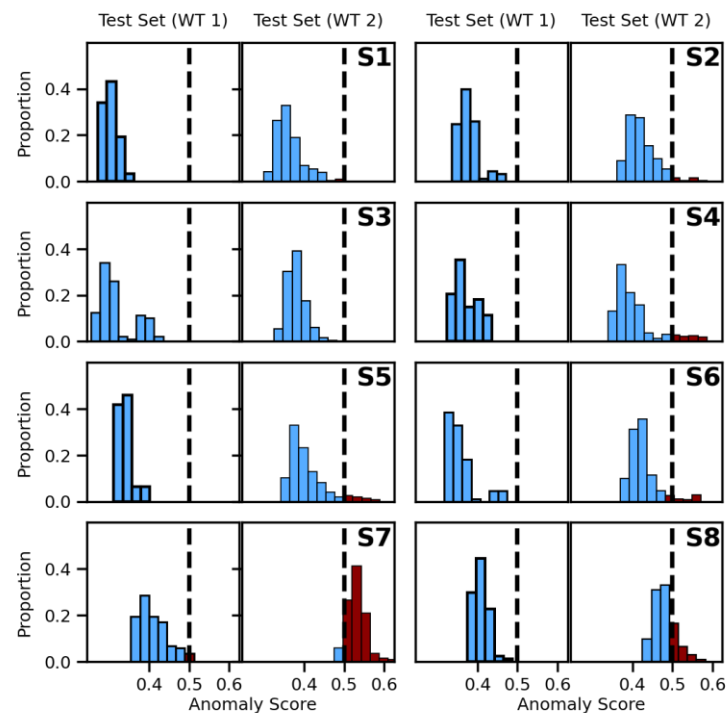


Figure 8. Fault detection based on isolation forest models. Visualization of the isolation forest anomaly scores for accelerometers S1–S8. In all cases, the vibration responses from WT1 were correctly classified as healthy based on their low anomaly score. Significantly larger anomaly scores were obtained for accelerometers S7 and S8 of WT2. This indicates a fault in the generator, which is consistent with the results of our reconstruction-error-based fault detection method (Figures 6 and 7).

We applied a rule to trigger a fault alarm if three consecutive anomalies occurred, as indicated by the shaded areas in Figure 9. No alarm was triggered for the components associated with S1–S6. A small number of spectrograms from other components exceeded the threshold of 0.5 according to the isolation forest models, as shown by sensors S2 and S4–S6 in Figure 8, but these did not result in persistently abnormal health scores, as shown in Figure 9. On the other hand, multiple persistent fault alarms were triggered for the generator coupling side towards the gearbox (S7, S8).




Figure 9. Fault detection based on isolation forest models. Health indices for the drivetrain components monitored with accelerometers S1–S8 of WT2. The red dashed line at 0.5 is used to separate normal vibration responses (health index $\lesssim 0.5$) from anomalous ones (health index > 0.5). We set the threshold value of 0.5 in accordance with [34]. Red shaded areas represent periods when the alarm notification criterion is active, following three consecutive crossings of the threshold.

After detecting the fault-affected vibration responses from sensors S7 and S8, a consultation of the WT2 operation logs confirmed that the generator coupling side towards the gearbox had suffered incipient damage. Thus, both presented fault detection methods successfully detected a fault in the generator from the spectral features learnt by the convolutional autoencoder.

Our study makes use of multiple model parameter values, such as the bottleneck size of the autoencoder, the number of trees in the isolation forests, or the number of maximal consecutive days with anomalies in the health index. We optimized the parameter values related to the autoencoder architecture and training, specifically the number of layers, their sizes, and learning parameters, based on a hyperparameter optimization approach as described above. Optimizing these values is possible because the autoencoder training can be considered a supervised task, in which the input spectrograms are taken to be the target outputs. Therefore, a loss value can be optimized by the choice of the autoencoder architecture and network parameters in the training process. On the other hand, the fault detection methods based on the spectrogram reconstruction error and based on the isolation forest perform unsupervised tasks because of the absence of a comprehensive set of labelled fault observations.

The optimal choice of parameters in the fault detection models is not within the scope of this study. Specifically, this includes the optimal number of trees in the isolation forest, the optimal contamination parameter in the isolation forest model, and the optimal reconstruction error threshold. We propose future studies to investigate in more detail the optimal choice of the parameter values based on additional fault observation datasets.

We also investigated the effectiveness of the presented fault detection methods for detecting known gearbox damages in one of the two 750 kW turbines from NREL (Section 3.2). Applying the same procedures as outlined for the commercial wind turbines above, we

found that both fault detection methods were successful in detecting the fault-affected components in the damaged NREL wind turbine from their vibration responses. Our results from the NREL wind turbines confirm the effectiveness of both the reconstruction-based and the isolation-forest-based fault detection approach.

5. Conclusions

Wind energy continues to expand strongly in Europe and around the world. The operation and maintenance costs of wind turbines account for a major fraction of the levelized cost of energy. Condition monitoring and artificial intelligence constitute powerful means for automating the early detection of incipient damage in wind turbines under various operating conditions. Machine learning methods enable early notification of wind turbine operators based on the vibration responses of the monitored components.

However, the existing vibration-based fault detection methods rely on the upfront definition of features in the frequency or time domains. This study has introduced two new fault detection methods for vibration-monitored parts that do not require any feature engineering. The proposed methods make use of convolutional autoencoders to detect unusual operation behaviour from the spectrograms of the monitored parts. The autoencoders learn and extract spectral features customized to the monitored components in an autonomous manner, without requiring any human assistance. In doing so, the autoencoders learn a spectral model of the component's normal behaviour from past accelerometer measurements.

We have demonstrated the new fault detection approaches—based on reconstruction errors and based on isolation forests—in four wind turbine drivetrains. We showed that both methods can successfully distinguish damaged from healthy vibration-monitored parts. First, we demonstrated their performances in detecting incipient generator damage from the vibration responses of the generators in two multi-MW onshore wind turbines. In addition, we confirmed the effectiveness of our presented fault detection methods in test rig measurements from NREL, where they successfully detected gearbox damage.

Comparing convolutional and dense autoencoders for feature extraction and reconstruction, we found that convolutional autoencoders can accomplish the spectrogram reconstructions with a drastically lower number of parameters as compared to dense autoencoders. Thus, the proposed convolutional autoencoders avoid overfitting and can generalize better to unseen data.

Importantly, both presented fault detection methods do not require any feature engineering. We discussed how convolutional autoencoders can autonomously extract the most relevant features from a continuous range of the spectrum. We demonstrated this for the range of $[0, 1000]$ Hz, because many characteristic frequencies fall in this range. In principle, however, the proposed method can also be applied to the full half spectrum. The autoencoders learn the normal vibration responses without requiring any upfront definition of features, thereby saving time and effort. An additional advantage of the presented methods is that a broad continuous range of the spectrum can be monitored instead of the usual focus on individual frequencies and harmonics.

The obtained results are very promising. While we demonstrated the new fault detection approach by means of gearbox and generator vibration responses, it can in principle also be applied to the structural health monitoring of other subsystems such as at the wind turbine tower. Further studies are needed for applications beyond the drivetrain to investigate the effect of variable operation conditions and how to account for them in the preprocessing and feature extraction in applications beyond the drivetrain. In the present study, we account for the effects of variable operation conditions by investigating the vibration responses at constant generator and drivetrain loads. The requirement to measure the vibration responses at constant loads can be easily accomplished in practical applications. For example, time-slice vibration measurements can be triggered whenever a specified load condition is met, e.g., by triggering the acquisition system at a certain rotational speed of the generator.

In addition to applications in other subsystems, future research should also investigate the proposed approach with more comprehensive datasets of vibration responses from damaged drivetrain components. It would be worthwhile to apply and investigate it for different damage types and intensities. It will also be interesting to study the temporal development of the proposed health indices in view of progressively increasing damage.

Author Contributions: Conceptualization, A.M.; Data curation, D.A.; Formal analysis, S.J.; Funding acquisition, B.B. and A.M.; Investigation, D.A.; Methodology, S.J. and A.M.; Project administration, B.B.; Resources, D.A.; Software, S.J.; Supervision, A.M.; Validation, S.J.; Visualization, S.J.; Writing—original draft, S.J. and A.M.; Writing—review and editing, D.A. and B.B. All authors have read and agreed to the published version of the manuscript.

Funding: The authors gratefully acknowledge that the present study was supported by a grant from the Swiss innovation agency, Innosuisse.

Data Availability Statement: The authors do not have permission to share the data used in this research.

Acknowledgments: The authors wish to thank Shawn Sheng from the National Renewable Energy Laboratory for providing test bench data from gearbox accelerometer measurements analysed in this study.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. International Energy Agency. *Global Energy Review 2021*; International Energy Agency: Paris, France, 2021.
2. International Energy Agency. *Renewables 2021*; International Energy Agency: Paris, France, 2021.
3. Tavner, P.J.; Xiang, J.; Spinato, F. Reliability Analysis for Wind Turbines. *Wind Energy* **2007**, *10*, 1–18. [[CrossRef](#)]
4. Faulstich, S.; Hahn, B.; Tavner, P.J. Wind Turbine Downtime and Its Importance for Offshore Deployment. *Wind Energy* **2011**, *14*, 327–337. [[CrossRef](#)]
5. Crabtree, C.J.; Zappalá, D.; Hogg, S.I. Wind Energy: UK Experiences and Offshore Operational Challenges. *Proc. Inst. Mech. Eng. Part J. Power Energy* **2015**, *229*, 727–746. [[CrossRef](#)]
6. Carroll, J.; McDonald, A.; McMillan, D. Failure Rate, Repair Time and Unscheduled O&M Cost Analysis of Offshore Wind Turbines: Reliability and Maintenance of Offshore Wind Turbines. *Wind Energy* **2016**, *19*, 1107–1119. [[CrossRef](#)]
7. Teng, W.; Ding, X.; Zhang, X.; Liu, Y.; Ma, Z. Multi-Fault Detection and Failure Analysis of Wind Turbine Gearbox Using Complex Wavelet Transform. *Renew. Energy* **2016**, *93*, 591–598. [[CrossRef](#)]
8. Feng, Z.; Qin, S.; Liang, M. Time–Frequency Analysis Based on Vold-Kalman Filter and Higher Order Energy Separation for Fault Diagnosis of Wind Turbine Planetary Gearbox under Nonstationary Conditions. *Renew. Energy* **2016**, *85*, 45–56. [[CrossRef](#)]
9. Yang, B.; Liu, R.; Chen, X. Sparse Time-Frequency Representation for Incipient Fault Diagnosis of Wind Turbine Drive Train. *IEEE Trans. Instrum. Meas.* **2018**, *67*, 2616–2627. [[CrossRef](#)]
10. Sun, R.; Yang, Z.; Chen, X.; Tian, S.; Xie, Y. Gear Fault Diagnosis Based on the Structured Sparsity Time-Frequency Analysis. *Mech. Syst. Signal Process.* **2018**, *102*, 346–363. [[CrossRef](#)]
11. Teng, W.; Ding, X.; Cheng, H.; Han, C.; Liu, Y.; Mu, H. Compound Faults Diagnosis and Analysis for a Wind Turbine Gearbox via a Novel Vibration Model and Empirical Wavelet Transform. *Renew. Energy* **2019**, *136*, 393–402. [[CrossRef](#)]
12. Abouel-seoud, S.A. Fault Detection Enhancement in Wind Turbine Planetary Gearbox via Stationary Vibration Waveform Data. *J. Low Freq. Noise Vib. Act. Control* **2018**, *37*, 477–494. [[CrossRef](#)]
13. Escaler, X.; Mebarki, T. Full-Scale Wind Turbine Vibration Signature Analysis. *Machines* **2018**, *6*, 63. [[CrossRef](#)]
14. Zimroz, R.; Bartelmus, W.; Barszcz, T.; Urbanek, J. Diagnostics of Bearings in Presence of Strong Operating Conditions Non-Stationarity—A Procedure of Load-Dependent Features Processing with Application to Wind Turbine Bearings. *Mech. Syst. Signal Process.* **2014**, *46*, 16–27. [[CrossRef](#)]
15. Tang, Z.; Wang, M.; Ouyang, T.; Che, F. A Wind Turbine Bearing Fault Diagnosis Method Based on Fused Depth Features in Time–Frequency Domain. *Energy Rep.* **2022**, *8*, 12727–12739. [[CrossRef](#)]
16. Mengjiao, W.; Zhenhao, T.; Bo, Z.; Yunfeng, H. Wind Turbine Bearing Fault Diagnosis Method Based on Multi-Domain Feature Extraction. In Proceedings of the 2022 4th International Conference on Power and Energy Technology (ICPET), Beijing, China, 27–30 July 2022; pp. 413–418.
17. Salameh, J.P.; Cauet, S.; Etien, E.; Sakout, A.; Rambault, L. Gearbox Condition Monitoring in Wind Turbines: A Review. *Mech. Syst. Signal Process.* **2018**, *111*, 251–264. [[CrossRef](#)]
18. Barszcz, T. *Vibration-Based Condition Monitoring of Wind Turbines*; Applied Condition Monitoring; Springer International Publishing: Cham, Switzerland, 2019; Volume 14, ISBN 978-3-030-05969-9.

19. Koukoura, S.; Carroll, J.; McDonald, A.; Weiss, S. Comparison of Wind Turbine Gearbox Vibration Analysis Algorithms Based on Feature Extraction and Classification. *IET Renew. Power Gener.* **2019**, *13*, 2549–2557. [[CrossRef](#)]
20. Li, Y.; Chen, P.; Wang, K.; Miao, Q. A Deep Convolutional Autoencoder for Wind Turbine Blades Health Condition Monitoring Based on FBG Strain Gauges Measurement. In Proceedings of the 7th Asia-Pacific Workshop on Structural Health Monitoring, APWSHM 2018, Hong Kong, China, 12–15 November 2018; pp. 973–981.
21. Yang, L.; Zhang, Z. A Conditional Convolutional Autoencoder-Based Method for Monitoring Wind Turbine Blade Breakages. *IEEE Trans. Ind. Inform.* **2021**, *17*, 6390–6398. [[CrossRef](#)]
22. Li, Y.; Jiang, W.; Zhang, G.; Shu, L. Wind Turbine Fault Diagnosis Based on Transfer Learning and Convolutional Autoencoder with Small-Scale Data. *Renew. Energy* **2021**, *171*, 103–115. [[CrossRef](#)]
23. Zaher, A.; McArthur, S.D.J.; Infield, D.G.; Patel, Y. Online Wind Turbine Fault Detection through Automated SCADA Data Analysis. *Wind Energy* **2009**, *12*, 574–593. [[CrossRef](#)]
24. Schlechtingen, M.; Santos, I.F.; Achiche, S. Wind Turbine Condition Monitoring Based on SCADA Data Using Normal Behavior Models. Part 1: System Description. *Appl. Soft Comput.* **2013**, *13*, 259–270. [[CrossRef](#)]
25. Tautz-Weinert, J.; Watson, S.J. Using SCADA Data for Wind Turbine Condition Monitoring—A Review. *IET Renew. Power Gener.* **2017**, *11*, 382–394. [[CrossRef](#)]
26. Meyer, A.; Brodbeck, B. Data-Driven Performance Fault Detection in Commercial Wind Turbines. In Proceedings of the PHM Society European Conference, Turin, Italy, 1–3 July 2020; Volume 5, p. 7.
27. Meyer, A. Multi-Target Normal Behaviour Models for Wind Farm Condition Monitoring. *Appl. Energy* **2021**, *300*, 117342. [[CrossRef](#)]
28. Bilendo, F.; Meyer, A.; Badihi, H.; Lu, N.; Cambron, P.; Jiang, B. Applications and Modeling Techniques of Wind Turbine Power Curve for Wind Farms—A Review. *Energies* **2023**, *16*, 180. [[CrossRef](#)]
29. Bourlard, H.; Kamp, Y. Auto-Association by Multilayer Perceptrons and Singular Value Decomposition. *Biol. Cybern.* **1988**, *59*, 291–294. [[CrossRef](#)] [[PubMed](#)]
30. Kramer, M.A. Nonlinear Principal Component Analysis Using Autoassociative Neural Networks. *AIChE J.* **1991**, *37*, 233–243. [[CrossRef](#)]
31. Goodfellow, I.; Bengio, Y.; Courville, A. *Deep Learning*; Adaptive Computation and Machine Learning; The MIT Press: Cambridge, MA, USA, 2016; ISBN 978-0-262-03561-3.
32. Ribeiro, M.; Lazzaretti, A.E.; Lopes, H.S. A Study of Deep Convolutional Auto-Encoders for Anomaly Detection in Videos. *Pattern Recognit. Lett.* **2018**, *105*, 13–22. [[CrossRef](#)]
33. Duman, E.; Erdem, O.A. Anomaly Detection in Videos Using Optical Flow and Convolutional Autoencoder. *IEEE Access* **2019**, *7*, 183914–183923. [[CrossRef](#)]
34. Liu, F.T.; Ting, K.M.; Zhou, Z.-H. Isolation Forest. In Proceedings of the 2008 Eighth IEEE International Conference on Data Mining, Pisa, Italy, 15–19 December 2008; pp. 413–422.
35. Meyer, A. Vibration Fault Diagnosis in Wind Turbines Based on Automated Feature Learning. *Energies* **2022**, *15*, 1514. [[CrossRef](#)]
36. Sheng, S. *Wind Turbine Gearbox Condition Monitoring Round Robin Study-Vibration Analysis*; National Renewable Energy Lab. (NREL): Golden, CO, USA, 2012.
37. Musial, W.; McNiff, B. *Wind Turbine Testing in the NREL Dynamometer Test Bed*; National Renewable Energy Lab. (NREL): Golden, CO, USA, 2000.
38. Allen, J. Short Term Spectral Analysis, Synthesis, and Modification by Discrete Fourier Transform. *IEEE Trans. Acoust. Speech Signal Process.* **1977**, *25*, 235–238. [[CrossRef](#)]
39. Li, L.; Jamieson, K.; DeSalvo, G.; Rostamizadeh, A.; Talwalkar, A. Hyperband: A Novel Bandit-Based Approach to Hyperparameter Optimization. *J. Mach. Learn. Res.* **2017**, *18*, 6765–6816.
40. Kingma, D.P.; Ba, J. Adam: A Method for Stochastic Optimization. ICLR. *arXiv* **2015**, arXiv:1412.6980.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.