

Research Article

Dictionary Learning Approach to Monitoring of Wind Turbine Drivetrain Bearings

Sergio Martin-del-Campo^{1,*}, Fredrik Sandin¹, Daniel Strömbergsson²

¹EISLAB, Luleå University of Technology (LTU), Universitetsvägen 1, Luleå, 97187, Sweden

²Division of Machine Elements, Luleå University of Technology (LTU), Universitetsvägen 1, Luleå, 97187, Sweden

ARTICLE INFO

Article History

Received 03 May 2020

Accepted 01 Nov 2020

Keywords

Wind turbine

Condition monitoring

Dictionary learning

Unsupervised learning

Feature extraction

Bearings

ABSTRACT

Condition monitoring is central to the efficient operation of wind farms due to the challenging operating conditions, rapid technology development, and a large number of aging wind turbines. In particular, predictive maintenance planning requires the early detection of faults with few false positives. Achieving this type of detection is a challenging problem due to the complex and weak signatures of some faults, particularly the faults that occur on the gearbox bearings of a turbine drivetrain. The results of former studies addressing condition-monitoring tasks using dictionary learning indicate that unsupervised feature learning is useful for diagnosis and anomaly detection purposes. However, these studies are based on small sets of labeled data from test rigs operating under controlled conditions that focus on classification tasks, which are useful for quantitative method comparisons but gives little insight into how useful these approaches are in practice or how can be used by existing condition-monitoring systems. Here, we investigate an unsupervised dictionary learning method for condition monitoring using vibration data recorded over 46 months under typical industrial operations. Thus, we contribute real-world industrial vibration data that are made publicly available and novel test results. In this study, dictionaries are learned from gearbox vibrations in six different turbines, and the dictionaries are subsequently propagated over a few years of monitoring data when faults are known to occur. We perform the experiment using two different sparse coding algorithms to investigate if the algorithm selected affects the features of abnormal conditions. We propose a dictionary distance metric derived from the dictionary learning process as a condition indicator and find the time periods of abnormal dictionary adaptation starting six months before a drivetrain bearing replacement and one year before the resulting gearbox replacement. In addition, we investigate the distance between dictionaries learned from geographically close turbines of the same type under healthy conditions. We find that the features learned are similar and that a dictionary learned from one turbine can be useful for monitoring a similar turbine.

© 2021 The Authors. Published by Atlantis Press B.V.

This is an open access article distributed under the CC BY-NC 4.0 license (<http://creativecommons.org/licenses/by-nc/4.0/>).

1. INTRODUCTION

Wind power is a renewable energy source that is growing rapidly and provides more than 11% of the electrical power in the European Union [1]. Wind power is harvested by wind farms, which typically include many similar wind turbines. Wind turbines are based on relatively new technology that has been scaled up from approximately 2MW to 10MW per turbine in one decade. The rapid development in combination with the challenging operating conditions of wind turbines over the typical 20-year service life implies that condition monitoring and predictive maintenance are central issues. When maintenance is needed, the cost of crane mobilization and energy production losses are high, and there are challenges acquiring spare parts in this rapidly expanding industry. The gearbox is a major component of a wind turbine, and the rolling element bearings that support the rotating components in the drivetrain are essential for reliable operation. Monitoring these bearings is an important and challenging issue given the predominance of bearing faults in wind turbines [2] and the complex and weak signatures of some faults

[3]. Because nearby turbines face similar environmental and operational conditions, methods can be adapted and validated with data from multiple machines.

Condition-based maintenance requires continuous monitoring of the machine to detect incipient failures so that the maintenance actions can be scheduled efficiently [4]. This procedure involves three stages: data acquisition, feature extraction, and diagnostics identification. The principle behind identification is that a “significant change [of a feature] is indicative of a developing failure” [5]. Feature selection and extraction is a key problem that typically determines the performance of decision support functions and thereby the efficiency of the condition-monitoring system. A feature is an individual measurable signal property or pattern that is characteristic of some particular type of source. The condition monitoring of wind turbines typically uses methods based on feature extraction with signals originating from vibration sensors mounted on the drivetrain. Hossain [4] describes common faults in wind turbines and the typical features used in their identification.

*Corresponding authors. Email: sergio.martindelcampo@ltu.se; sergio.mdc@gmail.com

The features used in the diagnostics of a wind turbine can be classified into three categories: time-domain, frequency-domain, and joint time-frequency domain. Time-domain features include traditional statistical parameters such as the root mean square (RMS), crest factor, and kurtosis. The trends of these parameters are used as indicators of the deterioration of the machine [6]. Frequency-domain features are typically derived by the conversion of the time-domain signal to the frequency domain using the fast Fourier transform (FFT). Kinematic data such as bearing defect frequencies are used to extract information in selected frequency bands [7], and thereby form a smaller set of features that can be trended and monitored. Analysis methods based on joint time-frequency domain features are also used, but these are more recent developments compared to the time-domain and frequency-domain methods. The wavelet transform is one example, which is useful for the analysis of nonstationary signals [5]. However, to date, such methods have not been widely used in industry because the analysis is more complex and requires trained experts to interpret the results [8]. Further information about these analysis methods can be found in, for example [9,10].

Features are typically manually selected by experts, which implies that the features are selected without explicit knowledge of the state of each machine. Furthermore, the dependence on experts is a bottleneck that limits the scalability of condition monitoring systems. In general, the detection, prediction, and diagnosis of faults in a rolling element bearing are challenging tasks due to the high number of variables that affect the operation. Thus, a machine learning approach can be useful in the development of more automated diagnosis and prognosis systems. Supervised machine learning is one approach, but it requires labeled data for training, which are difficult and expensive to generate [11]. An alternate approach is unsupervised learning methods, which, for example, can be used for feature learning and anomaly detection purposes. An unsupervised learning approach can assist in the analysis of the dark data, which is data acquired by sensors but it is not used for inference or decision-making process.

Here, we investigate an online feature learning approach based on dictionary learning that enables the optimization of the monitored feature set for each machine. In particular, we use dictionary learning to study signals recorded from vibration sensors installed on gearboxes in 2.5 MW turbines at a wind farm in northern Sweden. The learned features define a set of overcomplete and shift-invariant waveforms that are used to determine a sparse approximation of the corresponding vibration signal. We are interested in measures derived from the learning process that can be used to track the changes of such waveforms over time, and as a result, they can be used as key performance indicators in a condition monitoring process for the detection of abnormal changes.

Dictionary learning [12] and convolutional sparse coding [13] has attracted broad interest. Variations of the dictionary learning method have successfully been used in tasks such as signal compression, detection, separation, and denoising [14–16]. The methods developed here are based on the work by Smith and Lewicki [17,18], which is inspired by the earlier work of Olshausen and Field [19,20] in the area of sparse visual coding. The methodology includes a sparse regularization mechanism that reduces the influence of noise and some of the redundancy that is typically present in raw sensor signals. Here, the hypothesis is that the same general approach can

be used to characterize and analyze the signals generated by a rotating machine [21].

Liu *et al.* [22] were the first to apply dictionary learning to a dataset with bearing vibration signals. These authors trained dictionaries of waveforms of fixed length for different bearing conditions. The learned dictionaries were subsequently merged and used to classify the type of fault with a linear classifier. Furthermore, Martin-del-Campo *et al.* [23] showed that it is possible to distinguish different operational conditions through the learning of shift-invariant waveforms where the lengths of the waveforms are also optimized. Chen *et al.* [24] use a dictionary learning approach to detect a fault in a gearbox by the identification of impulse-like components in a vibration signal. Tang *et al.* [25] used shift-invariant sparse coding to generate a set of latent components that act as fault filters in a bearing or a gearbox. Moreover, studies by Ahmed *et al.* [26] and He *et al.* [27] proposed classification strategies that use the learned sparse representations on stacked autoencoders and large memory storage and retrieval neural networks, respectively. Further extensions of the work by Liu *et al.* [22] had been developed by Wang *et al.* [28] and Zhou *et al.* [29], who used the same dictionary learning method with different classification strategies.

Studies of dictionary learning for fault detection with bearing signals are based on simulated data and/or data from controlled experiments, where the faults are artificially introduced. Furthermore, most of these studies investigate either how the learned atoms can be used to classify faults or how to improve the accuracy of such a classifier. Thus, these results are not compatible with current approaches to health monitoring of wind turbines. Here, we extend the former studies with an investigation based on real-world vibration data collected from the vibration sensors on gearboxes in multiple wind turbines over an extended period of nearly four years. The data have been released publicly; see appendix for further information. The output shaft bearing and, subsequently, the gearbox were replaced in one of the turbines considered in this study. In addition to considering the fault detection problem, we investigate whether the dictionary of waveforms learned for one turbine is useful for the analysis of the corresponding signal in a nearby turbine of the same type. Furthermore, we study the possibility of using dictionary propagation and dictionary-based indicators to identify bad actors in a population of wind turbines, in a similar way that trend analysis is currently used to monitor the conditions of turbines. The experiments presented below also include a comparison of two different sparse coding algorithms.

The dictionary learning method and the proposed dictionary-based indicators are described in Section 2. The data used and the experiment procedure are described in Section 3.1. The results are presented in Section 4, followed by a discussion of the results in Section 5, and the conclusion in Section 6.

2. UNSUPERVISED FEATURE LEARNING METHOD

The unsupervised feature learning method used in this work is known as dictionary learning. Using this method, a vibration signal is decomposed into a sparse representation using a set of learned features known as atoms. The algorithm used to create the sparse

representation is known as a sparse coding algorithm, while the set of atoms is learned with the dictionary learning algorithm.

This section first introduces the sparse signal model used to represent the vibration signals. Then, the algorithms used to encode a sparse represented are presented, which is followed by a description of the dictionary learning algorithm used in the optimization process to learn a dictionary of shift-invariant atoms. This section concludes with the introduction of our proposed dictionary distance metric.

2.1. Sparse Signal Model

The signal $S(t)$ is modeled as a linear superposition of Gaussian noise and waveforms with compact support

$$S(t) = \sum_{i=1}^N a_i \phi_{m(i)}(t - \tau_i) + \epsilon(t). \quad (1)$$

The functions $\phi_m(t)$ are atoms that are learned from the signal, which we also refer to as features. An atom is an elementary waveform that describes a recurring feature of the signal. A set of atoms $\phi_m(t)$ defines a dictionary Φ that consists of M atoms

$$\Phi = \{\phi_1, \dots, \phi_M\}. \quad (2)$$

The triple $m(i), \tau_i, a_i$ defines one atom instance. We refer to the ratio between the total number of atom instances N and the total number of signal (segment) samples as the sparsity level. The temporal position and amplitude of the i -th instance of atom $\phi_m(t)$ are denoted by τ_i and a_i , respectively. The term $\epsilon(t)$ represents the model residual, which includes Gaussian noise.

The inverse problem defined by Eq. (1) is solved with an iterative two-step optimization process for each consecutive signal segment:

1. *Sparse coding*—While maintaining a fixed dictionary, determine the parameters $m(i), \tau_i$ and a_i of the N atom instances in Eq. (1) using the Matching Pursuit (MP) [30] or Orthogonal Matching Pursuit (OMP) algorithms [31].
2. *Dictionary update*—Given the set of atom instances and the residual $\epsilon(t)$, update the atoms in the dictionary Φ using a probabilistic gradient method (described below).

Step 1 is a convolutional sparse coding [32] process repeated until a stopping condition is reached. Typically, the stopping condition is defined in terms of the total number of terms N of the approximation, as in this case; alternatively, this condition can be defined in terms of the approximation error or signal-to-residual ratio. By defining the stopping condition in terms of the number of atom instances, the sparsity of the model in Eq. (1) is directly related to the number of iterations of the optimization process; hence, the computational requirements for online operation are well defined. This step tends to be computationally demanding but the computational cost will remain the same for each consecutive segment while the optimization parameters are not modified. Furthermore, real-time processing of signals might be possible with a new generation of processors, such as the Arm Ethos-N57 [33]. The Arm Ethos-N57 possess a performance of up to 1 TOP/s, which is at least one order magnitude above the number of operations carried out on the

scheme proposed here. Step 2, the dictionary update, is performed iteratively after the sparse coding of each signal segment. The process continues until there are no more signal segments. In this way, the online learning of atoms by processing consecutive signal segments is possible [21].

2.2. Signal Encoding Algorithm

The model described by Eq. (1) describes a continuous signal $S(t)$ as a linear combination of atoms. However, the problem of finding the optimal linear combination of an overcomplete set of atoms is an intractable (NP hard) problem. Therefore, several algorithms have been proposed to find approximate solutions to this problem. One set of algorithms are greedy algorithms, which rely on an iterative process to create the sparse representation, for example, MP [30], OMP [34], and Gradient Pursuit [35].

Here, we use the MP algorithm and the OMP algorithm to obtain a sparse approximation of each signal segment. Both algorithms are used to decompose the signal, given a dictionary of atoms. The algorithms operate on the residual of the signal, which initially is the signal segment to be decomposed: $R_0(t) = s_k(t)$. The cross-correlation between the residual and all of the elements of Φ is calculated in each iteration. An atom instance is defined by the atom with the maximum cross-correlation (inner product) across all possible timeshifts, and the amplitude a_i is defined by

$$a_i = \langle s_k(t) | \phi_{m(i)}(t - \tau_i) \rangle, \quad (3)$$

where the temporal position τ_i is

$$\tau_i = \arg \max_i \langle s_k(t) | \phi_{m(i)}(t - \tau_i) \rangle. \quad (4)$$

This process is repeated by determining a new atom instance for each iteration until the stopping condition is fulfilled. In each iteration, the atom instance $a_i \phi_{m(i)}(t - \tau_i)$ is subtracted from the residual to form a new residual to be used in the next iteration.

The MP and OMP methods have different residual update rules. In MP, the updated residual of the signal $R_i(t)$ after the i -th iteration is given by

$$R_i(t) = R_{i-1}(t) - a_i \phi_{m(i)}(t - \tau_i). \quad (5)$$

The OMP algorithm updates all coefficients a_i with an orthogonal projection of the signal segment onto the set of all previously selected atoms. Thus, the updated residual of the signal, $R_i(t)$, after the i -th iteration is

$$R_i(t) = R_{i-1}(t) - \Phi_i (\Phi_i^T \Phi_i)^{-1} \Phi_i^T R_{i-1}(t), \quad (6)$$

where Φ_i is the updated dictionary in the i -th iteration.

The iterative process continues until the stop condition is reached. This stop condition is defined based on the number of terms N relative to the number of samples of each signal segment $S(t)$. Moreover, the number of terms also determines the sparsity of the approximation. For further details about the MP method used here, see Ref. [23]; for the OMP method, see Ref. [34].

2.3. Learning of Shift-Invariant Dictionary

The next step in the iterative optimization process is to update the atoms in the dictionary Φ using the sparse approximation of the signal. The goal is to optimize the set of atomic waveforms ϕ_m in the dictionary Φ to minimize the residual of the sparse approximation. One solution to this problem can be obtained by rewriting Eq. (1) in probabilistic form:

$$p(s_k|\Phi) = \int p(s_k|a, \Phi)p(a)da \approx p(s_k|\hat{a}, \Phi)p(\hat{a}), \quad (7)$$

where \hat{a} is the maximum a posteriori (MAP) estimate of a [17],

$$\hat{a} = \arg \max_a p(a|s_k, \Phi) = \arg \max_a p(s_k|a, \Phi)p(a), \quad (8)$$

that is generated by the MP or orthogonal MP algorithms [18]. Furthermore, we assume that the noise term ϵ in Eq. (1) is Gaussian. Therefore, the data likelihood $p(s_k|a, \Phi)$ is also Gaussian and takes the form

$$p(s_k|a, \Phi) \approx \exp\left(-\frac{\|s_k - a\Phi\|^2}{2\sigma_\epsilon^2}\right), \quad (9)$$

where

$$\|s_k - a\Phi\|^2 = \left\| s_k - \sum_{i=1}^N a_i \phi_{m(i)}(t - \tau_i) \right\|^2, \quad (10)$$

and σ_ϵ^2 is the variance of the noise. Note that s_k , a and Φ are matrices in these probabilistic expressions and that the dictionary Φ includes all of the possible shifts of each atom ϕ_m .

Under these assumptions, the atoms in the dictionary can be optimized by performing gradient ascent on the approximate log data probability, thus, resulting in a gradient of Eq. (7) of the form

$$\frac{\partial}{\partial \phi_m} \log(p(s_k|\Phi)) = \frac{1}{\sigma_\epsilon^2} \sum_i a_i [s_k - \hat{s}_k]_{\tau_i}. \quad (11)$$

The term $[s_k - \hat{s}_k]_{\tau_i}$ represents the final residual of the sparse approximation (when the stop condition is met) that coincides with each atom at its respective temporal position τ_i identified by MP or OMP. Consequently, the gradient of each atom in the dictionary is proportional to the sum of the residuals within the support of each atom instance.

The use of the gradient for dictionary learning requires a step length parameter η that determines how much the atoms are updated. The resulting update rule for an atom is

$$\phi_m \rightarrow \phi_m + \frac{\eta}{\sigma_\epsilon^2} \sum_i a_i [s_k - \hat{s}_k]_{\tau_i}. \quad (12)$$

Therefore, the learning rate depends on how often atoms are selected during the sparse coding step, which implies that the learning rate of atoms can be different and that some atoms may not learn at all (see [36] for an alternative dictionary learning method where this is not the case). Furthermore, we zero-pad all atoms with ten elements at each tail and allow an atom to grow in length if the RMS of the tail exceeds 0.1 of the atom RMS, as described in [18].

Figure 1 presents an online monitoring scheme based on dictionary learning. The signal is divided into segments of equal length. The interval between the processed segments can be adapted to match either the processing capacity of the condition monitoring

system or the availability of the data communicated from the turbine (which is the case considered here). In the latter case, the interval between the segments can be up to hours or days due to the limitations of the communication network between the wind farm and the condition monitoring center. In an online processing scenario, the edge effects due to signal segmentation can be reduced by transferring the tail of the residual to the next segment to be processed. Martin-del-Campo *et al.* [37] describe this method for processing continuous signals.

The initial dictionary is either pseudorandomly generated (training stage) or copied from a repository that includes dictionaries learned from similar machines (monitoring stage) according to an experimental protocol such as the one defined in Section 3.2. Initially, the first segment is processed with the sparse coding algorithm (MP or OMP), and the resulting sparse representation is used to update the dictionary. Subsequently, the updated dictionary is used to process the next signal segment, which is a process known as dictionary propagation. The output of this process is the updated dictionary, the residual, and the coefficients and offsets of the selected atoms that define the sparse approximation of the signal. These parameters are used as features for monitoring the corresponding wind turbine.

Note that dictionary learning can be deactivated by setting the learning rate parameter η to zero. In that case, the dictionary is constant over time. However, the sparse representation of the signal can still be generated and analyzed for condition monitoring purposes.

2.4. Dictionary Distance

The dictionary is updated when each signal segment is processed (provided that the learning rate is nonzero). Therefore, it is interesting to quantify and monitor the difference between two dictionaries learned at different points in time, for example, by comparing the present dictionary to a baseline dictionary learned during a period when the machine was known to be in a healthy condition. Skretting and Engan [38] define the distance between two dictionaries Φ' and Φ as

$$\beta(\Phi, \Phi') = \frac{1}{2M} \left(\sum_{i=1}^M \beta(\Phi', \phi_i) + \sum_{j=1}^M \beta(\Phi, \phi'_j) \right), \quad (13)$$

where the two dictionaries have the same number of atoms, M , with memberships $\phi \in \Phi$ and $\phi' \in \Phi'$. The function β is the maximum similarity of an atom ϕ to the atoms in dictionary Φ ,

$$\beta(\Phi, \phi) = \arccos \mu(\Phi, \phi), \quad (14)$$

where μ is the mutual coherence, which is defined as follows [14] (here, the function μ should be a function of both Φ and ϕ according to how it is used in [14]):

$$\mu(\Phi, \phi_i) = \max_{\substack{\forall \phi_j \in \Phi \\ j \neq i}} \frac{|\phi_i^T \phi_j|}{\|\phi_i\|_2 \cdot \|\phi_j\|_2}. \quad (15)$$

The dictionary distance is measured in degrees and conceptually is a generalization of the conventional ‘‘cosine of angle’’ measure of similarity. For example, when $\beta(\Phi, \Phi') = 0$ the two dictionaries are equal.

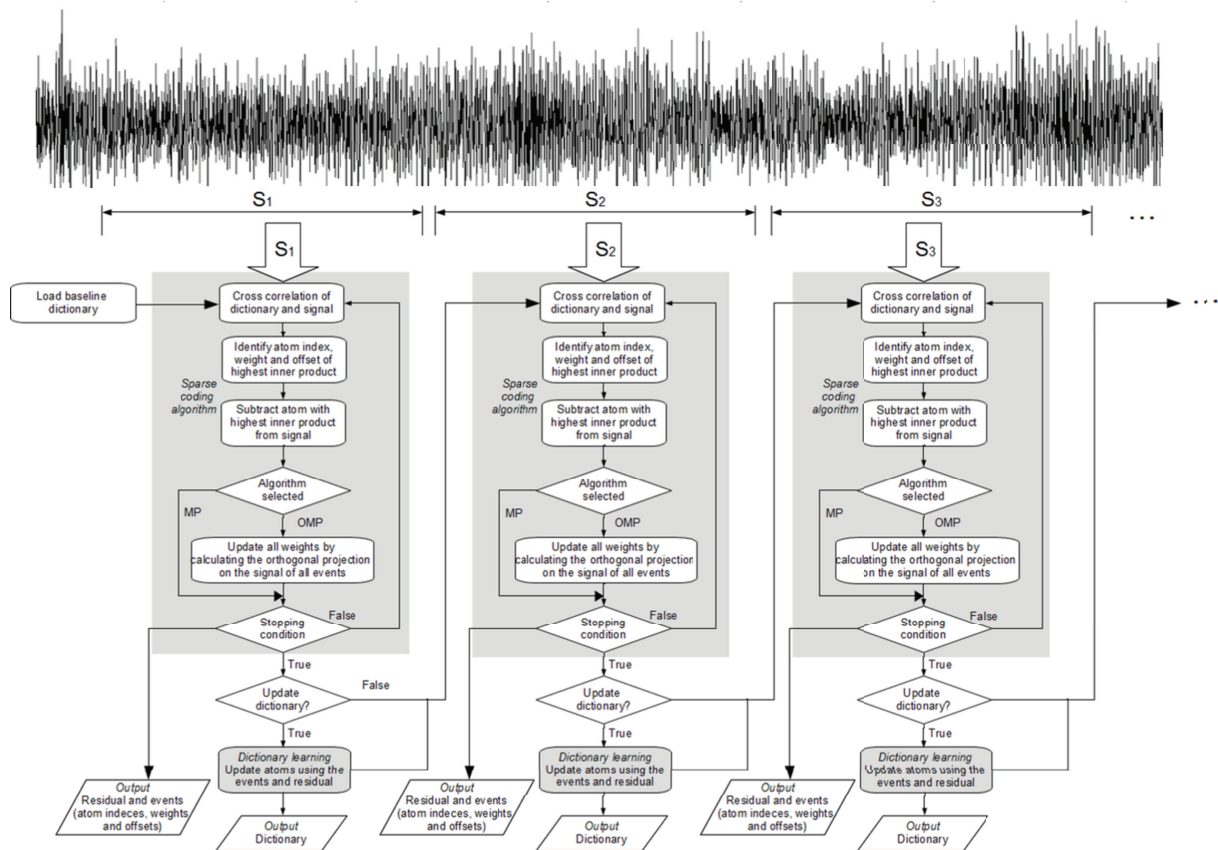


Figure 1 | Dictionary learning scheme for online condition monitoring.

The dictionary distance measure can be used to quantify the distance between one learned dictionary at two different points in time. We refer to this measure as the adaptation rate (of the dictionary) and define it as $\beta(\Phi^t, \Phi^{t-\delta})$, where t is the current time and δ refers to the same dictionary at some point in the past. The idea is that the adaptation rate can be used to quantify sudden abnormal changes in the signal, for example, due to a fault in the system. See also [21] where the rate of change of individual atoms after the introduction of a fault is investigated. In principle, the dictionary distance $\beta(\Phi^t, \Phi^{t-\delta})$ could be normalized with the time step δ to obtain a finite difference approximation of the “dictionary derivative” with respect to time.

3. METHODOLOGY

Firstly, we introduce the real-world bearing vibration data used on this work and which we are making publicly available. Afterward, we provide a description of the procedure followed during our numerical experiments.

3.1. Data Source

We aim to study the viability of a dictionary learning approach to condition monitoring using real-world data. As a result, we have no control over the operational and environmental conditions, which is in contrast to former studies based on data from controlled experiments. The data originate from a wind farm located in northern

Sweden. The wind turbines are the same model and have integrated condition monitoring systems that transmit data to a condition monitoring database; from this database, we can access the vibration data used in this study. Each wind turbine possesses a three-stage gearbox, including two sequential planetary gear stages followed by a helical gear stage. Each gearbox has four accelerometers located near the different gear stages. Figure 2 includes a schematic view of the gearbox and the locations of the accelerometers.

Raw time-domain vibration signals from six turbines within the same wind farm are considered in this study. All of the measurement data corresponds to the axial direction of the accelerometer, which is mounted on the housing of the output shaft bearing of each turbine. The sampling rate is 12.8 kHz, and each signal segment is 1.28 seconds long (16384 samples). The signal segments were recorded with an interval of approximately 12 hours over a period of 46 consecutive months in the last decade. In this period of time, five turbines remained healthy, which will henceforth be referred to as Turbine 1, Turbine 2, Turbine 3, Turbine 4, and Turbine 6. The other turbine, which we refer to as Turbine 5, had two bearing failures in this period. The locations of the defective bearings are highlighted in Figure 2. The descriptions of the failures are as follows:

- i. Inner raceway failure on a four-point ball bearing on the output shaft. The output shaft bearing was replaced after 1.2 years in operation.

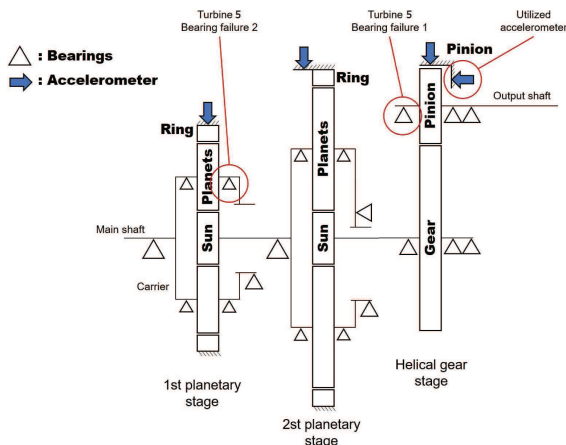


Figure 2 | Schematic view of the gearbox in a wind turbine. The components of each stage are shown, including the support bearings. Data from one wind turbine with two bearing failures are included in this study. The locations of the faulty bearings are highlighted in the figure. The measurement axes of the accelerometers are marked by arrows.

- ii. Inner raceway failure on one of the four cylindrical roller bearings supporting one of the planets in the first planetary gear. The entire gearbox was replaced after 2 years in operation. Figure 3 shows the result of this failure.

The dataset containing the raw time-domain vibration signals and speed measurements from the six turbines is publicly available; see [Appendix](#) for further information.

3.2. Outline of Numerical Experiments

In addition to investigating the bearing failures in Turbine 5 described above, we are interested in the similarities of the dictionaries learned from different wind turbines of the same type that are located in the same geographical area. Can the dictionary learned from one turbine be similar to the dictionary learned from a similar turbine that is subject to similar operational and environmental conditions? Furthermore, is a dictionary learned from one healthy turbine useful for monitoring of a similar turbine? To address these questions, we analyze the vibration signals described in Section 3.1 using the dictionary learning method described in Section 2. We process the data with our MATLAB/C++ implementations of MP and local OMP [34] as well as Smith and Lewicki’s dictionary learning algorithm [18]. There are two main stages in the experimental protocol we used: (1) learning of a baseline dictionary for each turbine with data from a period of nonfaulty operation (training stage) and (2) updates of each turbine dictionary using the successive recorded signal segments (monitoring stage). For each turbine, the baseline dictionary defines the first dictionary used in the monitoring stage.

3.2.1. Training stage

In the training stage, the aim is to learn a baseline dictionary for each turbine that corresponds to the signal recorded under

healthy/nonfaulty operational conditions. We use the signal segments recorded in the time period comprising approximately the period after the second year of operation of the turbines. Because we selected the period of time after the replacement of the gearbox in Turbine 5, this turbine also operated in a healthy condition during that time period. We use the same training period for the six turbines to ensure similar operational and environmental conditions during training. Signal segments with a vibration RMS above 0.5 G are included in the training process, while segments with a lower RMS are omitted from the analysis presented here. Figure 4 shows the RMS values of the signal segments *versus* the rotational speed. Below an RMS of 0.5 G, the turbines are sometimes unloaded and the corresponding signal segments are noisier (possibly due to the reduced load applied to the bearings). Thus, we introduce a threshold on the RMS to exclude signal segments recorded when the turbines/gearboxes are unloaded. Table 1 presents a summary of the signal segments available and considered in the training stage. Training is performed with 5000 signal blocks (used column) with a duration of one second (12800 samples). Each one-second block is randomly selected from within the signal segments with a duration of 1.28 seconds (16384 samples). The signal segments are randomly selected from the time period comprising the third year of operations, where the total number of available segments is shown in the column “available.” The considered column lists the number of segments after discarding segments with a vibration RMS below 0.5 G. Each block is preprocessed to have zero mean and unit variance. Both sparse coding algorithms are stopped at 90% sparsity, which means that each block of 12800 samples is modeled with 1280 atom instances. This means that the sparsity coefficient refers to the ratio of nonzero instances to segment length. We use a learning step size of $\eta = 10^{-6}$ for the dictionary update. A discussion on the effects of hyperparameter selection, such as sparsity level and learning step size is provided by Martin-del-Campo and Sandin [39].

Before the first signal block in the training stage is processed, the dictionary is initialized with a pseudorandom dictionary of eight atoms. Initially, the atoms are seventy elements long and are always defined in the same way at the start of the learning process. Each atom is generated from fifty elements sampled from a Gaussian distribution with zero mean that is zero-padded with ten samples at each tail. The atoms can grow in length during the learning process, and they are normalized after each update. Experiments with larger dictionaries show that additional atoms are rarely selected and do not adapt to the signal. The dictionary learned after the first 5000 iterations is henceforth referred to as a baseline dictionary, and it is used as the initial dictionary in the monitoring stage. This procedure is repeated for the six turbines at the beginning of their training stage using the same pseudorandom dictionary. Consequently, a baseline dictionary is generated for each turbine.

3.2.2. Monitoring stage

The monitoring stage is similar for the six turbines and is representative of the online monitoring scheme described in Figure 1. The learned dictionaries in the training stage, which are now known as baseline dictionaries, are used as input into the sparse coding with dictionary learning algorithm. In this second processing stage, we consider all of the available signal segments during the 46 consecutive months of data. However, as in the first learning stage, we use only segments with a vibration RMS value above 0.5 G. Table 2

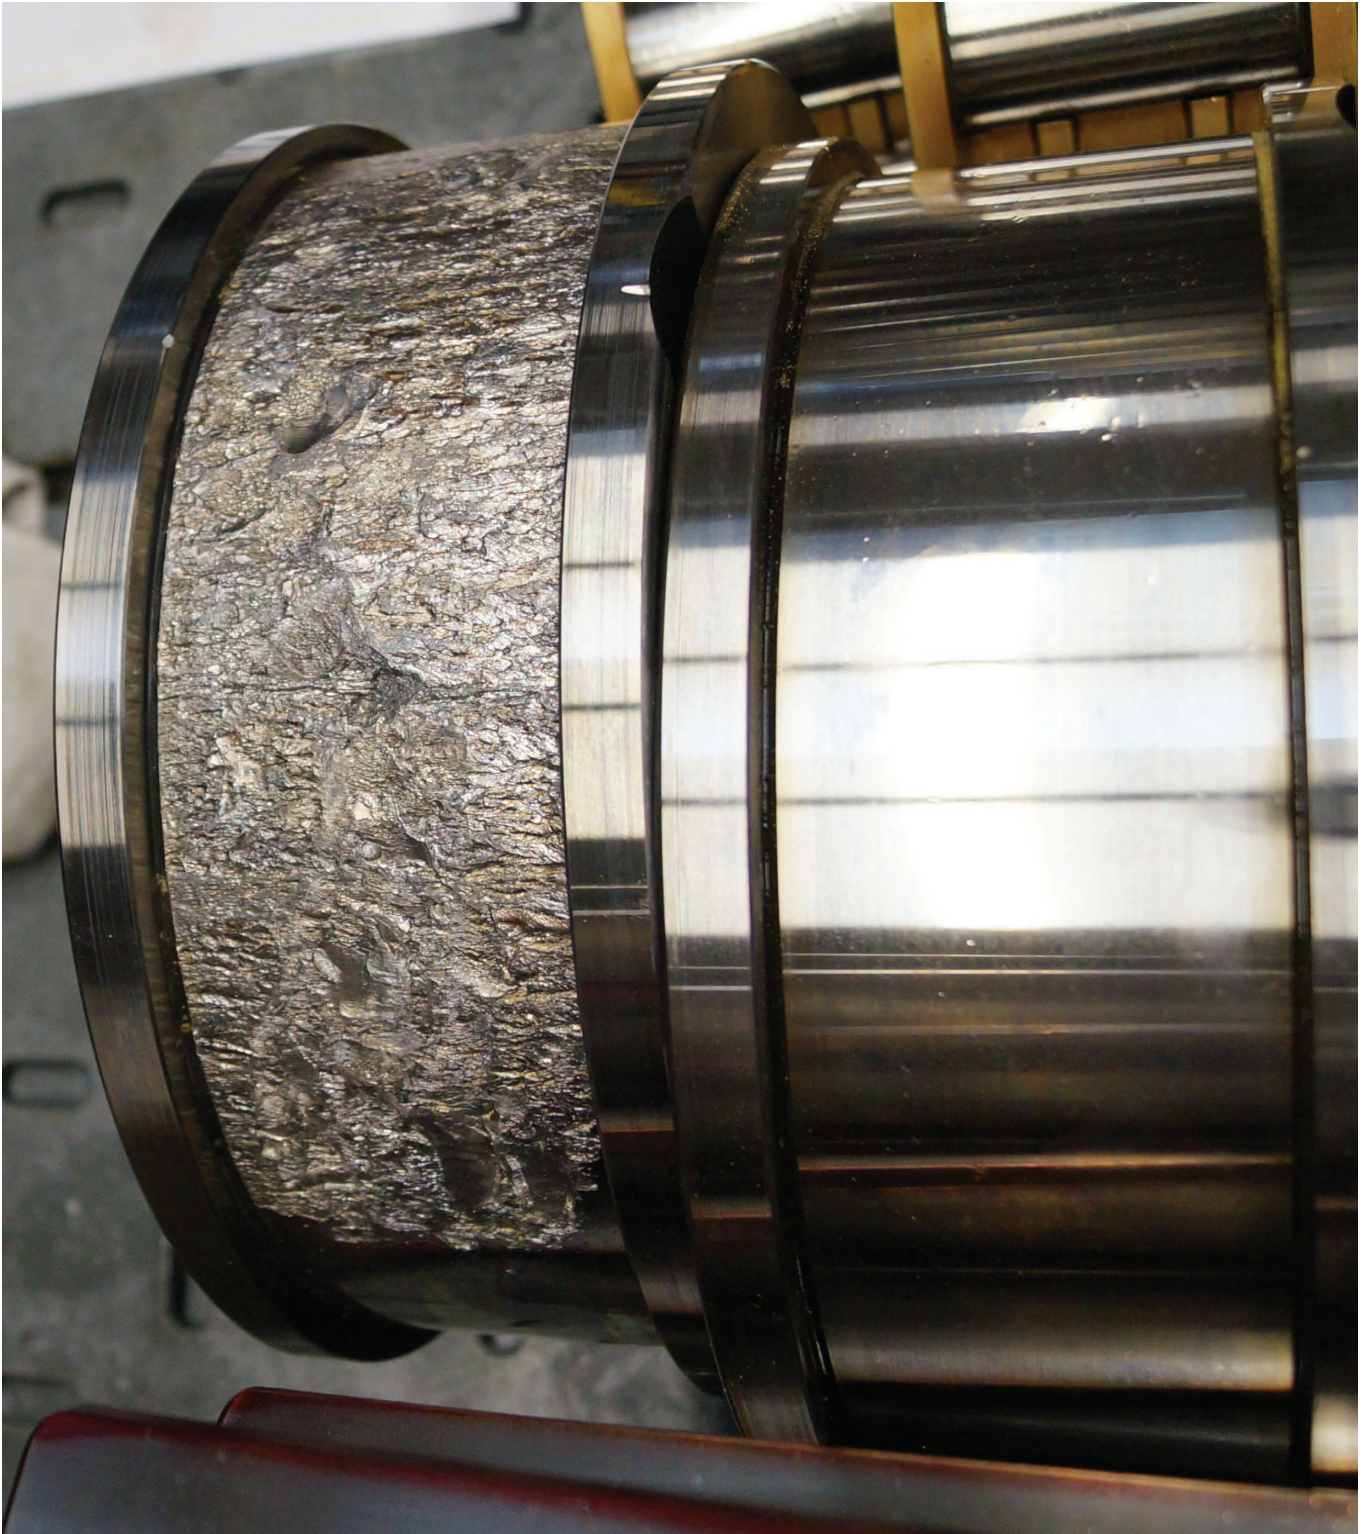


Figure 3 | End result of the inner raceway failure of the bearing supporting the gearbox planets in Turbine 5 (left). Inner raceway of a healthy bearing included for comparison (right).

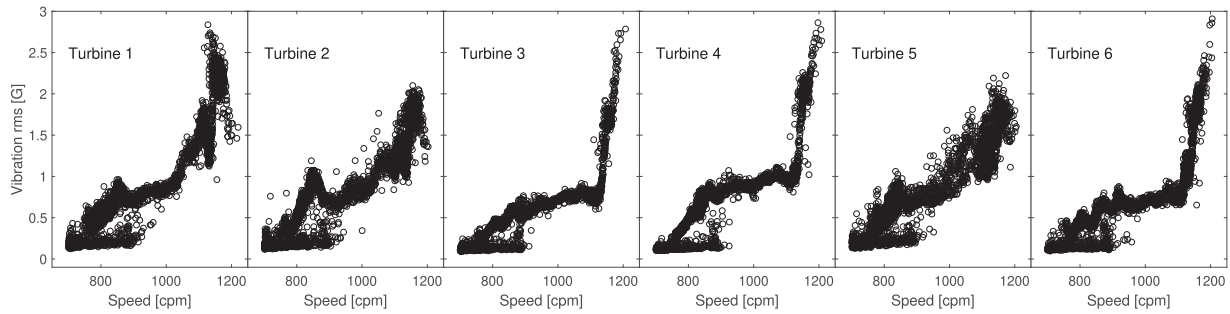


Figure 4 Scatter plot of vibration root mean square (RMS) of each recorded signal segment versus the speed in cycles-per-minute (CPM) for the six turbines. The data shown correspond to the full recording period of 46 months.

Table 1 Number of “Available” and “Selected” signal segments in the training stage. The processed “Used” blocks are one second long and are sampled at random offsets in the signal segments.

Case	Available	Considered	Used
Turbine 1	1212	859	5000
Turbine 2	1203	810	5000
Turbine 3	1243	527	5000
Turbine 4	1248	768	5000
Turbine 5	1237	803	5000
Turbine 6	1220	642	5000

Table 2 Number of “Available” and “Selected” signal segments in the monitoring stage.

Case	Available	Used
Turbine 1	2982	2078
Turbine 2	3005	2058
Turbine 3	2670	1135
Turbine 4	2667	1623
Turbine 5	2976	1907
Turbine 6	2953	1629

includes a summary of the number of used signal segments for each turbine. The available column lists all of the available segments for each turbine, and the used column lists the number of segments after discarding those with a vibration RMS below 0.5 G. Each segment is preprocessed to have zero mean and unit variance. As before, the MP and OMP algorithms are stopped at 90% sparsity, which means 1600 atom instances are used to model 16384 samples. The signal segments are analyzed in sequential order, as would be the case in an online monitoring situation. Therefore, the dictionary is said to be propagated over time, which means that it is gradually adapting to the structure of the signal. A method for edge effect reduction is not introduced in this stage because the learned translation-invariant atoms are about two orders of magnitude shorter than the processed signal window and the existing 12-hour gap between signal windows.

Two scenarios are considered in the monitoring stage. In the first scenario, the dictionary is propagated with a step size of $\eta = 10^{-6}$. In this scenario, we study the change in the propagated dictionary over time with respect to the baseline dictionary. In the second scenario, the dictionary remains constant by setting $\eta = 0$. In this scenario, we study the fidelity of the sparse representation model over time using the fixed baseline dictionary. In the second scenario, we

are interested in studying the effect of bearing failures on the fidelity of the sparse model.

The monitoring stage continues with the testing of two additional cases. These cases focus on the importance of the baseline dictionary by investigating the consequences of propagating a dictionary that is not optimized to the signal of the machine. Signal segmentation and preprocessing are performed in the same way as in the previous two monitoring cases. In one case, we use a baseline dictionary learned from each turbine to model and analyze the signals from the remaining five turbines. We repeat this procedure with the six turbines. In the other case, we use an arbitrary baseline dictionary learned from vibration signals obtained from the ball bearing data center at Case Western Reserve University (CWRU) [40]. In the latter case, the signals are generated by a rotating machine consisting of an electric motor, a torque transducer, a dynamometer, and a ball bearing supporting the motor shaft. An accelerometer located at the drive end of the motor is used to record the vibration data with a sampling rate of 12 kHz. We alternate between several recorded datasets from a healthy bearing to simulate a varying load between 0 HP and 3 HP. Thus, the dictionary used in this case does not encode information about the wind turbine signals and is not expected to result in particularly accurate sparse codes of the vibration signals.

4. RESULTS

This section presents the results of the numerical experiments. First, it provides an evaluation of the similarity of the learned dictionaries for the different turbines. This evaluation is followed by a study of the effects of selecting different baseline dictionaries. In this study, we use vibration data from the bearing data center [40] at CWRU. Next, we show how our proposed dictionary distance metric is used to identify anomalous assets and conclude by presenting the usefulness of this metric as a condition indicator.

4.1. Generalization Across Turbines

In the training stage, one baseline dictionary is learned for each turbine under healthy conditions. Here, we aim to investigate how similar the learned dictionaries are across different turbines. Thus, we process the signal from the same accelerometer location in the six turbines. Furthermore, we are interested in the differences that result from the use of two different sparse coding algorithms,

namely, MP and OMP. We use the same protocol and hyperparameters during learning with MP and OMP. The resulting dictionaries for three turbines using both algorithms are shown in Figure 5, and each dictionary includes eight updated atoms. The selected dictionaries include the following known cases: a healthy condition for Turbine 1, an electrical sensor failure for Turbine 2 and a gear-box failure for Turbine 5. However, all of the baseline dictionaries are trained during a period when the turbines are expected to be healthy (meaning that no faults were detected during or after that period of time). The dictionaries are obtained after learning from 5000 signal blocks, which correspond to approximately 83 minutes of vibration data and 64 million samples. The updated atoms have different length and (L2) normalized magnitudes, and all atoms are illustrated in the same scale in the two panels. The atoms corresponding to Turbine 1 are ordered by the ascending center frequency in both sparse coding algorithm cases. The center frequencies are calculated as the mean value of the power spectral density of each atom. The atoms of Turbine 2 and Turbine 3 are ordered in the corresponding way by maximizing the cross-correlation with each atom of Turbine 1. Table 3 summarizes the center frequencies of the atoms.

Some of the atoms learned from the three turbines appear to be similar, while a few are different, as shown in Figure 5. Furthermore, there are similarities between the atoms learned with the MP and OMP methods. For example, atoms one and two have sinusoidal components of relatively low frequency in all three cases regardless of the sparse coding method used. Atoms three and four are exchanged between the MP and OMP cases. Atom four with MP and atom three with OMP have a visible sinusoidal component. In contrast, atom three with MP has a smaller central frequency and a more noise-like appearance compared to atom four with OMP. Atoms five to eight are more noise-like and have higher center frequencies in both cases. Atoms six and seven are the most different across the three turbines when using MP or OMP. However, though they are different across the turbines, atoms six and seven are similar for MP and OMP, in particular, for Turbine 1 and Turbine 2 but not for Turbine 5.

Note that though all of the dictionaries are trained using healthy signal segments recorded under similar operational and environmental conditions, the bearings are not identical. The bearings in

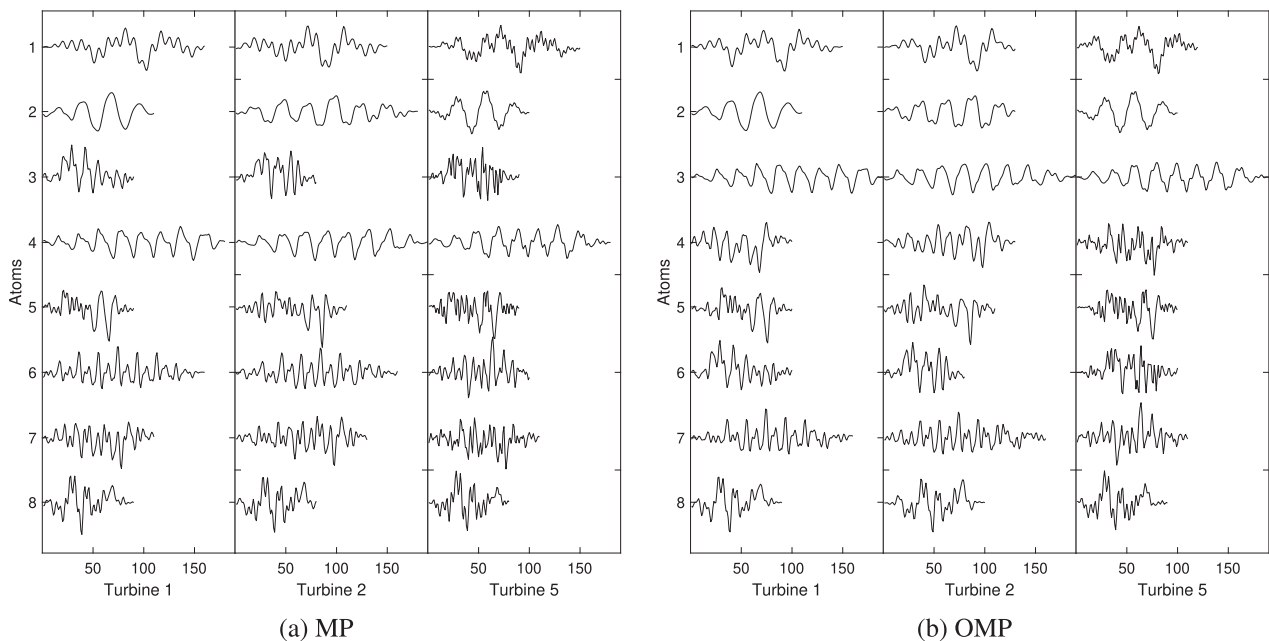


Figure 5 Atoms learned from vibration signals from the selected three turbines at the end of the training stage. Turbine 1 did not have any reported failures. Turbine 2 had an electrical sensor fault at the beginning of its operation. Turbine 5 had the faults described in Section 3.1. The atoms of Turbine 1 are ordered by the ascending center frequency. The atoms of Turbine 2 and Turbine 5 are ordered by maximizing the cross-correlation with respect to the atoms of Turbine 1. All atoms are normalized.

Table 3 Center frequencies of atoms in the baseline dictionary for MP (OMP).

Case	Center Frequency (kHz)							
	1	2	3	4	5	6	7	8
Turbine 1	0.28 (0.26)	0.47 (0.47)	0.64 (0.64)	0.64 (0.70)	0.92 (0.90)	1.36 (1.02)	1.40 (1.36)	1.49 (1.49)
Turbine 2	0.50 (0.29)	0.61 (0.49)	0.86 (0.60)	0.88 (1.28)	0.99 (1.45)	1.36 (1.04)	1.52 (1.36)	1.97 (0.70)
Turbine 3	0.14 (0.36)	0.56 (0.56)	0.76 (0.64)	0.96 (0.76)	1.00 (1.01)	1.11 (1.11)	1.28 (1.23)	1.60 (1.60)
Turbine 4	0.28 (0.57)	0.80 (0.75)	0.92 (0.80)	0.99 (0.93)	1.01 (1.05)	1.28 (1.22)	1.28 (1.46)	1.45 (1.66)
Turbine 5	0.26 (0.32)	0.45 (0.45)	0.64 (0.61)	0.92 (1.80)	1.60 (0.96)	1.60 (1.86)	1.85 (1.69)	2.68 (1.56)
Turbine 6	0.57 (0.32)	0.96 (0.60)	0.96 (0.69)	1.02 (1.09)	1.23 (1.28)	1.32 (1.28)	1.60 (1.66)	2.40 (2.38)

Turbine 5 are newer compared to the bearings in Turbine 1 and Turbine 2 due to the preceding gearbox replacement. The bearings in Turbine 1 and Turbine 2 have been in operation for more than two years. Thus, these bearings have been degraded compared to the new bearings in Turbine 5. Therefore, we cannot expect the updated atoms to be identical for the three turbines illustrated here.

4.2. Effect of Baseline Dictionary Selection

Next, we investigate the effects of selecting different initial dictionaries and dictionary learning step lengths. In a field implementation of dictionary learning, the baseline dictionary needs to be learned from the signal to be monitored to ensure that the model will have high fidelity and effectively separate the signal from the noise. The fidelity of the model in decibels is the ratio between the sparse approximation and the signal residual:

$$dB = 20 \log_{10} \left(\frac{\hat{s}_k}{\epsilon(t)} \right), \quad (16)$$

where dB is the fidelity in decibels, \hat{s}_k is the sparse approximation of a signal segment and $\epsilon(t)$ is the corresponding residual. Figure 6 shows the model fidelities for three different cases. In the first case, which corresponds to $\eta = 10^{-6}$, the dictionary at $t = 0$ is the baseline dictionary of each turbine; this dictionary is propagated over time with a finite dictionary update step length. In the second case, the dictionary update is set at $\eta = 0$, and the baseline dictionary is used as is without further modification. In the third case, which is

labeled *averaged*, the dictionary learned during the training stage of each individual turbine is used to model the signals of the remaining five turbines without further modification of the dictionary. In this case, the average fidelity of the resulting five models is shown. The fidelity is low-pass filtered with a first-order filter and a time constant of 15 days (30 signal segments) to improve the clarity of the plot. The $\eta = 10^{-6}$ and $\eta = 0$ cases result in similar fidelities, which indicates that the baseline dictionaries have converged and are not updated significantly by further learning. When we use the averaged baseline dictionary to model the signals from the other five turbines, the fidelities are slightly lower than when we use the correct baseline dictionaries. These results are observed independently of the sparse coding algorithm we used. MP has slightly lower fidelity than OMP in the three cases considered. The large decrease in the model fidelity for Turbine 5 at the beginning of year one corresponds to the period of time when the bearing was defective, and the two minima correspond to the replacement of the output shaft bearing and the gearbox. The low model fidelity for Turbine 2 at the beginning of year zero does not correspond to a previously known fault; this result is likely due to an electrical issue after the installation of the sensor, cables, or monitoring unit.

Next, two additional cases are considered where the baseline dictionary is learned from another machine. In these cases, the vibration data were obtained from the ball-bearing data center [40] at CWRU. The fidelity of these two cases using the MP and OMP algorithms is shown in Figure 7. Each algorithm has two cases: $\eta = 10^{-6}$, indicating when the baseline dictionary is propagated, and $\eta = 0$,

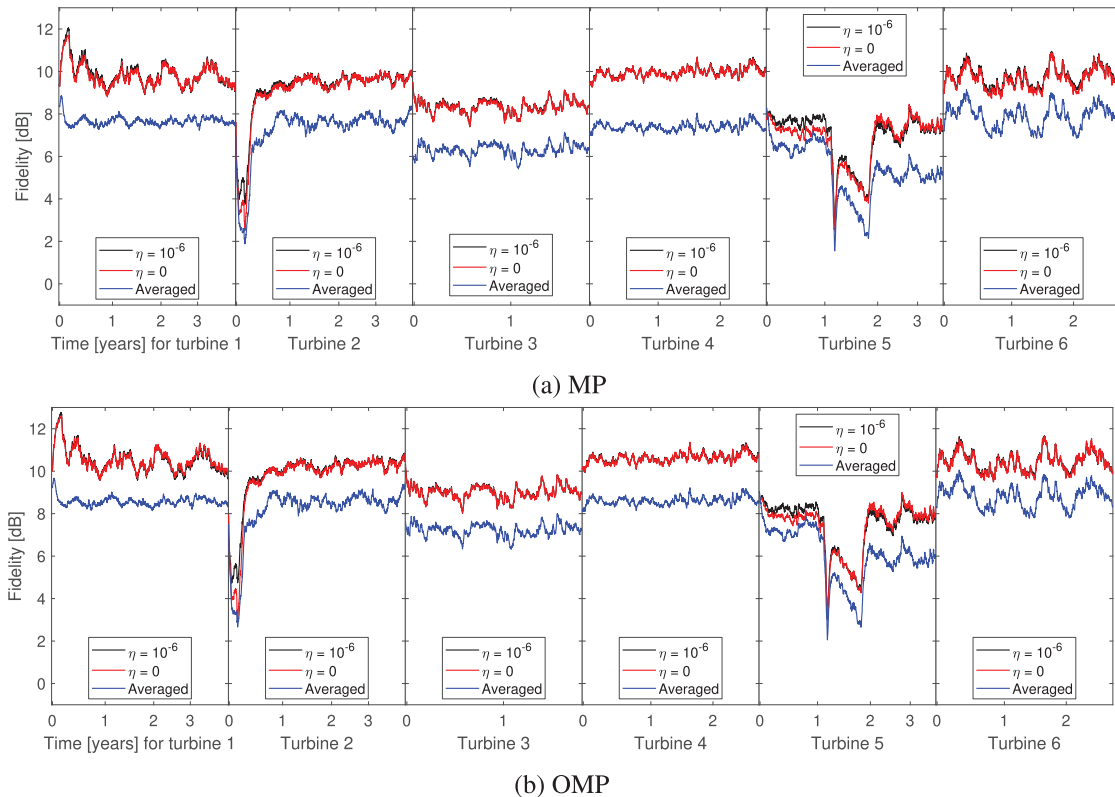


Figure 6 Fidelity of the vibration signal model for the baseline dictionary is used as is, without further adaptation. The averaged case corresponds to the fidelity we obtain when the baseline dictionary of each turbine is used without further modification to model the vibration signals of the other five turbines.

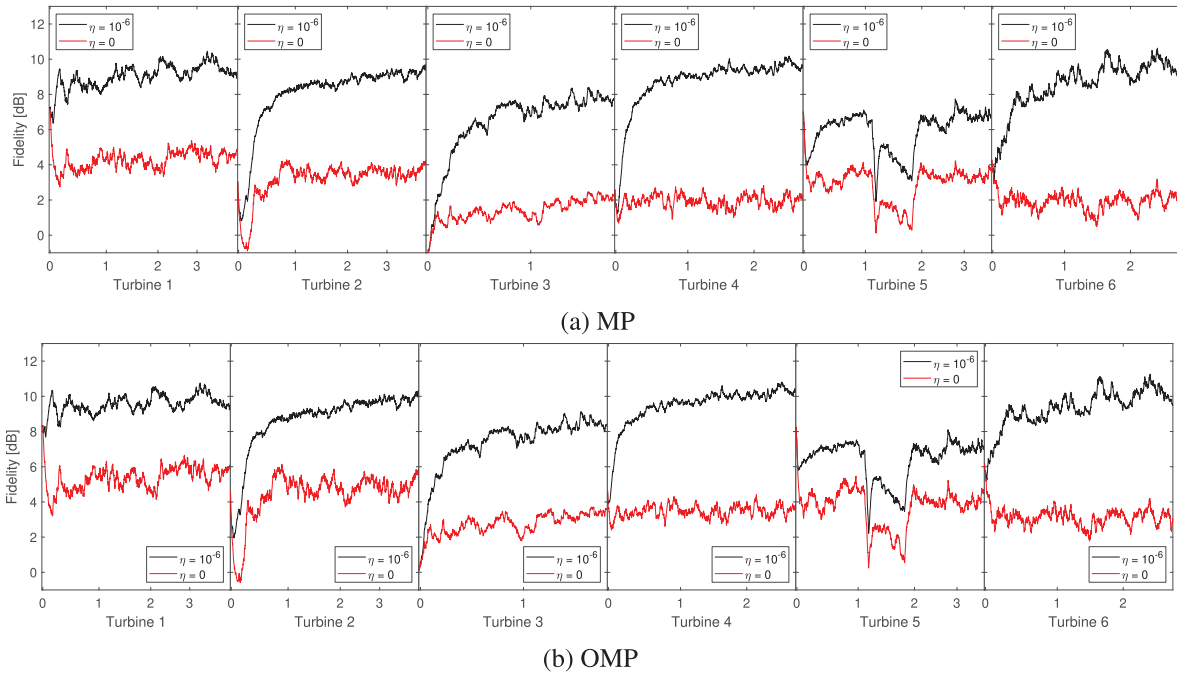


Figure 7 | Fidelity of the signal models based on the initial dictionaries learned from the Case Western Reserve University (CWRU) database for , the baseline dictionary learned from the CWRU database is used as is.

which occurs when the baseline dictionary is used as is and without further adaptation. Under the two algorithms, the fidelity in the $\eta = 10^{-6}$ case increases with time and reaches a similar level as in Figure 6 at the end of the time period, which is to be expected since the atoms in the dictionary are adapted to the signal. The CWRU baseline dictionaries cannot adequately model the wind turbine signals, resulting in a lower fidelity in the case of $\eta = 0$ compared to the fidelities in Figure 6. However, note that the replacement of the output shaft bearing and gearbox in Turbine 5 are still associated with abnormally low values of the fidelity.

4.3. Distance to the Baseline Dictionary

The dictionary distance defined by Eq. (13) quantifies the difference between two dictionaries. This distance can be used to detect a gradual change in a propagated dictionary by determining the distance between the propagated dictionary and either the baseline dictionary or a set of baseline dictionaries. Thus, faults that appear after a long period of degradation could possibly be detected by monitoring the distance between the propagated dictionary and the baseline. In principle, the baseline dictionary is defined by atoms learned in normal states of operation; see Section 3.2 for further details.

Figure 8 shows the corresponding dictionary distances for the six turbines using the two sparse coding algorithms. As described above, the baseline dictionaries are trained with signal blocks recorded when the turbines are operating in healthy conditions (after the gearbox replacement in Turbine 5). The resulting curve trends for both sparse coding algorithms describe similar behavior. Turbine 1, Turbine 3, Turbine 4, and Turbine 6 show an increase in the dictionary distance when the dictionary is propagated over time. For Turbine 2, there is a relatively fast initial increase in the distance, which after some time stabilizes and becomes similar to

the distance for Turbine 1. The rapid increase at the start is in agreement with the results presented in Figure 6; this figure shows that the model fidelity is initially low, which is most likely due to an electrical fault in the measurement system connected to the accelerometer. In contrast, the dictionary distance for Turbine 5 increases more quickly than the distances we determined for all of the other turbines. The dictionary distance for Turbine 5 is approximately two to three times higher than the distances for the other turbines at the point in time when the output shaft bearing is replaced in the gearbox. After the bearing replacement, the distance is approximately stationary until it reaches another peak just before the gearbox is replaced. After the replacement of the gearbox, the dictionary distance decreases and approaches the distance for Turbine 1, which indicates the return to a normal condition. The dictionary distances calculated with the OMP algorithm have a larger spread than the distances calculated with the MP algorithm. While the dictionary distance with the MP algorithm covers a 7-degree spread at the end of the recording period, the dictionary distance with the OMP algorithm covers a 10-degree spread.

The dictionary distance is an indicator that can be used for outlier detection in a population of monitored wind turbines. However, it can be challenging to interpret the dictionary distance value by itself. An alternative is to estimate the median absolute deviation (MAD) values for each turbine with respect to the entire turbine population. MAD is a measure that enables to identify the variability of an individual sample with respect to the population. The MAD values is estimated as follows:

$$MAD = \text{median} (|\beta_j - \tilde{\beta}|), \quad (17)$$

where β is the dictionary distance values of each turbine j as calculated for each time instance.

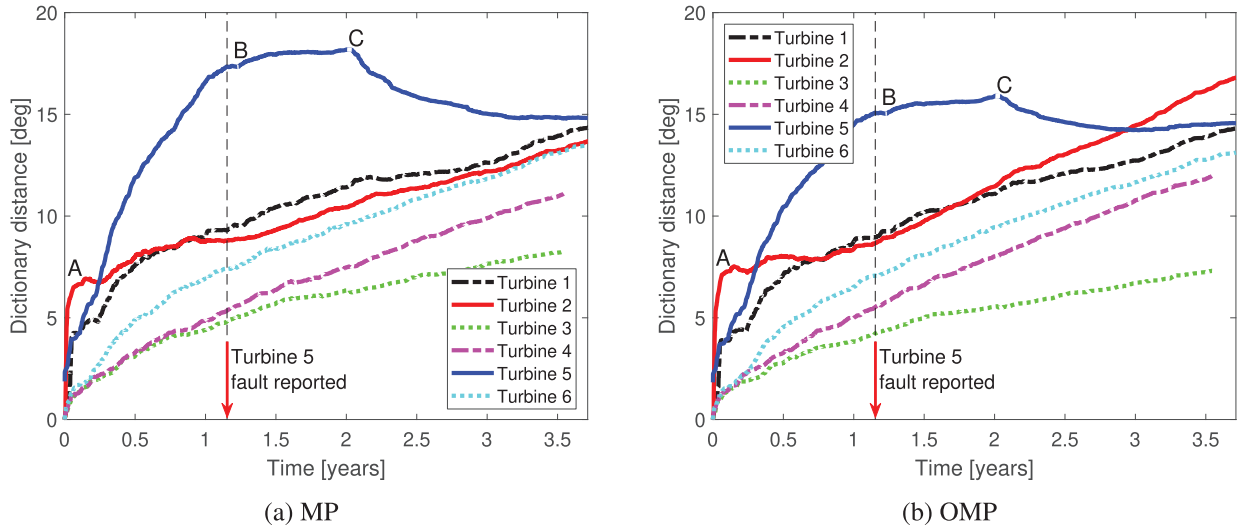


Figure 8 | Distance of a propagated dictionary with respect to the baseline dictionary versus time using the (a) matching pursuit (MP) and (b) orthogonal matching pursuit (OMP) algorithms for the six turbines. Label A indicates the end of a time period with a possible electrical fault in the data acquisition system of Turbine 2. In addition, label B indicates the time period when the output shaft bearing was replaced in Turbine 5, and label C indicates the subsequent gearbox replacement.

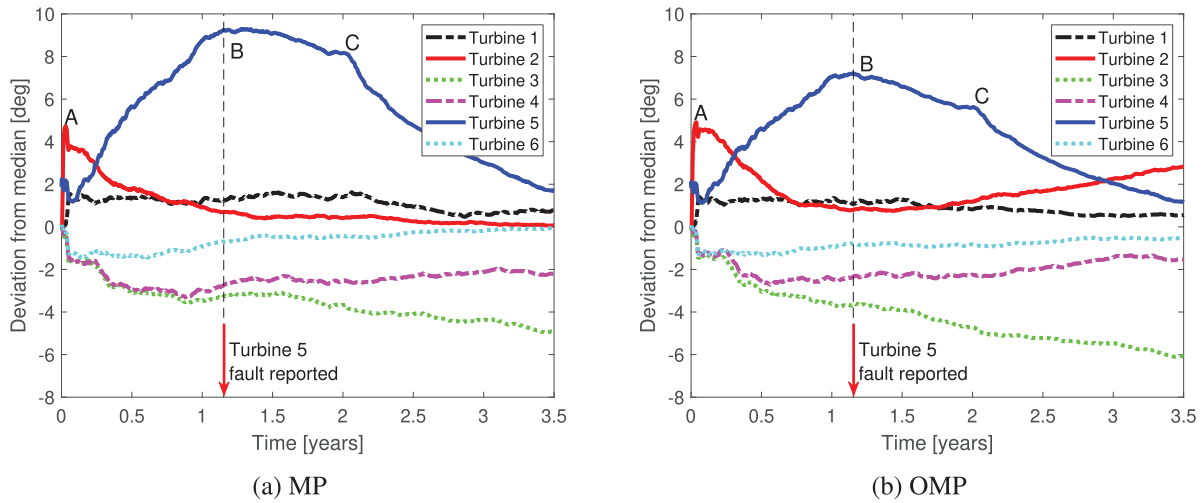


Figure 9 | Median absolute deviation (MAD) of the dictionary distance versus time based on the (a) matching pursuit (MP) and (b) orthogonal matching pursuit (OMP) algorithms. Label A indicates the end of a time period with a possible electrical fault in the data acquisition system of Turbine 2. In addition, label B indicates the time period when the output shaft bearing was replaced in Turbine 5, and label C indicates the subsequent gearbox replacement.

Figure 9 shows the MAD values of the dictionary distance shown in Figure 8. The MAD highlights a significant deviation of Turbine 5 in a period that extends approximately eight months earlier than the date when a fault report was filed for the turbine. The fault report for Turbine 5 was filed 1.2 years after the start of the vibration data recording. The trend of Turbine 3 deviates from the other turbines at the end of the recorded time period for both sparse coding algorithms. This deviation may be due to the geographical location of Turbine 3, which is the farthest away from the other turbines. The MAD values at times of interest when Turbine 2 had the electrical fault (point A) and Turbine 5 had the HSS bearing replaced (point B) and the gearbox replaced (point C) are shown in Table 4 for the MP case and Table 5 for the OMP case. The largest values are highlighted in bold and correspond to the periods of interest.

Table 4 | Median Absolute Deviation (MAD) values at times of interest for Matching Pursuit (MP) algorithm. All values in degrees.

Case	MP Algorithm		
	A	B	C
Turbine 1	1.49	1.20	1.44
Turbine 2	3.62	0.70	0.46
Turbine 3	-1.70	-3.28	-3.67
Turbine 4	-1.69	-2.74	-2.53
Turbine 5	1.92	9.25	8.21
Turbine 6	-1.49	-0.70	-0.46

An adequate selection of the initial dictionary under an online monitoring scheme is important to constrain the dictionary distance values and the rate of change. Using the baseline dictionaries

Table 5 | Median Absolute Deviation (MAD) values at times of interest for orthogonal Matching Pursuit (OMP) algorithm. All values in degrees.

Case	OMP Algorithm		
	A	B	C
Turbine 1	1.35	1.09	0.89
Turbine 2	4.50	0.80	1.16
Turbine 3	-1.45	-3.63	-4.68
Turbine 4	-1.36	-2.33	-2.26
Turbine 5	1.79	7.21	5.63
Turbine 6	-1.35	-0.80	-0.89

trained on signals from the turbines, the dictionary distance for Turbine 5 is roughly twice as large as the corresponding distance of the healthy turbines before the replacement of the faulty bearing. The two sparse coding algorithms we investigated result in similar dictionary distance trends. Note that in these numerical experiments, we use a sparse dataset with approximately 2.56 seconds of recorded signal per day. In an online monitoring implementation of this method, there would be significantly more data per time unit and a faster effective learning rate. Thus, to achieve online processing of the step length parameter η , Eq. (12) should be lowered to avoid short-term overfitting of the propagated dictionary to different healthy operational states.

4.4. ROC Analysis

We perform a basic receiver operating characteristic (ROC) analysis to study the usefulness of the dictionary distance measure as a condition indicator. ROC curves are commonly used to assess the efficiency of condition indicators for diagnostics purposes [41], and they are used more generally as a method for classifier evaluation and selection. A ROC curve illustrates the relationship between the true positive rate (TPR) and the false-positive rate (FPR). Each point on the curve corresponds to different parameters of the classifier model, for example, a threshold value of a condition indicator. This threshold value can be set up by an expert or an autonomous agent based on the performance of a population of turbines. Thus, the ROC curve describes the expected TPR and FPR for varying threshold values.

Figure 10 shows the ROC curves for two indicators based on the dictionary distance of the two sparse coding algorithms presented in Figure 8. One of the ROC curves is based on the slope of the dictionary distance *versus* time, and the other ROC curve is based on the minimum difference between the distance for one turbine compared to the distances for the other five turbines. The slope-based indicator gives a balanced ratio of the number of true positives to false positives. The minimum difference indicator is slightly skewed, which means that the indicator makes positive classifications with weak confidence since it classifies all positives correctly for false-positive rates above approximately one half. The indicators based on the MP algorithm cover a larger area than the results obtained with the OMP algorithm. This behavior corresponds to higher classification accuracy at a lower computational cost. Label B in Figure 8 marks the time of the bearing replacement in Turbine 5, while label A indicates the resolution of a suspected electrical issue that was introduced during the installation of the sensor system in Turbine 2. In the ROC analysis, we consider the data from Turbine

1, Turbine 3, Turbine 4, Turbine 6, and the data from Turbine 2 after time A as data that correspond to healthy states of operation. Furthermore, the data from Turbine 5 after the gearbox replacement at C are also considered as data that correspond to healthy states of operation. The data from Turbine 5 before C and the data from Turbine 2 before A are considered as data that correspond to faulty states of operation. The classification results used in the calculation of the TPR and FPR are defined by varying the threshold values of the slope-based and minimum difference indicators.

5. DISCUSSION

In this section, we provide a discussion of the results described in Section 4. We explain how a dictionary learned under certain conditions can be used by assets that experience similar conditions, which we follow by describing the effects of baseline dictionary selection. Afterward, we evaluate our proposal of using dictionary distance as a condition indicator. Finally, we provide considerations on the selection of the sparse coding algorithm.

5.1. Generalization of a Learned Dictionary

A working hypothesis that motivated this study is that condition monitoring signals from healthy turbines of the same type that operate in similar environments with nearby locations should be similar to some degree. Thus, a baseline signal model learned from one healthy turbine or a set of turbines could be useful for monitoring other turbines. We investigate this idea with vibration data recorded from accelerometers located at the same position in six wind turbines. We find that the dictionaries learned for the different turbines have remarkable similarities, see Figure 5. Furthermore, a dictionary learned from one turbine is successfully applied to the other two turbines, see Figure 6. This application cannot be made when one is using an arbitrary dictionary learned from another bearing vibration dataset; see Figure 7. Thus, we conclude that a dictionary-based sparse signal model learned from a sensor in one turbine can be generalized to the corresponding signal in another turbine, which permits further studies in this direction.

For example, using data from a larger population of turbines, it is possible to investigate whether there are some signal components or atoms that are common for all healthy turbines in a wind farm. By learning dictionaries for a larger population of wind turbines, one could also create a repository of dictionaries and dictionary elements, which would enable comparisons of dictionaries learned from similar turbines in different wind farms. In this way, the initial dictionary implemented in the condition monitoring systems for new wind turbines could be selected based on the typical features of healthy turbines. Similarly, dictionaries and atoms learned from confirmed faulty turbines can be stored and potentially used for the diagnosis of similar faults in other turbines of the same type.

5.2. Selection of a Baseline Dictionary

The baseline dictionary for each turbine considered here is learned from vibration data recorded during a period when there is no known fault. However, when taking a new turbine into service, it is not certain that the turbine will operate without problems or that no

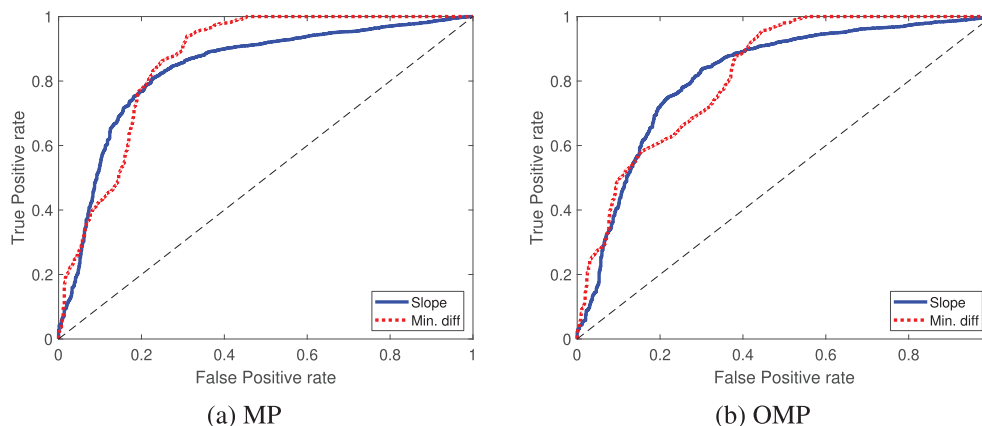


Figure 10 | Receiver operating characteristic (ROC) curves based on the dictionary distances between the baseline dictionaries and propagated dictionaries for the six turbines using the (a) matching pursuit (MP) and (b) orthogonal matching pursuit (OMP) sparse coding algorithms. The curves result from a threshold on the rate of change of the distance over time (blue) and the minimum difference of the dictionary distances in the population (red).

faults were introduced during the installation of the wind turbine components and condition monitoring system. Thus, if the baseline dictionary is trained by starting from a randomized dictionary, it may not be possible to identify a fault that is already present in the turbine from its initial operation.

The results presented in Figure 6 show that the difference in fidelity when using a baseline dictionary learned from the turbine itself, or a baseline dictionary learned from a similar turbine is small. Thus, an alternative to using a randomized initial dictionary is to further investigate the possibility of using a baseline dictionary learned from similar turbines that are known to operate in healthy conditions. Using such a baseline dictionary, it could be possible to identify abnormal conditions that appear when a turbine is taken into service.

5.3. Selection of the Condition Indicator

In general, a condition indicator is a quantitative measure of the performance or operational condition of a machine, and a feature is a measurable characteristic used when modeling a signal or set of data. Conventional condition indicators, such as the RMS and the energy within certain (kinematically determined) frequency bands, are successfully used as features both in conventional and machine learning approaches to condition monitoring. However, this approach typically requires human expertise to select and customize indicators for each particular application and machine type, which is costly. Furthermore, faults with unexpected characteristics can be difficult to detect with indicators engineered for particular purposes. These facts motivate investigations of unsupervised learning approaches as in the present study, which can complement and potentially replace manually defined condition indicators in some applications.

For example, the absolute value or the rate of change of the dictionary distance could be used as a condition indicator with a threshold level defining the allowed drift away from a baseline dictionary. However, these methods require further testing with data from a larger population of wind turbines to determine the appropriate

threshold value(s) and the expected true- and false-positive rates. Alternatively, dictionary distances can be used as scores in an unsupervised anomaly detection or ranking algorithm. Thus, unsupervised feature learning methods such as dictionary learning could pave the way for the development of unsupervised anomaly detection systems.

5.4. Selection of Sparse Coding Algorithm

In this work, the sparse representations are generated with MP and OMP, which are both greedy sparse coding algorithms. In both cases, the same probabilistic gradient ascent algorithm is used for dictionary learning. The OMP algorithm enforces the orthogonality of the terms in Eq. (1) by updating the weights of the previously selected terms for each consecutive term added to the sparse signal representation. This procedure results in higher computational cost and fidelity (given one particular dictionary) compared to the MP algorithm. However, when they are used in combination with dictionary learning, the fidelities achieved with the two methods can be comparable.

Figure 6 shows that the difference in fidelity obtained with the two algorithms is negligible. A decrease in the fidelity of the signal model for Turbine 5 when the HSS bearing is replaced at the beginning of the first year of operation is observed in both cases. The update of the propagated dictionary depends on the selected sparse coding algorithm, which consequently affects the calculated dictionary distance. Figure 8 shows that the dictionary distance distribution is wider under normal operating conditions in the case of the OMP algorithm compared to the MP algorithm. Considering the higher computational cost of the OMP algorithm, these results indicate that there is no evident benefit to using OMP for the calculation of dictionary-learning-based condition indicators and that MP-based indicators (for unknown reasons) may be beneficial for anomaly detection.

The computational cost of the sparse coding algorithm is significant, even if the implementations of the MP and OMP algorithms used here have been optimized for efficiency. Thus, an interesting

direction for future work is to investigate alternative methods for unsupervised feature learning, such as learning of co-sparse analysis operators [42], where the inverse problem addressed here is avoided.

6. CONCLUSION

This work focuses on the monitoring of rolling element bearings in wind turbines using an unsupervised dictionary learning approach and real-world wind turbine vibration data that have been made publicly available. The results presented above demonstrate that a condition indicator based on the dictionary distance metric is useful for the monitoring of drivetrain gearbox bearings of wind turbines and serves as a complement to currently used methods. In the case of Turbine 5 considered above, it is not known when the issue(s) leading to the bearing and gearbox replacements first appeared. The motivation for the bearing replacement was a sudden increase in the enveloped signal from the HSS axial sensor after approximately one year of operation. This replacement was followed by a gearbox replacement approximately one year later. The results shown in Figure 8 suggests that the abnormal behavior of Turbine 5 could have been detected several months earlier using dictionary learning. Earlier detection of a fault in the turbines represents an improvement in terms of maintenance planning and reducing the risk of costly failures. However, further tests are required to understand the strengths and weaknesses of a dictionary learning approach in realistic large-scale monitoring environments. For example, it is not understood whether the long-term drifts away from the baseline dictionaries observed in Figure 8 are related to mechanical wear of the turbines or the greedy approximation of the NP-hard dictionary learning problem. Further testing requires the acquisition and processing of condition monitoring data from a larger population of turbines, including documented faults and maintenance activities, which is the next step but is beyond the scope of the project reported here.

CONFLICTS OF INTEREST

The authors declare no conflict of interest.

AUTHORS' CONTRIBUTIONS

S. Martin-del Campo co-developed the framework, conducted the experiments and wrote the article. F. Sandin initiated the research, co-developed the framework and provided recommendations to this article. D. Strömbergsson collected and provided the data used on the experiments.

ACKNOWLEDGMENTS

We thank Per-Erik Larsson, Stephan Schnabel, Allan Thomson, and Joe Erskine for discussions that helped us improve the manuscript, and we thank Karl Skretting for contributing the idea to use a unique seed dictionary. This work is partially supported by SKF through its University Technology Center at LTU. The contribution of Fredrik Sandin is funded by the Kempe Foundations under contract Gunnar Öquist Fellow. The research

leading to these results has received funding from the People Program (Marie Curie Actions) of the European Union's Seventh Framework Program FP7/2007-2013/ under REA grant agreement number 612603 and the Swedish Foundation for International Cooperation in Research and Higher Education (STINT), grant number IG2011-2025.

REFERENCES

- [1] Global Wind Energy Council, Global wind Report - Annual Market Update 2017, Global Wind Energy Council, Brussels, Belgium, 2018.
- [2] A. Crowther, V. Ramakrishnan, N.A. Zaidi, C. Halse, Sources of time-varying contact stress and misalignments in wind turbine planetary sets, *Wind Energy*. 14 (2011), 637–651.
- [3] W. Liu, A review on wind turbine noise mechanism and denoising techniques, *Renew. Energy*. 108 (2017), 311–320.
- [4] M.L. Hossain, A. Abu-Siada, S.M. Mueen, Methods for advanced wind turbine condition monitoring and early diagnosis: a literature review, *Energies*. 11 (2018), 1–14.
- [5] F.P. Garcia-Marquez, A.M. Tobias, J.M.P. Perez, M. Papaalias, Condition monitoring of wind turbines: techniques and methods, *Renew. Energy*. 46 (2012), 169–178.
- [6] S. Lacey, An overview of bearing vibration analysis, *Maintenance Asset Manag.* 23 (2008), 32–42.
- [7] A.K. Jardine, D. Lin, D. Banjevic, A review on machinery diagnostics and prognostics implementing condition-based maintenance, *Mech. Syst. Signal Process.* 20 (2006), 1483–1510.
- [8] W.P. Tse, T.J. Leung, Advanced system for automatically detecting faults occurring in bearings, in: L.M. Simon (Ed.), *Fault Detection: Theory, Methods and Systems*, Nova Science Publishers, New York, NY, USA, 2011, pp. 1–67.
- [9] H. Dias Machado de Azevedo, A.M. Araujo, N. Bouchonneau, A review of wind turbine bearing condition monitoring: state of the art and challenges, *Renew. Sustain. Energy Rev.* 56 (2016), 368–379.
- [10] A. Rai, S. Upadhyay, A review on signal processing techniques utilized in the fault diagnosis of rolling element bearings, *Tribol. Int.* 96 (2016), 289–306.
- [11] R.B. Randall, J. Antoni, Rolling element bearing diagnostics—a tutorial, *Mech. Syst. Signal Process.* 25 (2011), 485–520.
- [12] S. Mallat, *A Wavelet Tour of Signal Processing: The Sparse Way*, third ed., Academic Press, Burlington, MA, USA, 2009.
- [13] V. Pappas, Y. Romano, J. Sulam, M. Elad, Theoretical foundations of deep learning via sparse representations: a multilayer sparse model and its connection to convolutional neural networks, *IEEE Signal Process. Mag.* 35 (2018), 72–89.
- [14] M. Elad, *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing*, Springer, New York, NY, USA, 2010.
- [15] A. Bruckstein, D. Donoho, M. Elad, From sparse solutions of systems of equations to sparse modeling of signals and images, *SIAM Rev.* 51 (2009), 34–81.
- [16] J.-L. Starck, F. Murtagh, J. Fadili, *Sparse Image and Signal Processing: Wavelets, Curvelets, Morphological Diversity*, Cambridge University Press, New York, NY, USA, 2010.
- [17] E. Smith, M.S. Lewicki, Efficient coding of time-relative structure using spikes, *Neural Comput.* 17 (2005), 19–45.

- [18] E. Smith, M.S. Lewicki, Efficient auditory coding, *Nature*. 439 (2006), 978–982.
- [19] B. Olshausen, D. Field, Emergence of simple-cell receptive field properties by learning a sparse code for natural images, *Nature*. 381 (1996), 607–609.
- [20] B.A. Olshausen, D.J. Field, Sparse coding with an overcomplete basis set: a strategy employed by V1?, *Vision Res.* 37 (1997), 3311–3325.
- [21] S. Martin-del-Campo, F. Sandin, Towards zero-configuration condition monitoring based on dictionary learning, in 23rd European Signal Processing Conference (EUSIPCO), Nice, France, 2015, pp. 1306–1310.
- [22] H. Liu, C. Liu, Y. Huang, Adaptive feature extraction using sparse coding for machinery fault diagnosis, *Mech. Syst. Signal Process.* 25 (2011), 558–574.
- [23] S. Martin-del-Campo, K. Albertsson, J. Nilsson, J. Eliasson, F. Sandin, FPGA prototype of machine learning analog-to-feature converter for event-based succinct representation of signals, in Machine Learning for Signal Processing, IEEE International Workshop, Southampton, UK, 2013.
- [24] X. Chen, Z. Du, J. Li, X. Li, H. Zhang, Compressed sensing based on dictionary learning for extracting impulse components, *Signal Process.* 96 (2014), 94–109.
- [25] H. Tang, J. Chen, G. Dong, Sparse representation based latent components analysis for machinery weak fault detection, *Mech. Syst. Signal Process.* 46 (2014), 373–388.
- [26] H. Ahmed, M. Wong, A. Nandi, Intelligent condition monitoring method for bearing faults from highly compressed measurements using sparse over-complete features, *Mech. Syst. Signal Process.* 99 (2018), 459–477.
- [27] M. He, D. He, J. Yoon, T.J. Nostrand, J. Zhu, E. Bechhoefer, Wind turbine planetary gearbox feature extraction and fault diagnosis using a deep-learning-based approach, *Proc. Inst. Mech. Eng. Part O J. Risk Reliab.* 233 (2019), 303–316.
- [28] X. Wang, H.Z. Zhu, T. Rui, Y. Li, T. Liu, M. Liu, Shift invariant sparse coding ensemble and its application in rolling bearing fault diagnosis, *J. Vibroeng.* 17 (2015), 1837–1848. <https://www.jvejournal.com/article/15739>
- [29] H. Zhou, J. Chen, G. Dong, R. Wang, Detection and diagnosis of bearing faults using shift-invariant dictionary learning and hidden markov model, *Mech. Syst. Signal Process.* 72–73 (2016), 65–79.
- [30] S. Mallat, Z. Zhang, Matching pursuits with time-frequency dictionaries, *IEEE Trans. Signal Process.* 41 (1993), 3397–3415.
- [31] Y.C. Pati, R. Rezahfar, P.S. Krishnaprasad, Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition, in The Twenty-Seventh Asilomar Conference on Signals, Systems and Computers, Pacific Grove, CA, USA, 1993, pp. 40–44.
- [32] A. Szlam, K. Kavukcuoglu, Y. LeCun, Convolutional matching pursuit and dictionary training, ArXiv e-prints, 2010.
- [33] Ethos-N Processor Series: Ethos-N57, Arm, 2020. <https://armkeil.blob.core.windows.net/developer/Files/pdf/ethos/arm-ethos-n57-product-brief.pdf>
- [34] B. Mailhe, R. Gribonval, F. Bimbot, P. Vandergheynst, A low complexity orthogonal matching pursuit for sparse signal approximation with shift-invariant dictionaries, in IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2009), Taipei, Taiwan, 2009, pp. 3445–3448.
- [35] T. Blumensath, M. Davies, Gradient pursuits, *IEEE Trans. Signal Process.* 56 (2008), 2370–2382.
- [36] F. Sandin, S. Martin-del-Campo, Dictionary Learning with Equiprobable Matching Pursuit, in International Joint Conference on Neural Networks (IJCNN), Anchorage, Alaska, 2017.
- [37] S. Martin-del-Campo, F. Sandin, S. Schnabel, P. Marklund, J. Delsing, Exploratory analysis of acoustic emissions in steel using dictionary learning, in 2016 IEEE International Ultrasonics Symposium (IUS), Tours, France, 2016, pp. 1–4.
- [38] K. Skretting, K. Engan, Learned dictionaries for sparse image representation: properties and results, in SPIE Proceedings, San Diego, CA, USA, 2011, vol. 8138.
- [39] S. Martin-del-Campo, F. Sandin, Online feature learning for condition monitoring of rotating machinery, *Eng. Appl. Artif. Intell.* 64 (2017), 187–196.
- [40] K. Loparo, Bearing Vibration Data Set, Case Western Reserve University, 2003. <http://csegroups.case.edu/bearingdatacenter/>
- [41] T. Fawcett, An introduction to ROC analysis, *Pattern Recognit. Lett.* 27 (2006), 861–874.
- [42] M. Seibert, J. Wörmann, R. Gribonval, M. Kleinstueber, Learning co-sparse analysis operators with separable structures, *IEEE Trans. Signal Process.* 64 (2016), 120–130.

APPENDIX

The raw time-domain vibration signals and speed data from the six wind turbines that are analyzed for the first time in this article are hosted by the library of LTU and publicly available at the permanent link <http://urn.kb.se/resolve?urn=urn:nbn:se:ltu:diva-70730>. The sampling rate is 12.8 kHz, and the signal segments are 1.28 seconds long (16384 samples per segment). Signal segments are recorded with an interval of approximately 12 hours over a period of 46 consecutive months for each turbine. Over this time period, bearing and gearbox faults appeared in one of the six turbines, as described in Sections 3 and 4.