Kinematic Frequencies of Rotating Equipment Identified with Sparse Coding and Dictionary Learning

Sergio Martin-del-Campo[1], Fredrik Sandin[2], and Stephan Schnabel[3]

[1] *EISLAB, Luleå University of Technology, Luleå, 97187, Sweden*
SKF-LTU University Technology Center, Luleå University of Technology, Luleå, 97187, Sweden
sergio.martindelcampo@ltu.se

[2] *EISLAB, Luleå University of Technology, Luleå, 97187, Sweden*
fredrik.sandin@ltu.se

[3] *SKF, Research and Technology Development - Diagnostics and Prognostics, Luleå, 97775, Sweden*
stephan.schnabel@skf.com

Abstract

The detection of faults and operational abnormalities in rotating machine elements like rolling element bearings and gears requires information about kinematic properties, such as ball-pass and gear mesh frequencies. Typically, condition-monitoring experts obtain such information from the manufacturers for diagnostics purposes. However, the reliability of such information can be compromised during installation and maintenance, for example, if components are replaced and do not match the documented specifications. Thus, methods enabling verification and online extraction of such kinematic properties are needed to improve diagnostic reliability. Unsupervised machine learning methods, like sparse coding with dictionary learning, enable automatic modeling and characterization of repeating signal structures in the time domain, which are naturally generated by rotating equipment. Sparse coding with dictionary learning represents a vibration signal as a linear superposition of noise and atomic waveforms. The activation rate of the atomic waveforms typically possesses a cyclic nature in rotating environments, similar to how bearing kinematic frequencies correlate with faults in a rolling element bearing. However, there is no explicit relationship between the activation rates of the atoms and the bearing kinematic frequencies. This motivates this investigation of the possibility to extract bearing kinematic frequencies from sparse representations. Former work describes the use of dictionary learning for the detection of anomalies in rolling element bearings. In this paper, we describe how a similar unsupervised machine learning method can be used to extract kinematic frequencies of bearings and gears, for example for anomaly detection purposes and comparisons with an expected signature. We study the activation rates and changes of atoms learned from vibration signals in two case studies. The first case is based on data from a well-known controlled experiment with faults seeded in the bearings. The second case is based on a public dataset recorded from the high-speed shaft of a wind turbine with a bearing failure. Furthermore, we compare the activation rates and weights of the atoms to the bearing kinematic frequencies and harmonics. Sparse coding with dictionary learning offers a possibility for self-learning of the kinematic frequencies of a bearing, which can be useful for the further improvement of automated anomaly detection methods in condition monitoring.

1. Introduction

Rotating machinery depends on the structural integrity of many machine elements to enable the relative motion between the moving parts. Rolling element bearings and gears are among the most common of these components. Condition monitoring is often applied to these components for early fault identification and prediction of fault conditions, thereby extending the operational life of the machine through condition-based maintenance. However, fault detection and prediction is a challenging task because of the large number of factors that affect the performance and behavior of the machines.

Typically, experts rely on kinematic properties of the machine elements when performing vibration analysis in order to identify defect frequencies that are expected for the typical failure modes. The precise mathematical relationship between the kinematic properties of rotating components and frequency content is well documented in the literature by Randall & Antoni (2011). Frequency bands, known as fault frequencies, in the spectra of velocity, acceleration or enveloped vibration signals are trended over time. The amplitude of those frequency bands typically remains stationary under healthy conditions, while the magnitude increases when a fault develops in the component. The faults frequencies are typically associated with periodic events related to the (Patidar & Soni, 2013):

- motor shaft speed,
- ball pass frequency, inner race (BPFI),
- ball pass frequency, outer race (BPFO),
- ball (roller) spin frequency (BSF),
- fundamental train frequency – cage speed (FTF),
- gear mesh frequency.

However, by themselves, these frequencies are not sufficient to enable early identification of the wide range of faults that can occur in a rolling element bearing or gear. Increased vibration analysis robustness and reliability is achieved by taking into consideration also the harmonics and sidebands of the fault frequencies. Typically, a vibration analysis expert would investigate when one of the defect frequencies or the related harmonics reach some relative threshold, and thereby determine if the component needs to be replaced/maintained or not.

The current interest for machine learning in condition monitoring is motivated by the need to automate the vibration analysis process described above and infer more accurate predictions. The number of works that consider machine learning methods to enable early fault diagnosis in rotating machinery is increasing, see Wei et al. (2019) for a recent review. Many of the strategies proposed, use the information of the fault frequencies and their harmonics as inputs to the proposed methods. These models require that the fault frequency information is up to date after each maintenance action, which can involve component replacements that render historical kinematic information incorrect. Thus, an unsupervised machine learning approach to condition monitoring with reduced intervention of human experts requires a method that enables automatic identification of kinematic information, possibly using a catalog listing the available options.

Former work in the area of unsupervised machine learning approaches to condition monitoring has focused on the identification of faults in vibration signals. In these approaches, the fault identification is treated as a classification problem where a fault is present or not, or as a localization problem where the fault is found, for example, in the inner race, outer race or ball. Yiakopoulos et al. (2011) use the K-means clustering algorithm on frequency spectra of raw and enveloped vibration signals, and among the considered features they introduce the fault frequencies. A deep learning approach has been proposed by Jia et al. (2016). They use stacked autoencoders, which are a type of feed-forward neural network, to learn a non-linear projection of the signal spectra to differentiate between labeled health conditions. Another unsupervised machine learning approach is sparse coding. B. Liu et al. (2002) uses the matching pursuit sparse coding algorithm on vibration data from test rigs to identify the presence of localized faults. H. Liu et al. (2011) builds on the previous work by introducing an adaptive scheme called dictionary learning where several basis functions are learned and used as features into a multiclass linear discriminant classifier to identify faults in the signals.

In this paper, we apply an unsupervised feature learning method called convolutional sparse coding with dictionary learning on vibration signals recorded from a controlled experiment and sensors installed on a gearbox of a wind turbine. With this method, the vibration signal is modeled as a linear superposition of noise and atomic waveforms. Our interest here focuses on the resulting sparse representation and the possibility to infer the fault properties given the cyclic nature of a rotating machine. Sparse coding with dictionary learning is useful for online monitoring (Martin-del-Campo et al., 2013) and anomaly detection (Martin-del Campo et al., 2019). The work presented here is novel because it focuses on the learning of fault frequencies of a rolling element bearing, without prior kinematic information. We observe clear differences in the weight distribution of the sparse representation of signals corresponding to healthy and faulty conditions. Under the presence of faults in a rolling element bearing, spikes in the weight distribution of certain learned waveforms appear at the fault frequencies, which is not the case for healthy signals. These results indicate that convolutional sparse coding with dictionary learning is useful for the extraction of kinematic information about machine elements in rotating machines.

2. CONVOLUTIONAL SPARSE CODING WITH DICTIONARY LEARNING

Convolutional sparse coding with dictionary learning produces succinct representations of signals, which means that the resulting representation occupies a minimum of space but it is still useful and informative for analysis. The model we use here was developed by Smith & Lewicki (2006) and it is inspired by earlier work of Olshausen & Field (1997). The work by Smith & Lewicki (2006) describes how learned waveforms, known as *atoms*, from speech data resemble cochlear impulse response functions. The working hypothesis behind our work is that features that characterize rotating machines can be learned in a similar manner. The sparse coding model decomposes a signal $x(t)$ as a linear superposition of atomic waveforms with compact support and noise

$$x(t) = \epsilon(t) + \sum_i a_i \phi_{m(i)}(t - \tau_i). \tag{1}$$

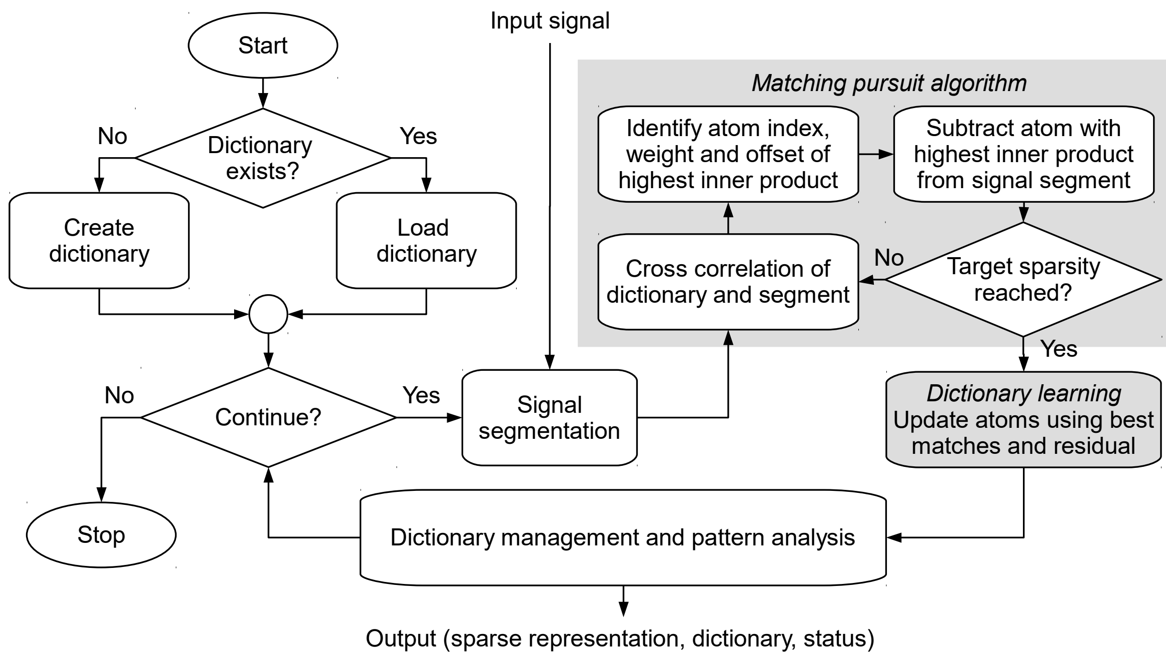

Figure 1. Convolutional sparse coding with dictionary learning method using the matching pursuit algorithm.

The functions $\phi_m(t)$ are the atoms, which represent shift-invariant morphological features of the signal, while τ_i and a_i indicate the temporal position (shift) and amplitude of each atomic *event*. The values of τ_i and a_i are calculated with the matching pursuit (MP) algorithm (Mallat & Zhang, 1993) and the triple $m(i), \tau_i, a_i$ represents one atomic event. A collection of atoms forms a *dictionary* as

$$\Phi = \{\phi_1, \cdots, \phi_M\}. \tag{2}$$

where M indicates the total number of atoms.

The optimization of the atoms follow an unsupervised approach consisting on gradient ascent of the approximate log data probability (Smith & Lewicki, 2006)

$$\frac{\partial}{\partial \phi_m} \log\left[p(x \mid \Phi)\right] = \frac{1}{\sigma_\epsilon^2} \sum_i a_i (x - \hat{x})_{\tau_i}, \tag{3}$$

where $(x - \hat{x})_{\tau_i}$ is the residual of the matching pursuit algorithm over the support of atom ϕ_m at time τ_i with an atom amplitude of a_i. This means that the shape and length of each atom are adapted from a weighted average of the residuals of the matches identified by the matching pursuit algorithm. The stop condition of the matching pursuit algorithm determines the sparseness of the representation. Note that the resulting representation is not a linear function of the input signal because the matching pursuit is non-linear.

The dictionary Φ is optimized in an iterative manner. The first step is the initialization of the dictionary. Here, we set the initial length of each atom in the dictionary to fifty and sample the initial amplitudes from a Gaussian distribution. The matching pursuit algorithm provides with the cross-correlation of the vibration signal (residual) with all atoms in the dictionary. The maximum cross-correlation defines one event, $m(i), \tau_i, a_i$, which is subtracted from the signal by subtracting the corresponding waveform, $a_i \phi_{m(i)}(t - \tau_i)$. The resulting residual is used as input to the next iteration of the matching pursuit algorithm. This process continues until the stop condition is reached. The stop condition used in this work is sparsity, which is the number of events per signal sample but the stop condition can be defined in terms of the signal-to-residual ratio as well.

The main challenge and area of opportunity of this approach is to learn the dictionary Φ, which makes it fundamentally different from other condition monitoring approaches like Fourier and wavelet analysis. We look for a dictionary of atoms Φ that maximizes the expectation of the log data probability

$$\Phi = \arg\max_\Phi \langle \log\left[p(x \mid \Phi)\right] \rangle, \tag{4}$$

where

$$p(x \mid \Phi) = \int p(x \mid a, \Phi) p(a) da. \tag{5}$$

The prior of the amplitude, $p(a)$, is defined to promote a sparse representation in terms of statistically independent atoms (Olshausen & Field, 1997). The integral is approximated with the maximum a posteriori estimate resulting from the matching pursuit algorithm. This results in a learning algorithm that involves gradient ascent on the approximate log data probability defined by Eq. (3). The gradient of each atom in the dictionary is proportional to the sum of residuals corresponding to the activation of that atom. The prefactor, $1/\sigma_e^2$, is the inverse variance of the residual that remains after matching pursuit. Additionally, we introduce a *learning rate* parameter η so that Eq. (3) is modified to

$$\Delta\phi_m = \frac{\eta}{\sigma_e^2}\sum_{i\,:\,m=m(i)} a_i(x-\hat{x})_{\tau_i}. \tag{6}$$

The resulting adaptation rates of each atom depends on the matching-pursuit activation rate, which implies that some atoms may adapt slowly or not at all. Our approach is comparable to that used by H. Liu et al. (2011) and is motivated by the relatively low complexity and simplicity of the algorithm.

Figure 1 presents a block diagram of our proposed approach, which describes the complete convolutional sparse coding with dictionary learning method. The diagram describes the stages responsible for generating the sparse representation using the matching pursuit algorithm and dictionary learning. The output includes the sparse representation of the input segment, which is used to identify the fault frequencies of rotating equipment.

3. CASE STUDIES

We apply convolutional sparse coding with dictionary learning to vibration data from two case studies. The first case study is a controlled experiment with the data taken from the bearing data center at Case Western Reserve University (Loparo, 2003). This experiment consisted of a motor, a torque transducer, and a dynamometer. The evaluated ball bearings were mounted on the motor shaft and data consists of vibration signals recorded near the drive end of the motor. The sampling frequency is 12 kHz and faults were manually introduced at the inner raceway (we consider a selected part of the data set). The second case study is based on data from a real-world condition monitoring system. The data originates from a database that collects information from condition monitoring systems installed in wind turbines located in northern Sweden. The wind turbine has a three-stage gearbox that includes two sequential planetary stages, followed by a helical gear stage. Raw time-domain vibration signals are measured by an accelerometer mounted on the housing of the output shaft bearing in the axial direction. The sampling frequency is 12.8 kHz and it is recorded in segments 1.28 seconds long with an interval of approximately 12 hours over a period of 46 consecutive months. The data is made publicly available by the Luleå University of Technology (Martin-del Campo et al., 2018).

The vibration data is processed with our MATLAB/C++ implementation of matching pursuit and the algorithm for dictionary learning by (Smith & Lewicki, 2006). In both case studies, signal segments are preprocessed to have zero mean and unit variance. The stop condition of the matching pursuit is 90% sparsity, which is comparable to a compression ratio of 0.1. In the controlled experiment case study, we use a dictionary with 16 atoms and in the wind turbine case study the dictionary contains eight atoms. During dictionary update, we use a step size of $\eta = 10^{-6}$ and the length of atoms are optimized using the method presented in (Smith & Lewicki, 2006). The atoms are normalized after each learning iteration. Further details of the testing setup of the controlled experiment are provided by Martin-del-Campo & Sandin (2017). Details of the wind turbine dataset and data collection conditions, as well as method evaluation setup is available in (Martin-del Campo et al., 2019).

3.1. Controlled Experiment

The processing of the vibration signals in the bearing data center data is carried out using a signal window of 5 seconds duration (60000 samples). The windows are randomly sampled from the different load and speed cases to simulate a time-varying load on the rotating machine. Two datasets are considered in this test, which intends to mimic the appearance and growth of a defect in the bearing, thereby simulating the evolution from a healthy state of operation to a faulty state. First, sparse coding with dictionary learning is applied to a healthy state of operation represented by 5 hours of vibration data (3600 segments of 5-seconds duration). This is referred to as the baseline (BL) case. Next, the atoms were adapted to 5 hours of data corresponding to a faulty bearing with a 7 mils (0.18 mm) diameter fault on the inner race. This case is henceforth referred to as IR7. Figure 2 shows the spectra of a vibration segment originated from the IR7 case. The dotted line in the figure points to the BPFI fault frequency. The complete list of the fault frequencies in this experiment is shown in Table 1.

Table 1. Fault frequencies in the controlled experiment.

Fault characteristic	Frequency
Motor shaft speed	30 Hz
BPFI	162 Hz
BPFO	107 Hz
BSF	70 Hz
FTF	12 Hz

In Figure 3a, we present the histogram of the cumulative weight of four atoms of the sparse representation at the end of the BL case. Frequencies are estimated from the inter spike interval (ISI) of consecutive atomic events per each atom in the dictionary. The distribution of the weights follows a normal profile that leans towards lower frequencies with a mean value centered around the motor shaft speed frequency, which

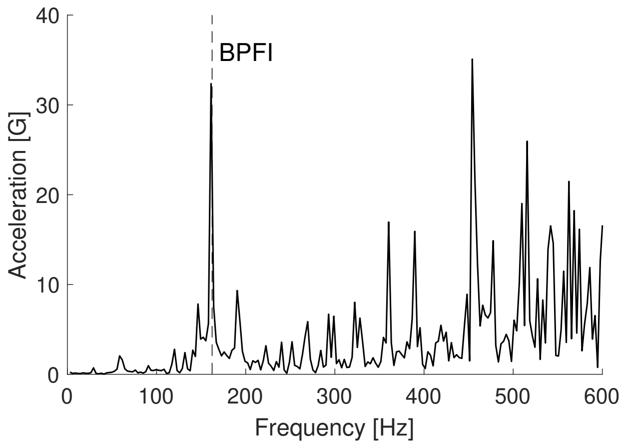

Figure 2. Frequency spectrum of a vibration signal from a ball bearing with an inner race fault of 7 mils diameter.

is represented by a dotted line. The waveform of each of these atoms is shown in Figure 3b. The learned atoms are normalized and are arranged in decreasing order of center frequency.

After the BL case, we studied the sparse representation of the IR7 case. The histograms of the cumulative weight of the sparse representation at the end of the IR7 case are shown in Figure 4a. The same four atoms described in Figure 3 are presented in Figure 4 with the intention to showcase the evolution of the histograms and atom waveform with the introduction of an inner race fault. A closer look at the histograms shows the appearance of a spike in the weight distribution at the BPFI fault frequency in atoms 15 and 16. Simultaneously, atom 14 has a spike at the motor shaft speed and atom 11 has equal height spikes at both, the shaft speed and BPFI frequencies. Notice the increase of one order of magnitude in the weights histogram amplitude between Figure 4 and Figure 3. Additionally, a visual inspection of the atoms Figure 4b shows how the atoms continued to adapt with the introduction of the fault on the inner raceway of the bearing. The remaining 12 atoms in the dictionary not shown in these figures did not present significant changes in their weight distribution between the BL and IR7 cases. Further information about this experiment, the remaining atoms not shown here, and the evolution of the waveforms is found in (Martin-del-Campo & Sandin, 2017).

3.2. Wind Turbine Data

The vibration signals used in the wind turbine case study originates from a wind farm located in Northern Sweden. The particular turbine used in this study had two bearing failures in the period of data made available to this investigation. The first fault was an inner raceway failure on a four-point ball bearing on the output shaft that resulted in the replacement of the bearing after 1.2 years of operation. The second failure occurred in one of the cylindrical roller bearings supporting

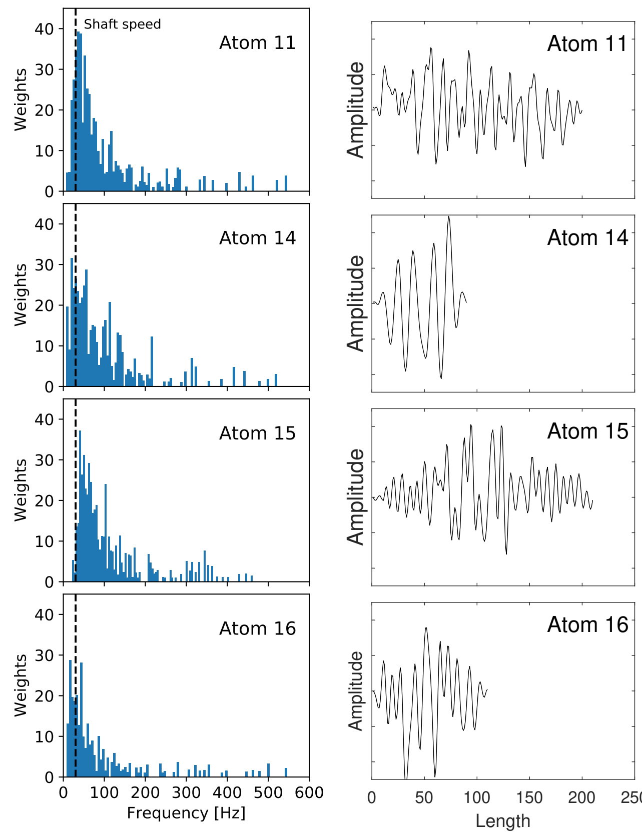

(a) Histogram (b) Waveform

Figure 3. Histograms of cumulative weight versus atom repetition frequency (a) and the learned atomic waveforms (b) for the vibration signals corresponding to the BL case. The dashed vertical line in the histogram corresponds to the shaft speed of the motor [Hz].

one of the planets in the first planetary gear of the gearbox. This failure required the replacement of the entire gearbox after two years of operation of the wind turbine. Processing of the vibration signals with the sparse coding with dictionary learning method required filtering signal segments that corresponded to an unloaded condition of the wind turbine. Afterward, a baseline dictionary was trained using 5000 segments of one-second duration (12800 samples) from a period of time where the wind turbine operated in healthy conditions. After the baseline dictionary was learned, it was propagated over time using all signal segments where the wind turbine was loaded (1907 segments total). Each signal segment is 16384 samples long and it was modeled using 1600 atomic events, which corresponds to 90% sparsity. Table 2 provides a summary of the fault characteristics of the wind turbine bearing evaluated in this study in terms of order, which is the ratio between frequency and shaft speed.

Figure 5 shows the histograms of the cumulative weight of three atoms of the sparse representation at three different

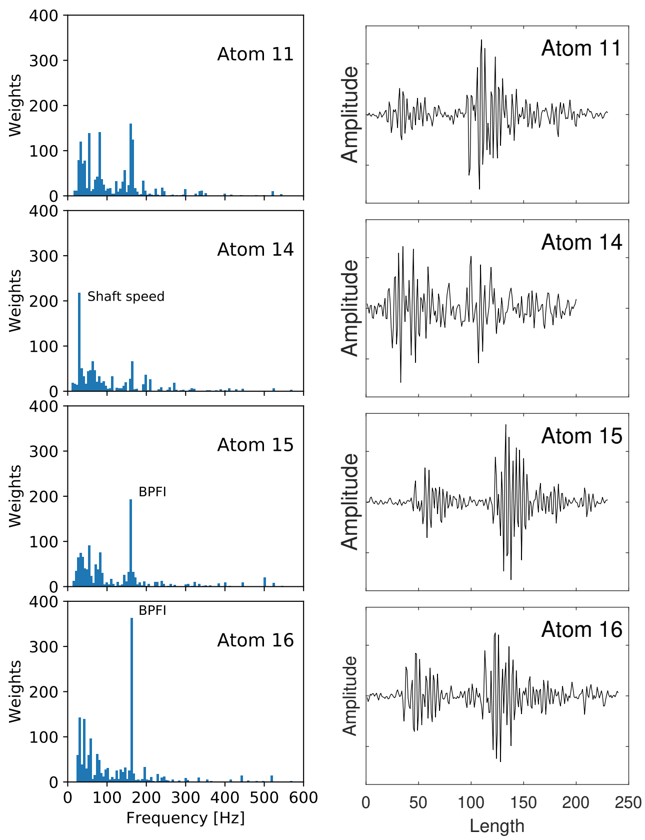

(a) Histogram (b) Waveform

Figure 4. Histograms of cumulative weight versus atom repetition frequency (a) and the learned atomic waveforms (b) for the vibration signal corresponding to the IR7 case. Note the impulse-like shape of these atoms, unlike the shape of the atoms learned in the BL case. The spikes on the histograms indicate the shaft speed of the motor and BPFI.

points in time during the operation of the wind turbine. The plots in row A correspond to a signal segment collected approximately two months prior to the replacement of the ball bearing in failure one, while row B corresponds to a segment collected a few days before the replacement of the bearing. Row C describes a healthy signal segment collected approximately six months after gearbox replacement described in failure two. Similarly to the controlled experiment case study, the order is estimated from the inter spike interval (ISI) of consecutive atomic events per each atom in the dictionary divided by the mean rotational speed at the particular signal segment. The dotted line present in all plots at order 1 corresponds to the shaft speed of the wind turbine. Atoms 1 and 2 in rows A and B have a spike in the histogram at the BPFI fault frequency that is represented by the dashed line. Simultaneously, atom 3 captures the gear mesh frequency at order 35. Once the wind turbine operates under healthy conditions, the spikes at the fault frequencies decrease as shown in row C. Notice the difference in the y-axis magnitude of the

Table 2. Fault frequencies in wind turbine.

Fault characteristic	**Order**
Motor shaft speed	1.0
BPFI	9.6
BPFO	7.4
BSF	3.7
FTF	0.4
Gear mesh	35

histogram between the faulty operational conditions and the healthy conditions. This change in the histogram magnitude is similar to the results described in the controlled experiment case study. A visual inspection of the atoms Figure 6 shows the continuous adaptation of atom 1 over time. Simultaneously, atoms 2 and 3 appear to had converged and further adaptation is not evident. The remaining five atoms in the dictionary not shown in this picture do not present significant changes between healthy and faulty conditions. Further information about this case study, together with a detailed description of gearbox schematics and the evaluation protocol used is provided by Martin-del Campo et al. (2019).

4. DISCUSSION

We study the possibility to extract the fault frequencies of rotating machine elements using unsupervised learning. We find that the distribution of the atom weights of the sparse representation of a vibration signal change between healthy and faulty operational conditions. In the presence of a fault in a rolling element bearing the weight distributions of some atoms exhibit spikes at the corresponding fault frequency, with an order of magnitude increase in the amplitude compared to the healthy case. This tendency is observed in two case studies, including real-world vibration signals collected in a condition monitoring system for wind turbines. These results motivate further improvements of the method and experiments, including different types of faults in the bearings and studies of the harmonics of the fault frequencies. An investigation of the effects of varying operational and application conditions would be beneficial to understand the capabilities and limitations of this approach. Convolutional sparse coding with dictionary learning is an interesting approach to condition monitoring automation, which requires few assumptions about the machine and the expected structure of the signal. Further work is required to integrate this approach with an anomaly detection framework that enables automatic early detection of faults in machine elements of rotating machines with reduced human intervention. Databases with kinematic information and atoms learned from signals of a population of similar machine elements can potentially be used to simplify the unsupervised optimization problem and improve the accuracy of the result.

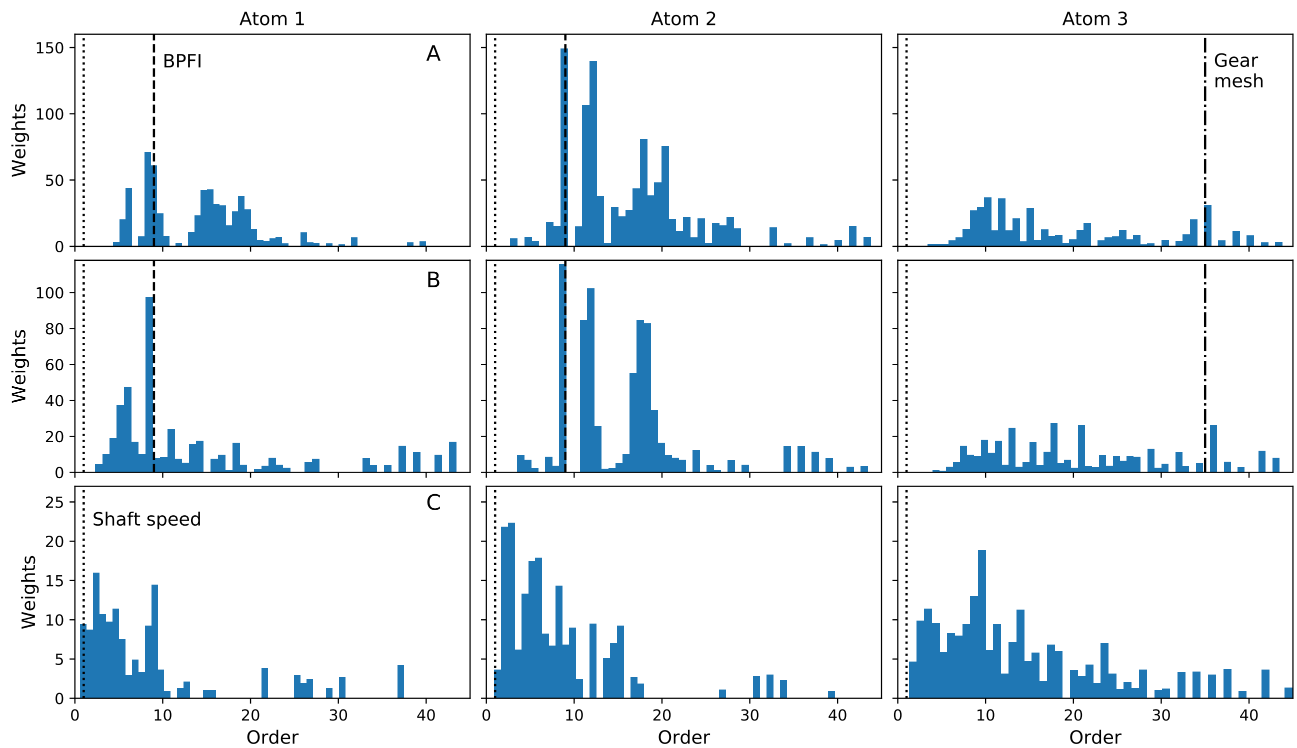

Figure 5. Histogram of cumulative weights of the sparse representation of a vibration signal from a wind turbine. Row A corresponds to a signal 3 months before bearing replacement, row B corresponds to a signal immediately before bearing replacement and row C corresponds to a signal in a healthy condition 6 months after gearbox replacement. Note the different scales on the vertical axes in the three cases.

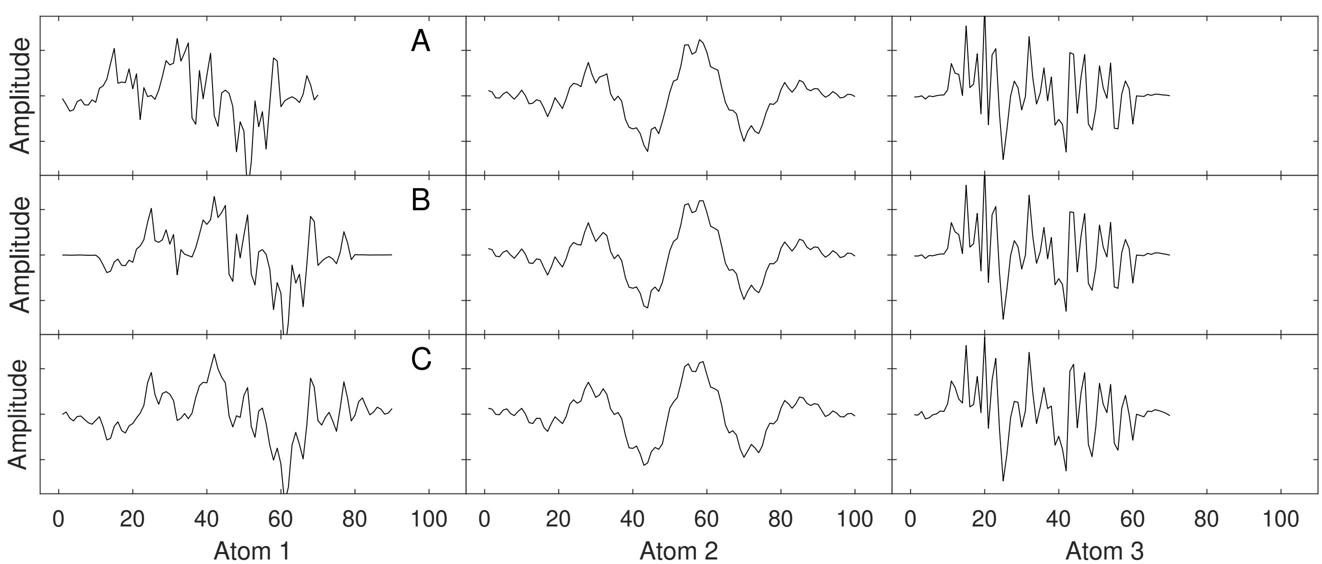

Figure 6. Learned atomic waveforms from a wind turbine. Row A corresponds to a signal 3 months before bearing replacement, row B corresponds to a signal immediately before bearing replacement and row C originates from a signal in a healthy condition 6 months after gearbox replacement.

ACKNOWLEDGMENT

This work is supported by SKF via their University Technology Center at LTU, and the Kempe Foundations.

REFERENCES

Jia, F., Lei, Y., Lin, J., Zhou, X., & Lu, N. (2016). Deep neural networks: A promising tool for fault characteristic mining and intelligent diagnosis of rotating machinery with massive data. *Mechanical Systems and Signal Processing*, *72–73*, 303–315.

Liu, B., Ling, S., & Gribonval, R. (2002). Bearing failure detection using matching pursuit. *NDT & E International*, *35*(4), 255 - 262.

Liu, H., Liu, C., & Huang, Y. (2011). Adaptive feature extraction using sparse coding for machinery fault diagnosis. *Mechanical Systems and Signal Processing*, *25*(2), 558-574.

Loparo, K. (2003). *Bearing vibration data set.* Case Western Reserve University. (http://csegroups.case.edu/bearingdatacenter/)

Mallat, S. G., & Zhang, Z. (1993, December). Matching pursuits with time-frequency dictionaries. *IEEE Transactions on Signal Processing*, *41*(12), 3397-3415.

Martin-del-Campo, S., Albertsson, K., Nilsson, J., Eliasson, J., & Sandin, F. (2013, September). 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.

Martin-del-Campo, S., & Sandin, F. (2017). Online feature learning for condition monitoring of rotating machinery. *Engineering Applications of Artificial Intelligence*, *64*, 187-196.

Martin-del Campo, S., Sandin, F., & Strömbergsson, D. (2018). *Dataset concerning the vibration signals from wind turbines in northern Sweden.* Luleå University of Technology. (http://urn.kb.se/resolve?urn=urn:nbn:se:ltu:diva-70730)

Martin-del Campo, S., Sandin, F., & Strömbergsson, D. (2019). Dictionary learning approach to monitoring of wind turbine drivetrain bearings. *arXiv preprint arXiv:1902.01426*.

Olshausen, B. A., & Field, D. J. (1997). Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vision Research*, *37*(23), 3311 - 3325.

Patidar, S., & Soni, P. K. (2013, May). An overview on vibration analysis techniques for diagnosis of rolling element bearings faults. *International Journal of Engineering Trends and Technology (IJETT)*, *4*(5), 1804-1809.

Randall, R. B., & Antoni, J. (2011). Rolling element bearing diagnostics–a tutorial. *Mechanical Systems and Signal Processing*, *25*(2), 485 - 520.

Smith, E. C., & Lewicki, M. S. (2006, February). Efficient auditory coding. *Nature*, *439*, 978-981.

Wei, Y., Li, Y., Xu, M., & Huang, W. (2019). A review of early fault diagnosis approaches and their applications in rotating machinery. *Entropy*, *21*(4).

Yiakopoulos, C., Gryllias, K., & Antoniadis, I. (2011). Rolling element bearing fault detection in industrial environments based on a k-means clustering approach. *Expert Systems with Applications*, *38*(3), 2888 - 2911.

BIOGRAPHIES

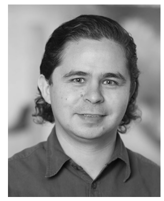

Sergio Martin-del-Campo received the B.Sc. degree in Mechatronic Engineering from Monterrey Institute of Technology and Higher Education (ITESM), Guadalajara, Mexico in 2007, the M.Sc. degree in Space Science and Technology from Luleå University of Technology (LTU), Sweden and Julius Maximilians University of Würzburg, Germany in 2012, and the Ph.D. degree in Industrial Electronics from LTU in 2017.

Currently, he is a post-doc in the Department of Computer Science, Electrical and Space Engineering at LTU, working on the use of machine learning for condition monitoring. His research interests include sparse decomposition of signals and time series analysis.

Fredrik Sandin received the M.Sc. degree in Engineering Physics from the Luleå University of Technology (LTU), Sweden, in 2001 and the Ph.D. degree in Theoretical Physics from LTU and the Swedish Graduate School of Space Technology in 2007.

He received two post-doctoral scholarships, one in Theoretical Physics (08–09) and one in Braininspired Computation (10–11), and a New-Talents award for an original work in theoretical physics from the International School of Subnuclear Physics in 2004 and the Gunnar Öquist Fellowship award from the Kempe foundations in 2014. He is currently an associate professor in Industrial Electronics at LTU and his research interests include machine learning and neuromorphic engineering.

Stephan Schnabel received the Dipl-Ing (FH) degree in Mechatronics from Ostbayerischer University of Technology (OTH), Regensburg, Germany in 2010 and the Ph.D. degree in Machine Elements from Luleå University of Technology (LTU), Sweden in 2016. He worked as researcher at LTU and guest researcher at Cornell University.

Since 2017, he is a research specialist with focus on diagnostics and prognostics of rotating machinery at SKF and his research interests include tribology, vibration-based condition monitoring and bearing dynamics.