A general anomaly detection framework for fleet-based condition monitoring of machines

Kilian Hendrickx^{a,b}, Wannes Meert^b, Yves Mollet^{a,c}, Johan Gyselinck^c, Bram Cornelis^a, Konstantinos Gryllias^{d,e}, Jesse Davis^b

^aSiemens Digital Industries Software, Interleuvenlaan 68, 3001 Leuven, Belgium

^bKU Leuven, Department of Computer Science, Celestijnenlaan 200A box 2402, 3001 Leuven, Belgium

^cUniversité Libre de Bruxelles, BEAMS, Avenue Franklin Roosevelt 50 (CP165/52), 1050 Brussels, Belgium

^dKU Leuven, Department of Mechanical Engineering, Celestijnenlaan 300, 3001 Leuven, Belgium

^eDynamics of Mechanical and Mechatronic Systems, Flanders Make, Belgium

Abstract

Machine failures decrease up-time and can lead to extra repair costs or even to human casualties and environmental pollution. Recent condition monitoring techniques use artificial intelligence in an effort to avoid time-consuming manual analysis and handcrafted feature extraction. Many of these only analyze a single machine and require a large historical data set. In practice, this can be difficult and expensive to collect. However, some industrial condition monitoring applications involve a fleet of similar operating machines. In most of these applications, it is safe to assume healthy conditions for the majority of machines. Deviating machine behavior is then an indicator for a machine fault.

This work proposes an unsupervised, generic, anomaly detection framework for fleet-based condition monitoring. It uses generic building blocks and offers three key advantages. First, a historical data set is not required due to online fleet-based comparisons. Second, it allows incorporating domain expertise by user-defined comparison measures. Finally, contrary to most black-box artificial intelligence techniques, easy interpretability allows a domain expert to validate the predictions made by the framework.

Two use-cases on an electrical machine fleet demonstrate the applicability of the framework to detect a voltage unbalance by means of electrical and vibration signatures.

Keywords: Anomaly detection, Clustering, Fleet monitoring, Condition monitoring, Electrical motors

Highlights

1. A framework is proposed for fleet-based condition monitoring
2. Online fleet-based comparisons avoid the need for a historical data set
3. It is easy for a domain expert to interact with this framework
4. A fleet of electrical drivetrains is used to experimentally validate the framework
5. The framework is benchmarked against a classic signal processing approach

1. Introduction

1.1. Context

The performance and reliability of machines are important for several industries, as in certain circumstances failures or malfunctions may lead to high production losses, severe injuries or even loss of lives. Condition monitoring allows for the early detection of machine failures, which helps to avoid further damage

Email address: kilian.hendrickx@siemens.com (Kilian Hendrickx)

and/or to reduce downtime. Many condition monitoring approaches are limited to analyzing a single machine [1]. However, some real-world industrial applications could benefit from monitoring multiple similar machines or entities (i.e., a fleet) simultaneously [2, 3]. This need arises in wind turbines [4], production lines, and aerospace engines (which often operate in pairs) [5]. Moreover, in fleet settings like monitoring the components of a train or a conveyor belt (e.g. the traction motors, bearings and bogies [6, 7]), the machines in the fleet are operating under similar conditions. A fleet-based approach leverages the fact that multiple entities are analyzed. In each of the previous examples, the different entities have comparable signature characteristics which deviate when a fault appears. A fault indicator can use this deviation.

1.2. Need for AI-based fleet monitoring approaches

The traditional condition monitoring approach involves handcrafting an indicator of the phenomena of interest, which requires significant expertise in machine signature analysis and machine usage. In deployment, the current value of the indicator is compared to a pre-selected threshold to decide on the health condition of a machine. In practice, this approach has several drawbacks. First, operational conditions can influence machine behavior, which requires either setting different thresholds [4, 8], identifying dedicated operational conditions for fault detection (e.g. ISO 10816-3:2009 [9]) or normalizing the fault indicator [10]. This approach is not scalable to a large number of operational conditions. Second, a machine can be sensitive to multiple fault types. Thus, a robust condition monitoring approach requires a manual tuning procedure for each expected machine fault and their combinations. For complex systems, it is often not possible to know all these upfront. For this reason, an alternative approach compares machine data with a simulation model of healthy behavior [11]. However, this assumes that every considered input parameter and the machine behavior does not vary over time.

Recent condition monitoring studies often suggest employing techniques from artificial intelligence as an alternative to these handcrafted solutions. In particular, the two most commonly mentioned subfields of AI are machine learning and data mining [4, 12, 13, 14]. Typically, researchers consider supervised, semi-supervised and unsupervised techniques from these two subfields.

Supervised approaches to fault detection such as Artificial Neural Networks (ANN) [15, 16, 17], Deep Learning (DL) [14, 18, 19, 20] and Support Vector Machine (SVM) [21, 22, 23] require large historical labeled data sets. That is, the data must be annotated with the true machine health condition. Moreover, they often assume the data contains examples of every possible fault type [24]. However, such data sets are not always available in industrial environments [25, 26]. In those applications, even a limited historical data set such as the [27] proposes can be challenging to obtain. As an alternative, a physical model can simulate faults [23, 28], but this is often very challenging in practice as not every possible fault can be accurately modeled.

To alleviate the need to collect data about faults, another approach is to consider semi-supervised techniques, like novelty detection [29, 30]. These techniques simply use a historical data set containing only healthy machines build a model of healthy behavior. When deployed, they detect behaviors that deviate compared to the model of normal behavior [31, 32, 33]. Consequently, it is very important that data from a faulty machine is not present in the training data. Unfortunately, this is hard to guarantee in reality as machines degrade over time. Moreover, this assumes that normal behavior does not change over time, which may not be the case. An alternative approach is to cluster historical data and assign a label to each cluster. This allows classifying new machine measurements by assigning it the label of its nearest cluster. However, this approach requires expert knowledge to provide the labels [34].

Finally, unsupervised anomaly detection can address the drawbacks of the supervised and semi-supervised approaches [35, 36, 37, 38]. These techniques require neither labeled data nor solely healthy behavior. Instead, unsupervised techniques assume that the majority of the data set belongs to the normal (healthy) class and identifies rare or infrequent behaviors as abnormal (faulty). Still, using these techniques poses two challenges in practice. First, they require selecting a threshold to determine the boundary between normal and anomalous behavior, which is challenging as these behaviors may overlap due to noise and uncertainties. Setting this parameter requires reasoning about the cost of misclassifying a faulty machine as healthy (and vice versa) and should thus be done by a domain expert. Second, in a condition monitoring context, it

is often necessary to collect an anomaly detection data set of machines in similar operational conditions, especially when a historical data set is not available.

A drawback shared by many of the existing approaches is that they employ black-box models, and hence it is difficult to gain insight into why a model made a certain prediction. Consequently, domain experts may not trust the predictions that the model is making. Moreover, it is non-trivial to incorporate expert knowledge into models such as ANNs.

1.3. Our Contribution

This paper’s contribution is an artificial intelligence based general condition monitoring approach. The framework simultaneously analyzes a fleet of machines under the assumption that each machine is in a comparable operational condition. To perform the analysis, it uses three inter-related generic blocks: comparing pairs of machines, clustering the entire fleet of machines, and identifying anomalous machines. Each block should have a meaningful outcome that can be visualized to give a domain expert an overview related to the observed machine behavior. First, we discuss several different ways to compare pairs of machines. Most interestingly, we consider a distance measure based on the amount of warping that occurs in a dynamic time warping. While measures based on Dynamic Time Warping are common in the AI literature, these receive less attention in the context of mechanics and condition monitoring. Second, we employ clustering which is a core artificial intelligence technique [39]. While the clustering can be done with a variety of algorithms, we recommend using a hierarchical method because of its simplicity and interpretability. We then propose an approach to partition a hierarchical clustering into a flat one that is novel within the context of machinery diagnostics. Finally, to identify potentially anomalous machines, we exploit the reasonable assumption that the majority of the machines in the fleet are healthy. Consequently, a good indication of a fault is when a machine’s current behavior deviates from the behavior currently exhibited by the majority of the machines. Like all condition monitoring and anomaly detection applications, this assumes that a fault alters the data generated by a machine [40]. Empirically, we consider machinery diagnostics tasks for a fleet of electrical drive trains. We examine two use-cases (electrical and vibration signatures), each with several variants that use differing amounts of domain knowledge. Applying fleet-based condition monitoring to electrical signature analysis is novel, particularly in terms of the fact that we analyze the raw time-series data. We find that our AI-based framework outperforms the more tradition condition monitoring approach the vast majority of the time.

The proposed methodology offers four advantages over existing approaches. First, unlike other approaches, the framework does not require any historical data. Instead, it makes use of online comparisons among machines. However, if they are available, recent machine measurements can improve performance by extending these comparisons. Second, the framework enables interpretable visualizations of the inter-related generic blocks. These allow a domain expert to gain insight into the framework’s predictions. Hence, this work considers relatively simple but intuitive techniques. Third, exploiting relevant domain knowledge allows tailoring the framework towards the specific use-case. For example, a measure to compare a machine pair’s behavior can be based on mechanical knowledge. This is not possible with techniques such as neural networks and deep learning. Finally, the framework can be applied directly on time-series data (i.e. raw measurements) and does not require a feature-based representation. However, the framework can use these features if available.

The rest of the paper is organized as follows. Section 2 proposes the fleet-based condition monitoring framework. Multiple possible implementations are discussed for each of the building blocks. Next, the framework is applied and evaluated on two use-cases considering a fleet of electrical drive trains. Section 4 processes the electrical signatures in three different signal domains, while section 5 analyzes vibration data. Both experimental evaluations are compared to handcrafted signal processing techniques for condition monitoring as a benchmark. Finally, section 6 discusses some general insights and conclusions.

2. General fleet-based condition monitoring framework

An anomaly detection framework for fleet condition monitoring of similar machines is proposed, which detects deviating, thus probably faulty, machine behavior. This approach can detect various faults, even

Figure 1: Interaction of the four blocks. Visualization can aid in interpreting the predictions. A user can select the machine comparison and anomaly score blocks based on domain knowledge, if available.

those not considered a priori, and does not require a training data set. The fault detection process relies on three key assumptions. First, the majority of the machines are healthy, machines whose behavior deviates from the majority are considered faulty. Second, all machines are operating in a similar operational condition, having an identical machine signature (e.g., speed). Third, the machines are subject to similar relevant environmental parameters influencing the machines’ behavior. If this were not the case, deviating behavior could arise due to a different environment instead of a machine’s health status. Practical examples where all of these assumptions hold are applications such as monitoring a conveyor belt or railway components. In both of these cases, a deviating machine is an indication of a machine fault.

Thoughtful preprocessing could allow utilizing the proposed framework in many more applications by considering external or contextual variables as well as domain knowledge. First, a heterogeneous fleet can cause machine diversity. However, expert knowledge such as an ontology could help to select comparable subsets of machines [41, 42, 43]. The framework could then assess each of these subsets individually. Second, the machine fleet can operate in dissimilar operational conditions. In such a case, two strategies could allow machine comparisons: normalization and a short-term historical database. The former exploits domain knowledge. For example, if the physics of a machine is available, then it is possible to normalize the measurements with respect to the working conditions such as by removing the effect of speed using angular resampling [10]. The latter contains recent machine measurements in various operational conditions. Clustering operational parameters could provide relevant data samples of machines in comparable operational conditions, which the framework would then use to assess their current health status’. Third, the machine fleet might be subject to different environmental parameters. Clustering machines based on these parameters could result in identifying subsets of comparable machines.

Four building blocks interact in the proposed framework: machine comparison, fleet clustering, anomaly detection, and visualization (fig. 1). The design and implementation choices of each block can be adapted to the application’s need, based on expert knowledge. The first block measures the similarity between two machine behaviors, where a variety of measures are possible. While this work discusses several concrete possibilities, a user can select an appropriate one for the application at hand based on domain knowledge. The second block uses a clustering algorithm with the selected measure to group machines that are behaving similarly. The third block uses the discovered clusters to assign an anomaly score to each machine. In a fleet where the majority of the machines are healthy, a faulty machine could occur in any cluster that consists of less than half the machines. The fourth block provides a deeper understanding of the framework by visualizing the result of the other blocks to help guide a domain expert to analyze a specific deviating machine.

The framework provides an online fleet analysis by considering small analysis windows. The required window size depends on the implementation of the first block, for example when evaluating in the frequency domain. The robustness of the prediction can be improved by considering a machine faulty based on n consecutive analysis windows.

2.1. Building block 1: Comparing two machines

Clustering uses the similarity of entities to infer a number of groups (i.e., clusters). Deciding how to compare a pair of machines (X, Y) is thus an important step in the proposed framework. Optionally, it allows exploiting engineering knowledge to guide the clustering. This block consists of two actions. First, the machine data is transformed into the desired signal domain. Second, the machines are compared in this new domain. Sequential data is represented as $(x_1, x_2, \dots, x_n)$ where x_i is one data point and n is the length of the sequence. Each machine in the fleet is assumed to have the same total number of data points n .

2.1.1. Domain specific preprocessing

Data about machines is generated by sensors that capture measurements in the time domain. Transforming the raw data to another signal domain can emphasize certain physical phenomena. The framework can support any type of transformation. This work focuses on four output domains: time, frequency, time/frequency, and a feature space.

For all domains, it is important to normalize the data before comparing machines to help remove differences that simply arise due to the unique characteristics of each machine and cope with the fact that the measurement values can originate from different mechanical units. Without normalization, a comparison between machines would be dominated by the unit with the largest range.

Time domain. Waveforms can be used to analyze time domain measurements, as they contain information about amplitude, shape, and existing symmetries. Time domain analysis is particularly interesting to detect trends and patterns over time.

Time series normalization is used to scale all machine measurements to the same range. Methods such as min-max scaling, $(X - \min(X))/(\max(X) - \min(X))$, and z-score scaling, $(X - \bar{X})/X_\sigma$, are popular, but sensitive to outliers and should be used with care. If outliers are expected, more robust estimators like percentile scaling, $(X - \text{median}(X))/(\%ile_{75}(X) - \%ile_{25}(X))$, are preferred.

Frequency domain. A Fast Fourier Transform (FFT) can be used to convert time domain data to the frequency domain, which can yield insights that are hidden in the time domain. However, it is not possible to track trends and patterns over time. Techniques such as cyclostationary analysis [44] can help to analyze rotating machinery.

The previously mentioned normalization techniques are applicable in this domain as well. To remove differences due to the unique machine characteristics, mechanical insights could be of use. For example, frequency spectra can be scaled by a specific frequency corresponding with a mechanical phenomenon or use signal processing techniques to obtain equal Parseval energy.

Time/frequency domain. The time/frequency domain detects frequency trends and patterns over time and can be shown in a spectrogram. Different normalization procedures can be applied, depending on the application's need. One can normalize each of the obtained spectra using the techniques described in the frequency domain. Another option is to scale the complete spectrogram using the techniques from the time domain.

Feature space. Various feature extraction methodologies, based on statistics, signal processing or physics, can extract indicators for fault diagnostics from any signal domain. Examples are amplitude in the time domain, order extraction or more complex signal processing techniques such as envelope analysis. The framework can use any of these indicators as a representation of machine behavior. In this case, normalization is required to equally weigh all features.

2.1.2. Dissimilarity measure

A measure¹ $s(X, Y)$ is needed to compare the distance or dissimilarity of a machine pair (X, Y) , represented by sequences $(x_1, x_2, \dots, x_n)$ and $(y_1, y_2, \dots, y_n)$ with length n . A data point x_i or y_i can be

¹For the proposed framework implementation, a measure suffices. The triangle inequality, which must hold for a metric, is not required.

1-dimensional, such as for the time domain, or m -dimensional, for example when representing a spectrogram or a feature space with m features.

Various functions to measure dissimilarity between sequences can be considered. An existing one can be selected using engineering knowledge or one can be designed specifically for the problem at hand. Two popular measures are:

1. *Euclidean distance.* The Euclidean distance is the straight-line distance between two sequences of length n , given by:

$$s(X, Y) = \sqrt{\sum_{i=1}^n c(x_i, y_i)^2} \quad (1)$$

with $c(x_i, y_i)$ being the Euclidean norm between two m -dimensional data points:

$$c(x_i, y_i) = \|x_i\|_2 - \|y_i\|_2 = \sqrt{x_{i,1}^2 + \dots + x_{i,m}^2} - \sqrt{y_{i,1}^2 + \dots + y_{i,m}^2} \quad (2)$$

A disadvantage of the Euclidean distance is its sensitivity to subtle shifts in the time series. This is solved by Dynamic Time Warping.

2. *Dynamic Time Warping.* Dynamic Time Warping (DTW) is a successful algorithm in time series analysis which aligns in a nonlinear fashion a sequence representing time series [45, 46]. The goal is to find an alignment of sequences X and Y that has the minimal total distance. The dissimilarity between X and Y is then expressed as:

$$s(X, Y) = \min_p \sqrt{\sum_{\theta=1}^L c(x_{p_{\theta,0}}, y_{p_{\theta,1}})^2} \quad (3)$$

Where $p = (p_1, p_2, \dots, p_L)$ expresses an optimal alignment with length L , also referred to as the warping path. Each p_θ represents a tuple where the first element refers to the element in the x series on position $p_{\theta,0}$ and the second element refers to the element in the y series on the $p_{\theta,1}$ position. It satisfies three conditions:

1. $p_1 = (1, 1)$ and $p_L = (n, n)$ (boundary condition)
2. $p_{1,0} \leq p_{2,0} \leq \dots \leq p_{L,0}$ and $p_{1,1} \leq p_{2,1} \leq \dots \leq p_{L,1}$ (monotonicity condition)
3. $p_{\theta+1} - p_\theta \in \{(1, 0), (0, 1), (1, 1)\}$ for $1 \leq \theta \leq L - 1$ (step size condition).

The dissimilarity of individual data points x_i and y_j is defined by a cost function $c(x_i, y_j)$, the Euclidean norm (eq. (2)) by default. The sequences should be normalized before applying DTW, to avoid suboptimal alignment due to differences in amplitude. Figure 2 shows an example in which two arbitrary time series are aligned. Note that the Euclidean distance is considerably larger compared to the DTW dissimilarity even though the series are quite similar.

A variant to DTW used in this work is Ψ -DTW. Periodic signals can be misaligned due to a phase offset, which can result in DTW introducing undesired warping at the beginning and end (fig. 3a). This additional warping can lead to an unwanted increase in the DTW measure. Ψ -DTW [47] addresses this issue and avoids the undesired warping (fig. 3b). It does this by relaxing the DTW boundary condition by $p_1 = ([1, \Psi], 1)$ or $p_1 = (1, [1, \Psi])$, and $p_L = ([n - \Psi, n], n)$ or $p_L = (n, [n - \Psi, n])$, where $\Psi \in \mathbb{N}$ is a user-defined parameter. Apart from this, it operates in a similar way as the classic DTW. The example in fig. 3 shows that for periodic signals, Ψ -DTW leads to a better match.

DTW warping amount (Ψ)-DTW is known to work well when it is important to compare time series based on their shape. However, in some cases, DTW removes the difference between healthy and faulty signals. One example is in electrical current waveforms which are considered in one of the use-cases. To overcome this challenge, a measure that considers the amount of warping that occurs in Ψ -DTW is used [48]:

(a) Original time series.

(b) Optimal aligned time series. The X-axis is hidden, as the time domain is warped.

Figure 2: DTW finds an optimal alignment between two time series (blue) by means of a warping path (orange). Euclidean distance in (a) is 7.0, the non-linear DTW in (b) has a dissimilarity of 1.0.

(a) Classic DTW. Dissimilarity: 1.8

(b) Ψ -DTW. Dissimilarity: 0.0

Figure 3: DTW on sine waves with a phase shift. Warping paths are shown in orange. While identical, the sine waves are considered dissimilar by classic DTW due to additional warping in the beginning and end. Ψ -DTW removes this effect and correctly identifies the sine waves as identical.

Figure 4: A dendrogram representing a cluster hierarchy. Agglomerative (bottom-up) clustering begins by assigning each item (blue dot) to its own cluster and iteratively merges the two most similar clusters until only one cluster is left. In this example, the horizontal distance between the items is considered being their distance.

$$s(X, Y) = \sum_{\theta=1}^{L-1} [p_{\theta+1} - p_{\theta} \neq (1, 1)]. \quad (4)$$

It is advised to use a normalized form, $s(X, Y)/L$, when using Ψ -DTW, as the length of the paths can vary.

2.2. Building block 2: Clustering the machines in the fleet

This block identifies similarly behaving machines by employing a clustering algorithm that uses the selected measure $s(X, Y)$ to compare pairs of machines. There are two important points to consider when selecting a clustering algorithm. First, the selected clustering algorithm should work for an arbitrary measure $s(X, Y)$. Some algorithms only work with certain measures (e.g., *K-means* requires a space where the mean is defined). Therefore, it is preferable to use an algorithm that represents clusters by actual data points and is hence more generic, such as K-medoids [49, 50], spectral clustering [51], hierarchical clustering [52] and DB-scan [53]. Second, the number of clusters is unknown in anomaly detection. If all the machines are healthy, a single cluster is expected whereas multiple clusters are needed if there are faulty machines. Some methods, like K-medoids, require the user to specify the number of clusters k in advance. In contrast, approaches like hierarchical clustering and DB-scan derive the number of clusters from the structure of the data, often by using a parameter that describes how tight a cluster should be.

Hierarchical clustering is proposed as the default implementation for this block. It offers the benefits of being relatively simple, producing results that can be easily visualized, and not having to set k upfront.

Hierarchical clustering. Hierarchical clustering builds a hierarchy of clusters. This can be done in either an agglomerative (bottom-up) or divisive (top-down) manner. However, agglomerative clustering is more common. This approach begins by assigning each item to its own subcluster. Then, the algorithm iteratively merges the two nearest (or most similar) subclusters until there is only one cluster left. The function to measure the distance between two subclusters is a parameter of hierarchical clustering. Two common choices are: *single linkage*, which considers the minimum distance between the elements of two subclusters and *complete linkage*, which considers the maximum. Other alternatives are *average linkage* and Ward’s method [54].

The result of the clustering can be represented in a tree-like structure called a dendrogram (fig. 4). It is relatively easy to interpret this structure, which is an additional benefit over other clustering techniques. At the lowest level, each cluster contains a single item. Going from the bottom-up, each subcluster is merged with its most similar subcluster until a single cluster is obtained. The vertical axis shows the distance between each of these subclusters. The higher this merge occurs, the less similar the subclusters are.

The proposed anomaly detection block requires that each machine is assigned to only one cluster. The hierarchical structure of the clustering needs to be partitioned to satisfy this constraint. A naive approach is to cut the hierarchy by using a threshold on the distance between subclusters. Selecting an appropriate threshold depends on $s(X, Y)$ (e.g., its range) and does not take the underlying cluster structure into account. To overcome this challenge, a threshold on the cophenetic correlation thr_{cc} is proposed. This offers the advantages of being independent of the range of $s(X, Y)$ and considering the structure of the data.

The cophenetic correlation measures how well a clustering preserves the original pairwise distances [55], that is, how well the clustering reflects the original structure of the data. A high cophenetic correlation indicates that two or more cluster partitions are present in the data, while a low cophenetic correlation

Figure 5: Upper: 10 data points form one cluster. Single linkage hierarchical cluster generates the corresponding dendrogram. There is a clear difference between pairwise and dendrogrammic distances, shown in triangular matrices with a green (low) to yellow (high) color scale. This difference leads to a low correlation. Lower: a second cluster (blue) of 3 additional data points. In this situation, the difference in pairwise and dendrogrammic distances is relatively small. This leads to a higher correlation, indicating the presence of multiple clusters.

suggests that only a single cluster partition is present in the data. Formally, the cophenetic correlation is defined as eq. (5):

$$\text{cophenetic_correlation} = \frac{\sum_{(X,Y) \in P} (s(X,Y) - \bar{s})(t(X,Y) - \bar{t})}{\sqrt{[\sum_{(X,Y) \in P} (s(X,Y) - \bar{s})^2][\sum_{(X,Y) \in P} (t(X,Y) - \bar{t})^2]}}. \quad (5)$$

where P is the set of all machine pairs (X, Y) , $s(X, Y)$ is the measure used in the clustering, and $t(X, Y)$ is the dendrogrammic distance, which is the height at which machines X and Y are first joined in the dendrogram.² Finally, $\bar{s}$ and $\bar{t}$ are the average pairwise dissimilarities and dendrogrammic distances. A pseudocode for splitting a dendrogram is given in algorithm 1.

The intuition behind the cophenetic correlation procedure is illustrated with a simple example (fig. 5). First, the upper part considers a single cluster (shown in red). There is a difference between the pairwise and dendrogrammic distances. The correlation between them is rather low. Second, the lower part of the figure shows a situation where a second cluster is present (shown in blue). Now, the distances within a cluster are relatively small. In this situation, the larger distances between clusters dominate the correlation, which is now close to 1.

Algorithm 1 Hierarchical cluster partitioning

```

1: procedure PARTITION(dendrogram, thrcc)
2:   cc = cophenetic_correlation(dendrogram) (eq. (5))
3:   if cc > thrcc then
4:     d1, d2 = cut_at_highest_level(dendrogram)
5:     return Partition(d1, thrcc), Partition(d2, thrcc)
6:   else
7:     return d

```

The cophenetic correlation threshold thr_{cc} is the most important parameter in the default implementation, as it controls the cluster partitioning. It is set depending on the cost of not identifying a faulty machine (i.e., a false negative, missed detection) or incorrectly labeling a healthy machine as faulty (i.e., a false positive, false alarm). Moreover, it determines when an alarm occurs when monitoring slowly degrading machines. If thr_{cc} is set too high, a faulty machine might be partitioned together with machines considered to be healthy. The anomaly detection procedure would then report faulty machines as being

²Note that the dendrogrammic distance depends on the selected linkage function.

healthy. Multiple small partitions containing machines considered to be healthy would be formed if thr_{cc} is set too low. Anomaly detection will then report these machines as being faulty. The value of this threshold is set depending on the application.

2.3. Building block 3: Anomaly detection

Block 3 uses the discovered clustering to assign an anomaly score to each machine. Machines with similar behavior are expected to be in the same cluster. The larger a cluster, the more likely it is to indicate healthy behavior. Each machine’s anomaly score is simply the fraction of total machines that are not assigned to its cluster, hence an anomaly score close to one indicates a machine that is likely to be faulty. Thresholding this score to classify each machine as either healthy or faulty requires reasoning about the size of the fleet and the expected proportion of faulty machines.

More sophisticated unsupervised anomaly and outlier detection techniques can be used. These can be based on clustering, neighborhood, or statistics [29, 35, 56]. Cluster-based techniques are preferred for fleet-condition monitoring applications. This class of techniques identifies anomalies as small clusters or unclustered data points. Neighborhood-based techniques like k-th Nearest Neighbor (kNN) or Local Outlier Factor (LOF) detect anomalies based on respectively the distance to the K-th nearest neighbor or the density of its neighborhood. However, the accuracy of these methods depends on the number of analyzed machines, which can be limited. Statistical-based techniques assume that anomalies arise in low-probability regions. However, selecting an appropriate hypothesis test for high-dimensional data is non-trivial. Moreover, these tests require making assumptions about the underlying data distribution, which might not hold in practice.

2.4. Building block 4: Visualization

Visualizing the results of each block can yield insights into the machine health conditions and the framework’s predictions. A simple visualization shows results of each block (e.g. fig. 6). The use-cases show more examples in figs. 12, 13 and 16.

First, both parts of the machine comparison block can give detailed insights. The template has space for preprocessed machine data, to give the user a more detailed view (fig. 6a.). A dissimilarity matrix (fig. 6b.) shows the result of pairwise machine comparisons. A color map helps to see relative differences, as absolute values depend on the selected comparison method. Second, clustering gives a more global overview of the machine fleet. A dendrogram can, for example, show a cluster hierarchy, to detect divergent machines (fig. 6c.). Moreover, colors can indicate cluster partitions (fig. 6d.). This helps to tune the threshold thr_{cc} . Finally, the anomaly scores correspond to the predicted health conditions (fig. 6e.).

3. Application and evaluation of the proposed framework on experimental use-cases

The proposed framework is validated in two different use-cases. Both involve the detection of a voltage unbalance in a fleet of electrical drivetrains. This results in early machine failure and is often caused by a grid fault or a bad wire connection. Early detection will extend the lifetime of a machine and will minimize wear due to increased temperature and additional vibrations.

Each use-case includes signal processing techniques for condition monitoring. These are used as a benchmark for the framework. Its broad applicability is shown by different variants, using different degrees of domain knowledge. Each variant makes use of different implementations for machine comparison but uses the default fleet clustering, anomaly detection, and visualization blocks. It analyzes and predicts machine status for 0.5s non-overlapping analysis windows. A machine is considered to be faulty after being identified as anomalous for five consecutive windows. This reduces misclassifications due to signal noise.

Three performance metrics consider the two types of misclassifications: false negatives (fn) and false positive (fp) (table 1). Precision (eq. (6a)) and Recall (eq. (6b)) help the user to take the cost into account of labeling a healthy machine as being faulty (precision, false alarm) versus missing a faulty machine (recall, missed detection). The F1-score (eq. (6c)) considers both faults equally important and is used as a single performance measure.

Figure 6: An example of the visualizations provided by the framework, which are analyzed in detail in section 4.3. In this example, drivetrains D1_2 and D2_10 are faulty. Signal data (a.) gives the user a detailed view. In this example, min-max normalization has been applied to the waveforms. A global overview of the fleet is given by pairwise comparisons (b.) and a dendrogram (c.). Finally, the cluster partitions (d.), obtained with $thr_{cc} = 0.9$, and their corresponding anomaly scores (e.) indicate the prediction for a machine’s health condition. In this example, the framework correctly identifies D1_2 and D2_10 as being anomalous.

The evaluation procedure considers both stationary as well as dynamic operational conditions. The former allows analyzing the consistency of the framework, the latter its generality. Within these operational conditions, the health status of machines D1_2 and D2_10 is varied.

Table 1: Possible prediction mistakes

		True condition		
		Faulty	Healthy	
Predicted condition	Faulty	True positive (tp)	False positive (fp)	Precision (eq. (6a))
	Healthy	False negative (fn)	True negative (tn)	
		Recall (eq. (6b))		F1-score (eq. (6c))

$$Precision = \frac{tp}{tp + fp} \quad (6a)$$

$$Recall = \frac{tp}{tp + fn} \quad (6b)$$

$$F1 - score = 2 \frac{Precision \cdot Recall}{Precision + Recall} \quad (6c)$$

4. Use-case 1: Voltage unbalance – Electrical signatures

This use-case involves electrical measurements of an electrical machine fleet. The task is to detect a voltage unbalance fault based on current signals. Three variants show the general applicability of the framework in different signal domains. The first variant uses time domain data and demonstrates a generic approach for waveform analysis. The second variant is an example on how to include domain knowledge. Finally, the third variant is a hybrid method and analyzes time/frequency data.

4.1. Experimental fleet test rig and data set

A single phase current of ten electrical drivetrains (fig. 7) is measured at 25600 Hz. Each drivetrain consists of an electrical motor pair, connected by a flexible jaw coupling. A 3-phase Squirrel cage induction motor (SCIM)³ drives the shaft. Ten ABB⁴ drive controllers manage the drive sides with internal closed-loop direct-torque control. All drivetrains are set up to always have similar speeds. The other motor acts as load and is either a 3 kW Direct Current (DC) motor (drivetrains D1_1 – D1_5) or a 3 kW Wound Rotor Synchronous Motor (WRSM) (drivetrains D2_6 – D2_10). The specifications of each machine are shown in table 2. Resistors generate load while keeping the current excitation constant. The resulting load torque is proportional to speed and corresponds to the rated load at rated speed.

A 3Ω external resistor R_{add} emulates a voltage unbalance in drivetrains D1_2 and D2_10. This resistor is inserted between the drive controller and the motor (fig. 7b). A higher resistance would trigger the internal safeties of the ABB drive controller.

The data set includes both stationary and dynamic system behavior under loaded and unloaded operating conditions. For the stationary case, each drivetrain is running at a fixed speed (820 RPM or 1500 RPM). On the other hand, during the dynamic operational condition, all machines perform a run-up from 0 RPM to 1200 RPM.

(a) Motor fleet.

(b) Schematic of drivetrain D1.2. The analyzed current sensor (green) measures phase B.

Figure 7: The measurement setup consists of ten drivetrains, having a 3kW SCIM (drive, inner part of the fleet) coupled to a 3kW WRSM (left) or DC (right) motor (load). Drive and load side motors are connected with a rubber coupling.

Table 2: Rated parameters of motors

	Power (kW)	Phase-to-phase voltage (V)	Current (A)	Frequency (Hz)	Pole pairs	Speed
SCIM	3	400	6.2	50	2	1385
WRSM	3	400	4.3	50	2	1500
DC	3	270	11.2	DC	1	1500

4.2. Classic voltage unbalance diagnosis techniques

Motor Current Signature Analysis (MCSA) is a popular class of techniques for detecting motor faults that apply signal processing methods to current signals. It can be used to detect a voltage unbalance [57, 58]. Some of these methods require measuring multiple motor current signals. For example, the phase shift between the three current signals of a healthy system is 120 degrees. A voltage unbalance results in changes to this angle [59]. Another example measures an estimator for I_{RMS} based on two⁵ current signals. A fluctuating component at twice the fundamental frequency is observed due to the presence of inverse sequence current in the stator [58].

³Siemens 1AV3105B 1LE10031AB522AA4

⁴ABB ACS880-01-09A4-3

⁵The third current signal can be inferred from the measured signals.

Figure 8: Voltage unbalance detection based on amplitude is not possible in a fleet context. Machine D2_10 is indicated in orange. The variance within the healthy fleet is larger than its amplitude increase due to a voltage unbalance. All machines are running stationary at 820 RPM (loaded).

Figure 9: Current waveform analysis can detect a voltage unbalance. The fault affects the current waveform of a faulty drivetrain (red). The exact shape is however dependent on the operational speed. Above 1385 RPM, the fault effect disappears due to an increase of the flux. A current signal of a healthy drivetrain (blue) is shown for reference.

A single current sensor is sensitive to a secondary effect of the voltage unbalance [58]. The oscillations of the shaft lead to an increase in the third harmonic. This work compares the framework with three approaches that make use of this signal. First, careful time domain analysis reveals an increase in the current amplitude. However, a fleet-based approach cannot use this current amplitude as the difference due to a fault is lower than the variability within the fleet (fig. 8). Moreover, the current amplitude highly depends on the operational conditions (i.e. rotational speed and load). Second, the fault affects the current waveform. However, the shape of the waveform depends on the operational speed (fig. 9). Any type of deviation should thus trigger a machine fault prediction. Finally, the frequency domain shows an increase of the third harmonic amplitude (fig. 10). The harmonic amplitude is not sensitive to changes in speed but does depend on the load. Different load conditions could thus require different thresholds.

In each of these approaches, the machine is affected up to a certain speed (~ 1400 RPM). This is demonstrated in fig. 10, where the amplitude of the third harmonic decreases at that point. The hypothesis is that at higher speeds, a flux weakening region is entered. For example at the rated load above 1385 RPM, the stator frequency exceeds the rated 50 Hz. Therefore, the flux is decreased in the machine to keep the stator voltage under its limit. As a consequence, the magnetic saturation level decreases and the third harmonic in the current becomes less visible.

The remainder of this use-case utilizes a simple condition monitoring method as a benchmark to validate the proposed framework. This method considers a machine as healthy when its third harmonic amplitude is within σ standard deviations of the fleet's mean. Note that this is a simplified fleet approach, which avoids setting thresholds manually. Table 3 shows the fault detection performance for this methodology with different values of σ , Figure 10 illustrates the specific case for $\sigma = 2$.

Figure 10: The third harmonic can be used to detect a voltage unbalance. Faulty machines D1.2 (red) and D2.10 (orange) show a clear increase compared to the healthy (blue) machines between 300 and 1420 RPM. The orange band indicates the range in which a machine is considered healthy, being 2 standard deviations around the mean. This correctly detects the machine health condition in most of the cases (table 3).

Table 3: Performance of classic voltage unbalance detection using the third current harmonic. Scores corresponding to a scenario's highest F1-score are indicated in bold.

		σ								
		Metric	0.5	1	1.5	2	2.5	3	3.5	4
820 RPM - load	Precision	0.174	0.236	0.289	0.323	0.334	0.334	0.334	0.334	
	Recall	0.857	0.857	0.845	0.793	0.758	0.695	0.641	0.596	
	F1	0.289	0.371	0.431	0.459	0.464	0.451	0.440	0.428	
820 RPM - no load	Precision	0.157	0.213	0.259	0.285	0.297	0.297	0.297	0.297	
	Recall	0.855	0.855	0.855	0.800	0.770	0.706	0.652	0.605	
	F1	0.265	0.341	0.397	0.421	0.429	0.418	0.408	0.399	
1500 RPM - load	Precision	0.145	0.164	0.172	0.172	0.172	0.172	0.172	0.172	
	Recall	0.623	0.523	0.422	0.337	0.307	0.281	0.259	0.241	
	F1	0.235	0.250	0.244	0.228	0.220	0.213	0.207	0.201	
1500 RPM - no load	Precision	0.142	0.190	0.236	0.273	0.283	0.283	0.283	0.283	
	Recall	0.855	0.855	0.855	0.832	0.796	0.730	0.673	0.625	
	F1	0.243	0.311	0.370	0.411	0.418	0.408	0.399	0.390	
Run up - load	Precision	0.245	0.325	0.390	0.444	0.465	0.465	0.465	0.465	
	Recall	0.833	0.827	0.824	0.822	0.815	0.747	0.689	0.640	
	F1	0.378	0.467	0.530	0.577	0.592	0.573	0.556	0.539	
Run up - no load	Precision	0.232	0.302	0.361	0.410	0.427	0.427	0.427	0.427	
	Recall	0.821	0.786	0.768	0.757	0.737	0.676	0.624	0.579	
	F1	0.361	0.436	0.491	0.532	0.540	0.523	0.507	0.491

Figure 11: DTW of a healthy (blue) and faulty (red) machine. While the measured current signals have a clear difference in shape (a), DTW can non-linearly align them so the signals become similar (b).

4.3. Variant 1: Time domain

This variant evaluates the different operational conditions in the time domain representation. The goal is to detect any deviation on the current waveform with a generic comparison measure.

4.3.1. Block 1: Machine comparison

Due to the fact that the differences appearing in current amplitudes cannot be used in a fleet-based approach (fig. 8), min/max normalization is performed. Moreover, the current signals are downsampled to 50 samples per period. This removes the effect of signal noise and allows the analysis to focus on the waveforms.

The waveform of a machine affected by a voltage unbalance depends on the operational speed. A robust technique is thus required that detects any difference in this shape. DTW is a popular technique for pattern recognition, but the classic implementation considers waveforms of healthy and faulty machines as similar (fig. 11). However, aligning healthy and faulty signals involves significant amounts of warping. In contrast, little warping occurs when comparing two healthy machines. Therefore, the amount of warping (eq. (4)) can be used as a powerful indicator for waveform deviations, regardless of the exact shape. The Ψ -DTW extension avoids unnecessary warping in the boundary points. Unaligned start- and endpoints would otherwise dominate the DTW warping amount measure. The used value of Ψ is corresponding to half a period of the current signal, allowing to align any phase offset.

4.3.2. Block 2: Fleet clustering

Hierarchical clustering, combined with a cophenetic correlation partition procedure, generates machine clusters. First, hierarchical clustering uses pairwise machine comparisons and single linkage to construct a cluster hierarchy. Single linkage considers the minimum distance as the similarity of two subclusters. The cophenetic correlation procedure partitions this hierarchy with recursive top-down evaluations, if the cophenetic correlation is larger than thr_{cc} (algorithm 1).

The value of thr_{cc} affects the effectiveness and the performance of the method but its optimal value is application dependent. Therefore, the framework is evaluated for various values of thr_{cc} to show the parameter's impact on predictive performance. A lower value results in more clusters, which leads to more machines incorrectly labeled as faulty (low precision, higher false alarm rate). A higher value, on the other hand, leads to more misses of faulty machines (low recall, higher missed detection rate).

4.3.3. Block 3: Anomaly detection

The anomaly score is defined as the fraction of machines outside each cluster. This is a simple scoring mechanism to demonstrate the concept of anomaly detection. At least $2/3$ of the machines are assumed to be healthy. The anomaly threshold thr_{ad} , above which a machine is predicted as faulty, is thus set equal to $2/3$. If required, more advanced anomaly detection techniques can be used.

4.3.4. Block 4: Visualization

Visualizations allow a domain expert to interpret the framework’s predictions. By showing the results of each block, an expert can confirm the predictions and avoid taking unneeded actions. All use-cases use the default implementation of section 2.4, being a (normalized) signal domain representation, machine dissimilarity, clustering, and anomaly scoring.

4.3.5. Results

The example of fig. 6 shows the evaluation of a single analysis window. A domain expert can use this visualization to validate the predictions or to tune the general framework. The time series show deviating behavior of machines D1_2 and D2_10 in each of the subfigures corresponding to the different building blocks.

The evaluation procedure uses performance metrics to evaluate the framework’s ability to detect a deviating waveform (table 4). Both the benchmark (table 3) and this implementation are able to detect an unbalance, except at 1500 RPM in a loaded condition due to the decrease of the flux. In other stationary scenarios, this implementation outperforms the benchmark method. Similar results are obtained in dynamic run-up scenarios. However, this implementation of the framework offers the advantage that it detects any deviations in the current waveforms, not only caused by an increase of the third harmonic.

Table 4: Performance of the proposed framework for fault detection using electrical signature analysis via waveform comparison. Scores corresponding to a scenario’s highest F1-score are indicated in bold. For reference, the last column shows the performance of the classic diagnostics approach (table 3) with the highest F1-score. In five out of the six cases, the proposed approach obtains the best F1 score.

		thr_{cc}						Classic diagnostics approach
	Metric	0.5	0.7	0.8	0.85	0.9	0.95	
820 RPM - load	Precision	0.218	0.542	0.696	0.932	1.000	1.000	<i>0.334</i>
	Recall	0.857	0.857	0.857	0.857	0.857	0.812	<i>0.758</i>
	F1	0.348	0.664	0.768	0.893	0.923	0.897	<i>0.464</i>
820 RPM - no load	Precision	0.189	0.573	0.922	1.000	1.000	1.000	<i>0.297</i>
	Recall	0.855	0.855	0.855	0.855	0.855	0.855	<i>0.770</i>
	F1	0.309	0.686	0.887	0.922	0.922	0.922	<i>0.429</i>
1500 RPM - load	Precision	0.135	0.043	0.000	1.000	1.000	1.000	<i>0.164</i>
	Recall	0.527	0.027	0.000	0.000	0.000	0.000	<i>0.523</i>
	F1	0.215	0.034	0.000	0.000	0.000	0.000	<i>0.250</i>
1500 RPM - no load	Precision	0.382	0.959	1.000	1.000	1.000	1.000	<i>0.283</i>
	Recall	0.855	0.855	0.855	0.855	0.818	0.245	<i>0.796</i>
	F1	0.528	0.904	0.922	0.922	0.900	0.394	<i>0.418</i>
Run up - load	Precision	0.479	0.697	0.821	1.000	1.000	1.000	<i>0.465</i>
	Recall	0.852	0.852	0.852	0.852	0.852	0.778	<i>0.815</i>
	F1	0.613	0.767	0.836	0.920	0.920	0.875	<i>0.592</i>
Run up - no load	Precision	0.118	0.686	1.000	1.000	1.000	1.000	<i>0.427</i>
	Recall	0.857	0.857	0.679	0.643	0.643	0.643	<i>0.737</i>
	F1	0.207	0.762	0.809	0.783	0.783	0.783	<i>0.540</i>

4.4. Variant 2: Frequency domain

This variant demonstrates the case in which detailed domain knowledge is included. A voltage unbalance affects the third harmonic of the fundamental frequency, which is used in this variant to compare the machines. In general, this is very similar to the benchmark method, as both depend on domain knowledge.

4.4.1. Block 1: Machine comparison

A FFT converts time series within an analysis window to their frequency data, which is converted to a log scale and normalized using min/max scaling. No outliers are present in the considered data set, other cases might need more robust techniques. Machines are represented in the feature space, by the amplitude or their current’s third harmonic. This increases if one of the machines is faulty, as shown in section 4.2.

The periodic nature of current signals allow for estimating the fundamental frequency. This corresponds to the highest peak in the frequency spectrum. The third harmonic is found at three times this frequency. The maximum value in a small window of $\pm 5\text{Hz}$ around the estimated frequency is used as its amplitude.

This variant compares a machine pair by the difference in their third harmonic amplitudes (eq. (7)).

$$s(X, Y) = |\text{harmonic}(3, X) - \text{harmonic}(3, Y)| \quad (7)$$

4.4.2. Blocks 2 – 4: Fleet clustering, Anomaly detection & Visualization

Clustering, anomaly detection, and visualization make use of the default implementations. Hierarchical clustering combined with cophenetic correlation clusters the fleet. Therefore, the framework is evaluated for various values of thr_{cc} to show the parameter’s impact on predictive performance. Anomaly scores use these clusters and are defined as the fraction of machines outside each cluster. The implementation considers a score above $2/3$ as anomalous behavior. The visualization shows the frequency spectra enriched with the harmonic amplitudes windows.

Figure 12: Visualization of frequency domain signals (a.), machine similarities (b.), clustering (c. & d.) and anomaly scores (e.). The red vertical line in (a.) indicate the small windows of 10Hz around the fundamental and third harmonic frequencies. Its amplitudes are determined by each window's maximum amplitude. Clusters are partitioned with $thr_{cc} = 0.9$. Each machine is running stationary at 820 RPM. The framework correctly identifies D1_2 and D2_10 as anomalous.

4.4.3. Results

Figure 12 shows the evaluation of a single analysis window. The anomalous behavior of D1_2 and D2_10 is successfully detected.

The evaluation procedure uses performance metrics to evaluate the framework's ability to detect deviating third harmonic amplitudes (table 5). This framework implementation has similar performance compared to the previous variant (table 4). It outperforms the benchmark method (table 3) in almost all stationary operational conditions. At 1500 RPM in loaded condition, both methods show poor behavior. In the dynamic run-up scenarios, on the other hand, the benchmark method is very competitive. It shows a higher recall (less missed detections), but its precision is slightly lower (more false alarms). This results in F1-scores in favor of the benchmark method.

Table 5: Performance of the proposed framework for fault detection using electrical signature analysis via harmonic amplitudes. Scores corresponding to a scenario’s highest F1-score are indicated in bold. For reference, the last column shows the performance of the classic diagnostics approach (table 3) with the highest F1-score. In five out of the six cases, the proposed approach obtains the best F1 score.

		thr_{cc}						Classic diagnostis approach
Metric		0.5	0.7	0.8	0.85	0.9	0.95	
820 RPM - load	Precision	0.097	0.287	0.865	1.000	1.000	1.000	<i>0.334</i>
	Recall	0.857	0.857	0.857	0.857	0.643	0.045	<i>0.758</i>
	F1	0.174	0.430	0.861	0.923	0.783	0.085	<i>0.464</i>
820 RPM - no load	Precision	0.087	0.307	0.895	1.000	1.000	1.000	<i>0.297</i>
	Recall	0.855	0.855	0.855	0.855	0.627	0.009	<i>0.770</i>
	F1	0.158	0.452	0.874	0.922	0.771	0.018	<i>0.429</i>
1500 RPM - load	Precision	0.098	0.118	0.250	1.000	1.000	1.000	<i>0.164</i>
	Recall	0.855	0.309	0.009	0.000	0.000	0.000	<i>0.523</i>
	F1	0.176	0.171	0.018	0.000	0.000	0.000	<i>0.250</i>
1500 RPM - no load	Precision	0.097	0.250	0.712	0.887	0.979	1.000	<i>0.283</i>
	Recall	0.855	0.855	0.855	0.855	0.855	0.855	<i>0.796</i>
	F1	0.175	0.387	0.777	0.870	0.913	0.922	<i>0.418</i>
Run up - load	Precision	0.100	0.535	1.000	1.000	1.000	1.000	<i>0.465</i>
	Recall	0.852	0.852	0.741	0.741	0.444	0.259	<i>0.815</i>
	F1	0.179	0.657	0.851	0.851	0.615	0.412	<i>0.592</i>
Run up - no load	Precision	0.100	0.218	1.000	1.000	1.000	1.000	<i>0.427</i>
	Recall	0.857	0.857	0.607	0.429	0.429	0.000	<i>0.737</i>
	F1	0.179	0.348	0.756	0.600	0.600	0.000	<i>0.540</i>

4.5. Variant 3: Time/Frequency domain

This variant demonstrates a hybrid approach, in which a domain expert expects a fault to manifest itself in the frequency content, without specifying the exact frequency. It enables fault detection in a large range of working conditions. In this case, the framework is implemented to detect deviations in the spectrogram, a time/frequency domain representation.

4.5.1. Block 1: Machine comparison

Each machine is represented by its spectrogram, a sequence of FFTs obtained from consecutive windows. This implementation converts the data on a log scale and normalizes using min/max scaling FFT to have a similar value ranges for each machine. A low-pass filter at 200 Hz removes unrelated high frequency content.

DTW (eq. (3)) is preferred as dissimilarity measure over Euclidean distance. If the drivetrain’s speed fluctuates, DTW will optimally align the spectrograms. In that case, these are considered as n -dimensional time series, with n the number of frequency bins. Each of these bins is equally weighted in this implementation. If desired, a user could opt to have different weighting for frequency bins of interest.

A high frequency resolution can result in false predictions due to the curse of dimensionality [60]. In that case, noise over unaffected frequencies would be larger than the amplitude increase in a single frequency bin. The task is however only to perform fault detection. Since the exact frequency at which a deviation occurs is not of interest, it is safe to have a low frequency resolution. In this variant, the selected FFT window size is 0.05s.

4.5.2. Blocks 2 – 4: Fleet clustering, Anomaly detection & Visualization

This variant reused the default clustering, anomaly detection, and visualization implementations. The first block uses hierarchical clustering and cophenetic correlation. Therefore, the framework is evaluated for various values of thr_{cc} to show the parameter’s impact on predictive performance. The second considers the fraction of machines outside each cluster as anomaly scores. A score above $2/3$ is considered as anomalous behavior. Finally, the visualization presents the signal data as a spectrogram.

4.5.3. Results

Figure 13 shows the evaluation of a single analysis window. A low-resolution spectrogram shows different machine behavior for faulty machines D1_2 and D2_10. Moreover, the dendrogram suggests using a lower value for thr_{cc} , as the healthy machine D2_7 appears in its own cluster.

The evaluation procedure uses performance metrics to evaluate the framework’s ability to detect deviating spectrograms (table 6). It outperforms the benchmark method (table 3) in most stationary scenarios. None of the methods detect an unbalance at 1500 RPM in a loaded condition. In the dynamic run-up scenarios, on the other hand, both methods have similar results. In general, this variant offers the advantage that it does not require domain knowledge or handcrafted rules.

Figure 13: Visualization of time/frequency domain signals (a.), machine similarities (b.), clustering (c. & d.) and anomaly scores (e.). The frequency resolution in (a.) is coarse, to avoid the curse of dimensionality while clustering. Clusters are partitioned with $thr_{cc} = 0.9$. Each machine is running stationary at 820 RPM. The framework correctly identifies D1.2 and D2.10 as anomalous and incorrectly partitions D2.7 by itself. This suggests that the choice of thr_{cc} might not be optimal.

Table 6: Performance of the proposed framework for fault detection using electrical signature analysis via spectrogram analysis. Scores corresponding to a scenario's highest F1-score are indicated in bold. For reference, the last column shows the performance for the classic diagnostics approach (table 3) with the highest F1-score. In five out of the six cases, the proposed approach has the highest F1 score.

		thr_{cc}						Classic diagnostic approach
Metric		0.5	0.7	0.8	0.85	0.9	0.95	
820 RPM - load	Precision	0.157	0.505	0.738	0.906	0.990	1.000	0.334
	Recall	0.857	0.857	0.857	0.857	0.857	0.768	0.758
	F1	0.265	0.636	0.793	0.881	0.919	0.869	0.464
820 RPM - no load	Precision	0.236	0.913	1.000	1.000	1.000	1.000	0.297
	Recall	0.855	0.855	0.809	0.809	0.618	0.091	0.770
	F1	0.370	0.883	0.894	0.894	0.764	0.167	0.429
1500 RPM - load	Precision	0.106	0.198	0.000	1.000	1.000	1.000	0.164
	Recall	0.627	0.155	0.000	0.000	0.000	0.000	0.523
	F1	0.182	0.173	0.000	0.000	0.000	0.000	0.250
1500 RPM - no load	Precision	0.208	0.839	1.000	1.000	1.000	1.000	0.283
	Recall	0.855	0.855	0.855	0.636	0.118	0.000	0.796
	F1	0.335	0.847	0.922	0.778	0.211	0.000	0.418
Run up - load	Precision	0.178	0.487	0.864	1.000	1.000	1.000	0.465
	Recall	0.852	0.704	0.704	0.630	0.593	0.333	0.815
	F1	0.295	0.576	0.776	0.773	0.744	0.500	0.592
Run up - no load	Precision	0.221	0.900	0.900	1.000	1.000	1.000	0.427
	Recall	0.679	0.643	0.643	0.643	0.393	0.000	0.737
	F1	0.333	0.750	0.750	0.783	0.564	0.000	0.540

Figure 14: A 3D accelerometer (indicated in red) is mounted on the drive side motor and is sampled at 12800Hz.

5. Use-case 2: Voltage unbalance – Vibration signatures

This use-case involves vibration measurements of an electrical machine fleet, similar to the previous use-case. The task is to detect a voltage unbalance fault based on accelerometer signals. Accelerometers are often already present for condition monitoring of other (mechanical) components, hence it is valuable if they can also be used for voltage unbalance detection. This section illustrates that it is indeed feasible and moreover shows how the framework can use multiple sensor channels and indicators of interest.

5.1. Experimental fleet test rig and data set

This use-case considers a setup of ten electrical drivetrains. A voltage unbalance is again emulated with an external resistor. PCB ICP AC accelerometers measure vibrations. One sensor is mounted on the side of the drive side motor (fig. 14) and is sampled at 12800 Hz. Operational speed is measured with a single tachometer mounted on D1.2, triggered once every rotation. This is representative for the full fleet, as all drivetrains are set up to always have the same speed, controlled by an ABB drive controller. One exception is machine D2.8, whose run-up parameters are set incorrectly. For this reason, this drivetrain is not used in the evaluation procedures when considering a run-up.

5.2. Classic voltage unbalance diagnosis techniques

For a healthy, symmetrical, motor, electrodynamic forces have no alternating components. In the balanced case, the forces are rotating with the flux. In the case of unbalance they have a component rotating in the opposite direction. This results in a pulsation of the amplitude of the forces at twice the stator frequency. This causes radial vibrations in the induction motor [58, 61, 62]. The forces have a tangential component (creating the torque) but also a radial component (the stator flux attracts the rotor like an electromagnet) generating the vibrations.

Detecting voltage unbalance based on vibration signatures is an example of a condition monitoring case that requires different thresholds, as operational conditions affect vibration amplitudes. For example, additional load or operating at a resonance frequency leads to an increase, while lower operating speeds result in lower vibration levels (fig. 15). Moreover, many frequencies and sensor directions can be considered. Detailed machine knowledge is required to know these frequencies a priori.

The machine response varies for D1.2 and D2.10, the faulty drive trains (fig. 15). Their resonance frequencies differ and in loaded conditions, D2.10 is significantly more affected by the fault compared to D1.2. The optimal indicator depends thus not only on operational conditions but also on the evaluated machine.

The remainder of this use-case utilizes a simple condition monitoring method as a benchmark to validate the proposed framework. This method considers a machine as healthy when the amplitude of all its harmonics 3 – 6 is within σ standard deviations of the fleet's mean. If one harmonic is outside this range, the machine is predicted as being faulty. Note that this is a simplified fleet approach, which avoids setting threshold manually. Table 7 shows the fault detection performance for this methodology with different values of σ , Figure 15 illustrates the specific case for $\sigma = 2$.

Figure 15: Harmonics 1 (top) – 6 (bottom) of the X-axis accelerometer in a run-up. The faulty drivetrains are shown in blue (D1_2) and red (D2_10). The optimal indicator to detect the voltage unbalance depends on the machine of interest and the operational speed. The orange band indicates the range in which a machine is considered healthy, being 2 standard deviations around the mean.

Table 7: Performance of classic voltage unbalance detection using accelerometer harmonics 3 – 6. Scores corresponding to a scenario's highest F1-score are indicated in bold.

	Metric	σ							
		0.5	1	1.5	2	2.5	3	3.5	4
820 RPM - load	Precision	0.110	0.143	0.168	0.191	0.201	0.201	0.201	0.201
	Recall	0.871	0.871	0.817	0.784	0.772	0.708	0.653	0.607
	F1	0.195	0.246	0.279	0.307	0.319	0.314	0.308	0.302
820 RPM - no load	Precision	0.104	0.116	0.125	0.132	0.133	0.133	0.133	0.133
	Recall	0.871	0.753	0.648	0.569	0.526	0.483	0.445	0.414
	F1	0.185	0.201	0.210	0.214	0.212	0.208	0.205	0.201
1500 RPM - load	Precision	0.098	0.101	0.107	0.107	0.106	0.106	0.106	0.106
	Recall	0.847	0.730	0.623	0.510	0.463	0.425	0.392	0.364
	F1	0.176	0.177	0.183	0.176	0.173	0.170	0.167	0.165
1500 RPM - no load	Precision	0.098	0.107	0.115	0.119	0.119	0.119	0.119	0.119
	Recall	0.871	0.837	0.736	0.620	0.564	0.517	0.477	0.443
	F1	0.176	0.189	0.199	0.199	0.196	0.193	0.190	0.187
Run up - load	Precision	0.246	0.301	0.349	0.367	0.368	0.368	0.368	0.368
	Recall	0.817	0.756	0.705	0.611	0.558	0.512	0.473	0.439
	F1	0.378	0.431	0.467	0.458	0.444	0.428	0.414	0.400
Run up - no load	Precision	0.249	0.295	0.322	0.332	0.332	0.332	0.332	0.332
	Recall	0.848	0.789	0.670	0.562	0.511	0.469	0.433	0.402
	F1	0.385	0.429	0.434	0.418	0.403	0.389	0.376	0.364

5.3. Variant 1: Frequency domain

This variant demonstrates the use of multiple sensors. The implementation combines each of the accelerometer directions (X, Y, Z) in the comparison measure. Moreover, it considers domain knowledge by multiple harmonic frequencies for each of these directions. A data-driven approach automatically selects the best indicator value.

5.3.1. Block 1: Machine comparison

Machines are represented by harmonic frequencies 3 – 6 of each sensor direction. The implementation uses the tachometer to define the fundamental frequency. A FFT converts the time domain measurements to frequency spectra, in which the maximum value within a window of $\pm 5\text{Hz}$ around the estimated frequencies is considered as their amplitude. Together, these form a sequence with length 1 and 12 dimensions (3 sensor directions, 4 harmonics).

Accelerometer signals can be noisy and require a robust normalization procedure. This variant uses percentile scaling, as data can be affected by outliers. In this case, normalization also makes sure to equally weight every harmonic. Otherwise, a harmonic with a larger range could dominate the clustering.

The Euclidean distance measures the differences between a machine pair (eq. (1)).

5.3.2. Blocks 2 – 4: Fleet clustering, anomaly detection & visualization

This variant reuses the default clustering, anomaly detection, and visualization blocks. The first combines hierarchical clustering with a cophenetic correlation procedure. Therefore, the framework is evaluated for various values of thr_{cc} to show the parameter’s impact on predictive performance. The second uses the fraction of machines in each cluster as the anomaly score. A score above $2/3$ is considered as anomalous behavior. Finally, the frequency spectra of each drivetrain are shown combined with the considered harmonic windows.

5.3.3. Results

Figure 12 shows the evaluation of a single analysis window. The anomalous behavior of D1.2 and D2.10 is successfully detected in each block.

The evaluation procedure uses performance metrics to evaluate the framework’s ability to detect faults based on multiple sensor directions and harmonics (table 8). In loaded scenarios, it outperforms the benchmark method (table 7). Both approaches show poor results in unloaded scenarios, which can be explained by the reduced vibrations in the absence of a load. At 1500 RPM, none of the methods detect the fault due to a decrease of the flux. Lower general performance (compared to electrical signature analysis, section 4) in dynamic run-up scenarios is caused by the fault not manifesting itself in the vibration signature at every operational speed. In general, this implementation offers the advantage that it can consider multiple harmonics, without needing individual threshold values. This is an advantage especially when the optimal indicator varies with the operational conditions.

Figure 16: Visualization showing the three accelerometer frequency domain signals (a.), machine similarities (b.), clustering (c. & d.) and anomaly scores (e.). The red vertical lines in (a.) indicate the used harmonics 3 – 6. Clusters are partitioned with $thr_{cc} = 0.9$. Each machine is running stationary at 820 RPM, with D1.2 and D2.10 having a voltage unbalance. D1.2 is less affected by the voltage unbalance compared to D2.10, which is visible in the dendrogram.

Table 8: Performance of the proposed framework for fault detection using vibration signature analysis via harmonic amplitudes. For reference, the last column shows the performance for the classic diagnostics approach (table 7) with the highest F1-score. In five out of the six cases, the proposed approach has the highest F1 score.

		thr_{cc}						Classic diagnostics approach
Metric		0.5	0.7	0.8	0.85	0.9	0.95	
820 RPM - load	Precision	0.121	0.462	0.545	0.565	0.610	0.908	0.201
	Recall	0.871	0.871	0.871	0.871	0.871	0.871	0.772
	F1	0.213	0.603	0.671	0.686	0.718	0.889	0.319
820 RPM - no load	Precision	0.097	0.153	0.185	0.196	0.205	0.163	0.132
	Recall	0.629	0.452	0.435	0.435	0.435	0.266	0.569
	F1	0.168	0.229	0.260	0.270	0.278	0.202	0.214
1500 RPM - load	Precision	0.091	0.004	0.000	0.000	0.000	0.000	0.107
	Recall	0.508	0.008	0.000	0.000	0.000	0.000	0.623
	F1	0.154	0.005	0.000	0.000	0.000	0.000	0.183
1500 RPM - no load	Precision	0.105	0.133	0.178	0.172	0.116	0.000	0.119
	Recall	0.702	0.435	0.379	0.258	0.129	0.000	0.620
	F1	0.183	0.204	0.242	0.206	0.122	0.000	0.199
Run up - load	Precision	0.345	0.710	1.000	1.000	1.000	1.000	0.349
	Recall	0.857	0.786	0.696	0.536	0.536	0.429	0.705
	F1	0.492	0.746	0.821	0.698	0.698	0.600	0.467
Run up - no load	Precision	0.240	0.508	0.469	0.469	0.200	0.143	0.322
	Recall	0.857	0.589	0.411	0.411	0.107	0.071	0.670
	F1	0.375	0.545	0.438	0.438	0.140	0.095	0.434

6. Conclusion

Machine fleet condition monitoring offers several advantages over both traditional condition monitoring approaches and supervised machine learning techniques. First, it removes the need to know all potential machine faults a priori, as the fleet-based approach can detect any deviation in machine behavior. However, the framework can incorporate domain knowledge if available. Second, fleet monitoring does not require a high-quality historical (labeled) data set for training. Both traditional signal processing approaches as supervised machine learning need this data to learn faulty machine behavior. Moreover, these require ground-truth knowledge about the machine’s actual health status. This is not required for fleet monitoring, as it assumes the majority of the machines to be healthy. Third, fleet monitoring allows analysis in dynamic operational conditions. Online comparisons allow detecting deviating machine behavior even in unconsidered operational conditions. Finally, the framework offers a high level of interpretability. Visualizations allow a domain expert to get insights into the predictions and gain confidence in the methodology. This is especially in contrast with black box machine learning techniques, whose models are very hard to interpret.

Acknowledgments

The authors acknowledge the financial support of VLAIO (Flemish Innovation & Entrepreneurship) through the Baekeland PhD mandate [nr. HBC.2017.0226]; the O&O project REFLEXION [nr. IWT. 150334]. Jesse Davis is partially supported by the KU Leuven research funds [C14/17/070], Research Foundation - Flanders [EOS No. 30992574], and the Flanders AI Impulse Program.

References

- [1] M.L.D. Wong, L.B. Jack, and A.K. Nandi. Modified self-organising map for automated novelty detection applied to vibration signal monitoring. *Mechanical Systems and Signal Processing*, 20(3):593–610, apr 2006.
- [2] Jay Lee, Hung-An Kao, and Shanhu Yang. Service Innovation and Smart Analytics for Industry 4.0 and Big Data Environment. In *6th Conference on Industrial Product-Service Systems*, volume 16, pages 3–8, Windsor, Ontario, Canada, 2014. Elsevier B.V.
- [3] Bryan L Matthews, Santanu Das, Kanishka Bhaduri, Kamalika Das, Rodney Martin, and Nikunj Oza. Discovering Anomalous Aviation Safety Events Using Scalable Data Mining Algorithms. *Journal of Aerospace Information Systems*, 10(10):467–475, 2013.
- [4] David Siegel. *Prognostics and Health Assessment of a Multi-Regime System using a Residual Clustering Health Monitoring Approach*. PhD thesis, University of Cincinnati, 2013.
- [5] Will Jacobs, Huw Edwards, Visakan Kadirkamanathan, and Andrew R Mills. Inter-engine variation analysis for health monitoring of aerospace gas turbine engines. *European PHM*, 2018.
- [6] Victoria J. Hodge, Simon O’Keefe, Michael Weeks, and Anthony Moulds. Wireless Sensor Networks for Condition Monitoring in the Railway Industry: A Survey. *IEEE Transactions on Intelligent Transportation Systems*, 16(3):1088–1106, jun 2015.
- [7] Emanuele Fumeo, Luca Oneto, and Davide Anguita. Condition Based Maintenance in Railway Transportation Systems Based on Big Data Streaming Analysis. *Procedia Computer Science*, 53(1):437–446, 2015.
- [8] Fabien Meinguet, Eric Semail, Xavier Kestelyn, Yves Mollet, and Johan Gyselinck. Change-detection algorithm for short-circuit fault detection in closed-loop AC drives. *IET Electric Power Applications*, 8(5):165–177, may 2014.
- [9] International Organization for Standardization. ISO 10816-3:2009 – Mechanical vibration – Evaluation of machine vibration by measurements on non-rotating parts – Part 3: Industrial machines with nominal power above 15 kW and nominal speeds between 120 r/min and 15 000 r/min when measured in situ, 2009.
- [10] Charles R C.R. Farrar, Hoon Sohn, and Keith Worden. *Data Normalization : A Key For Structural Health Monitoring*, 2001.
- [11] Xiaodong Jia, Chao Jin, Matt Buzza, Yuan Di, David Siegel, and Jay Lee. A deviation based assessment methodology for multiple machine health patterns classification and fault detection. *Mechanical Systems and Signal Processing*, 99(June):244–261, jan 2018.
- [12] Jay Lee, Fangji Wu, Wenyu Zhao, Masoud Ghaffari, Linxia Liao, and David Siegel. Prognostics and health management design for rotary machinery systems Reviews, methodology and applications. *Mechanical Systems and Signal Processing*, 42(1-2):314–334, jan 2014.
- [13] Ruonan Liu, Boyuan Yang, Enrico Zio, and Xuefeng Chen. Artificial intelligence for fault diagnosis of rotating machinery: A review. *Mechanical Systems and Signal Processing*, 108:33–47, aug 2018.
- [14] Rui Zhao, Ruqiang Yan, Zhenghua Chen, Kezhi Mao, Peng Wang, and Robert X. Gao. Deep learning and its applications to machine health monitoring. *Mechanical Systems and Signal Processing*, 115:213–237, jan 2019.

- [15] B.S. Yang, T. Han, and J.L. An. ARTKOHONEN neural network for fault diagnosis of rotating machinery. *Mechanical Systems and Signal Processing*, 18(3):645–657, may 2004.
- [16] Muhammet Unal, Mustafa Onat, Mustafa Demetgul, and Haluk Kucuk. Fault diagnosis of rolling bearings using a genetic algorithm optimized neural network. *Measurement*, 58:187–196, dec 2014.
- [17] Meik Schlechtingen and Ilmar Ferreira Santos. Comparative analysis of neural network and regression based condition monitoring approaches for wind turbine fault detection. *Mechanical Systems and Signal Processing*, 25(5):1849–1875, jul 2011.
- [18] Feng Jia, Yaguo Lei, Jing Lin, Xin Zhou, and Na Lu. 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, 2016.
- [19] Samir Khan and Takehisa Yairi. A review on the application of deep learning in system health management. *Mechanical Systems and Signal Processing*, 107:241–265, 2018.
- [20] Zhu Yun Chen, Konstantinos Gryllias, and Weihua Li. Intelligent Fault Diagnosis for Rotary Machinery Using Transferable Convolutional Neural Network. *IEEE Transactions on Industrial Informatics*, pages 1–1, 2019.
- [21] Ming Ge, R. Du, Guicai Zhang, and Yangsheng Xu. Fault diagnosis using support vector machine with an application in sheet metal stamping operations. *Mechanical Systems and Signal Processing*, 18(1):143–159, jan 2004.
- [22] Achmad Widodo and Bo-Suk Yang. Support vector machine in machine condition monitoring and fault diagnosis. *Mechanical Systems and Signal Processing*, 21(6):2560–2574, aug 2007.
- [23] K.C. Gryllias and I.A. Antoniadis. A Support Vector Machine approach based on physical model training for rolling element bearing fault detection in industrial environments. *Engineering Applications of Artificial Intelligence*, 25(2):326–344, mar 2012.
- [24] Alexandre Mauricio, Carina Freitas, Jaques Cuenca, Konstantinos Gryllias, Bram Cornelis, Karl Janssens, and Kilian Hendrickx. Condition monitoring of gears under medium rotational speed. In *24th International Congress on Sound and Vibration*, page 8, London, 2017.
- [25] L. Bull, K. Worden, G. Manson, and N. Dervilis. Active learning for semi-supervised structural health monitoring. *Journal of Sound and Vibration*, 437:373–388, dec 2018.
- [26] T.J. Rogers, K. Worden, R. Fuentes, N. Dervilis, U.T. Tygesen, and E.J. Cross. A Bayesian non-parametric clustering approach for semi-supervised Structural Health Monitoring. *Mechanical Systems and Signal Processing*, 119:100–119, mar 2019.
- [27] Lei Zhang, Qixin Cao, and Jay Lee. Performance Assessment for a Fleet of Machines Using a Combined Method of Ant-Based Clustering and CMAC. *Advances in Mechanical Engineering*, 5:603071, jan 2013.
- [28] Cameron Sobie, Carina Freitas, and Mike Nicolai. Simulation-driven machine learning: Bearing fault classification. *Mechanical Systems and Signal Processing*, 99:403–419, 2018.
- [29] Marco A.F. Pimentel, David A. Clifton, Lei Clifton, and Lionel Tarassenko. A review of novelty detection. *Signal Processing*, 99:215–249, jun 2014.
- [30] Konstantinos Gryllias, Ioannis Antoniadis, and Christos Yiakopoulos. A novel semi-supervised mathematical morphology-based fault detection & classification method for rolling element bearings. In *The 22nd International Congress on Sound and Vibration*, page 8, Florence, 2015.
- [31] Ziad Hasan Abu-El-Zeet and Vijaykumar Chhaganlal Patel. METHOD OF CONDITION MONITORING (Patent US 7275018 B2), 2006.
- [32] G.M. West, S.D.J. McArthur, and D. Towle. Industrial implementation of intelligent system techniques for nuclear power plant condition monitoring. *Expert Systems with Applications*, 39(8):7432–7440, jun 2012.
- [33] Peng Li, Jens Eickmeyer, and Oliver Niggemann. Data Driven Condition Monitoring of Wind Power Plants Using Cluster Analysis. In *2015 International Conference on Cyber-Enabled Distributed Computing and Knowledge Discovery*, pages 131–136. IEEE, sep 2015.
- [34] Sameer Al-Dahidi, Francesco Di Maio, Piero Baraldi, Enrico Zio, and Redouane Seraoui. A framework for reconciling data clusters from a fleet of nuclear power plants turbines for fault diagnosis. *Applied Soft Computing*, 69:213–231, aug 2018.
- [35] Varun Chandola, Arindam Banerjee, and Vipin Kumar. Anomaly detection: A survey. *ACM Computing Surveys*, 41(3):1–58, jul 2009.
- [36] Niels Henrik Pontoppidan, Sigurdur Sigurdsson, and Jan Larsen. Condition monitoring with mean field independent components analysis. *Mechanical Systems and Signal Processing*, 19(6):1337–1347, nov 2005.
- [37] Afroz Purarjomandlangrudi, Amir Hossein Ghapanchi, and Mohammad Esmalifalak. A data mining approach for fault diagnosis: An application of anomaly detection algorithm. *Measurement*, 55:343–352, sep 2014.
- [38] Vincent Vercruyssen, Wannes Meert, Gust Verbruggen, Koen Maes, Ruben Baumer, and Jesse Davis. Semi-Supervised Anomaly Detection with an Application to Water Analytics. In *2018 IEEE International Conference on Data Mining (ICDM)*, volume 2018-Novem, pages 527–536. IEEE, nov 2018.
- [39] Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, mar 2009.
- [40] Charu C. Aggarwal and Saket Sathe. *Outlier Ensembles*. Springer International Publishing, Cham, 2017.
- [41] Gabriela Medina-Oliva, Alexandre Voisin, Maxime Monnin, Flavien Peysson, and Jean-Baptiste Léger. Prognostics assessment using fleet-wide ontology. In *Annual Conference of the Prognostics and Health Management Society 2012, PHM Conference 2012*, Minneapolis, Minnesota, 2012.
- [42] Gabriela Medina-Oliva, Alexandre Voisin, Maxime Monnin, and Jean Baptiste Léger. Predictive diagnosis based on a fleet-wide ontology approach. *Knowledge-Based Systems*, 68:40–57, 2014.
- [43] Alexandre Voisin, Gabriela Medina-Oliva, Maxime Monnin, Jean-Baptiste Léger, and Benoît Iung. Fleet-wide Diagnostic

- and Prognostic Assessment. In *Annual Conference of the Prognostics and Health Management Society*, volume 3, 2011.
- [44] Robert B. Randall and Jérôme Antoni. Rolling element bearing diagnosticsA tutorial. *Mechanical Systems and Signal Processing*, 25(2):485–520, feb 2011.
 - [45] Meinard Müller. Dynamic Time Warping. In *Information Retrieval for Music and Motion*, pages 69–84. Springer-Verlag Berlin Heidelberg, 1 edition, 2007.
 - [46] Donald J Bemdt and James Clifford. Using dynamic time warping to find patterns in time series. In *AAAIWS’94 Proceedings of the 3rd International Conference on Knowledge Discovery and Data Mining*, pages 359–370, Seattle, 1994. AAAI Press.
 - [47] Diego Furtado Silva, Gustavo E. A. P. A. Batista, and Eamonn Keogh. On the Effect of Endpoints on Dynamic Time Warping. In *SIGKDD MiLeTS’16*, page 10, San Francisco, California, 2016. ACM.
 - [48] Kilian Hendrickx, Wannes Meert, Bram Cornelis, Karl Janssens, Konstantinos Gryllias, and Jesse Davis. A Fleet-Wide Approach for Condition Monitoring of Similar Machines Using Time-Series Clustering. In Springer, editor, *Proceedings of the 6th International Conference on Condition Monitoring of Machinery in Non-Stationary Operations, CMMNO’2018*, pages 101–110. Springer, Cham, Santander, Spain, jun 2019.
 - [49] Leonard Kaufman and Peter J. Rousseeuw. *Finding Groups in Data: An Introduction to Cluster Analysis*. Wiley, 1990.
 - [50] Hae Sang Park and Chi Hyuck Jun. A simple and fast algorithm for K-medoids clustering. *Expert Systems with Applications*, 36(2 PART 2):3336–3341, 2009.
 - [51] Ulrike von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, dec 2007.
 - [52] F. Murtagh. A Survey of Recent Advances in Hierarchical Clustering Algorithms. *The Computer Journal*, 26(4):354–359, nov 1983.
 - [53] Martin Ester, Hans-Peter Kriegel, Sander Jorg, and Xiaowei Xu. A Density-Based Clustering Algorithms for Discovering Clusters. In *KDD’96 Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, pages 226–231, Portland, Oregon, 1996.
 - [54] Joe H. Ward. Hierarchical Grouping to Optimize an Objective Function. *Journal of the American Statistical Association*, 58(301):236–244, mar 1963.
 - [55] V. Parker Lessig. Comparing Cluster Analyses with Cophenetic Correlation. *Journal of Marketing Research*, 9(1):82, feb 1972.
 - [56] Shikha Agrawal and Jitendra Agrawal. Survey on Anomaly Detection using Data Mining Techniques. *Procedia Computer Science*, 60(1):708–713, 2015.
 - [57] Rasool Sharifi and Mohammad Ebrahimi. Detection of stator winding faults in induction motors using three-phase current monitoring. *ISA Transactions*, 50(1):14–20, 2011.
 - [58] Yves Mollet, Matteo Pergolesi, Mathieu Sarrazin, Karl Janssens, Herman Van der Auweraer, Paolo Chiariotti, Paolo Castellini, and Johan Gyselinck. Multi-Physical Signature Analysis of Induction Machines under Unbalanced Supply Voltage. In *2018 XIII International Conference on Electrical Machines (ICEM)*, volume 32, pages 2378–2384, Alexandroupoli, Greece, sep 2018. IEEE.
 - [59] Sang-Joon Lee, Jun-Koo Kang, and Seung-Ki Sul. A new phase detecting method for power conversion systems considering distorted conditions in power system. In *Conference Record of the 1999 IEEE Industry Applications Conference. Thirty-Forth IAS Annual Meeting (Cat. No.99CH36370)*, volume 4, pages 2167–2172, Phoenix, AZ, USA, 1999. IEEE.
 - [60] Jerome H. Friedman. On bias, variance, 0/1-loss, and the curse-of-dimensionality. *Data Mining and Knowledge Discovery*, 1(1):55–77, 1997.
 - [61] G.S Maruthi and K. Panduranga Vittal. Electrical Fault Detection in Three Phase Squirrel Cage Induction Motor by Vibration Analysis using MEMS Accelerometer. In *2005 International Conference on Power Electronics and Drives Systems*, volume 2, pages 838–843. IEEE, 2006.
 - [62] P.J. Tavner. Review of condition monitoring of rotating electrical machines. *IET Electric Power Applications*, 2(4):215, 2008.