



Novel features for the detection of bearing faults in railway vehicles

Matthias Kreuzer¹
Multimedia Communications and Signal Processing
Friedrich-Alexander University Erlangen-Nürnberg
Cauerstr. 7, 91058 Erlangen, Germany

Alexander Schmidt²
Multimedia Communications and Signal Processing
Friedrich-Alexander University Erlangen-Nürnberg
Cauerstr. 7, 91058 Erlangen, Germany

Walter Kellermann³
Multimedia Communications and Signal Processing
Friedrich-Alexander University Erlangen-Nürnberg
Cauerstr. 7, 91058 Erlangen, Germany

ABSTRACT

In this paper, we address the challenging problem of detecting bearing faults from vibration signals. For this, several time- and frequency-domain features have been proposed already in the past. However, these features are usually evaluated on data originating from relatively simple scenarios and a significant performance loss can be observed if more realistic scenarios are considered. To overcome this, we introduce Mel Frequency Cepstral Coefficients (MFCCs) and features extracted from the Amplitude Modulation Spectrogram (AMS) as features for the detection of bearing faults. Both AMS and MFCCs were originally introduced in the context of audio signal processing but it is demonstrated that a significantly improved classification performance can be obtained by using these features. Furthermore, to tackle the characteristic data imbalance problem in the context of bearing fault detection, i.e., typically much more data from healthy bearings than from damaged bearings is available, we propose to train a One-class Support Vector Machine (SVM) with data from healthy bearings only. Bearing faults are then classified by the detection of outliers. Our approach is evaluated with data measured in a highly challenging scenario comprising a state-of-the-art commuter railway engine which is supplied by an industrial power converter and coupled to a load machine.

1. INTRODUCTION

Bearings represent crucial components in rotating machinery, e.g., wind turbines and railway vehicles. As a bearing failure may cause substantial additional costs and also damages to other components,

¹matthias.kreuzer@fau.de

²alexander.as.schmidt@fau.de

³walter.kellermann@fau.de

an early and reliable detection of bearing faults is of major importance. Since these components are often difficult to access, manual inspection is time-consuming and expensive. However, damages to the bearing alter the vibration signature of the machine, which can be captured by acceleration sensors that are installed on the housing of the machine. Consequently, faults can be identified by analyzing the vibration signal.

For this, various techniques have been developed over the recent decades. The analysis of the vibration signals is typically performed in the time domain or the frequency domain. In this context, especially the so-called fault frequencies play a crucial role: If one of the four major bearing components, i.e., the inner race, the outer race, the cage and the rolling elements, is damaged, impulses with component-specific frequencies can be observed in the vibration signal. Many methods analyze the spectrum of the envelope of the vibration signal as the fault frequencies become more apparent after demodulating the signal [1]. Yet, the signal components that are related to the bearing fault are often buried in environmental noise caused by other machine components, which necessitates the application of signal enhancement algorithms. In order to separate the fault-related components from the environmental noise, several signal decomposition techniques, e.g., the Empirical Mode Decomposition [2] and Wavelet Transforms [3], and adaptive filtering techniques [4] have been applied to this problem. Detailed reviews for signal processing techniques applied to the topic of bearing fault detection can be found in [3].

As the vibration pattern changes with the appearance of a fault, statistical features, e.g., the variance, the kurtosis etc., of the vibration signal also change which can be used to detect a fault. Consequently, data-driven machine learning classification algorithms, e.g., linear discriminant analysis, k-nearest neighbors, SVMs and neural networks, can be trained on features extracted from the time- and frequency-domain representations of the vibration signal [5]. Naturally, the performance of these classifiers is strongly dependent on the extracted features. Thus, determining the best-suited features is crucial for this kind of approach. Recently, also end-to-end deep learning-based approaches have gained wide popularity [6]. In contrast to the previous classification algorithms, a manual feature selection step is not required as the features are extracted by the network itself. A comprehensive overview for a wide variety of deep learning architectures including autoencoders, convolutive neural networks and recurrent neural networks can be found in [6]. Although the manual feature selection step can be omitted, most of these approaches suffer from the fact that they require a large amount of labeled training data for both healthy and damaged bearings in order to show satisfactory results. However, the available data is often unbalanced as data for damaged bearings in realistic scenarios is hard to obtain, e.g. due to safety reasons, whereas data for healthy bearings is usually available in abundance. Furthermore, the labeling process for naturally developed damages is difficult as it is hard to determine precisely the point in time when the first signs of a fault begin to show. Hence, approaches that are not reliant on data from damaged bearings or only require a very little amount, are especially well-suited for the classification of bearing faults. This type of classification problem is referred to as outlier or anomaly detection in the literature.

In this paper, we evaluate and compare various features from the time and frequency domain for the detection of bearing faults in a highly challenging scenario. The experimental setup consists of a state-of-the-art commuter railway engine coupled to a load machine. Both engine and load are supplied by an industrial power converter. In contrast to most of the publicly available datasets the rotational frequency and the applied torque vary strongly and the power converter strongly affects the measured vibration patterns. In addition to the well-studied features from literature, we also investigate novel features from the domain of audio signal processing and automatic speech recognition (ASR), namely Amplitude Modulation Spectrogram (AMS) and Mel Frequency Cepstral Coefficients (MFCCs). It is shown that these features are highly effective and outperform conventional time- and frequency-domain features. Further, in order to address the commonly encountered data imbalance problem we only use features obtained from healthy bearings to train a One-Class Support Vector Machine (SVM) and treat the classification of bearing faults as an anomaly detection task. It is shown that even with

this rather simple classifier and an appropriate set of features, classification accuracies over 98 % can be achieved.

This article is organized as follows. In Section 2.1, the underlying signal model is introduced before conventional features for bearing fault detection from time and frequency domain are briefly reviewed in Section 2.2. After this, the MFCCs and the AMS are discussed in more detail in Section 2.3 and Section 2.4, respectively. In Section 2.5, the One-Class SVM classification method is described. In Section 3, the experimental setup (cf. Section 3.1), the data generation and the training process along with used evaluation metrics (cf. Section 3.2) are described, before the experimental results are discussed in detail in Section 3.3. A summary in Section 4 concludes the paper.

2. BEARING FAULT CLASSIFICATION

In this section, the investigated features, most importantly the proposed AMS and MFCCs, are introduced and the classification method is described.

2.1. Signal Model

In the following, the discrete time-domain vibration signal will be denoted as $x[k]$ with sample index k , which is obtained after sampling the continuous-time domain vibration signal with sampling frequency f_s . Given K samples of the discrete-time signal $x[k]$, the Discrete Fourier Transform (DFT) can be defined as follows

$$X[\mu] = \text{DFT}_M\{x[k]\} = \sum_{k=0}^{K-1} x[k]e^{-j\frac{2\pi}{M}k\mu}, \quad (1)$$

where μ and M denote the frequency bin index and the length of the DFT, respectively. The Short-Time Fourier Transform (STFT) of $x[k]$ for a segment of length K is then given by

$$X[\mu, n] = \sum_{k=n}^{n+K-1} x[k]w_M[k-n]e^{-j\frac{2\pi}{M}k\mu}, \quad (2)$$

with time frame index n . $w_M[k]$ is a window function of length M [7].

2.2. Conventional Time- and Frequency-domain Features

In this section, we briefly review conventional features from the time domain and the frequency domain. Subsequently, the proposed features, i.e., the MFCCs and AMS, are introduced in Sections 2.3 and 2.4, respectively. The selected features from the time domain include average, variance, root mean square value (RMS), kurtosis, skewness, amplitude range and peak-to-RMS ratio of $x[k]$ and according estimators are summarized in Table 5. These features are well-established for the detection for bearing faults and used for example used in [8–11]. In the following, the set of time-domain features is abbreviated with TD . Analogously, statistical features can be derived from the frequency-domain vibration signal $X[\mu]$. We consider spectral centroid, spectral spread, spectral kurtosis, spectral entropy, spectral crest and roll-off point (cf. Table 6). These features have been successfully applied for bearing fault detection in [9, 12, 13]. Henceforth, the set of features listed in Table 6 is denoted as SD , i.e., the spectral descriptors. As localized faults can be linked to characteristic frequencies, the amplitudes at multiples of the characteristic fault frequencies in the envelope magnitude spectrum are also considered as features. These features are computed as

$$AMP_{fault} = \sum_{i=1}^3 |X_{env}[i \cdot \mu_{fault}]| \quad (3)$$

with $\mu_{fault} \in \{\mu_{BPFO}, \mu_{BPFI}, \mu_{CA}, \mu_{RE}\}$,

where $X_{env}[\mu]$ denotes the envelope spectrum (cf. [1]), which is obtained by applying the Hilbert transform [7] and μ_{fault} is the frequency bin index corresponding to the characteristic fault frequencies f_{BPFO} , f_{BPFI} , f_{CA} and f_{RE} for a localized fault in the outer race, the inner race, the rolling cage and the rolling elements, respectively. Equations to determine the characteristic fault frequencies can be found in [1, 2].

Table 1: Overview of considered time-domain features.

Feature	Formula
Average \bar{x}	$\frac{1}{K-1} \sum_{k=0}^{K-1} x[k]$
Variance σ_x^2	$\frac{1}{K-1} \sum_{k=0}^{K-1} (x[k] - \bar{x})^2$
Root Mean Square (RMS) x_{RMS}	$\sqrt{\frac{1}{K-1} \sum_{k=0}^{K-1} x[k]^2}$
Kurtosis	$\frac{\frac{1}{K-1} \sum_{k=0}^{K-1} (x[k] - \bar{x})^4}{(\sigma_x^2)^2}$
Skewness	$\frac{\frac{1}{K-1} \sum_{k=0}^{K-1} (x[k] - \bar{x})^3}{(\sqrt{\sigma_x^2})^3}$
Amplitude Range	$\max(x[k]) - \min(x[k])$
Peak-to-RMS ratio	$\frac{\max(x[k])}{x_{RMS}}$

2.3. Mel Frequency Cepstral Coefficients

MFCCs are state-of-the-art features for automatic speech recognition tasks and acoustic scene classification. As the name suggests, the MFCCs combine the cepstrum with a Mel-frequency scaling which allows for a compact representation of the frequency spectrum of a signal. In speech analysis, the real-valued cepstrum, i.e., the Inverse Discrete Fourier Transform (IDFT) of the log magnitude of the Discrete-Time Fourier Transform (DTFT) of a signal, is widely popular for various applications such as, e.g., pitch detection and acoustic scene classification [14]. The MFCC feature extraction scheme can be summarized as follows:

1. The time domain signal is divided into frames by windowing.
2. An N -point DFT is computed of the windowed input samples.
3. The resulting power magnitude spectrum is then filtered using a Mel-filterbank, which is approximated, e.g., by applying triangular windows in the DFT domain [14].
4. Finally, the MFCCs are obtained by computing the Discrete Cosine Transform (DCT) of the logarithm of summed filter bank energies.

Table 2: Overview of considered frequency-domain features. $\mu_1 = 0$ and $\mu_2 = K - 1$ denote the lower and upper frequency bin index, respectively, f_μ denotes the frequency corresponding to frequency bin μ . κ is set to 0.95.

Feature	Formula
Spectral Centroid (SC)	$\frac{\sum_{\mu=\mu_1}^{\mu_2} f_\mu X[\mu] }{\sum_{\mu=\mu_1}^{\mu_2} X[\mu] }$
Spectral Spread (SSpr)	$\sqrt{\frac{\sum_{\mu=\mu_1}^{\mu_2} (f_\mu - \text{SC})^2 X[\mu] }{\sum_{\mu=\mu_1}^{\mu_2} X[\mu] }}$
Spectral Kurtosis	$\frac{\sum_{\mu=\mu_1}^{\mu_2} (f_\mu - \text{SC})^4 X[\mu] }{(\text{SSpr})^4 \sum_{\mu=\mu_1}^{\mu_2} X[\mu] }$
Spectral Entropy	$\frac{-\sum_{\mu=\mu_1}^{\mu_2} X[\mu] \log(X[\mu])}{\log(\mu_2 - \mu_1)}$
Spectral Crest	$\frac{\max(X[\mu])}{\frac{1}{\mu_2 - \mu_1 + 1} \sum_{\mu=\mu_1}^{\mu_2} X[\mu] }$
Roll-off point	$\sum_{\mu=\mu_1}^i X[\mu] = \kappa \sum_{\mu=\mu_1}^{\mu_2} X[\mu] $

For time-frame n , the spectral energies $X_i[n]$ of time frame i with $i = 0, \dots, K - 1$ can be computed as

$$X_i[n] = \sum_{v=0}^{N-1} g_{iv} \left| \sum_{k=0}^{N-1} x[n-k] w[k] e^{-\frac{2\pi k v}{N}} \right|^2, \quad (4)$$

where $x[k]$ are in the input samples, $w[k]$ is a window function and g_{iv} is a triangular window function of a Mel-filterbank. The Mel filters mimic the human auditory system which has a finer resolution for lower frequency regions than for higher frequency regions. The bandwidths of the practically used Mel filters are constant for center frequencies below 500 Hz and increase approximately exponentially for higher center frequencies. The number of subbands, i.e., number of filters, varies typically from $K = 24$ to $K = 40$ subbands. After filtering, the MFCCs $c[\mu]$ are obtained after applying a DCT

$$c[\mu] = \sum_{i=1}^K \log X_i[n] \cos\left(\frac{\pi(2i-1)\mu}{2K}\right), \quad \mu = 1, \dots, K. \quad (5)$$

which remove the correlation between the cepstral coefficients introduced by the overlapping filterbanks.

2.4. Amplitude Modulation Spectrogram

The AMS is a popular feature in hearing-related signal processing and speech recognition and is known for its usefulness for the detection of envelope fluctuations of speech signals [15]. The signal processing steps that are required for computing the AMS are illustrated in the block diagram in Figure 2.4. First, the time-domain signal $x[k]$ is transformed to the STFT domain. In the next step, a second STFT of the squared magnitude of the complex-valued STFT coefficients for every frequency




Figure 1: Block diagram for AMS.

subband is computed. The subband signals after the squaring operation can be interpreted as the envelope signal of the spectral amplitudes for every subband. This second STFT can be understood as computing a Fourier series of the temporal envelope signal for a given frequency bin, and thereby detecting periodic components (modulation frequencies) in the given frequency subband. Finally, the AMS is obtained by compressing the magnitude logarithmically. Sometimes the number of subbands is reduced before the second STFT by applying a Mel- or Bark-filterbank [15, 16]. In this article the filtering step is omitted and the parameters listed in Table 2.4 are utilized for the computation of the two consecutive STFTs.

To illustrate that the AMS can be a valuable feature for the detection of bearing faults, we consider Figures 2 and 3. Figure 2 shows the spectrogram (top) and the AMS (bottom) for a vibration signal excerpt of 2 s for a bearing with a localized fault in the outer race and a rotational frequency of 500 rpm (≈ 8.33 Hz) and a sampling frequency f_s of 51.2 kHz. It can be noticed that the spectrogram exhibits distinct periodical vertical lines. The corresponding AMS is shown in the lower subfigure of Figure 2 where the subbands of the first STFT are given along the y-axis, and the modulation frequencies are illustrated along the x-axis. It can be observed that for center frequencies above 20 kHz strong modulations are present. This is, however, not the case for the AMS of the vibration signals of a healthy bearing operated under identical operational conditions (cf. Figure 3). The aforementioned distinct vertical lines of the spectrogram are especially noticeable in the high frequencies above 15 kHz (highlighted by dotted red lines) and are well captured by the energy in the AMS. One might argue that inspecting the spectrograms in Figure 2 and Figure 3 is sufficient for detecting the fault. However, the AMS provides a larger energy difference in the considered frequency range. In order to quantify these differences with a single scalar value, we propose to sum up the amplitudes for center frequencies above 20 kHz and modulation frequencies below 80 Hz.

Table 3: Parameters for the computation of the two STFTs. The number of samples is given for the sampling frequency of $f_s = 51.2$ kHz.

Parameter	1. STFT	2. STFT
Window length	25 ms (1280 samples)	512 ms (128 samples)
Step size	4 ms (205 samples)	256 ms (64 samples)
N-FFT	512	256

2.5. One-Class SVM

The One-Class SVM represents a flexible and effective method for the detection of outliers. It was originally introduced by Schölkopf et al. [17] in the context of novelty detection and aims at estimating a boundary function which returns positive values for data points belonging to a given class. For data points of other classes and outliers, the function returns negative values.

In the following, we will briefly review the fundamental idea of One-Class SVMs. For this, we consider a training data set $\{\mathbf{x}_n\}_{n=1,\dots,N}$ belonging to one class, where N is the overall number of training samples. The boundary function is estimated in a so-called kernel-induced feature space where the




Figure 2: Spectrogram (top) and AMS (bottom) for a damaged bearing with a fault in the outer race.




Figure 3: Spectrogram (top) and AMS (bottom) for a healthy bearing.

training data is assumed to be linearly separable from the origin. An according decision boundary is then given by

$$f(\mathbf{x}) = \mathbf{w}^T \phi(\mathbf{x}) + b, \quad (6)$$

where b and \mathbf{w} denote the distance from the origin and the normal vector, respectively. $\phi(\cdot)$ is the mapping function to the kernel-induced feature space. Similar to conventional SVMs, $\phi(\cdot)$ does not have to be known explicitly what is generally referred to as kernel-trick [18]. The parameters of $f(\cdot)$ must be optimized such that the training data is separated from the origin with maximal distance and $f(\mathbf{x}_n) \geq 0$ holds for all n . Considering that the Euclidean distance from any arbitrary point \mathbf{x} to $f(\cdot)$ is given by $f(\mathbf{x})/\|\mathbf{w}\|_2$ and ensuring that $\phi(0) = 0$ holds, a boundary with maximal distance from the origin can be obtained by solving the following optimization problem:

$$\begin{aligned} & \underset{\mathbf{w}, b}{\text{minimize}} && \|\mathbf{w}\|_2 - b \\ & \text{subject to} && \mathbf{w}^T \phi(\mathbf{x}_n) + b \geq 0, \forall n. \end{aligned} \quad (7)$$

Equation 7 can be solved using methods from constrained optimization, e.g., solving the Lagrangian [18]. Similar to conventional SVMs, the hyperplane $f(\cdot)$ will be defined by only some data points

from the training set, the so-called support vectors. Furthermore, it is shown in [17] that the obtained decision boundary is an estimator for the support of a density that generated the training data points. Note that Equation 7 can be extended by slack variables to trade-off the width of the margin and the number of misclassified data points, cf. [17].

As mentioned above, the implicit mapping of \mathbf{x} by $\phi(\cdot)$ does not have to be known explicitly as we are only interested in the inner products of the mapped inputs, which can be computed using a kernel. Usually, One-Class SVMs operate with a Gaussian kernel $e^{-\gamma\|\mathbf{x}_i-\mathbf{x}_j\|}$ for data points \mathbf{x}_i and \mathbf{x}_j . The Gaussian kernel has only one free parameter γ which is referred to as the *width* of the kernel. Note that the Gaussian kernel ensures the existence of a hyperplane defined by the One-Class SVM optimization problem [17].

3. EVALUATION

In the following, the experimental setup (cf. Section 3.1), the data generation and training process as well as the considered evaluation metrics (cf. Section 3.2) are described, before the obtained results are thoroughly discussed in Section 3.3.

3.1. Experimental Setup

In the experimental study, a state-of-the-art commuter rail induction engine with a nominal capacity of 405 kW, a nominal load of 1800 Nm and a nominal rotation speed of 2150 rpm is studied. In total, 3 different bearings are investigated, including a healthy bearing and two bearings with different types of damages. One of the bearings exhibits an artificially introduced damage in the outer race with a fault diameter of approximately 3.5 mm and a depth of approximately 1.5 mm. The width of the outer race is 22 mm. This can be considered as a localized fault in a rather early stage. The second bearing returned from actual usage in the field in a railway vehicle after a run-time of approximately 1.5 million km, which exhibits distributed faults at multiple bearing components including the outer race, the rolling elements, the cage and the rolling elements. This bearing represents a naturally developed bearing damage in an advanced development stage. In the following these two damage patterns are referred to as artificially and naturally developed damage, respectively. The investigated bearings (type 6016) are installed at the drive-end position of the induction engine.

The experimental setup, which is shown in Figure 4, further includes a second induction engine of identical type that is used as a load machine in a back-to-back setup. In total, two acceleration sensors (Sensor A and Sensor B, $f_s = 51,2$ kHz) are mounted on the housing of the induction engine at the drive-end and the non-drive-end position, respectively. An additional sensor is used to track the rotational speed. Further, a railway traction converter is used in the experiments, which strongly influences the vibration pattern of the induction engine. The applied torque lies in the interval -2900 Nm and 2900 Nm with four different torque levels (0 %, 33 %, 66 %, 100 %) for accelerating and braking. The applied rotational speed follows a steplike characteristic with 8 distinct rotational speed levels: 500 rpm, 750 rpm, 1000 rpm, 1500 rpm, 2000 rpm, 2500 rpm, 3000 rpm and 3500 rpm. The rotational speed curve is shown in Figure 5. It can be observed that the rotational speed was kept constant for a period of approximately 60 s for every rotational speed. Overall, the examined scenario is highly challenging and close to the operating conditions in the field as the rotational speed and the torque are strongly varying, which is not the case for publicly available datasets as, e.g., the Case Western Reserve University (CWRU) dataset [6].

3.2. Data Generation, Training, and Evaluation Metrics

For generating the training set and the test set, the vibration signals were divided into non-overlapping frames with a length of 2 s each. After segmentation, frames containing transitions of the rotational frequency were dismissed. For both sensors, this resulted in 1558 frames for the healthy

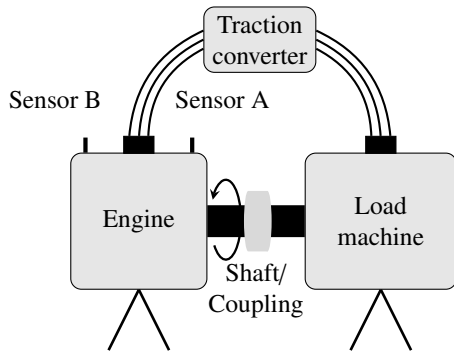


Figure 4: Experimental setup.




Figure 5: Rotational speed curve.

Predicted Class	P	TP	FP
	N	FN	TN
		P	N
		True Class	

Figure 6: Confusion matrix for a binary classification problem.

Table 4: Evaluation metrics.

Name	Formula
TPR (True Positive Rate)	$\frac{TP}{TP+FN}$
TNR (True Negative Rate)	$\frac{TN}{TN+FP}$
BA (Balanced Accuracy)	$\frac{TPR+TNR}{2}$

measurement, 2512 frames for the artificially damaged bearing and 2728 frames for the naturally developed damage. For training the One-Class SVM, 500 samples obtained from the healthy bearing were randomly chosen for the training set. Another small dataset containing 50 additional samples from the measurements of the healthy bearing was used as an evaluation dataset for optimizing the parameters of the One-Class SVM using a grid-search procedure.

In the following, the evaluation metrics that are considered for rating the classification performance are introduced. We consider a binary classification problem where we distinguish between the class P (positive, healthy bearing) and the class N (negative, damaged bearing). The performance of a classifier can be assessed in an illustrative manner with a so-called confusion matrix (CM) (cf. Figure 6). In a CM, the predicted class labels (y-axis) are compared to the true class labels (x-axis). Hence, the diagonal entries, i.e., True Positive (TP) and True Negative (TN), show the number of correctly classified test instances, whereas the off-diagonal entries, i.e., False Negative (FN) and False Positive (FP), represent the false predictions. Applied to the classification of bearing faults, the FN value corresponds to the number of healthy bearings that were falsely classified as damaged and vice-versa for FP. By construction our test dataset is highly imbalanced as the number of negative test instances is much greater than the number of positive test instances. Consequently, metrics such as the accuracy or the F1-score are not suited and we consider the True Positive Rate (TPR) and the True Negative Rate (TNR), which specify how well the positive class and the negative class were predicted, respectively, instead. The arithmetic mean of the TPR and the TNR leads to the Balanced Accuracy (BA). The formulae for computing TPR, TNR and BA are summarized in Table 4. Further, the False Negative Rate (FNR) and the False Positive Rate (FPR) denote the fraction of the test instances that were misclassified as P or N and are given as $1 - \text{FNR}$ and $1 - \text{FPR}$.

3.3. Results

For evaluation, the One-Class SVM was trained using the five different feature sets that were introduced in Sections 2.2-2.3: the time-domain features (*TD*), the spectral descriptors (*SD*), the amplitudes at the characteristic fault frequencies in the envelope spectrum (*ENV-AMP*), the scalar

feature derived from the Amplitude Modulation Spectrogram (*AMS*) and the first 13 Mel Frequency Spectral Coefficients (*MFCC (13)*). Further, to investigate whether the addition of the rotational frequency f_r is beneficial, all feature sets were again evaluated after appending f_r to the feature vectors. Training and testing were repeated 10 times using different realizations of the training and test set. Averaged results for TPR, TNR and BA are summarized in Table 5 for Sensor A. Overall, it can be stated that the feature set *MFCC (13)* leads to the best classification results with a BA of 98.79 % and 98.91 % after appending the rotational frequency to the feature set. BAs over 90 % are only achieved for the *ENV AMP + f_r* and *AMS+f_r* with 91.58 % and 94.81 %, respectively, whereas the *TD* and *SD* feature performs worst. Surprisingly, a large performance drop can be witnessed when f_r is added to *SD* (from 88.33 % to 82.81 % in terms of BA). The reasons for this remains to be investigated. However, for the other feature sets, the addition of f_r to the feature vector improves the classification results always at least slightly. Yet, the TPR does not always improve, which calls for further investigation. The most noticeable difference can be observed for *AMS* where the BA improves from 85.13 % to 94.81 %. This steep rise in performance can be explained by inspecting Figure 7 where the *AMS* feature is plotted for different f_r and healthy (blue), artificially damaged (red) and naturally damaged (yellow) bearings. A clear margin between the artificially damaged and the healthy bearing can be observed. The *AMS* feature values for the naturally damaged bearing populate the region between the two other classes. However, for $f_r > 30$ Hz also a clear distinction can be made between the naturally damaged and the healthy bearing. The observation that the *AMS* feature is well-suited for the classification of the artificially damaged bearing is further validated by a BA of 99.65 %, when only test instances from the artificially damaged bearing are considered in the test set. However, the classification performance increases also significantly for all other feature sets in this scenario.

Inspecting the results for Sensor B (cf. Table 6), which is located at the opposite end of the induction engine, it becomes evident that the larger distance between bearing and sensor heavily impacts the classification results. From Table 6, a significant drop in performance by approximately 10 percentage points for *TD* and *SD* and 7 percentage points for *ENV AMP* and *AMS* in terms of BA can be observed. The TNR drops significantly, i.e., many bearing faults remain undetected and are misclassified as healthy. Solely the MFCCs and to some extent the *AMS + f_r* feature sets yield high TNRs. As the MFCCs have proven to be the most effective features, it was further investigated how the performance changes with a varying number of MFCCs. Accuracy, FPR, FNR and BA in dependence of the number of used MFCCs for Sensor B are shown in Figure 8. In Figure 8 a significant performance gain can be observed for Sensor B when the number of included MFCCs rises above 23 and the FNR, the number of healthy bearings that are misclassified, drops noticeably. Yet, for Sensor A, very good classification results can already be obtained after using only the first five MFCCs. A similar behavior can also be observed for Sensor B, when a few instances of damaged bearings are also included in the evaluation dataset that is used for optimizing the SVM parameters. The use of such an evaluation dataset did not, however, improve the classification performance for the other feature sets to the same extent.

4. SUMMARY

In this article, well-established features for the detection of bearing faults together with state-of-the-art features from the audio signal processing domain were investigated. In our approach, bearing faults were classified following an anomaly detection approach in which a One-Class SVM was trained solely with data obtained from healthy bearings. Extensive experiments, conducted on a state-of-the-art commuter railway induction engine in a challenging scenario, proved that the MFCCs and the *AMS* feature can be highly effective features for the detection of bearing faults.

Table 5: Classification results for Sensor A.

Feature Set	TPR	TNR	BA
TD	97.81	80.51	89.16
TD + f_r	97.22	82.58	89.90
SD	98.53	78.13	88.33
SD + f_r	99.34	66.29	82.81
ENV AMP	85.55	95.32	90.44
ENV AMP + f_r	86.73	96.44	91.58
AMS	97.35	72.92	85.13
AMS + f_r	99.33	90.29	94.81
MFCC	98.23	99.34	98.79
MFCC + f_r	98.20	99.62	98.91

Table 6: Classification results for Sensor B.

Feature Set	TPR	TNR	BA
TD	96.67	60.91	78.79
TD + f_r	97.57	62.11	79.84
SD	97.66	65.37	81.51
SD + f_r	99.34	66.29	82.81
ENV AMP	96.65	71.99	84.32
ENV AMP + f_r	97.96	71.24	84.60
AMS	94.19	50.37	72.28
AMS + f_r	91.15	82.95	87.05
MFCC	98.30	94.39	96.35
MFCC + f_r	98.25	95.32	96.79




Figure 7: AMS feature in dependency of the rotational frequency for healthy (blue), artificially (red) and naturally damaged (yellow) bearing.




Figure 8: Accuracy, FPR, FNR and BA in dependency of the number of MFCCs for Sensor B

REFERENCES

- [1] Randall, R.B. and Antoni, J. Rolling element bearing diagnostics—a tutorial. *Mechanical Systems and Signal Processing*, 25(2):485–520, 2011.
- [2] Zhang, Z., Entezami, M., Stewart, E., and Roberts, C. Enhanced fault diagnosis of roller bearing elements using a combination of empirical mode decomposition and minimum entropy deconvolution. *Proc. Institution of Mechanical Engineers, Part C: Journal of Mechanical*

Engineering Science, 231(4):655–671, 2017.

- [3] Wei, Y., Li, Y., Xu, M., and Huang, W. A review of early fault diagnosis approaches and their applications in rotating machinery. *Entropy*, 21(4):409–435, Apr 2019.
- [4] Elasha, F., Ruiz-Carcel, C., Mba, D., and Chandra, P. A comparative study of the effectiveness of adaptive filter algorithms, spectral kurtosis and linear prediction in detection of a naturally degraded bearing in a gearbox. *Journal of Failure Analysis and Prevention*, 14(5):623–636, Jan 2014.
- [5] Hamadache, M., Jung, J. H., Park, J., and Youn, B. D. A comprehensive review of artificial intelligence-based approaches for rolling element bearing PHM: shallow and deep learning. *JMST Advances*, pages 1–27, 2019.
- [6] Neupane, D. and Seok, J. Bearing Fault Detection and Diagnosis using Case Western Reserve University Dataset with Deep Learning Approaches: A Review. *IEEE Access*, pages 1–26, 2020.
- [7] Oppenheim, A. V. and Schaffer, R. W. *Discrete-Time Signal Processing*. Prentice Hall Press, USA, 3rd edition, 2009.
- [8] Vakharia, V., Gupta, V. K., and Kankar, P. K. A comparison of feature ranking techniques for fault diagnosis of ball bearing. *Soft Computing*, 20(4):1601–1619, April 2016.
- [9] Vargas-Machuca, J., Garcia, F., and Coronado, A. M. Detailed Comparison of Methods for Classifying Bearing Failures Using Noisy Measurements. *Journal of Failure Analysis and Prevention*, 20(3):744–754, June 2020.
- [10] Nayana, B. R. and Geethanjali, P. Analysis of Statistical Time-Domain Features Effectiveness in Identification of Bearing Faults From Vibration Signal. *IEEE Sensors Journal*, 17(17):5618–5625, September 2017.
- [11] BR, N. and Geethanjali, P. Feature Extraction for Bearing Fault Diagnosis in Noisy Environment: A Study. In *2019 Innovations in Power and Advanced Computing Technologies (i-PACT)*, volume 1, pages 1–5, March 2019.
- [12] Arun, P., Abraham Lincon, S., and Prabhakaran, N. Detection and characterization of bearing faults from the frequency domain features of vibration. *IETE Journal of Research*, 64(5):634–647, 2018.
- [13] Yuan, Y., Chen, C., and Yuan, Y. Fault detection of rolling bearing based on principal component analysis and empirical mode decomposition. *AIMS Mathematics*, 5(6):5916–5938, 2020.
- [14] Lawrence, R. W. and Schaffer, R. W. *Introduction to Digital Speech Processing*. Now Publishers Inc., MA, USA, 2007.
- [15] Kollmeier, B. and Koch, R. Speech enhancement based on physiological and psychoacoustical models of modulation perception and binaural interaction. *The Journal of the Acoustical Society of America*, 95(3):1593–1602, March 1994.
- [16] Tchorz, J. Combination of amplitude modulation spectrogram features and MFCCs for acoustic scene classification. Technical report, DCASE2018 Challenge, September 2018.
- [17] Schölkopf, B., Platt, J. C., Shawe-Taylor, J., Smola, A. J., and Williamson, R. C. Estimating the support of a high-dimensional distribution. *Neural Computation*, 13(7):1443–1471, 2001.
- [18] Bishop, C. M. *Pattern Recognition and Machine Learning*. Springer, Berlin, DE, 2006.