Using Optimized Time-Frequency Representations for Acoustic Quality Control of Motors

Christoph Heitz
Centre for Data Analysis and Model Building
University Freiburg
Albertstr. 26-28
D-79104 Freiburg i. Br.

January 1995

Abstract

In this article a problem of industrial quality control is investigated. Electric motors have to be tested for proper working by analyzing their emitted noise signal in regular running. The noise signal being (approximately) periodic, it can be seen as a sequence of single signal segments $p_i(t), i = 1, \ldots, N;\ t = 1, \ldots, L_i$ with approximately equal length L_i. These segments are represented in the time-frequency space by their time-frequency representations (TFR) $C_i(t, \nu; \Phi)$. For each of the two considered fault classes "beating" and "gear noise" a fault parameter is defined based on the sequence of periods. This parameter implements a coarse knowledge of the corresponding physical fault.

Since the representation of the signal segments depends explicitely on the kernel function Φ of the TFR, the ability of the fault parameter to discriminate between good and defective motors can be optimized by varying Φ. The resulting optimum kernel functions enables a reliable classification.

The computational efficiency of the classification algorithm is sufficiently high to make possible its application in a realistic industrial production environment.

1 Introduction

In the present study a problem of industrial quality control is investigated: After the assembly of small electric motors with an attached gear cog-wheel, the motor-gear-unit (which will be referred to by “motor” in the following) has to be checked on proper working. This is done by an aural check: human testers listen to the emitted noise in regular running, which renders a lot of information about the state of the motor. Thus, several fault classes can be detected and distinguished.

In the present article we demonstrate that with optimized time-frequency representations (TFRs) the faulty motors can be detected by an automatic algorithm. We examine the two fault classes “beating” and “gear noise”. The other fault classes (e.g. “too loud”) can be detected by simple methods and are not regarded here.

The structure of the noise signal under consideration can be derived from the physics of noise production: The motor-gear unit possesses two different axles: first the motor axle, rotating with approximately 100 Hz, and, second, the gear axle whose rotation velocity is reduced by the ratio 35:1 with respect to the motor axle. Thus, for the analysis of these signals it is natural to divide the signal $f(t)$ into succeeding segments, each segment consisting of the signal emitted in a time periode which is needed for the motor axle to rotate by an angle of 2π. Since the rotation frequency of the motor axle is not constant, the segments can have different lengths.

After a complete rotation, the motor axle is in the same physical state than before and the gear axle has rotated by an angle of $\frac{2\pi}{35}$. Thus we expect that the emitted noise in the next cycle is similar to that of the considered cycle, i.e. the different segments should be similar one to the other. The exactly identical physical state of the motor-gear unit is restored after 35 complete rotations of the motor axle. Thus, the segments 1,36,71,... should produce, apart from stochastical fluctuations, the same signal.

For evaluating the similarity between different segments it is convenient to perform a time normalization: each segment is normalized to a common length L. Thus, the signal segments can be compared by standard techniques.

The central idea is to replace the signal $f(t)$ by its TFR $C_f(t, \nu; \Phi)$ which can be viewed as another representation of the signal $f(t)$. The TFR can be segmented and time normalized as before. Thus, we get a sequence of time-frequency pieces $C_i(t, \nu), t = 1, \ldots, L$, which yields a description of the signal in terms of elements which are well adopted to its internal structure. The segments C_i being similar to each other, we can regard them as different realizations of an underlying random variable. For motors with gear noise it will be shown that the variance of this variable is larger than the variance for good motors. For beating motors the expectation value for one of the 35 segments needed for a full rotation of the gear axle, differs significantly from the expectation values of the other segments.

The advantage of using a TFR instead of the signal directly is that the representation of a signal in time-frequency space depends on an arbitrary kernel function $\Phi(\xi, \tau)$, defining the TFR. The freedom in chosing an adequate kernel can be taken advantage of for the analysis task to be solved: Given a test quantity which should serve to distinguish between good and defective motors, the discrimination performance of this quantity can be maximized by varying the kernel function Φ. Thus, for given signal classes to be distinguished, an optimum kernel, i.e. an optimum TFR, can be found which is best fitted to discriminate between the classes by means of the considered test quantity. Frequently, in the time-frequency formulation a much superior discrimination can be achieved than in the time or, respectively, frequency formulation.

In the present article we show how this idea can be used to construct an algorithm for detecting defective motors by using optimized TFRs. The underlying theory of optimized TFRs is described in detail in [3].

The article is organized as follows: In Section 2 the data set under consideration is presented. In Section 3 the general framework of Cohen's class of TFRs is introduced and applied to the specific structure of our data. The algorithms for the segmentation of the signal into succeeding periods and for the time-normalization are described in Section 4. In Sections 5 and 6 two different fault classes are considered. It is shown how, in each case, an optimum kernel function for the TFR can be found such that the resulting representation of the signal yields a good separation between the defective motors and the good ones. The performance of the classifier is examined in more detail in Section 7. Finally, in Section 8 some remarks on the required computing time for the proposed classification scheme are made.

2 The data

For the present study 200 motor noise signals have been available. The signals had been recorded by air pressure microphones in a noisy environment. The length of each signal is 1.66 s, sampled at a sample rate of 24 kHz, corresponding to 40000 data points. Depending on the exact motor frequency, this length corresponds to 160 to 200 rotations of the motor axle (see Fig. 1).

Four human testers had previously jugded each motor. For the two features "beating" and "gear noise" a number from 1 (not present) to 5 (highly present) had to be assigned to the signals noise, the discriminating value between "good" and "defective" motors being fixed at 3. Additionally, the testing persons had to assign a general goodness value for each motor, regarding also additional sound features which are irrelevant for our purposes. This number also had a range from 1 (good) to 5 (defective).

Each test person had to judge each signal two times with a time distance of two weeks. Thus, for each motor 8 judgements had been given which, however,

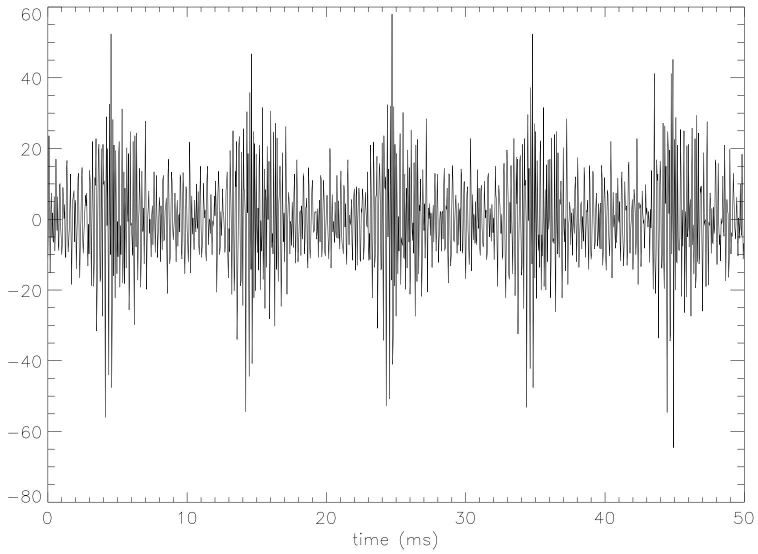

Figure 1: Example of the acoustic signal of a motor. The periodicity of about 1 ms, corresponding to a complete rotation of the motor axle, can be seen clearly.

show a relatively large variability[1]. Thus, for most of the motors it cannot be decided precisely if they are faulty or not. For our evaluations a motor has been regarded to be faulty (in the corresponding fault class) if at least 5 judgements of the corresponding feature yielded a value larger than 3. A motor was regarded to be good if at least 6 judgements of the goodness number were less than 3. In this way we got 32 signals assumed to be good, 9 signals of beating motors and 11 signals for the gear-noise class.

3 Time-frequency representations

A member of Cohen's class of TFRs $C_f(t,\nu;\Phi)$ of a real signal $f(t), t = 1,\ldots,L$ is defined by [1]

$$C_f(t,\nu;\Phi) = \sum_{\xi=-L/2}^{L/2} \sum_{\tau=-\hat{\tau}}^{\hat{\tau}} A_f(\xi,\tau)\Phi(\xi,\tau)\, e^{-2\pi i(\frac{\xi t}{L}+\frac{\nu\tau}{\hat{\tau}})}, \quad t = 1,\ldots,L, \quad \nu = 0,\ldots,\hat{\tau}$$

where

$$A_f(\xi,\tau) = \frac{1}{L}\sum_{\mu=1}^{N} f(\mu+\tau)f(\mu-\tau)\, e^{2\pi i\xi\frac{\mu}{L}}, \qquad \tau = -\hat{\tau},\ldots,\hat{\tau}, \quad \xi = -\frac{L}{2},\ldots,\frac{L}{2}$$

[1]A previous study showed a variability of 28% for the general goodness value [4].

is called the *ambiguity function* of f. Each TFR is defined by its kernel function $\Phi(\xi, \tau)$ which, in principle, can be chosen arbitrarily. Here, however, we assumed that $\Phi(\xi, \tau) = 0$ for $|\tau| > \hat{\tau}$. Furthermore we admit only real kernels (see [3]).

For a subsequent optimization, the kernel must be parametrized. We chose the following parametrization

$$\Phi(\xi, \tau) = \begin{cases} e^{-(\frac{\xi}{\xi_0})^2} e^{-(\frac{\tau}{\tau_0})^2} & \text{for } |\tau| < \hat{\tau} \\ 0 & \text{else} \end{cases} \tag{1}$$

with $\hat{\tau} = 20$. The resulting TFR is known as Smoothed Pseudo-Wigner Distribution [1].

For the analysis of the motor signals, instead of the original signal $f(t)$ we use the TFR $C_f(t, \nu; \Phi)$, where the kernel Φ is to be optimized by searching for the optimum parameter values ξ_0 and τ_0.

This TFR is now segmentated (see Section 4) and, for each frequency ν, time normalized to a length $L_2 = 256$. This leads to a sequence of $\tilde{C}_i(t, \nu; \Phi), t = 1, \dots, L_2,\ \nu = 0, \dots, \hat{\tau}$. Each $\tilde{C}_i$ can be viewed as a vector in a $L_2 \cdot (\hat{\tau} + 1)$-dimensional linear space Ω. In this linear space a vector product and a norm can be defined as usual (see [3]). By normalizing each segment we arrive at the sequence

$$C_i = (\tilde{C}_i)^N \equiv \frac{\tilde{C}_i}{\|\tilde{C}_i\|}$$

Here $(\cdot)^N$ means "normalized".

In the following we regard only vectors in Ω with unit norm. The distance between two such vectors is given by $D(C_1, C_2) = 1 - < C_1, C_2 >= \frac{1}{2}\|C_1 - C_2\|^2$. For comparison we also regard a distance measure in the original time representation, defined by $d(f_1, f_2) = 1 - < \tilde{f}_1, \tilde{f}_2 >= \frac{1}{2}\|\tilde{f}_1 - \tilde{f}_2\|^2$, where $\tilde{f}(t)$ is the signal $f(t)$ after time normalization to a length of L_2 and variance normalization to unit norm. Here the norm is the usual Euclidian norm in the L_2-dimensional linear space.

In the succeeding sections we will interpret the C_i as realizations of a random variable $\mathbf{C}$. We will use the estimate of its expectation value which is defined as follows: for a given set of $C_i, i = 1, \dots, n$ we look for the vector C^0 with unit norm for which the mean distance $\frac{1}{n}\sum_{i=1}^{n} D(C_i, C^0)$ is minimum. It can be seen easily that C^0 is given by the normalized mean of the C_i:

$$C^0 = \overline{C}^N = \left(\frac{1}{n}\sum_{i=1}^{n} C_i\right)^N \tag{2}$$

4 Segmentation and time normalization

In order to segmentate the signals the following algorithm was used: For the time points $\hat{t} = 1000, 1400, 1800, \dots$ an estimate of the local (time dependent) segment

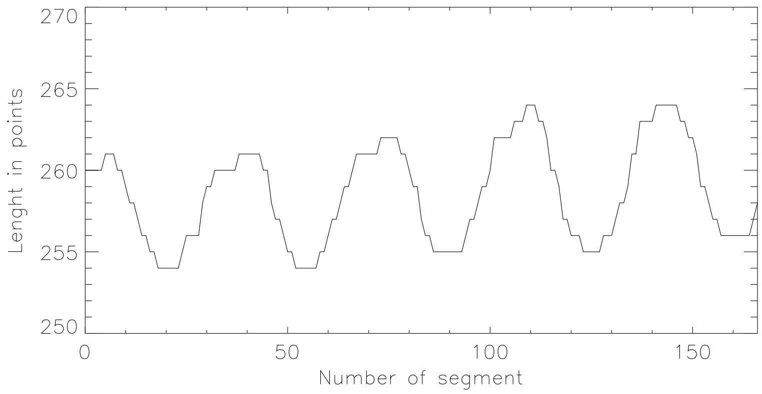

Figure 2: Estimated lengths of the succeeding segments of one motor signal. The underlying periodicity of 35 segments, corresponding to a complete rotation of the gear axle, can be well seen in this example since a periodic change of the motor frequency is present. This is not the case for all motors.

length $\tilde{L}(\hat{t})$ was obtained by multiplying the signal with a rectangular window of length 2000, centered at $\hat{t}$, and searching the maximum of the autocorrelation function for the windowed signal in the lag interval $\tau = 200, \ldots, 300$. For the leaving time points $t \neq \hat{t}$ the segment length was assumed to be $\tilde{L}(\hat{t}')$, $\hat{t}'$ being the nearest considered $\hat{t}$. Thus, for each time point, a local segment length was estimated.

Beginning with an arbitrary starting point t_1 for the first segment, the starting point t_2 for the subsequent segment can be determined by adding $\tilde{L}(t_1)$. In this way, a sequence of starting points is constructed[2], defining segment i to be the interval $f(t), t \in [t_i, t_{i+1} - 1]$. Its length, corresponding to the local rotation frequency of the motor axle for this segment, is given simply by $t_{i+1} - t_i - 1$. In Fig. 2 the obtained lengths of the succeeding segments for one signal is shown.

The time normalization for the time-frequency segments was performed, seperately for each frequency ν, by a simple scaling algorithm, which bases on the following idea: If a signal $g(t), t = 1, \ldots, L_1$ is a segment from a signal which is periodical with period L_1 (an assumption which is reasonable in our situation, not only for the signal segment directly, but also for the functions $C_f(t, \nu)$ for fixed frequency ν) then Fourier transformation

$$\tilde{g}_\nu = \sum_{t=1}^{L_1} g(t)\, e^{2\pi i t \nu}, \qquad \nu = -\frac{L_1/2}{L_1}, -\frac{L_1/2 - 1}{L_1}, \ldots, \frac{L_1/2 - 1}{L_1}, \frac{L_1/2}{L_1}$$

yields the Fourier coefficients $\tilde{g}_\nu$ of $g(t)$. These coefficients are now interpreted

[2]Its has been verified that the estimation errors do not cumulate but cancel each other. A systematic shift of the estimated starting points has not been observed.

as Fourier coefficients from a signal of length L_2. To construct this signal by an inverse Fourier transformation, we need exactly L_2 coefficents. Therefore, for $L_2 > L_1$ we pad with zeros to produce the leaving $(L_2 - L_1)$ coefficients. If, on the other hand, $L_1 > L_2$, the coefficients for the highest frequencies are ignored for the inverse Fourier transformation.

For computational efficiency we do not perform a FFT with a length L_1. Instead $g(t)$ is padded with zeros to a length L_1' which is chosen to be a power of 2. The values $\tilde{g}_\nu$ for $\nu = -\frac{L_1/2}{L_1}, -\frac{L_1/2-1}{L_1}, \ldots, \frac{L_1/2-1}{L_1}, \frac{L_1/2}{L_1}$ are estimated by linear extrapolation of the neighbouring, actually calculated, $\tilde{g}_\nu$ (with $\nu = -\frac{L_1'/2}{L_1'}, \ldots, \frac{L_1'/2}{L_1'}$). For $L_1' \to \infty$ this yields the correct values. For our purposes a value of $L_1' \geq 4L_1$ has been found to be sufficient.

5 Gear noise

One of the faults to be detected is specified by the aural impression "gear noise". It has its origin in a not exactly known defect in the gear unit. Human beings describe the resulting motor sound as "rough".

In our framework, this impression could be represented by two different signal features. First, the expectation value C^0 for the signal periods of the defective motors could differ from the analogous expectation value for the good motors. Thus, the change in sound would correspond to a change of the mean of the random variable $\mathbf{C}$. In contrast, it could as well be some sort of irregularity, expressed in the *variance* of the underlying random variable $\mathbf{C}$, which leads to the "rough" sound impression, without significant difference of the means of the two variables.

The first possibility could not be verified: nor the time signals $p_i(t)$ for the single motor cycles nor their spectra nor their TFRs C_Φ with various kernels showed a significant difference between good and defective motors.

Instead, the second alternative was found to be adequate. In order to quantify the irregularity, the following test quantity was constructed: All C_i with $i \bmod 35 = k, k = 0, \ldots, 34$ were viewed as realizations of the random variables $\mathbf{C}_k$. Their expectation values are given by $C_k^0 = \left(\sum_{i \bmod 35 = k} C_i\right)^N$. Let ϵ_k be the deviation of the C_k^0 from the overall expectation value $C^0 = \left(\sum_i C_i\right)^N$:

$$\epsilon_k = ||C_k^0 - C^0||^2$$

Then the sum

$$S = \sum_{k=0}^{34} \epsilon_k$$

is a measure of the irregularity of the sequence of the segments. The time period for one segment being about 10 ms, the irregularities cannot be resolved by the human ear (unlike in the following example of beating motors) but are perceived as roughness.

For each motor a different S-value is obtained. Thus, we can interpret the calculated S-value as a realization of a random variable $\mathbf{S}$. Precisely speaking, we assume the good motors being described by a random variable $\mathbf{S}_1$, whereas the gear noise motors are described by $\mathbf{S}_2$. The statistical properties of the two random variables have to be estimated based on the available realizations. If the means differ sufficiently regarding the variances, a discrimination of the two classes can be obtained by comparing the calculated S-value for a given motor with a threshold γ.

In fact it is the case, that, for most kernel parameters, the S-values of the gear noise motors are larger than that of the good motors. The discriminating power of the S-value is estimated by the Fisher discrimination measure

$$F = \frac{\overline{S}_1 - \overline{S}_2}{\sqrt{\text{var } S_1} + \sqrt{\text{var } S_2}}$$

where $\overline{S}_i$ is the empirical mean and var S_i the empirical variance of the considered random variable.

Since we used TFRs, S and hence F are depend on Φ: $F = F_\Phi$. Maximizing F_Φ yields the optimum kernel parameters, i.e. the optimum TFR, for the discrimination of the two classes. In Fig. 3 the F-values for a region in the kernel parameter space are shown. The optimum parameters are found to be $(\xi_0, \tau_0) = (1.0, 1.0)$.

In [3] it has been shown that for kernels with $\tau_0 = 0$ or $\xi_0 = 0$ not the complete TFR theory would be needed but, equivalently, smoothed periodograms (i.e. $|\tilde{f}(\nu)|^2$) or smoothed intensities (i.e. $|f(t)|^2$) could be used for the evaluation of the ϵ_k and hence S. Note, however, that neither the kernels with $\tau_0 = 0$ nor the kernels with $\xi_0 = 0$ yield a good F-value (i.e. a good discrimination between the classes). Therefore, in our case, this simplification is not possible.

The cumulative distributions of the S_Φ-values for the optimum kernel parameters is shown in Fig. 4 (a). The maximum distance between the two cumulatives is the non-parametric classification rate, if the threshold γ is chosen to be the S-value where the maximum occurs. It can be seen that the two distributions differ sufficiently to ensure a good separation of the classes. A discimination rate of 96.7 % has been obtained.

Instead of using the C_i as representations of the single signal segments, one could as well try to use the $p_i(t)$ directly. After time normalization a sequence of equally long signals is obtained which could be used to calculate an irregularity measure S which is defined quite analogously. Note, however, that no kernel function appears. Fig. 4(b) shows the empirical distributions of the S-values obtained in this way. Appearently no distinction between the classes is possible.

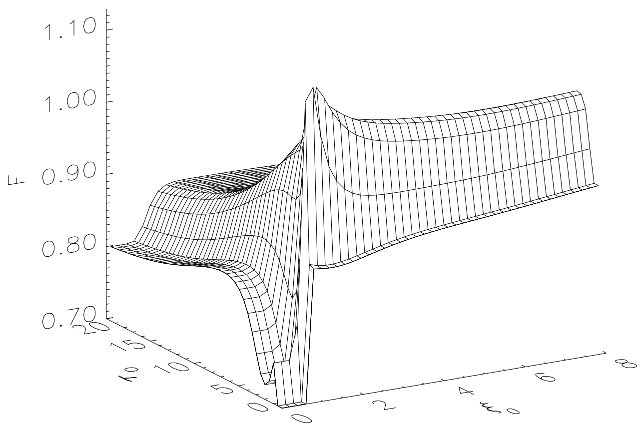

Figure 3: The found F_Φ-values for the kernel parameters in the region $\tau_0 = 0, \ldots, 20, \xi_0 = 0, \ldots, 8$. The maximum value of F is $F = 1.13$ with $\xi_0 = 1.0$ and $\tau_0 = 1.0$.

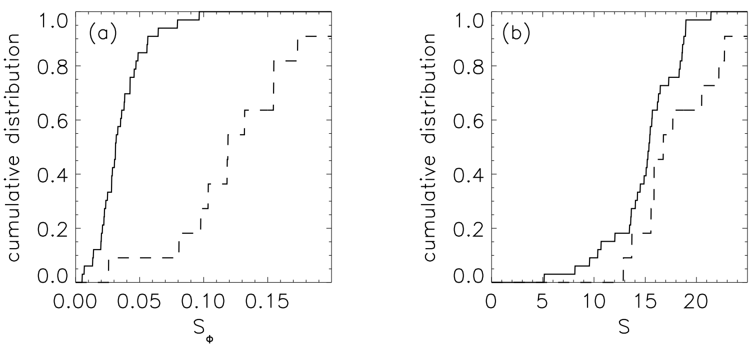

Figure 4: Gear noise motors: (a) Empirical cumulative distributions of the S_Φ-values for the good motors (solid line) and the defective ones (dashed line), where the optimum kernel Φ as stated in text was used. (b) Analogous as (a), but S being evaluated with the orginal time segments p_i instead of the optimum TFRs C_i.

6 Beating motors

Another class of faults is "beating". In regular time intervals of about 0.3 seconds a short "tick" can be heard in the motor's noise. The origin of these beats is some fault in the gear unit: Each time the gear axle passes a certain rotation angle, a beat is produced. Thus, within every 35th rotation of the motor axle, an additional short sound event is produced.

Although these beats can be heard rather well, the total power output of the beats is very low. Simple methods like estimating the local power or spectral methods fail in detecting these beats.

However, by optimizing a TFR to this kind of problem, the beating motors could be detected. The idea is the following: For good motors, the segments should be similar to each other, but, for beating motors, each 35th segment (which includes the additional sound event of the gear fault) should differ significantly from the other segments. The segments are represented by their TFRs C_i, the mutual distance measured by D_Φ. For corrupted segments the distance to the regular segments should be large, whereas for good motors all segments should have a little distance from each other. As explained before, the distance is kernel dependent and can thus be optimized in order to render an optimum distinction between the two motor classes.

Let's assume for the moment that the corrupted segments are the segments number 35,70,105,..., i.e. the segments with $i \bmod 35 = 0$. We regard these segments to be realizations of a random variable $\mathbf{C}_2$, whereas the other segments being realizations of another random variable $\mathbf{C}_1$. We have to test if the two random variables differ, i.e. if the realizations of $\mathbf{C}_2$ differ significantly from the mean of $\mathbf{C}_1$, compared with the mean inner class distance within $\mathbf{C}_1$. A measure of this difference is given by [3]

$$m_\Phi = \frac{\overline{D_\Phi(C_1, C_2)}}{\overline{D_\Phi(C_1, C_1')}}$$

Here $\overline{D_\Phi(C_1, C_2)}$ denotes the mean distance between two members of class 1 and 2, respectively, and $\overline{D_\Phi(C_1, C_1')}$ is the mean inner class distance of class 1. It can be shown [3] that this reduces to

$$m_\Phi = \frac{1 - < \overline{C_1}, \overline{C_2} >}{\frac{n_1}{n_1 - 1}(1 - \|\overline{C_1}\|^2)}$$

where n_1 is the number of the available realizations of $\mathbf{C}_1$, and $\overline{C_1} = \frac{1}{n_1}\sum_{k=1}^{n_1} C_1^k$.

A large m_Φ means a large distance of the expectation values of the random variables. Thus, beating motors should have a larger m_Φ than good ones.

Remember that we assumed the segments with $i \bmod 35 = 0$ to be corrupted. For the 35 sets of (eventually corrupted) segments given by $i \bmod 35 = k, k = 0, \ldots, 34$ we get 35 different m_Φ-values. The maximum of these values is called $\hat{m}_\Phi$. A large $\hat{m}_\Phi$ means that at least one set of segments is corrupted.

Until now nothing has been said about the choice of the kernel though, apparently, $\hat{m}_\Phi$ is kernel dependent. One could imagine that it would be best to use the kernel which maximizes $\hat{m}_\Phi$. Then the corrupted segments could be detected optimally thus indicating a beating motor. For the kernel parametrization as in

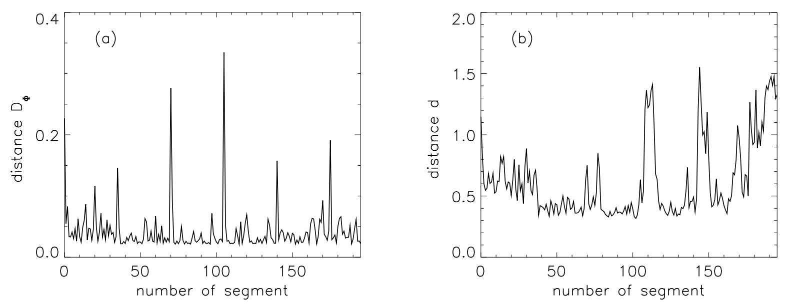

Figure 5: Example of a beating motor: distances of the single segments to the mean of class 1. (a) Distance evaluated with D_Φ for the TFR with maximum $\hat{m}_\Phi$. The corrupted segments which have a distance of 35 segments from one another can be detected easily. (b) Distance evaluated with the the ordinary distance d for the time segments directly, without reference to a TFR.

Eq. 1, Fig. 5(a) gives an example of the result of such an optimization which rendered $\hat{m}_\Phi = 7.6$. When, on the other hand, using the original time representation $p_i(t)$ of the segments and the distance d as in Section 3, a $\hat{m}_\Phi$-value of only 1.2 is achieved (see Fig. 5(b)) .

However, this leads to different kernels for different motors. For the detection of the beating motors, in contrast, we want to find a single TFR for the analysis of all signals. In addition, not only $\hat{m}_\Phi$ must be large for the beating motors, but also $\hat{m}_\Phi$ must be small for the good motors.

For the optimization of the kernel Φ the empirical cumulative distributions of the $\hat{m}_\Phi$-value for the two classes were calculated. The maximum value of the difference between the two cumulatives is the non-parametric classification rate. The optimum kernel was estimated by maximizing this classification rate.

For the found optimum kernel with the parameters $(\xi_0, \tau_0) = (8.0, 2.2)$ a classification rate of 96.7% was achieved. The corresponding empirical cumulative distributions are shown in Fig. 6(a).

The analogous cumulatives for the $\hat{m}$-values evaluated directly on the time segments p_i (instead of C_i) is shown in Fig. 6(b). It can be seen that again no class separation is possible in this case.

7 Determination of the classifier performance

In Sections 5 and 6 it was shown how good motors can be separated from defective ones by postulating a fault parameter and optimizing the kernel of the TFR in order to get maximum discrimination.

The kernel parameters have been estimated on the base of a training set of

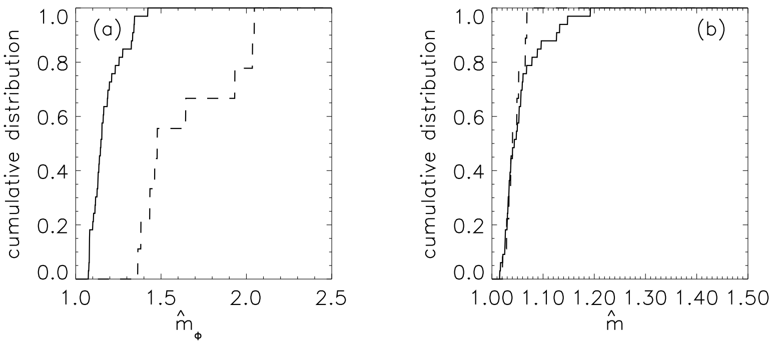

Figure 6: Beating motors: a) Empirical cumulative distributions of the $\hat{m}_\Phi$-values for the good motors (solid line) and the beating motors (dashed line), where the optimum kernel Φ as stated in text was used. (b) Analogous as (a), but here $\hat{m}$ being evaluated with the orginal time segments p_i instead of the optimum TFRs C_i.

good and faulty motors. In order to determine the performance of a classifier, however, it must be tested on an independent test data set. For small data sets, frequently resampling methods are used (see e.g. [2]).

We used the leaving-one-out method to test the two fault detectors: One of the data samples is removed as test data. The optimum kernel parameters are estimated with the other signals as described in Sections 6 and 5. The resulting cumulative distributions of the two classes are calculated and the classification threshold γ is determined by searching the maximum distance of the cumulatives. Then, the test data is classified by calculating its $\hat{m}_\Phi$- or S_Φ-value, respectively. This procedure is repeated for each of the available data. By couting the correct classifications, the classification rate can be estimated.

Applying this method to our data sets, regarding the the discrimination good/beating, two good motors and two beating motors were misclassified, corresponding a classification rate of about 90% for

For the good/gear noise case three of the good motors were classified to be defective, and one of the faulty motors was classified to be good. This corresponds to a classification rate of about 93%.

Given the optimum kernels Φ_1 and Φ_2, each motor can be represented as a point in a two-dimensional space which is spanned by the variables $\hat{m}_{\Phi_1}$ and S_{Φ_2}. Note that for the calculation of the $\hat{m}_{\Phi_1}$ and S_{Φ_2} values different kernel parameters have to be used. Fig. 7 shows the distribution of the data in this feature space with the kernel parameters of Sections 5 and 6. The three classes are located in different areas of this feature space and are relatively well clustered. The only outlier is one of the gear-noise motors which has a very untypical S_Φ-value.

It should be emphasized that the above given values for the classifier perfor-

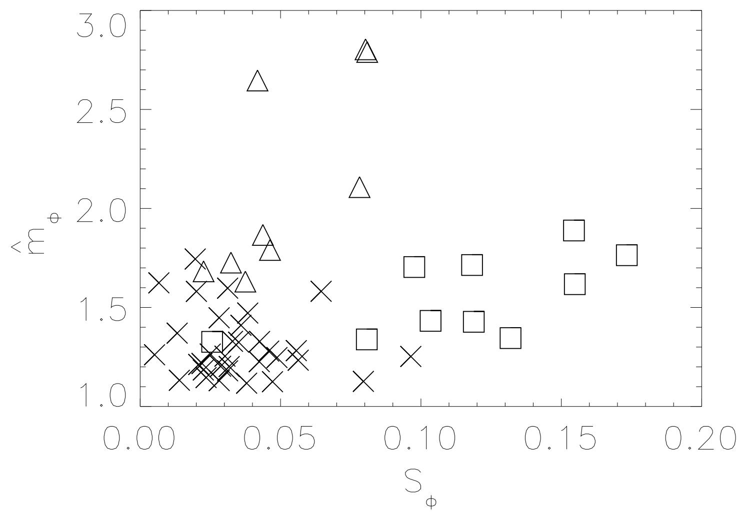

Figure 7: Distribution of the three considered classes of motors (gear noise: □, beating: △, good: ×) in the two-dimensional space which is spanned by $\hat{m}_{\Phi_1}$ and S_{Φ_2}. One of the gear noise motors is outside the plot region because of a S_Φ-value which is too large.

mance are to be taken not too literal, since the data base is rather small. For a better statistic more data should be used which, however, were not available to the author. Fig. 7 gives nevertheless an impression of the algorithm's behaviour.

8 Computational Efficiency

An important feature for an algorithm is its computational efficiency. When dealing with TFRs, frequently the computational costs are very high. However, this depends strongly on the used kernel.

In our case, the situation is fairly good: since the algorithm relies only on inner products and means of TFRs, instead of the TFR $C_f(t,\nu)$ its Fourier transform with respect to ν, denoted by $Y(t,\tau)$, can be used equivalently.

For product kernels $\Phi(\xi,\tau) = g_1(\xi)\, g_2(\tau)$, Y is calculated by

1. calculating the array $X(t,\tau) = f(t+\tau)f(t-\tau)\, g_2(\tau)$

2. smoothing $X(t,\tau)$ by Fourier transformation of $X(t,\tau)$ with respect to t, multiplying with $g_1(\xi)$ and backward Fourier transformation, resulting in the final $Y(t,\tau)$.

When using a smoothed Pseudo-Wigner distribution with the parametrization as proposed, $g_2(\tau)$ can be set to 0 for $\tau > 4\tau_0$. In our case we got $\tau_0 = 1.0$

(gear noise) and $\tau_0 = 2.2$ (beating motors) from the optimization procedure as explained in Sections 5 and 6, so $Y(t, \tau)$ must be calculated only for $\tau = 0, \ldots, 9$. Thus, the proposed algorithm is extremely fast compared with, e.g., algorithms using the Wigner distribution[3].

With an IBM RISC-6000, model 250, workstation the following run times were obtained:

- starting program, reading data, etc. : 1.0 sec
- segmentation of the signal: 3.9 sec
- Calculation of the ambiguity functions for each segment (based on the time normalized TFR of the corresponding signal segment): 9.7 sec
- Calculation of the values $\hat{m}_{\Phi_1}$ and S_{Φ_2} with the two optimized kernels Φ_1 and Φ_2 and final classification: 0.5 sec

Thus, with a total computing time of 15.1 seconds a motor's signal could be classified. It must be noted, however, that the author payed little attention on optimization of the program's code. With a more efficient code a further reduction of computing time can be expected.

It should be stressed that the obtained run times are not too far away from the requirements of realistic industrial production lines. Thus, an application of the algorithm should be possible.

9 Conclusions

In the present article we showed how optimized TFRs can be used to detect defective motors from their emitted noise signals in regular running. The sound which is produced for each of the succeeding rotations of the motor axle is represented by a time-normalized TFR array. Based on the sequence of these arrays, for the two fault classes under consideration test quantities could be derived from an intuitive unterstanding of the physics of the motor defect. Note that, for doing this, neither an exact knowledge of the fault mechanism nor a complicated physical model for the sound production is needed.

In the case of gear noise the test quantity corresponds to the degree of the "roughness" of sound, for the "beating" motors it corresponds to the extent of the existence of outliers. Good motors are represented by small number of both feature values, whereas a faulty motor shows a large value of the corresponding quantity.

[3]In general, of course, this depends on the used kernel parameters. The larger τ_0, the higher are the computational costs. The value of ξ_0, in contrast, does not relevantly influence the computing time.

It has been shown how the discrimination power of the features can be optimized by varying the kernel function $\Phi(\xi, \tau)$ of the TFR. For the present study, a Smoothed Pseudo-Wigner distribution parametrization was chosen. It should be noted, however, that this is not necessary. Other parametrizations could be chosen as well.

We have demonstrated furthermore that, with the optimum kernel parameters, the resulting feature values yield a representation of the signals in a two-dimensional feature space which separates the good from the defective motors as well as the two fault classes one from the other.

The computational efficiency of the used algorithm has been quantified. The required computation time is still too long to allow a real-time application, but by optimization of the program's code an application in a realistic industrial production environment should be possible.

It should be emphasized that the proposed classification scheme is not restricted to the special case regarded in this article. The main ideas can be transferred to any other periodic or quasi-periodic signals: Let every signal period be represented by its (eventuelly time-normalized) TFR, construct plausible characteristic features and optimize the kernel parameters in such a way that they discriminate as well as possible between the different classes. A wide field of possible applications can thus be covered. It is the author's experience that, due to the flexibility of the TFRs, a large variety of signal classes can be treated in this way (see also [3]).

10 Acknowledgements

The author likes to thank Prof. Dr. J. Honerkamp for his stimulating interest in this work. Many thanks go to Dr. K. Heldmann and Dr. W. Keiper from the Bosch GmbH, Stuttgart, Germany, for fruitful discussions and comments and for making their data available.

References

[1] L. Cohen. Time-frequency–distributions – a review. *Proc. IEEE*, 77(7):941–981, July 1989.

[2] Richard O. Duda and Peter E Hart. *Pattern Classification and Scene Analysis.* John Wiley & Sons, 1973.

[3] C. Heitz. Optimized time-frequency representations for the classification and detection of signals. *Applied Signal Processing.* submitted for publication.

[4] Klaus Heldmann. *Wahrnehmung, gehörgerechte Analyse und Merkmalsextraktion technischer Schalle.* Number 109, Reihe 17. VDI-Verlag, 1994.