ICSV13 - Vienna

The Thirteenth International Congress on Sound and Vibration

Vienna, Austria, July 2-6, 2006

Time-Frequency Segmentation for Engine Speed Monitoring

Fabien Millioz* and Nadine Martin

Laboratory of Image and Signal (LIS)
961 rue de la Houille Blanche, BP 46
38402 Saint Martin d'Hères Cedex, France
email: 1stname.2ndname@lis.inpg.fr

Abstract

This paper presents a segmentation algorithm of a Time-Frequency Representation, which automatically selects time-frequency patterns containing signal of interest. Considering a deterministic non-stationary signal embedded in a white Gaussian noise, we know that the real and imaginary parts of the Short Time Fourier Transform coefficients (STFT) have a Gaussian distribution. We already proposed an unsupervised segmentation based on the spectrogram, the squared modulus of the STFT, whose coefficients have a non-central chi-square distribution. In order to keep simple Gaussian distributions and the phase information, we consider here only real and imaginary part of the STFT.

We first highlight the difference existing between the variance of these real and imaginary parts. The ratio of these variances is a function of the nature and length of the analysis window.

Estimated on a cell around each time-frequency location, these variances are used to determine if a time-frequency location contains deterministic signal or noise only. Given that the noise variance is unknown, an iterative algorithm is proposed.

In addition of the STFT's parameters, which are the size and shape of the analysis window, the overlap between two consecutive windows, and the amount of zero padding, three different parameters control the segmentation. The most important one is the kurtosis of the distribution, calculated on time-frequency coefficients supposed to contain noise only. It is used to define a stop criterion of the segmentation and permits the monitoring of the signal pattern segmentation. The influence of the other parameters on the segmentation result is also discussed.

This tool is applied to monitor a three-phase AC induction motor on the test bench GOTIX of the laboratory. Sensors measure vibration, torque and phase tensions and intensities. Segmented patterns provide information about time evolution of the spectral energy, and permit the tracking of the engine speed.

INTRODUCTION

In nonstationnary signal analysis, Time-Frequency Representations (TFR) are efficient way of describing the spectral energy along time. In order to help in these analysis, we already proposed TFR segmentation algorithms, which automatically select time-frequency patterns containing signal of interest. A first one is based upon spectrogram ([1], [2]), and have successfully been applied to automatic speech recognition ([3], [4]).

More recently, we proposed a second algorithm based directly upon the Short Time Fourier Transform (STFT) [5]. Theses segmentations are performed by taking as a model of signal a deterministic signal $d[m]$ embedded in a white Gaussian noise $n[m]$

$$x[m] = d[m] + n[m]. \tag{1}$$

$d[m]$ is the signal of interest to be segmented, which contains all the nonstationnary parts of $x[m]$. We derive the TFR coefficient probability distributions from this model. At each iteration, the segmentation algorithm estimates the noise level, and then uses the statistical features of each time-frequency location to determine either the location contains a deterministic part or noise only. Locations containing deterministic part are then segmented into patterns.

In this paper, we focus on the segmentation algorithm based on the STFT. We firstly describe the distributions of the real and imaginary coefficients of the STFT, in order to introduce the algorithm in a second section. Three parameters control the segmentation in addition of the STFT parameters: a threshold on the kurtosis calculated on time-frequency points supposed to contain only noise; a false alarm probability on the local variance distribution of the locations supposed to contain only noise, which determines at a given iteration what points may contain deterministic signal; and the proportion of these points to be segmented before switching to the next iteration.

In the third section Monte Carlo simulations assess the influence of the three parameters of the segmentation algorithm. We finally apply this method to signals issued from a three-phase AC induction motor of the test bench GOTIX of the laboratory, in order to track the engine speed.

REAL AND IMAGINARY PART STFT DISTRIBUTIONS

The STFT $X_\phi[n,k]$ of a discrete time signal $x[m]$ is the succession of Fourier Transforms of N windowed overlapping segments centered on the time index n. Spectral content of $x[m]$ is estimated around each instant n. $X_\phi[n,k]$ is defined by:

$$X_\phi[n,k] = \sum_{m=n-\frac{M_\phi-1}{2}}^{n+\frac{M_\phi-1}{2}} x[m]\phi[m-n]e^{-2i\pi k\frac{m}{M_\phi+Z}}, \tag{2}$$

where ϕ is the M_ϕ-length window function, Z the zero padding and k the frequency index. $\phi[m]$ is normalized as

$$\sum \phi[m]^2 = 1. \tag{3}$$

Assuming model (1), real part $X^r_\phi[n,k]$ and imaginary part $X^i_\phi[n,k]$ are a sum of Gaussian variables, so they have a Gaussian distribution. It's easy to show that

$$X^r_\phi[n,k] \sim \mathcal{N}\left(D^r_\phi[n,k], \alpha[n,k]\sigma^2\right), \tag{4}$$

$$X^i_\phi[n,k] \sim \mathcal{N}\left(D^i_\phi[n,k], (1-\alpha[n,k])\sigma^2\right), \tag{5}$$

where σ^2 is the variance of the noise $n[m]$, $D^r_\phi[n,k]$ and $D^i_\phi[n,k]$ are the real and imaginary parts of the STFT of $d[m]$ respectively, and $\alpha[n,k]$ a function fully determined by the STFT parameters. $\alpha[n,k]$ is approximately equal to 0.5 except for high and low values of k [5]. We can define new random variables of constant variance σ^2 such as

$$X^r_\phi[n,k]' = \frac{X^r_\phi[n,k]}{\sqrt{\alpha[n,k]}}, \qquad X^i_\phi[n,k]' = \frac{X^i_\phi[n,k]}{\sqrt{1-\alpha[n,k]}}. \tag{6}$$

In summary, the choice of Gaussian model (1) leads to a simple Gaussian distribution in the TFR, where the mean depends only on the signal to segment $d[m]$, and the variance depends only on the unknown noise variance σ^2.

SEGMENTATION ALGORITHM

The algorithm has to determine either a time-frequency location contains a deterministic part or not, thus it has to determine if the time-frequency distribution of this location has a non-zero mean or not. To discriminate these locations, the noise level needs to be estimated. We show in [5] that estimating the noise level with points containing non-zero mean leads to an overestimation. We thus propose an iterative process: at each iteration, noise variance is estimated from non-segmented points, so each estimation leads to a less overestimated value.

At each iteration, we first determine what points called candidates are suitable for being segmented, and create or rise patterns with these candidates.

Determination of the candidates to the segmentation

To determine which points are supposed to contain a deterministic part, we estimate the variance of a (n,k) site on a small cell $C_{n,k}$ of P points centered on this site. Local variance estimator writes

$$\widehat{\sigma^2}[n,k] = \frac{1}{P}\sum \left(X^r_\phi[n,k]'\right)^2. \tag{7}$$

For time-frequency coefficients having a non deterministic part, we derive that $\widehat{\sigma^2}[n,k]$ has a χ^2 distribution

$$\widehat{\sigma^2}_{\text{noise only}}[n,k] \sim \frac{\sigma^2}{\delta}\chi^2_\delta \tag{8}$$

where δ is the degree of freedom, verifying $\delta < P$. δ is a deterministic parameter, depending on the STFT parameters only. The only unknown parameter of this distribution is the noise level σ^2, which is estimated at each iteration.

In order to choose the candidates to the segmentation, we define then a threshold t_{σ^2}

$$t_{\sigma^2} \quad / \quad \mathcal{P}rob\{\widehat{\sigma^2}_{\text{noise only}}[n,k] > t_{\sigma^2}\} = p_{fa}, \tag{9}$$

where p_{fa} is a given false alarm probability.

Candidates to the segmentation are the time-frequency locations whose local variance is higher than this threshold.

Pattern creation

In order to segment these candidates in spectral patterns, we enter in the propagation sub-loop of the algorithm. At each sub-iteration, the candidate with the highest variance is chosen as a "seed", associated with a label l. The candidates which are close to the "seed" in the time-frequency representation are contaminated with this label, and become new seeds, contaminating a new spectral pattern iteratively. Nevertheless, if patterns were already created at the previous iteration, these patterns are successively taken as seeds, before considering new seeds. Each propagation is validated with a test based on the Kolmogorov distance.

The sub-loop is stopped when a given proportion p_{cand} of the candidates is segmented.

End of the algorithm

Time-frequency coefficients are known to have a Gaussian distribution. In order to define a segmentation stop, a kurtosis criterion is defined as [6]

$$K = \frac{\mu_4}{(\sigma^2)^2} - 3, \tag{10}$$

where μ_4 is the fourth centered moment. The kurtosis of a Gaussian distribution is known to be null. When all (n,k) sites having a deterministic part are segmented, the remaining sites have a zero mean Gaussian distribution, thus the kurtosis on these points is zero.

When the absolute value of the kurtosis estimated on the non-segmented points is lower than a given threshold t_k, we assume that we have segmented all points having a deterministic part. There is thus no need to continue and the algorithm is stopped.

PARAMETER INFLUENCE

In order to study the different influences of the three parameters p_{fa}, p_{cand} and t_k, we use a synthetic signal made up of three sinusoids of amplitude 1, over three distinct time supports and frequencies, and embedded in a white Gaussian noise of known variance. Figure 1 shows the spectrogram of the signal without noise. The optimal segmentation of this synthetic signal results in three distinct patterns, one for each sinusoid.

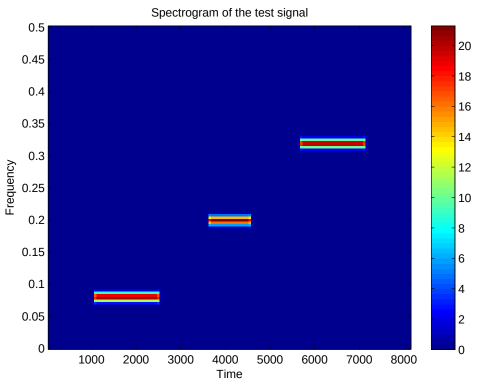

Figure 1: Spectrogram of the test signal, without additive noise.

Applying the time-frequency segmentation proposed in this paper, the influence of each parameter is studied for different local Signal-to-Noise Ratios (SNR) defined as

$$\mathrm{SNR} = 10\log_{10}\left(\frac{0.5}{\sigma^2}\right), \qquad (11)$$

the amplitude of the sinusoids being 1.

For each SNR, segmentations are performed with 100 realisations of noise. We focus on the number of *oversegmentations* and *undersegmentations*, *i.e.* the number of segmentation results which displays more and less than thee patterns respectively.

Influence of the probability of false alarm

As seen in equation (9), this probability of false alarm p_{fa} determines a threshold which set the time-frequency candidates to the segmentation. At a given iteration, there are less candidates as p_{fa} decreases. An oversegmentation may appear if there is not enough candidates: a single pattern will be split in several groups of candidates, each one being able to create an isolated pattern. On the contrary, too many candidates can merge different patterns into a single one, and lead to a undersegmentation.

Figures 2 and 3 show the percentage of oversegmentations and undersegmentations respectively, for different SNR and different values of p_{fa}, with $p_{cand} = 0.5$, $t_k = 0.1$.

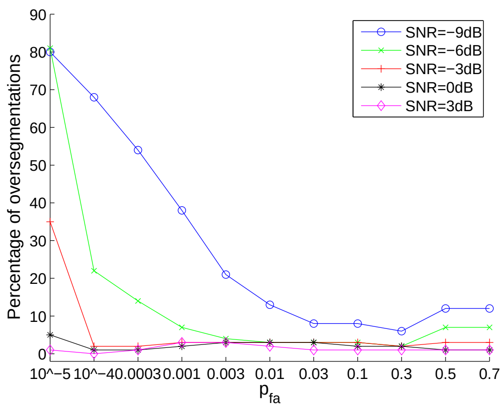

Figure 2: Percentage of oversegmentations for different SNR and p_{fa}

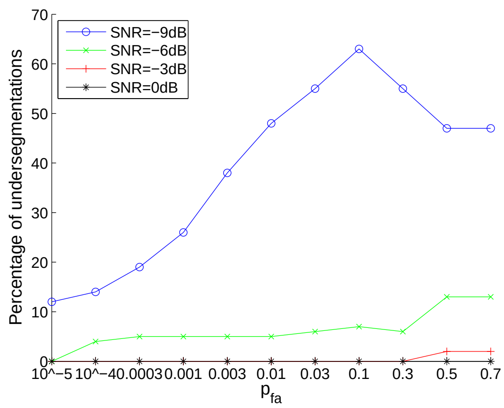

Figure 3: Percentage of undersegmentations for different SNR and p_{fa}

Influence of the proportion of candidates to be segmented

The parameter p_{cand} is the proportion of candidates to segment at each iteration. p_{cand} enhances the effect of p_{fa}: on the first hand, a value of 1 will create all the possible patterns among the candidates, and leads to an oversegmentation. On the other hand, a very low value segments only a few points at each iteration: the growth of existing patterns is privileged over new pattern creation. The choice of p_{cand} is thus a compromise between these two extreme cases.

Figures 4 and 5 show the percentage of oversegmentations and undersegmentations respectively, for different Signal to Noise Ratios (SNR) and different values of p_{cand}, with $p_{fa} = 0.01$ and $t_k = 0.1$. For this test signal, which has homogeneous patterns well separated in the TFR, high values of p_{cand} have no effect.

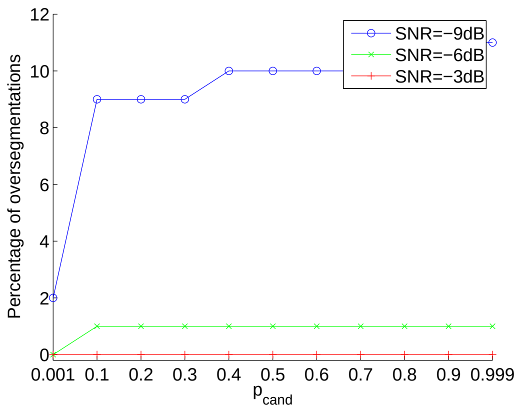

Figure 4: Percentage of oversegmentations for different SNR and p_{cand}

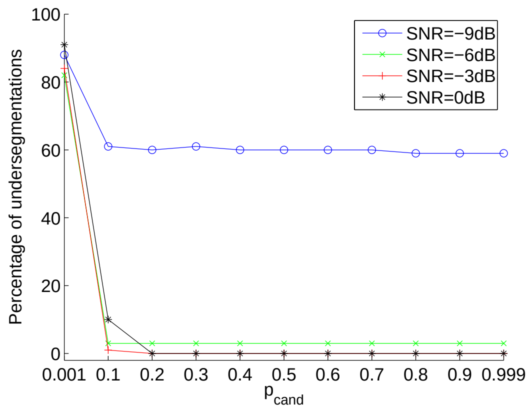

Figure 5: Percentage of undersegmentations for different SNR and p_{cand}

Influence of the threshold on the kurtosis

During the segmentation process, the distribution of the non-segmented points converges at a Gaussian distribution, and its kurtosis converges at zero. When all signal patterns are segmented, if the kurtosis on the non-segmented points is not null, the algorithm creates new patterns containing noise. We thus need to set a threshold t_k on the kurtosis in order to avoid oversegmentation. However, a too high t_k will stop the algorithm before all signal patterns are segmented.

Figures 6 and 7 show the percentage of oversegmentations and undersegmentations respectively, for different Signal to Noise Ratios (SNR) and different values of t_k, with $p_{fa} = 0.01$ and $p_{cand} = 0.5$.

Choice of a set of parameters

As seen on figure 2 to 7, a generic choice of $p_{fa} = 0.01$, $p_{cand} = 0.5$ and $t_k = 0.1$ results in a good compromise between few oversegmentations and undersegmentations. Moreover, we

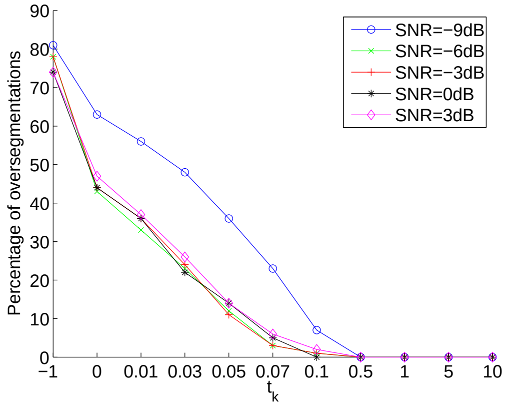

Figure 6: Percentage of oversegmentations for different SNR and t_k

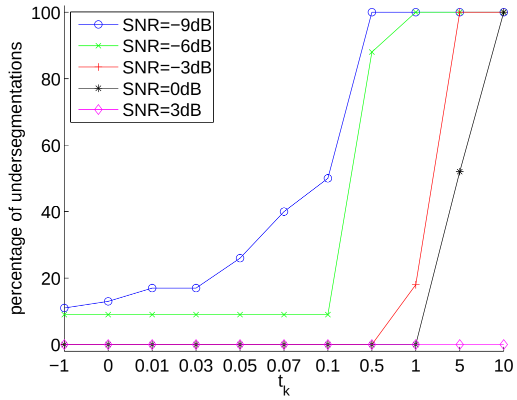

Figure 7: Percentage of undersegmentations for different SNR and t_k

notice that higher SNR causes both less oversegmentations and undersegmentations. The next section shows that another set of parameters allows the algorithm to fit to a wished goal.

APPLICATION

In order to monitor the three-phase AC induction motor of the test bench GOTIX of the laboratory, the segmentation is applied to a signal provided by an accelerometer located on the engine shaft, which thus measures its vibrations. Figure 8 and 9 show the spectrogram and a time-frequency zoom of the measure. The black rectangle in figure 8 defines a time-frequency region where the segmentation algorithm is applied. Indeed, whether the algorithm takes all the time-frequency coefficients or only a sample, the coefficients keep the same probability distributions.

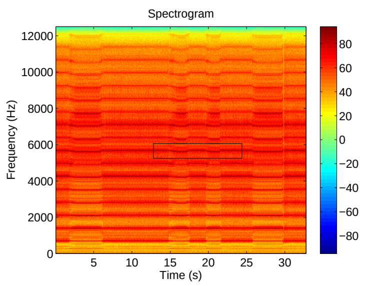

Figure 8: Spectrogram of the engine shaft's vibrations.

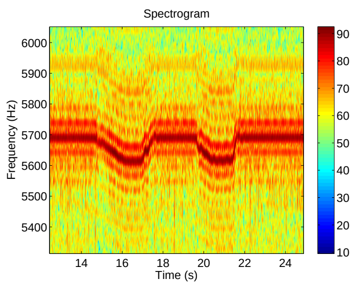

Figure 9: Zoom on the black rectangle of figure 8.

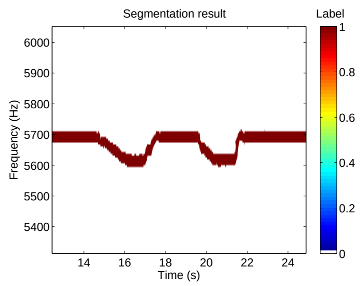

Figure 10: Segmentations results. p_{fa}=10^{-4}, p_{cand}=0.1, t_k=15

In this application, we are interested to segment only a single harmonic, in order to recover the engine speed. We thus choose a low p_{cand} to avoid oversegmentation; a high t_k to stop the algorithm as soon as the main pattern is segmented; and a low p_{fa} to avoid segmenting the nearest harmonic in the main pattern. Figure 10 shows the segmentation results for $p_{fa} = 10^{-4}$, $p_{cand} = 0.1$, $t_k = 15$.

The pattern obtained has the same wave shape than the command signal which was manually applied on the torque during the experiments.

CONCLUSIONS

Taking as a model of signal a deterministic part embedded in a white Gaussian noise, this paper describes the distributions of the STFT's real and imaginary parts. The distributions are then used in a segmentation algorithm, controlled by three parameters. The algorithm is applied on a synthetic signal with different realisations of noise, and informs about the influence of these parameters on the segmentation results. It helps to choose a good set of parameters to segment a vibration signal of an engine shaft, in order to get a simple pattern describing the frequency evolution.

A more complete characterisation of the algorithm will consider validations of patterns. The validation is currently based on a Kolmogorov distance. Works are in progress in this direction.

References

[1] C. Hory, N. Martin and A. Chehikian, *Spectrogram Segmentation by Means of Statistical Features of Non-stationnary Signal Interpretation*, IEEE Transactions on Signal Processing, 50(12), (2002), 2915–2925.

[2] N. Martin, C. Hory and M. Huchard, *Chi-square Law Mixture for an Unsupervised Time-Frequency Pattern Extraction*, in *ICSV10*, Stockholm, Sweden (2003).

[3] F. Romero Rodriguez, W.M. Liu, N.W.D. Evans and J.S.D. Mason, *Morphological Filtering of Speech Spectrograms in the Context of Additive Noise*, in *Procedings of Eurospeech*, Geneva, Switzerland (2003).

[4] W.M. Liu, V.J. Rivas Bastante, F. Romero Rodriguez, N.W.D. Evans and J.S.D. Mason, *Morphological Filtering of Spectrograms for Automatic Speech Recognition*, in *Proceedings of Visualization, Imaging, and Image Processing*, Marbella, Spain (2004).

[5] F. Millioz, J. Huillery and N. Martin, *Short Time Fourier Transform Probability Distribution for Time-Frequency Segmentation*, in *Proceedings of ICASSP*, Toulouse, France (2006).

[6] J. F. Kenney, E. S. Keeping, *Mathematics of Statistics*, Van Nostrand (1951).