



HAL
open science

Enhanced visual analysis of aircraft engines based on spectrograms

Julien Griffaton, José Picheral, Arthur Tenenhaus

► **To cite this version:**

Julien Griffaton, José Picheral, Arthur Tenenhaus. Enhanced visual analysis of aircraft engines based on spectrograms. ISMA2014, Sep 2014, Leuven, Belgium. pp.2809-2822. hal-01103775

HAL Id: hal-01103775

<https://centralesupelec.hal.science/hal-01103775v1>

Submitted on 15 Jan 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Enhanced visual analysis of aircraft engines based on spectrograms

J.Griffaton¹, J. Picheral², A. Tenenhaus²

¹SAFRAN, Snecma, Whole Engine Dynamics
77550 Moissy-Cramayel, France

²Supélec, Department Signals and Electronic Systems
Gif-sur-Yvette, 91192, France

Abstract

Detection of damaged bearings in aircraft engines by visual analysis of vibration spectrograms has proven its efficiency. Enhancement of this visual analysis can be achieved by data processing of the spectrograms. The data processing flow presented in this paper is composed of three successive steps: i) individual statistical normalization of the wideband noise (background noise of the spectrogram), ii) elimination of ‘sane’ narrowband vibration using a determinist parametrical model, and iii) elimination of residual ‘sane’ vibration using a statistical model whose parameters have been estimated from a database of ‘sane’ engine. This 3-step procedure has been tested on a set of vibration data from aircraft engines and the results are discussed.

1 Introduction

In the scope of aircraft engine health monitoring, analysis of vibrations can be very powerful despite the noisy environment. Among the many ways of identifying damages leading to impending failures of bearings, visual analysis of the vibration spectrograms has proven its efficiency. Therefore, there is a crucial need of automatic procedure able to help the diagnosis of failure from spectrograms. This paper presents new automatic strategy for improving the visual analysis of spectrograms and highlighting the main causes of failures.

Anomaly detection using vibration dataset has been investigated in [1], the key idea was to model the wideband noise by a Gamma distribution and to classify the pixel by thresholding their energy. This approach suffers of some drawbacks; mainly because it does not take into account the narrow band characteristics of the vibration, since the criterion is mainly based on the pixel energy. Another drawback is that the method has been designed to monitor a unique engine.

Another approach proposed in [2] is based on estimating distance between spectrogram of the inspected recording and the baseline; distances are then compared to the characteristic patterns of the bearing failure modes. Main drawback of this approach is that it assumes that the patterns of bearing failure are known.

The paper starts with a brief description of the data set and the features in the spectrograms of engines, with and without damaged bearings. The methodology for damage detection is explained step by step. Finally, the results of tests on the database of aircraft engines are discussed.

2 Data

2.1 Data set description

The data set is composed of 50 records of sane engines and 2 records of engines with a damaged bearing.

Those data are collected during an on-ground vibration survey: the engine is slowly accelerated from idle to max power in two minutes. Those recorded data are the rotating speeds N1 and N2 of the two main shafts of the engine shafts and the acceleration measured on the engine carcass.

The processed data used in this paper is a Short Time Fourier Transform (STFT) of the measured acceleration, where time samples have been resampled in order to assure a uniform angular sampling (of the axis where the N2 speed is measured). The spectrogram which results of the STFT is thus defined in the [frequency, N2] plane.

2.2 Features of a sane engine

An example of spectrogram of a sane engine is shown in Figure 1. The signals which compose this data can be roughly classified in two types. i) Narrow bands signal which is generally associated with natural activities from the rotating parts, for instance shafts unbalance, blades passing frequencies or gear mesh frequencies and their harmonics and modulations. ii) Wideband noise (or back ground noise) due to the natural vibration activity in the engine, for instance structural modes response or combustion noise.




Figure 1: Example of a spectrogram of a sane engine

2.3 Features of a damaged engine

An example of spectrogram of an engine with a damaged bearing is shown in

Figure 2. The signature of a damaged bearing in the spectrogram of an aircraft engine can be described a honeycomb of narrow bands with the following characteristics:

- The bands may be present in one or few ranges of rotating speeds. This is an effect of the load on the damaged that may be different for different rotating speeds
- The bands are approximately continuous when rotating speeds are varying
- The bands are regularly spaced. This can be explained by modulations around a main frequency with the rotating speeds of the bearing components (the inner race, the outer race and the cage).
- The frequency locations of the bands approximately follow orders of rotating speeds of the shaft(s)

There is therefore a great variety of signatures of damaged bearings, and each signature is relatively complex. The idea of looking for atypical narrow bands in the spectrograms is therefore a good way to catch any signature of damaged bearings.




Figure 2: an example of a whole spectrogram of an engine with a damaged bearing

3 Data processing steps

The study is performed in two phases. Firstly, the training phase aims to identify characteristics of the narrow band vibrations of sane engine and is partially supervised in order to be able to cancel these ‘normal’ narrowband vibrations. To this purpose a training dataset composed of 50 records from 50 sane engines is used. Characterization of the sane vibrations is provided by two sets: one of determinist parameter for N1 and N2 modes and another set of statistical parameters to characterize the remaining sane vibration signal. The second phase is the processing of new data in order to enhance the visualization. Main idea is to remove the sane vibration in order to exhibit potential vibration due to a damaged bearing. In order to cancel the sane vibration, the models obtained in the training phase are used.

3.1 Training phase

A general overview of training phase is presented in Figure 3 and can be summarized into four main steps:

- *Individual background normalization.* Each spectrogram is normalized by local Gamma laws. This step reduces the spectrogram background noise. The spectrogram background frequency content is not characteristic of a damaged bearing. The output of this step is a normalized spectrogram.
- *N1 and N2 orders identification.* Most of narrowband signal are modal vibrations ordered by N1 and N2. The aim of this step is to identify the order number of typical narrow band vibration. This step produces one order model for each spectrogram.
- *Statistical classification of narrow band signals.* Since all the narrow band signal cannot be modeled as order of N and N2, the purpose of this step is to use a supervised classification in order to identify remaining sane vibration. This step produces one statistical template for each spectrogram.
- *Model summary.* All individual models established for each spectrogram of the dataset are aggregated in a single order model and one single statistical template.




Figure 3: training phase, general overview of the data flow and data processing

3.1.1 Pre-processing – individual background normalization

This first step of data processing has been inspired by [1] that deals with the identification of local Gamma laws in the background of jet engines vibration spectrograms. The probability density function of a Gamma law is defined by two parameters, a shape parameter a and a scale parameter b .

$$f(x) = \frac{x^{a-1}}{\Gamma(a)b^a} \exp\left(-\frac{x}{b}\right) \quad (1)$$

The mean and the standard deviations for the Gamma law can be expressed with those two parameters.

$$\mu = ab \quad (2)$$

$$\sigma = \sqrt{ab} \quad (3)$$

The spectrogram is divided in a set of neighboring sub-matrices of 74x33 pixels, length 74 in N2 and length 33 in frequencies. The choice of a larger length in N2 is driven by the observation that the background noise remains approximately constant for different N2.




Figure 4: histogram of the vibration amplitude in a sub-matrix of a spectrogram

For each sub-matrix, a histogram of vibration magnitudes can be fitted by a first approximate local Gamma law using maximum likelihood estimator. A threshold is defined to segregate magnitudes that are in the tail distribution. Threshold is set to $3\hat{\sigma}$, where $\hat{\sigma}$ is the estimated standard deviation of the Gamma law. Once the tail distribution magnitudes have been removed, a second approximate Gamma law can be assessed with a better accuracy for the description of the background vibration. An example of the results is illustrated in figure 4.




Figure 5: relationship between standard deviation and mean of the Gamma law of all the sub-matrices of a spectrogram, red line is the linear approximation

The parameters distribution of the local Gama laws of a single spectrogram has been analyzed. The means and standard deviations of the local Gamma laws plotted in Figure 5 show that there is a strong linear relationship between the two parameters. Exploiting (2) and (3), it is straightforward that the linear relationship between the mean and the standard deviations implies that the shape parameter a is a constant in a spectrogram.




Figure 6: background normalization of a spectrogram

Moreover, the Gamma law has the following scaling property: if X is Gamma-distributed with shape a and scale b , then $\tilde{X} = cX$ is Gamma-distributed with shape $\tilde{a} = a$ and scale $\tilde{b} = cb$. Therefore, the normalized spectrogram is obtained by dividing each sub-matrix by its mean value $\mu = ab$. In consequence, each normalized sub-matrix follows a new Gamma law with a shape parameter $\tilde{a} = a$ and a scale parameter $\tilde{b} = \frac{b}{\mu} = 1/a$. Since the parameter a is roughly constant over the whole spectrogram, each normalized sub-matrix follows the same Gamma law. The result of such transformation is presented in Figure 6. The effect of background normalization can be easily seen: the smoothed patterns in the background are removed while narrow bands remain.

3.1.2 N1 and N2 orders identification

Spectrograms exhibit many narrow bands that are orders of N1 or N2. An effective way to identify those orders is the application of a Hough transform.

Let $S(f, N2)$ the spectrogram, $T_1(\alpha)$ and $T_2(\alpha)$ the Hough transforms of $S(f, N2)$ that have been applied separately to identify N1 and N2 orders. $T_1(\alpha)$ and $T_2(\alpha)$ are defined by :

$$T_1(\alpha) = \int_{N2_{\min}}^{N2_{\max}} S(N2, \alpha N1(N2)) dN2 \quad (4)$$

$$T_2(\alpha) = \int_{N2_{\min}}^{N2_{\max}} S(N2, \alpha N2) dN2 \quad (5)$$

The results of those transforms, for a single spectrogram, are plotted in Figure 7. As expected, a bunch of local maxima are exhibited both in N1 and in N2 orders.




Figure 7: examples of Hough transform

3.1.3 Order model summary

All individual Hough transforms results have been compiled in order to edit a z-score. The order model summary allows creating a template of the pixel location for both in N1 and N2 orders, of the maxima of the z-score. An illustration of the application of this order model on a spectrogram is shown in Figure 8.

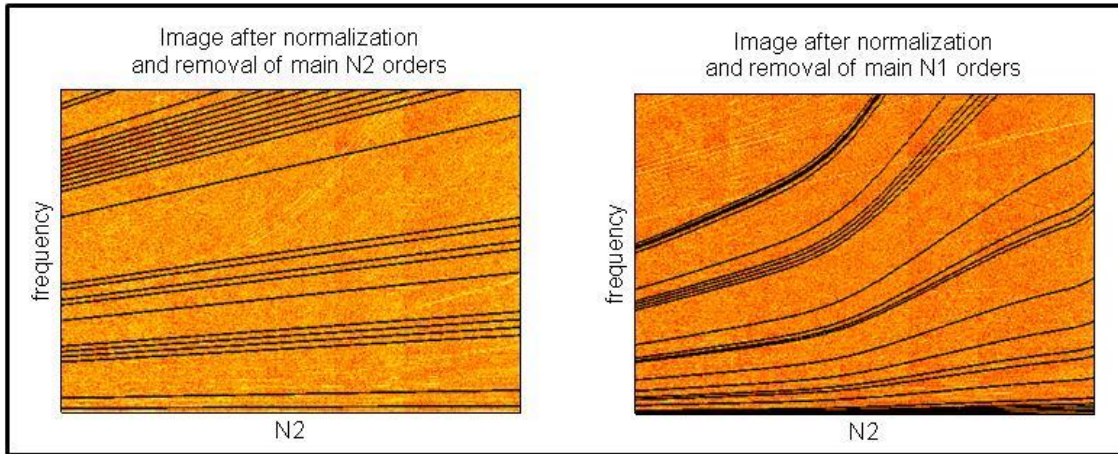


Figure 8: examples of parametric template applied on a spectrogram

3.1.4 Statistical classification of sub-matrices

This step consists in identifying abnormal vibration activity looking like local narrow bands in a spectrogram.

For this purpose each spectrogram is divided in small sub-matrices. Each sub-matrix will be labeled among 2 categories (binary classification): sub-matrices with a narrow band are labeled 1 and sub-matrices without a narrow band are labeled 0. Two sizes of sub-matrix ($N2 \times \text{frequency}$) length have been considered: 3×5 and 5×11 pixels.

A database of sub-matrices belonging either to categories 0 or 1 has been manually generated. Few examples of the result of the manual classification are shown in Figure 9.

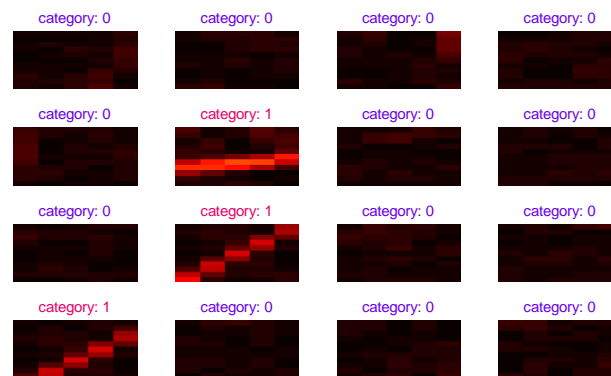


Figure 9: examples of manual classification of sub-matrices of size 5×11

Two state-of-the-art methods of classification have been tested on the database, namely Quadratic discriminant analysis (QDA) and Support Vector Machines (SVM) combined with either linear, quadratic or Gaussian kernel. We refer the interested reader to [3] for details.

For SVM, the regularization parameter C and the kernel parameter σ have been optimized by k -fold cross validation. More precisely:

- The learning dataset comprises $2/3$ of the sub-matrices, and the test set comprises the remaining part of the sub-matrices (3-fold cross-validation).
- Parameters C and σ respectively in the ranges of $[0.01 \ 3]$ and $[0.1 \ 30]$
- For each couple of parameters (C, σ), the error rate was calculated with 20 partitions of the dataset

An example of results of such cross-validation optimization is illustrated in Figure 10. Performances of the methods in term of sensitivity and specificity are reported in table 1.




Figure 10: result of an optimization of parameters C and σ for sub-matrices of size 3×5 ; the white dot is the optimized solution $C_{\text{optim}}=0.33$ and $\sigma_{\text{optim}}=1.33$, with an error rate of 9.1%

| Method <i>sub-matrices size</i> | Sensitivity (%) | | Specificity (%) | |
|------------------------------------|-----------------|-------------|-----------------|-------------|
| | <i>3x5</i> | <i>5x11</i> | <i>3x5</i> | <i>5x11</i> |
| QDA | 79 | 88 | 99 | 97 |
| SVM, linear kernel | 78 | 91 | 97 | 95 |
| SVM, quadratic kernel | 77 | 90 | 98 | 99 |
| SVM, Gaussian kernel | 79 | 94 | 98 | 97 |

Table 1: comparison of the classifiers

SVM with Gaussian kernel produces the classifier giving the best results as regards sensitivity and specificity. QDA and Gaussian-SVM give relatively similar results. All methods have a good specificity, which can be interpreted as a low rate of false alarm. Whatever the method, sensitivity is always higher with the sub-matrices of size 5×11 . From those conclusions, and considering that QDA has no parameter to optimize, QDA combined with sub-matrices of size 5×11 have been retained for the generation of a statistical template.

During the benchmark of the different classification methods, the threshold for the probability of a sub-matrix to belong either to category 0 or 1 (decision boundary) was set to 0.5. This threshold has been optimized in order to get a better compromise between sensitivity and specificity.




Figure 11: ROC curve with QDA and 5×11 , varying parameter is boundary decision threshold, the dot is the optimized solution with threshold= $1e-5$.

The Receiver Operating Characteristic (ROC) curve is shown in Figure 11. The area above the curve is 0.96. A good compromise between sensitivity and specificity can be achieved by finding the point of the curve that is the closest to (0, 1), which means finding the threshold that minimizes:

$$\sqrt{(1 - \text{sensitivity})^2 + (1 - \text{specificity})^2} \quad (6)$$

The associated threshold is 10^{-5} .

3.1.5 Statistical template summary

Finally, the statistical template is obtained using the following procedure: for each spectrogram, QDA is applied to each sub-matrix of size 5×11 . The resulting decision function which is expressed in terms of the probability of a sub-matrix to belong the first class is thresholded using a threshold of 10^{-5} . This classifier is then applied to the spectrogram of each sane engine of the dataset.

The resulting binary matrices are then combined to obtain the statistical template summary using the following rule: a pixel is affected to the first class in the statistical template whenever at least 2 spectrograms have affected the first class at the pixel location.

The statistical template is reported in Figure 12.




Figure 12: statistical template, black pixels are 0, white pixels are 1

3.2 Diagnostic phase

A general overview of diagnostic phase is presented in Figure 13. The data processing follows 4 main steps:

- *Background normalization.* This is the same step that was used in the training phase.
- *N1 and N2 orders removal.* The order model estimated during the training phase is used to force to zero the pixel where sane narrow band vibration is expected.
- *Statistical narrow band vibration removal.* As for the previous steps, pixels are forced to zero using the statistical template.
- *Detection of anomalies.* After the removal of the sane vibrations, the spectrogram is analyzed by the statistical classifier in order to enhance potential vibrations due to anomalies. It provides a score of anomaly and a picture of anomaly locations.




Figure 13: diagnostic phase, general overview of the data flow and data processing

Using the output of the classifier, an anomaly rate can be evaluated as the ratio between the number of abnormal pixels to the total number of pixels. The plot of classified pixel provides a clarified spectrogram where only vibrations associated to anomaly are represented. An example of the result of the diagnostic phase is shown in Figure 14 for a sane engine.




Figure 14: results for a sane engine, anomaly rate is 0.6%

4 Results

4.1 Cross validation

The database of the sane engines has been tested by the 'leave one-out' method. For each engine x:

- The training phase is performed with the database of the sane engines at the exception of engine x
- The diagnostic phase is then applied on engine x

The expected results are very low anomaly rate since the engines that are tested were tagged sane engines. The results of this cross validation are shown in table 2.




Table 2: results of the cross validation test on sane engine

Most anomaly rates are below 1%, at the exception of few engines.

Engine 8D_100026_Campbell_Ch3_Up has the highest rate. Its clarified spectrogram is shown in Figure 15. The record of this engine is atypical: it shows a honeycomb of wavy curves that can be seen in a small proportion of engines that have been confirmed to be sane. The score obtained with this engine is therefore not a surprise. A way to reduce the false alarm rate due to this kind of feature would require to feed the database of sane engines with those atypical engines.

8D 100026 Campbell Ch3 Up




Figure 15: results for an atypical sane engine, anomaly rate is 3.0%

4.2 Test with damaged engines

The results for the two engines that have a damaged bearing are shown in Figure 16 and Figure 17. The anomaly rates for those two engines are above the rates of almost all sane engines. Features that remain in the clarified spectrograms can be easily tagged and identified by an analyst.

8D 200001 Campbell Ch3 Up




Figure 16: results for an engine with a damaged bearing, anomaly rate is 2.4%

8D 200002 Campbell Ch3 Up




Figure 17: results for an engine with another damaged bearing, anomaly rate is 1.7%

5 Conclusion

Different methods for narrow bands recognition have been successfully tested and have proven their efficiency. Quadratic discriminant analysis has been elected because of the easiness of its parameter optimization.

Tests performed on real aircraft engines with and without damaged bearings succeeded in damage detection.

The initial intent was the detection of engines with damaged bearings, but a side effect is the detection of any engines that exhibit atypical narrow bands in its spectrogram. In such case, when a new record shows an atypical behavior and if the engine is considered as a sane engine by an analyst, this record shall be added in the database and the statistical template shall be up-dated.

The outputs of the data processing can efficiently help the work of an analyst:

- The engines with the lowest anomaly scores can be considered as sane engine and they can be filtered out
- The clarified spectrograms of the remaining engines make it easy to focus on the zones of interest in the original spectrogram

On the other hand, this algorithm on anomaly detection is limited to the detection of damages that generate narrow bands in a spectrogram.

Future improvements of the current algorithm are under study, mainly in the pre-processing phase, for continuity recognition, reconstruction and amplification.

Acknowledgements

The authors wish to thank the students of Supélec, and more precisely Baptiste MAUREL and Tangfei WANG, for their works in the analysis of spectrogram and the study of the flow of algorithms that are presented in this paper.

References

- [1] David A. Clifton and Lionel Tarassenko, *Novelty Detection in Jet Engine Vibration Spectra*, Institute of Biomedical Engineering, Department of Engineering Science, University of Oxford.
- [2] Klein, R., Masad, E., Rudyk, E., & Winkler, I. (2014). Bearing diagnostics using image processing methods. *Mechanical Systems and Signal Processing*, 45(1), 105-113.
- [3] Hastie, T., Tibshirani, R., Friedman, J., Hastie, T., Friedman, J., & Tibshirani, R. (2009). *The elements of statistical learning*. New York: Springer.