Novelty Detection in Jet Engine Vibration Spectra

David A. Clifton and Lionel Tarassenko

Institute of Biomedical Engineering
Department of Engineering Science, University of Oxford
Old Road Campus, Roosevelt Drive, Oxford, OX3 7DQ, UK
davidc@robots.ox.ac.uk

Abstract

This paper presents a principled method for detecting "abnormal" content in vibration spectra obtained from rotating machinery. We illustrate the use of the method in detecting abnormality in jet engine vibration spectra corresponding to unforeseen engine events. We take a novelty detection approach, in which a model of normality is constructed from the typically large numbers of examples of "normal" behaviour that exist when monitoring jet engines. "Abnormal" spectral content is then detected by comparing new vibration spectra to the model of normality. The use of novelty detection allows us to take an engine-specific approach to modelling, in which the engine-under-test becomes its own model, rather than relying on a model that is generic to a large population of engines. A probabilistic approach is taken that employs Extreme Value Theory to determine the boundaries of "normal" behaviour in a principled manner. We also describe a novel visualisation technique that highlights significant spectral content that would otherwise be too low in magnitude to see in a standard plot of spectral power.

1 Introduction

Vibration spectra obtained from rotating systems (such as gas-turbine engines, combustion engines, or machining tools), are characterised by peaks in spectral power at the fundamental frequency of rotation, and smaller peaks at harmonics of that fundamental frequency. In jet engine terminology, these peaks are conventionally called *tracked orders*. Methods exist [1, 2, 3] for the principled analysis of information pertaining to these tracked orders, such that precursors of system failure can be identified, and preventative maintenance action taken. However, many modes of failure manifest themselves as changes in vibration spectra that are not related to energy of the tracked orders.

An example of this is the failure of engine bearings, which are small ball-bearings enclosed within fixed cages such that they may rotate freely. These are used to form load-bearing contacts between the various rotating engine shafts, and they maintain the position of the shafts relative to one another. Damage to the surfaces of these bearings may result in previously-unobserved vibration energy at high frequencies, significantly removed from the narrow frequency bands of the tracked orders observed under "normal" conditions. Failure of the cages in which bearings are mounted can result in constant peaks in spectral energy at previously-unseen multiples of the fundamental tracked orders [4, 5]. The latter could be described as *novel*

tracked orders (NTOs) because they are peaks in vibration energy within narrow frequency bands, and are thus tracked orders, but occur at frequencies for which tracked orders are not observed under "normal" conditions.

This paper describes a method for identifying NTOs and other abnormal content in spectral data, allowing the identification of modes of failure that methods based on modelling of tracked orders cannot detect. Principled methods are used for modelling time-series of spectral data observed under "normal" conditions. The goal is to learn an engine-specific model of normality on-line, to provide sensitive novelty detection without the need for tuning heuristic parameters.

A model of normality is introduced in Section 2, and in Section 3, we discuss principled methods of identifying which components of a vibration spectrum are significant with respect to background noise. We use these models to transform the problem into probability space, in Section 4, describe how to perform novelty detection in Section 5, and present results from jet engine vibration data in Section 6.

Throughout this paper, absolute values of vibration amplitude and frequency are not given for purposes of commercial confidence, and units of measurement have been omitted from some figures.

2 Modelling Normality

Previous work [1, 2] has shown that the vibration amplitude of tracked orders can be characterised with respect to the operating point of the engine; i.e., the rotational speed of the engine shafts. We here extend this approach to examining whole vibration spectra with respect to engine speed, which we will use both for determining engine-specific noise-floor thresholds, and ultimately to construct a model of normality for use in novelty detection.

Figure 1 shows the mean vibration amplitude in each of the f_i, $i = 1 \ldots N_{\mathrm{FFT}}$ spectral bins computed using all data from an exemplar jet engine (here, $N_{\mathrm{FFT}} = 1024$), collected in 1% sub-ranges of shaft speed ω_{LP}, which is the rotational speed of the low-pressure (LP) shaft. Energy associated with the fundamental tracked orders appears in the lowest decile of the frequency range. Energy associated with the second harmonic tracked orders appears in the second decile of the frequency range (though at lower amplitudes than the fundamental tracked orders). The whole speed-frequency space shown in the figure has been partitioned into $N_\omega = 10$ bins on the speed axis, and $N_f = 10$ bins on the frequency axis, giving a 10×10 matrix representation of the speed-frequency space.

Characterising the spectral energy in each spectral bin across the full range of shaft speeds, as illustrated in Figure 1, provides a useful summary of engine vibration energy during the "training period" in which the data used to construct the summary, or model, are acquired. Thus, a representation of this form can be used for novelty detection, because we anticipate that "abnormal" vibration energy observed during testing will appear in some way abnormal with respect to this representation.

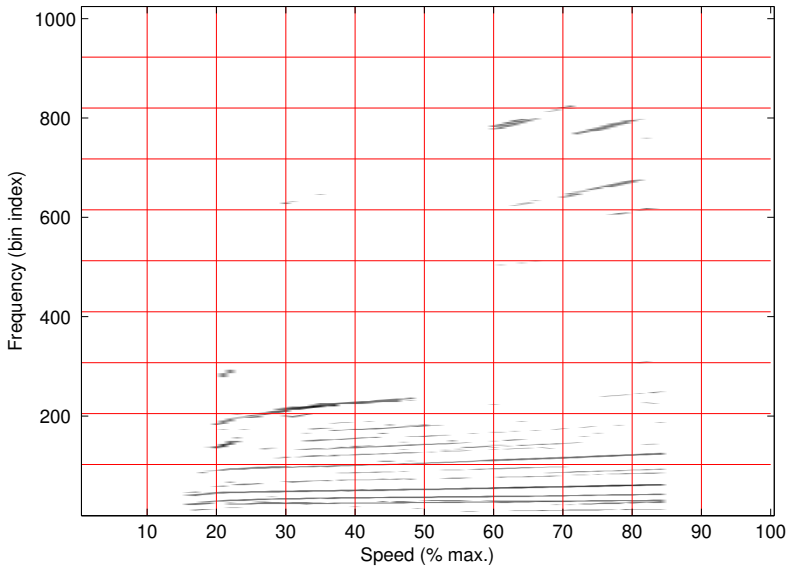

Figure 1: A speed-based representation of spectral data, showing the mean vibration amplitude from an exemplar jet engine collected in 1% sub-ranges of LP shaft speed. In this figure, the speed-frequency space is quantised into $N_\omega \times N_f = 10 \times 10$ bins, shown in red. The vertical axis shows spectral bin indices f_i ranging between $[0\ f_s/2]$.

3 Noise-Floor Estimation

In order to identify NTOs within time-series of spectra, we must first identify and disregard spectral components corresponding to background noise. In existing heuristic methods, this has been performed by defining a single "noise-floor" threshold on spectral energy, below which spectral components are deemed to be noise. A single value for this threshold has been applied to all gas-turbine engines of a similar class, using expert knowledge [6]. In order to avoid large numbers of false-positive classifications during novelty detection when using this single threshold with an entire class of engines, it must be set conservatively due to inter-engine variability within that class. However, this causes subsequent novelty detection to have low sensitivity, because the conservative threshold results in a high number of false-negative classifications.

Furthermore, as will be shown in this section, a constant noise-floor threshold applied to all frequencies of spectral data results in an inability to determine significant engine events that manifest as changes in spectral energy at frequencies much higher than the first few harmonics of fundamental tracked orders. In order to detect bearing failure and other events occurring at higher frequencies, a dynamic approach to noise-floor estimation is required.

3.1 Examining Noise Distributions

Experimental observation of jet engine vibration data shows that the vibration amplitude of noise generally decreases with increasing spectral frequency. In [7], the authors observed similar behaviour for background noise in acoustic signals, where

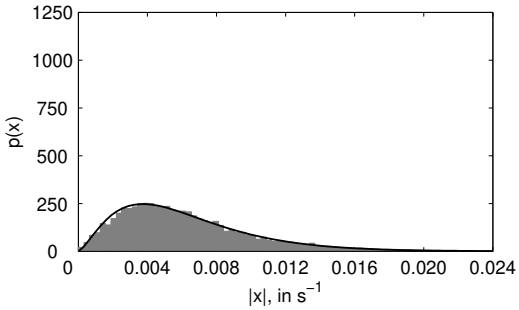

(a) Bin $b_\omega = 8, b_f = 1$

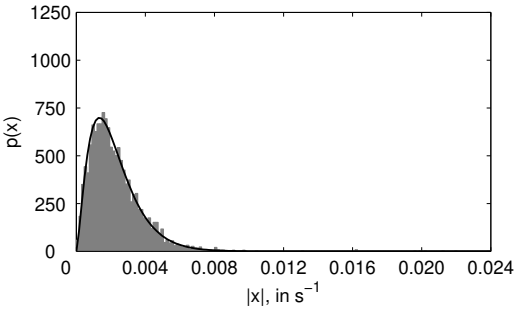

(b) Bin $b_\omega = 8, b_f = 5$

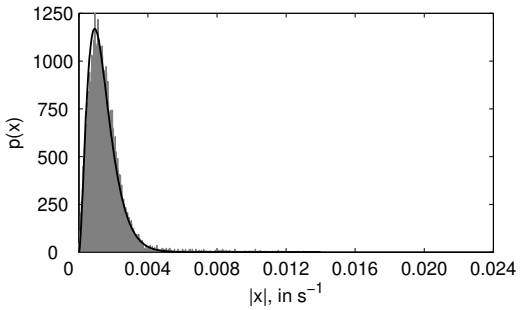

(c) Bin $b_\omega = 8, b_f = 10$

Figure 2: Histograms of vibration amplitudes falling within selected speed-frequency sub-ranges for vibration data observed from an exemplar jet engine. Histograms are normalised such that they integrate to unity and are approximately pdfs, $p(x)$. Maximum likelihood gamma distributions for each bin are shown in black.

narrow frequency bands used to transmit speech were seen to contain Gamma-distributed background noise.

Bin $(b_\omega = 8, b_f = 1)$, shown in Figure 2(a), covers the speed range $70\% \leq \omega_{\text{LP}} \leq 80\%$. This bin contains the highest vibration energy of all those shown. The figure shows that the distribution of vibration amplitudes within this bin is approximately Gamma, corresponding to background noise.

Bin $(b_\omega = 8, b_f = 5)$, shown in Figure 2(b), covers the same range of shaft speeds as the bin described above, but corresponds to a higher range of frequencies. The figure shows that the distribution of vibration amplitudes within this bin remains approximately Gamma, but is reduced in location and scale (where the *location* and *scale* of the Gamma distribution $\text{Gam}(x|a, b)$ are a and b, respectively) compared with that of the lower-frequency bin.

Bin $(b_\omega = 8, b_f = 10)$, shown in Figure 2(c), again covers the same range of shaft speeds, but corresponds to the highest frequencies in the FFT . The distribution of vibration amplitudes contained within this is approximately Gamma, but further reduced in location and scale compared with the lower-frequency bins.

3.2 Modelling Noise Distributions

In order to model the distribution of amplitudes $\mathbf{X}_{\omega,f} = \{x_1, \ldots, x_N\}$ within each of the $N_\omega \times N_f = 10 \times 10$ bins considered above, we find the Maximum Likelihood (ML)

Gamma distribution, $\mathrm{Gam}_{\mathrm{ML}}\left(x|\hat{a},\hat{b}\right)$ for which the log likelihood may be written

$$\ln \mathcal{L}(a,b) = Na \ln b - N \ln \Gamma(a) + (a-1)\sum_{i=1}^{N} \ln(x_i) - b\sum_{i=1}^{N} x_i \tag{1}$$

The maximum likelihood parameters $(\hat{a},\hat{b})$ are given by $\mathrm{argmax}_{a,b} \ln \mathcal{L}(a,b)$. Solutions for this maximisation do not exist in closed form, and we use the iterative methods described in [8].

3.3 Determining Noise-Floor Thresholds

Using the noise distributions, a threshold $x^{+}_{\omega,f}$ may now be defined for each of the 10×10 bins, below which vibration amplitudes will be deemed to be background noise. Thus, we wish to determine the upper limit of extrema generated from the Gamma distribution in each bin.

Extreme Value Theory (EVT) provides a means of estimating the maximum likely value to be generated from a given distribution F, which effectively models the tails of F [9, 10]. According to the Fisher-Tippett theorem [11] upon which EVT is based, the tails of the univariate Gamma distribution $F = \mathrm{Gam}(a,b)$ are modelled using the Gumbel distribution,

$$p_e(x|c,d) = \exp(-\exp(-y)) \tag{2}$$

where $y = (x-c)/d$ for Gumbel parameters (c,d), and where these latter may be estimated [12] from the Gamma parameters (a,b) and the number of data m drawn from the Gamma distribution,

$$c = \frac{1}{b} \tag{3}$$

$$d = \frac{1}{b}\Big(\ln m + (a-1)\ln\ln m - \ln\Gamma(a)\Big) \tag{4}$$

This Gumbel distribution describes where we expect the maximum of m values drawn from the Gamma distribution to lie. To estimate a noise-floor threshold, we wish to determine the most extreme value generated by noise under "normal" conditions, which is given by this Gumbel distribution. Given a spectrum of N_{FFT} components for each observation, each bin will contain N_{FFT}/N_f data. We thus define $m = \lceil N_{\mathrm{FFT}}/N_f \rceil$.

Let $P_e(x)$ be the cdf obtained by integrating the pdf $p_e(x)$ defined in equation (2). We can set a noise-floor threshold $x^{+}_{\omega,f}$ by setting this cdf equal to some suitable probability value; e.g., $P_e(x) = 0.99$.

Thus, for each of the $N_\omega \times N_f$ speed-frequency bins in our representation,

- we can find the ML Gamma distribution using the method described above;
- we find the Gumbel distribution defined in equation (2) that describes where we expect extrema generated from that distribution to lie; and thus,

- we can set a noise-floor threshold for this speed-frequency bin at some probability value, such as $P_e(x) = 0.99$.

Vibration data in each speed-frequency bin are compared to the corresponding noise-floor threshold determined using EVT. Data falling below the threshold are thus deemed to have been generated by the noise process, and discarded, while data exceeding the threshold are deemed to be significant spectral content which will be analysed further.

4 Analysing Spectral Data

This section considers how spectral data may be tested with respect to the model of normality.

Consider now a test sample which consists of $\mathbf{X} = \{x_1, \ldots, x_{N_{\mathrm{FFT}}}\}$ spectral components, observed at LP shaft speed ω_{LP}. Each spectral component x_i in $\mathbf{X}$ is evaluated with respect to the Gumbel distribution of the appropriate speed-frequency bin in the $N_\omega \times N_f$ quantisation. Figure 3(a) shows spectra of vibration amplitudes for part of a flight of the exemplar engine, in which the engine initially accelerates at a high rate from time index $t = 50$ to $t = 180$, and then more gradually until $t = 700$. Fundamental tracked orders may be seen in the lowest decile of the frequency range, varying in frequency proportional to engine speed. Between time indices $t = 100$ and $t = 150$, significant vibration energy may be seen to occur in the lower three deciles of the frequency range, associated with harmonic tracked-order vibration. Low-amplitude vibration associated with higher-order harmonic tracked orders may be seen throughout the example, appearing as a series of inclined lines extending across the frequency range. Horizontal bands of low-amplitude vibration noise may be seen extending across the frequency range throughout the example.

Figure 3(b) shows the result of determining cumulative Gumbel probabilities $P_e(x)$ for each spectral component with respect to the Gumbel distributions for each of the 10×10 bins considered previously. It may be seen that harmonic and fundamental tracked orders alike have $P_e(x) \approx 1$, as vibration amplitudes associated with them are considerably higher than the distributions of background noise used to construct the Gumbel distributions in the 10×10 bins. The horizontal bands of low-amplitude noise occurring across the frequency range generally have probabilities $P_e(x) \ll 1$.

For the purposes of visualisation, we define a novelty score

$$z(x) = -\log_{10}\{1 - P_e(x)\} \tag{5}$$

Figure 3(c) shows $z(x)$ for each spectral component in the example. The horizontal bands of low-amplitude vibration noise have been attenuated, while the fundamental and harmonic tracked orders appear clearly. Tracked orders can be seen throughout the frequency range as lines with positive gradient, increasing in frequency proportional to the rotational frequency of the engine, which is increasing during this interval. The 14HP harmonic tracked order is visible at the top of the frequency range, and is labelled in the figure.

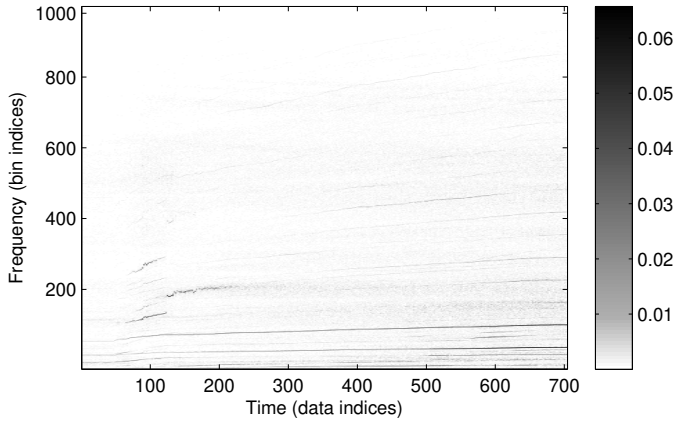

(a) Time-series of spectra

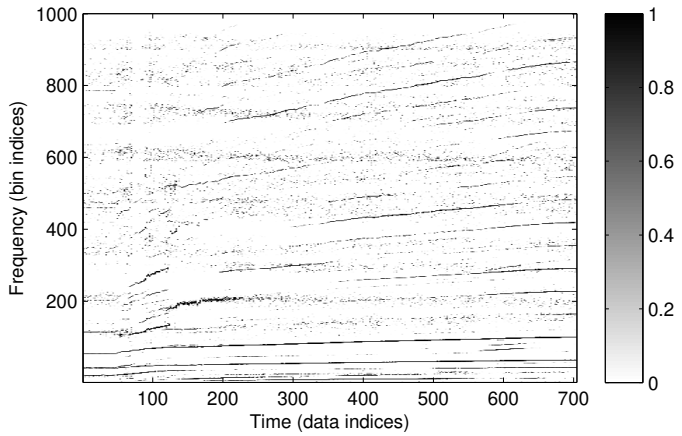

(b) Probabilities, $P_e(x)$

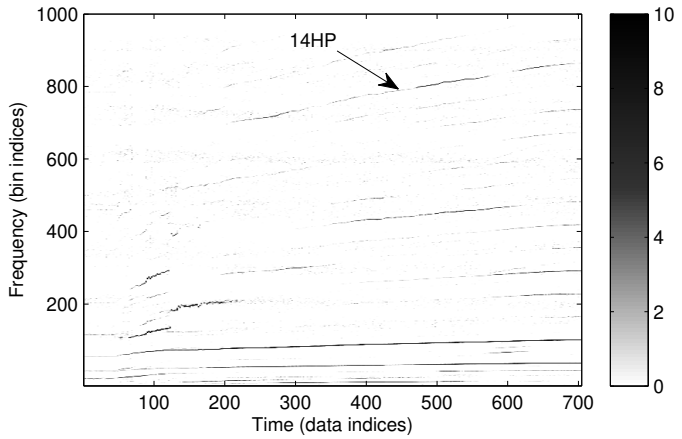

(c) Novelty scores, $z(x)$

Figure 3: Evaluating spectra from a test flight with respect to noise-floor estimates. (a) Vibration amplitudes for test flight of an exemplar engine, shown as a time-series of vibration spectra. (b) Probabilities $P_e(x)$ for each spectral component evaluated with respect to cumulative Gumbels $P_e(x)$ from noise distributions in 10×10 bins. (c) Novelty scores $z(x)$ as defined in (5) for each spectral component.

We note that energy associated with higher-order tracked orders is usually not visible in a standard plot of spectral power against speed-frequency, such as that in Figure 3(a), because of its low amplitude. The figure shows that our proposed probabilistic method can be used to highlight such spectral content because it is highly "novel" with respect to the background noise. Such spectral content, even though of very low absolute magnitude, thus takes high novelty scores $z(x)$, which causes it to be highlighted in the visualisation, as shown in Figure 3(c).

5 Identifying Novel Tracked Orders

Spectral components with vibration amplitudes above the probabilistic noise-floor thresholds $x^+_{\omega,f}$ have been identified, and it remains for them to be separated into those which correspond to spectral content previously observed within the training data, and those which are novel with respect to the training data.

In order to identify novel vibration energy in test spectra, we must record the location of vibration energy observed in the training data. Note that we here disregard all spectral components with vibration amplitudes deemed to be "background noise" (using the method described in Section 3), and hereafter only refer to those spectral components with vibration amplitudes above the appropriate noise-floor thresholds.

To record the location of vibration energy from the training data in speed-frequency space, we again partition that space into $N_\omega \times N_f = 10 \times 10$ bins. We define counting matrix $C_{i,j}$ such that element (i, j) is set to the number of FFT components in the training set that fall within the corresponding speed-frequency bin (i, j).

The C-matrix characterises the distribution of spectral energy throughout the speed-frequency space for the training data. We now define a quantity N_m such that any bin (i, j) with $C_{i,j} \geq N_m$ will be considered to contain "known tracked order" vibration energy. Any FFT component from test data (with vibration amplitude above the appropriate noise-floor threshold) falling into that bin (i, j) will be classified "normal"; i.e., it lies in a part of the speed-frequency space for which a sufficient number of FFT components above the noise-floor threshold were observed in the training data.

Conversely, if an FFT component from test data falls into a bin (i, j) for which $C_{i,j} < N_m$, it will be classified "abnormal" with respect to the training data; i.e., it lies in a part of the speed-frequency space where an insufficient number of FFT components above the noise-floor threshold were observed during training.

The selection of parameter N_m is performed using example datasets such that the rate of novelty detection is sufficiently high, while the false-alarm rate is acceptably low. The optimisation of N_m will vary depending on the nature of the system that is to be analysed, and on the quantisation $N_\omega \times N_f$. Experimental observation [13] has shown that $N_m = 3$ provides a suitable compromise between being able to detect example abnormal events in jet engine vibration data and in producing low numbers of false-positive novelty detections.

Table 1: Datasets used in the investigation described by this paper.

Dataset	Engine	Flights	Comments
$\mathcal{B}_1$	3-shaft class A, engine I	4	$\mathcal{B}_{1,3}, \mathcal{B}_{1,4}$ contain bearing failure
$\mathcal{B}_2$	3-shaft class B, engine I	4	$\mathcal{B}_{2,4}$ contains bearing-cage failure
$\mathcal{B}_3$	3-shaft class A, engine II	6	Normal engine data

6 Test Data

The investigation described by this paper uses three datasets of full spectral data, described in Table 1. The first three engine flights in datasets $\mathcal{B}_1$ and $\mathcal{B}_2$ were deemed to be "normal" by engine experts. Flights $\mathcal{B}_{1,3}$ and $\mathcal{B}_{1,4}$ (the final two flights in dataset $\mathcal{B}_1$) were deemed to contain evidence of a bearing cage event. Flight $\mathcal{B}_{2,4}$ (the fourth and final flight in dataset $\mathcal{B}_2$) was deemed to contain evidence of a bearing event. Dataset $\mathcal{B}_3$ contains data recorded from an engine with no abnormal events, which is used to investigate the performance of the method described by this chapter when presented with "normal" data.

For each dataset, an engine-specific model of normality was constructed using the techniques described in Sections 3, 4 and 5. Each model was trained using data from the first two flights of the engine. Data from subsequent flights in each dataset were used as test data, compared to the model of normality corresponding to that engine.

6.1 Test Set $\mathcal{B}_1$

Vibration spectra from part of flight $\mathcal{B}_{1,3}$ with a suspected bearing-cage event are shown in Figure 4(a). $P_e(x)$ for "abnormal" spectral data is shown in Figure 4(b). This shows the high-amplitude vibration energy at frequencies corresponding to the bearing-cage event, which is a series of novel tracked orders occurring between $t = 200$ and $t = 400$. This abnormal spectral energy is correctly identified using this speed-frequency quantisation and N_m value. Note that all fundamental and harmonic tracked orders corresponding to "normal" operation are not shown in Figure 4, which shows only "abnormal" spectral content.

6.2 Test Set $\mathcal{B}_2$

Vibration spectra from part of flight $\mathcal{B}_{2,4}$ are shown in Figure 5(a). $P_e(x)$ for "novel" spectral data is shown in Figure 5(b). This shows the novel tracked order corresponding to the bearing event at f_{780}, which is correctly identified using this speed-frequency resolution and N_m value.

6.3 Test Set $\mathcal{B}_3$

Table 2 shows the results of testing all six flights with respect to the noise-floor thresholds and C-matrix constructed using the first two runs for the "normal"

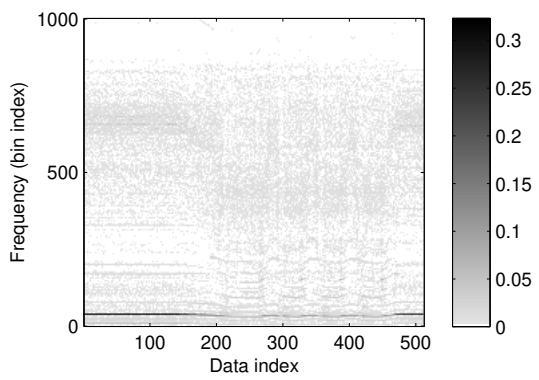

(a) Vibration spectra from flight $\mathcal{B}_{1,3}$

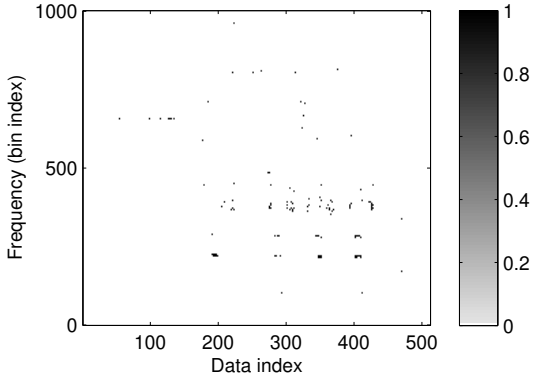

(b) $P_e(x)$ for flight $\mathcal{B}_{1,3}$

Figure 4: Vibration spectra for episodes from flights from dataset $\mathcal{B}_1$, with $P_e(x)$ for "abnormal" spectral data. (a) Vibration spectra from flight $\mathcal{B}_{1,3}$. (b) $P_e(x)$ for "abnormal" spectral data from flight $\mathcal{B}_{1,3}$

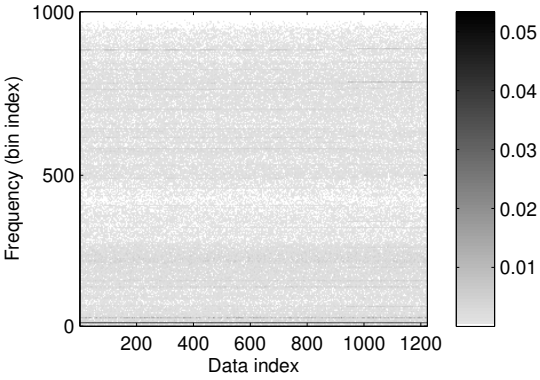

(a) Vibration spectra from flight $\mathcal{B}_{2,4}$

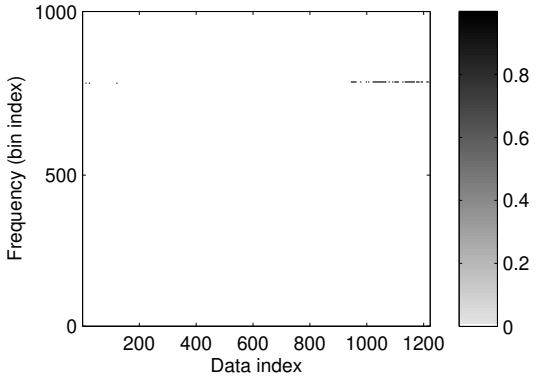

(b) $P_e(x)$ for flight $\mathcal{B}_{2,4}$

Figure 5: Vibration spectra for episodes from flight $\mathcal{B}_{2,4}$, with $P_e(x)$ for "novel" spectral data. (a) Vibration spectra from flight $\mathcal{B}_{2,4}$. (b) $P_e(x)$ for "novel" spectral data from flight $\mathcal{B}_{2,4}$.

dataset $\mathcal{B}_3$. Any "novel" classifications made are false positives. It can be seen from the table that the use of method results in a low number of false-positive novelty detections.

7 Conclusion

We have presented a method of constructing engine-specific models of normality that describe "normal" spectral content in vibration spectra. The method is based on a probabilistic methodology that first seeks to separate noise from significant spectral content, and then classifies any significant spectral content as either "normal" or "abnormal" with respect to the model. This method offers advantages over conventional noise-thresholding techniques, in that a separate noise threshold is identified for each quanta in the speed-frequency space. This causes low-amplitude spectral

Table 2: Proportion of spectral bins in each flight of dataset $\mathcal{B}_3$ that were classified "novel"

Flight	Proportion of bins classified "novel"
$\mathcal{B}_{3,1}$	0
$\mathcal{B}_{3,2}$	0
$\mathcal{B}_{3,3}$	0.10×10^{-4}
$\mathcal{B}_{3,4}$	0.21×10^{-4}
$\mathcal{B}_{3,5}$	0.11×10^{-4}
$\mathcal{B}_{3,6}$	0.05×10^{-4}

content to be retained, whereas conventional methods assume it to be noise, and thus reject it.

Results have been presented which show that the method identifies abnormal spectral behaviour corresponding to subtle bearing events in jet engine vibration data. These events are manifested as persistent peaks of low spectral power occurring high in the frequency range with respect to the fundamental orders of spectral power. These low spectral powers typically mean that conventional methods find such events difficult to detect. By transformation of the problem into a probabilistic domain, we have shown that such events of low absolute magnitude in power take high novelty scores when compared to an engine-specific model of normality. Novelty detection then takes place in the transformed probabilistic domain.

The resulting system has a low false-positive detection rate when tested using a large dataset of "normal" jet engine vibration data, in which no engine events were present. Though we have illustrated the method using jet engine vibration data, we anticipate that the method is applicable to rotating machinery from which vibration spectra may be obtained.

We have shown that the conventional spectral power plot can be transformed into a probabilistic "novelty" plot that highlights subtle, abnormal spectral content that would otherwise be undetectable. This visualisation method is intended to assist domain experts in the diagnosis of events that manifest themselves as subtle "abnormal" spectral energy.

Acknowledgements

DAC was supported by the NIHR Biomedical Research Centre, Oxford. Thanks to Robert Slater, Steve King, Paul Anuzis, and Mark Walters of Rolls-Royce PLC., and Dennis King, Ken MacDonald, Tom Collings, and Srini Sundaram of Oxford BioSignals Ltd. for useful discussions and in their implementation of this work. Thanks also to Mike Nicholls of J.A. Kemp & Co. Patent Attorneys for useful discussions in clarifying the technical detail described in this paper.

References

[1] L. Tarassenko, D.A. Clifton, P.R. Bannister, S. King, and D. King, "Novelty detection," *Encyclopaedia of Structural Health Monitoring*, 2009.

[2] D.A. Clifton, L.A. Clifton, P.R. Bannister, and L. Tarassenko, "Automated novelty detection in industrial systems," *Studies in Computational Intelligence*, vol. 116, pp. 269–296, 2008.

[3] D.A. Clifton, N. McGrogan, L. Tarassenko, S. King, P. Anuzis, and D. King, "Bayesian extreme value statistics for novelty detection in gas-turbine engines," in *Proceedings of IEEE Aerospace, Montana, USA*, 2008, pp. 1–11.

[4] T.I. Liu and N.R. Iyer, "Diagnosis of roller bearing defects using neural networks," *International Journal of Advanced Manufacturing Technology*, vol. 8, pp. 210–215, 1993.

[5] N. Tandon and A. Choudhury, "A review of vibration and acoustic measurement methods for the detection of defects in rolling element bearings," *Tribology International*, vol. 32, pp. 469–480, 1999.

[6] P. Hayton, S. Utete, and L. Tarassenko, "QUOTE project final report," Tech. Rep., University of Oxford, 2003.

[7] X. Zhang, "Narrow band noise reduction for speech enhancement," Patent, No. WO/2007/130766, 2007.

[8] S.C. Choi and R. Wette, "Maximum likelihood estimation of the parameters of the gamma distribution and their bias," *Technometrics*, vol. 11, no. 4, pp. 683–690, 1969.

[9] S.G Coles and J.A. Tawn, "Statistical methods for multivariate extremes: An application to structural design," *Applied Statistics*, vol. 43, no. 1, pp. 1–48, 1994.

[10] E. Castillo, A. S. Hadi, N. Balakrishnan, and J. M. Sarabia, *Extreme Value and Related Models with Applications in Engineering and Science*, John Wiley and Sons, New York, 2005.

[11] R. A. Fisher and L. H. C. Tippett, "Limiting Forms of the Frequency Distributions of the Largest or Smallest Members of a Sample," *Proceedings of the Cambridge Philosophical Society*, vol. 24, 1928.

[12] P. Embrechts, C. Kluppelberg, and T. Mikosch, *Modelling Extremal Events for Insurance and Finance*, Springer-Verlag, Berlin, 4 edition, 2008.

[13] D.A. Clifton, *Novelty Detection with Extreme Value Theory in Jet Engine Vibration Data*, Ph.D. thesis, University of Oxford, 2009.