

Application of an Intuitive Novelty Metric for Jet Engine Condition Monitoring

David A. Clifton^{1,2}, Peter R. Bannister¹, and Lionel Tarassenko¹

¹ Department of Engineering Science, Oxford University, UK,
{davidc, prb, lionel}@robots.ox.ac.uk

² Oxford BioSignals Ltd., Magdalen Centre, Oxford Science Park,
Oxford, OX4 4GA, UK

Abstract. Application of novelty detection to a new class of jet engine is considered within this paper, providing a worked example of the steps necessary for constructing a model of normality. Abnormal jet engine vibration signatures are automatically identified with respect to a training set of normal examples. Pre-processing steps suitable for this area of application are investigated. An intuitive metric for assigning novelty scores to patterns is introduced, with benefits for reducing model sensitivity to noise, and in pruning patterns from the model training set.

Keywords Application to Manufacturing; Novelty Detection; Jet Engine Application; Machine Learning.

1 Introduction

Novelty detection, defined to be the identification of departures from normal system behaviour, is appropriate for areas of application in which examples of normal behaviour greatly outnumber examples of abnormal behaviour. The identification of abnormal jet engine vibration data is one such example, in which the majority of available examples are “normal”. This paper describes the application of the *shape analysis* method of novelty detection [2] to a new engine example, investigating the necessary pre-processing steps required to provide accurate classification of abnormal examples from a new class of modern jet engine.

An intuitive novelty metric is introduced for assigning novelty scores to engine examples such that sensitivity of the model of normality to noise in example patterns may be reduced. The novelty metric is used as a basis for meaningful pruning of data from the training set such that a robust model of normality is constructed.

Results from the application of this novelty detection process to the new class of jet engine considered within this investigation are presented, providing evidence of the suitability of the approach in the identification of abnormal engine vibration signatures.

Table 1. Data classified according to maximum vibration amplitude

Sub-set	$ D_n $	Class membership criteria
D_1	29	Examples for which $\max\{a(s)\} < 0.9H$
D_2	15	Examples for which $0.9H \leq \max\{a(s)\} \leq H$
D_3	17	Examples for which $H < \max\{a(s)\}$
D_4	10	Examples in which weights are deliberately applied to the engine fan

2 Data Description

This investigation considers a modern civil jet engine, consisting of several rotating engine *shafts*, mounted concentrically. In normal operation, air enters the *low pressure* (“LP”) shaft for compression, passed to the *intermediate pressure* (“IP”) and *high pressure* (“HP”) shafts for further compression, before ultimately being mixed with aviation fuel for use within the combustion chamber. Vibration energy at the frequency of the fundamental harmonic of vibration associated with each shaft is referred to as the *tracked order* of each shaft.

During product testing, engines of the class considered within this investigation perform a controlled two-minute acceleration from idle to maximum speed, followed by a two-minute deceleration back to idle. Vibration amplitude levels of the fundamental tracked orders corresponding to each shaft are recorded during these tests, from which the speed-based vibration signature $a(s)$ is constructed for rotational speeds s of each engine shaft. If $a(s) \leq H$ for all speeds s , where H is a contractual vibration limit, then the engine may be released into service.

The data set used in the investigation described in this article consists of 71 engine examples, initially divided into four sub-sets $D_1 \dots D_4$, according to their maximum vibration amplitude $\max\{a(s)\}$ compared with the contractual vibration limit H , as shown in Table 1. Note that sub-set D_4 is formed from recordings of engines with seeded fault conditions, in which small masses are deliberately applied to the engine fan blades. These weights are deliberately applied to fan (and sometimes turbine) blades in order to correct any unbalances during rotation of those blades about the shaft to which they are connected. This unbalance is noted by engine manufacturers to result in very large vibration amplitude levels of the LP shaft, with approximately normal behaviour in other shafts. Novelty detection *applied to the IP and HP shafts* should therefore not identify examples from sub-set D_4 as being novel.

3 Quantisation of Vibration Signatures

Increasing dimensionality of data requires exponentially increasing numbers of patterns within the data set used to construct a general model; this is termed the *curse of dimensionality* [1]. In order to avoid this problem, each signature is summarised by a *shape vector* \mathbf{x} . This is performed by computing a weighted

average of the vibration amplitude values $a(s)$ over $D = 10$ speed sub-ranges [2]. The d^{th} dimension of shape vector \mathbf{x}^d , for $d = 1 \dots D$, is defined to be:

$$\mathbf{x}^d = \int_{s_{\min}}^{s_{\max}} a(s)\omega_d(s)ds \quad (1)$$

in which the vibration amplitude $a(s)$ is integrated over the speed range $s : [s_{\min} \ s_{\max}]$, using weighting functions $\omega_d(s)$, for $d : 1 \dots D$. The weighting functions used are overlapping trapezoids, as shown in Figure 1.

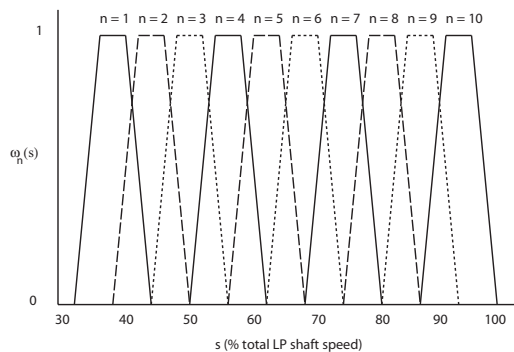


Fig. 1. Example of weighting functions $\omega_d, d : 1 \dots 10$, used to create shape vectors

4 Normalisation

In order to identify changes in the *shape* of 10-dimensional vectors derived from vibration signatures with respect to a population of normal patterns, regardless of absolute values of vibration amplitude, normalisation is applied. This pre-processing step removes dependence upon absolute amplitudes, whilst preserving information about the general shape of signatures.

Five normalisation methods were investigated for scaling the $d : 1 \dots 10$ dimensional shape vectors $\{\mathbf{x}_1^d, \mathbf{x}_2^d, \dots, \mathbf{x}_I^d\}$ constructed from vibration signatures. These normalisation methods are defined in Table 2.

We define the component-wise normalisation function $N(\mathbf{x}_i)$ to be a transformation of the d elements within pattern \mathbf{x}_i :

$$N(\mathbf{x}_i) = \frac{\mathbf{x}_i^d - \boldsymbol{\mu}^d}{\boldsymbol{\sigma}^d}, \quad d = 1 \dots D \quad (2)$$

where $(\boldsymbol{\mu}, \boldsymbol{\sigma})$ are vectors of D elements, computed component-wise across all $i = 1 \dots I$ patterns:

$$\boldsymbol{\mu}^d = \frac{1}{I} \sum_{i=1}^I \mathbf{x}_i^d \quad (3)$$

Table 2. Normalisation methods applied to engine shape vectors

Method	Definition
Energy-based normalisation	$\sum(\mathbf{x}^d)^2 = 1$
Amplitude normalisation	$\sum \mathbf{x}^d = 1$
Unit normalisation	$\max\{\mathbf{x}^d\} = 1$
Zero-mean, unit-variance normalisation	$\mu(\mathbf{x}^d) = 0, \sigma(\mathbf{x}^d) = 1$
Component-wise normalisation	See Eq.(2)

$$\sigma^d = \left(\frac{1}{I-1} \sum_{i=1}^I (\mathbf{x}_i^d - \boldsymbol{\mu}^d)^2 \right)^{\frac{1}{2}} \quad (4)$$

In order to examine the results of normalisation, the data set was visualised by projecting the set of 10-dimensional shape vectors into 2 dimensions.

Topographic projection is a transformation that attempts to best preserve, in the projected space of lower-dimensionality (*latent space*, \mathbb{R}^q), distances between data in their original high-dimensional space (*data space*, \mathbb{R}^d). The *Sammon stress metric*[5] is based upon the distance d_{ij} between points (x_i, x_j) in \mathbb{R}^d , and the distance d_{ij}^* between projected points (y_i, y_j) in \mathbb{R}^q :

$$E_{\text{sam}} = \sum_{i=1}^N \sum_{j>i}^N (d_{ij} - d_{ij}^*)^2 \quad (5)$$

in which the distance measure is typically Euclidean. The NeuroScale model [3, 4] trains a radial basis function (*RBF*) neural network to perform the mapping from \mathbb{R}^d to \mathbb{R}^q , in which E_{sam} is minimised; i.e. distances between points are best preserved after projection.

A NeuroScale network was used for projecting shape vectors derived from the example data set described previously, with $d = 10$ inputs (corresponding to the number of elements in each shape vector) and $q = 2$ outputs (for 2-dimensional projection).

The projection of the un-normalised vibration signatures from the HP shaft during acceleration manoeuvres, as generated by the NeuroScale network, is shown in Figure 2(a). A new NeuroScale mapping is generated for use with each normalisation scheme, by training a RBF network using data normalised by each method previously considered. The resulting projections from each NeuroScale network are shown in Figure 2(b)-(f). Note that the axes of NeuroScale projections have no units of measurement: they are arbitrary orthogonal axes that allow minimisation of E_{sam} .

The figures show that component-based normalisation provides increase in separation of “abnormal” patterns (sub-set D_3), while the other four normalisation methods result in significant overlap between “normal” and “abnormal” patterns. Thus, component-wise normalisation is used in pre-processing the 10-dimensional shape vectors for this engine data-set.

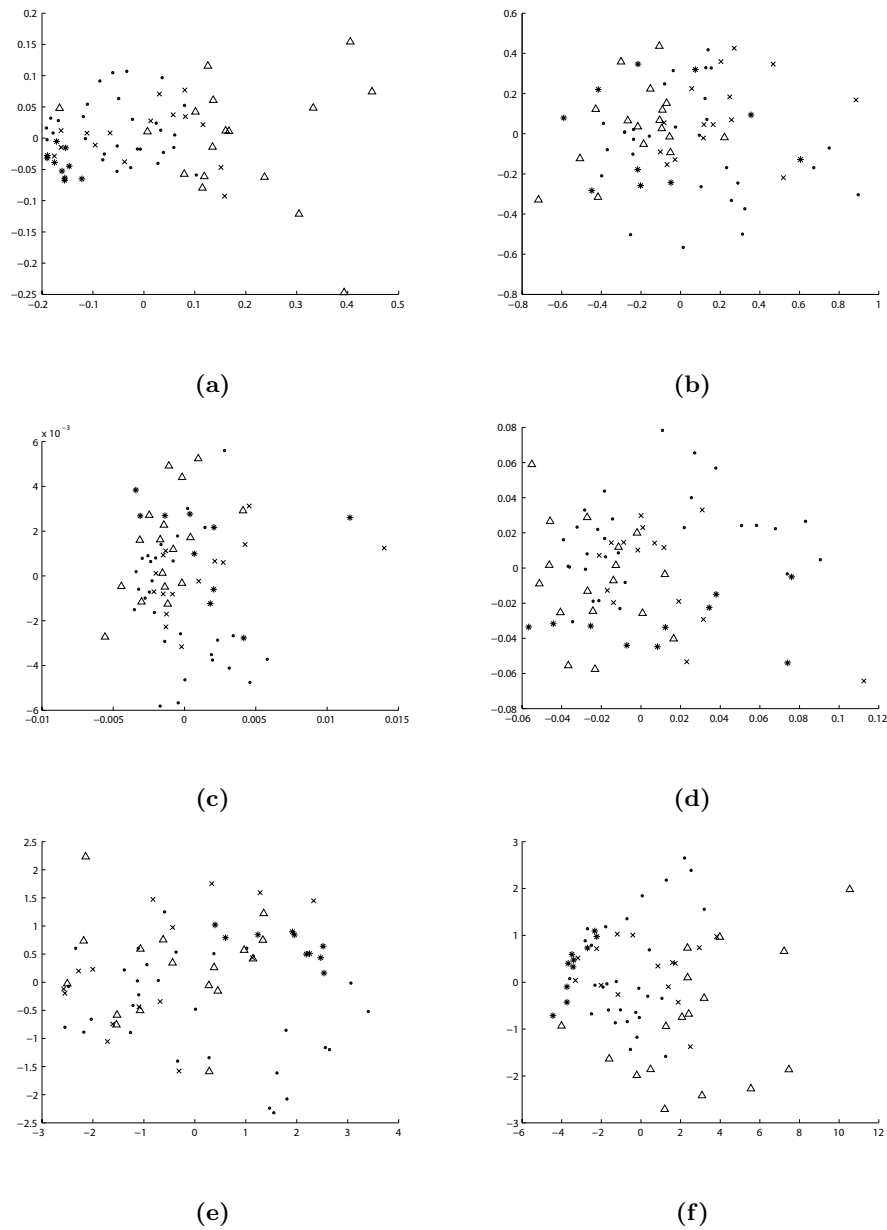


Fig. 2. Projections of HP Shaft Accelerations: (a) un-normalised; (b) unit normalisation; (c) amplitude normalisation; (d) energy-based normalisation; (e) zero-mean, unit-variance normalisation; (f) component-wise normalisation. Data sub-sets $\{D_1 \dots D_4\}$ are shown by $\{\bullet \times \triangle +\}$, respectively.

5 Modelling Normality

The k -means clustering algorithm was used, as described in [6], to construct a model of normality from “normal” patterns (i.e. those from sub-sets D_1, D_2). In this method, the distribution of “normal” patterns is defined by \mathbf{C}_K cluster centres in \mathbb{R}^{10} space, each with an associated *cluster width* σ_K .

A novelty score $z(\mathbf{x})$ may be computed for shape vector \mathbf{x} with respect to the K cluster centres:

$$z(\mathbf{x}) = \min_{k=1}^K \frac{d(\mathbf{x}, \mathbf{C}_k)}{\sigma_k} \quad (6)$$

where $d(\mathbf{x}, \mathbf{C}_k)$ is Euclidean distance. A threshold H_z is defined such that patterns with novelty scores $z(\mathbf{x}) \geq H_z$ are classified as “abnormal”.

The k -means model assigns individual patterns to the nearest \mathbf{C}_k cluster centre. We define the *population* γ_k of cluster \mathbf{C}_k to be:

$$\gamma_k = \left\{ i : [1 \dots I] \mid d(\mathbf{x}, \mathbf{C}_k) < d(\mathbf{x}, \mathbf{C}_j) \right\}, \quad \forall j : [1 \dots K], j \neq k. \quad (7)$$

A pattern is said to be outside a cluster if it lies beyond that cluster’s width, σ_k . Thus, an oft-used first approximation of cluster membership is to set the width σ_k of cluster \mathbf{C}_K according to:

$$\sigma_k = \max_{i \in \gamma_k} d(\mathbf{x}_i, \mathbf{C}_k), \quad (8)$$

effectively placing K hyperspherical clusters into data space \mathbb{R}^d , each corresponding to the membership region of a cluster.

We propose a new definition of width σ_k :

$$\sigma'_k = \sqrt{\frac{1}{|\gamma_k|} \sum_{i \in \gamma_k} d(\mathbf{x}_i, \mathbf{C}_k)^2}, \quad (9)$$

effectively setting width σ'_k to be the standard deviation of intra-cluster distances for each cluster’s population γ_k . Thus, the novelty score $z(\mathbf{x})$ in (6) may be interpreted as the number of standard deviations that \mathbf{x} lies from its nearest cluster centre.

5.1 Sensitivity to Noise

The use of σ_k from (8) causes cluster widths to be dependent on the location of a single pattern in \mathbb{R}^d , i.e. that pattern from population γ_k which is furthest from centre \mathbf{C}_k . Small variations in the placement of that pattern relative to \mathbf{C}_k can result in considerable variations in σ_k , resulting in poor robustness to noisy data. Furthermore, the cluster width σ_k does not represent the *distribution* of its population of patterns γ_k .

This sensitivity to single patterns can be avoided by:

- i. placing cluster centres \mathbf{C}_k using the k -means algorithm, as before;




Fig. 3. Cluster centre widths calculating using: (a) σ_k ; (b) σ'_k . Training patterns and cluster centres are shown by $\{\bullet +\}$, respectively. Cluster boundaries are shown by dashed lines. Outlying pattern shown by a large \bullet .

- ii. computing σ'_k using (9);
- iii. computing $z(\mathbf{x})$ for all patterns using (6);
- iv. re-allocating patterns to cluster populations γ_k according to minimum novelty score using:

$$\gamma_k = \left\{ i : [1 \dots I] \mid \frac{d(\mathbf{x}, \mathbf{C}_k)}{\sigma'_k} < \frac{d(\mathbf{x}, \mathbf{C}_j)}{\sigma'_k} \right\}, \quad \forall j : [1 \dots K], j \neq k; \quad (10)$$

- v. re-computing the final σ'_k using γ_k from the previous step.

This is illustrated in Figure 3, using example data drawn from two circular distributions in \mathbb{R}^2 , and a single outlying pattern. We refer to these two distributions as p_1 and p_2 , appearing in the lower-left and upper-right of Figure 3, respectively. The single outlying pattern, indicated by a large \bullet , lies between p_1 and p_2 , and is slightly closer to the centre of the former.

Figure 3(a) shows widths σ_k computed using (8). The outlying pattern has been assigned to \mathbf{C}_1 , i.e. the cluster centre placed by the k -means algorithms to represent data drawn from p_1 . Thus, σ_1 is placed to include all of its population γ_1 , and does not closely represent the distribution of the data in γ_1 .

Figure 3(b) shows widths σ'_k after the single outlying pattern has been allocated to γ_2 , using (10). Thus, $(\mathbf{C}_1, \sigma'_1)$ closely represents distribution p_1 , and $(\mathbf{C}_2, \sigma'_2)$ has extended slightly to cover the single outlying pattern.

Though an extreme example, this illustrates the improvement in model-fit achieved with (9), in which the resulting two cluster centres closely model the underlying distribution from which the example data were drawn. This is beneficial in application to engine vibration data, in which patterns (derived from engine vibration signatures) are subject to noise.




Fig. 4. Projection of all LP shaft patterns. Patterns classified as “abnormal” are circled. Data sub-sets $\{D_1 \dots D_4\}$ are shown by $\{\bullet \times \triangle *\}$, respectively.

5.2 Improved Pruning of Training Data

The use of (9) to compute σ'_k results in novelty scores $z(\mathbf{x})$, computed using (6), which may be interpreted as the number of standard deviations that \mathbf{x} lies from its closest cluster centre, using normalised distance $d(\mathbf{x}, \mathbf{C}_k)/\sigma'_k$.

This interpretation allowed by the proposed definition of the novelty score $z(\mathbf{x})$ allows initial pruning of patterns from the training set, in order to discard outliers unrepresentative of normality, and which should not be included in the training set. This can be achieved using:

- i-v.** as before, in 5.1;
- vi.** discarding all patterns \mathbf{x} for which $z(\mathbf{x}) > H'_z$.

Thus, if, for example $H'_z = 3.0$, all patterns which lie more than three standard deviations from their nearest cluster centre (using normalised distance) are discarded. This corresponds to the intuitive notion of discarding outliers based on their novelty relative to the distribution of the remainder of the data.

6 Results of Application to Example Engine Data

The results of novelty detection using models trained on patterns from each engine shaft may be visualised using NeuroScale projection from \mathbb{R}^{10} onto \mathbb{R}^2 . Results for the LP and IP shafts are presented.

6.1 LP Shaft Results

Engines that were deliberately unbalanced by attachment of weights to the fan exhibit extremely high vibration amplitude within the LP shaft, as discussed in Section 2. Figure 4 shows a NeuroScale projection of all patterns generated from LP shaft vibration signatures.

6 patterns derived from unbalanced engines (i.e. drawn from sub-set D_4) are contained within the LP shaft data set, all of which are correctly classified “abnormal” by the novelty detection process. The figure shows that they are clearly




Fig. 5. Projection of all IP shaft patterns. Patterns classified as “abnormal” are circled. Data sub-sets $\{D_1 \dots D_4\}$ are shown by $\{\bullet \times \triangle *\}$, respectively.

separated from the cluster formed by normal training patterns, supporting the assumption that unbalanced patterns should appear highly abnormal in the LP shaft.

4 of the 5 patterns that form the set of abnormal example patterns for the LP shaft (i.e. drawn from sub-set D_3) are classified “abnormal” with respect to the model. As can be seen from the figure, these 4 patterns classified “abnormal” are clearly separated from the normal patterns.

The remaining pattern of this abnormal class lies in the centre of the normal cluster of patterns. This engine was released into service following pass-off testing (source A), but exceeded the contractual vibration limit for that shaft, placing it within sub-set D_3 (as defined in Section 2). This provides support for the claim that comparison of maximum vibration amplitude to a fixed contractual limit H_a does not provide a reliable indication of abnormality, and that this shape analysis technique can provide better discrimination between “normal” and “abnormal” engines.

6.2 IP Shaft Results

All abnormal patterns (from sub-set D_3) derived from IP shaft vibration signatures are correctly classified by the novelty detection process, as shown in Figure 5. Obvious separation between “abnormal” and “normal” classes can be seen in the figure.

A single pattern from the “normal” class is classified as “abnormal”. This pattern was pruned from the training set prior to training, due to its outlying position relative to the rest of the “normal” patterns.

The figure shows that patterns from sub-set D_4 are not classified as “abnormal” within the IP shaft, following the observation that changes in vibration levels caused by unbalance conditions are localised within LP shaft vibration signatures.

7 Conclusions

Results from the use of the novelty detection procedure described within this investigation showed that abnormal examples of jet engine vibration signatures could be automatically identified, with a very low false-positive rate.

Appropriate pre-processing of vibration signatures was investigated, with component-wise normalisation proving most effective in increasing separation between normal and abnormal engine examples.

The pruning method resulting from application of the novelty metric described in this investigation results in appropriate removal of patterns that are significantly novel with respect to the majority of training patterns. This allows a model of normality to be constructed that closely models vibration signatures taken from normal engines, such that abnormal engine examples are identified.

Acknowledgements The authors gratefully acknowledge the support of the EPSRC; Dennis King and Mark Walters of Rolls-Royce Plc.; and Nicholas McGrogan of Oxford BioSignals Ltd.

References

1. Bishop, C.: Neural networks for pattern recognition. Oxford University Press (1995)
2. Nairac, A., Townsend N., Carr, R., King, S., Cowley, P., Tarassenko, L.: A system for the analysis of jet engine vibration data. Integrated Computer-Aided Engineering (1999)
3. Lowe, D., Tipping, M.E.: Feed-forward neural networks and topographic mappings for exploratory data analysis. Neural Computing and Applications (1996)
4. Nabney, I.T.: NETLAB, algorithms for pattern recognition. Springer (2002)
5. Sammon, J.W.: A non-linear mapping for data structure analysis. IEEE Transactions on Computers (1969)
6. Nairac, A., Corbett-Clark, T.A., Ripley, R.M., Townsend, N.W., Tarassenko, L.: Choosing an appropriate model for novelty detection. IEE 5th International Conference on Artificial Neural Networks (1997)