Learning Shape for Jet Engine Novelty Detection

David A. Clifton[1,2], Peter R. Bannister[1], and Lionel Tarassenko[1]

[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 a neural network approach to data exploration and the generation of a model of system normality is described for use in novelty detection of vibration characteristics of a modern jet engine. The analysis of the *shape* of engine vibration signatures is shown to improve upon existing methods of engine vibration testing, in which engine vibrations are conventionally compared with a fixed vibration threshold. A refinement of the concept of "novelty scoring" in this approach is also presented.

1 Introduction

Novelty detection, defined to be the identification of departures from a model of "normal" behaviour of a system, is particularly suited to the monitoring of jet engines, in which examples of normal engine behaviour greatly outnumber examples of abnormal behaviour (in common with many high-integrity systems).

Current engine testing techniques depend on the comparison of engine vibration to a simple maximum vibration threshold. This paper presents a method for novelty detection of engine vibration characteristics developed for use with a modern civil jet engine to improve these existing methods, and provides a worked example of the use of neural networks in developing a model of normality for use within novelty detection.

Existing work in the analysis of the *shape* of engine vibration characteristics is extended for use with the new engine type, in which *novelty scores* are computed for each pattern and compared to a decision threshold. This paper presents a refinement of the score definition, such that novelty scores have an intuitive interpretation. Results of applying this technique are shown to emphasise the benefits over existing methods, and the exploration of data and resultant models of normality using neural network methods are presented.

2 Example Data

This investigation considers a modern civil jet engine, consisting of several rotating engine *shafts*, mounted concentrically. In normal operation, air enters the

J. Wang et al. (Eds.): ISNN 2006, LNCS 3973, pp. 828–835, 2006.

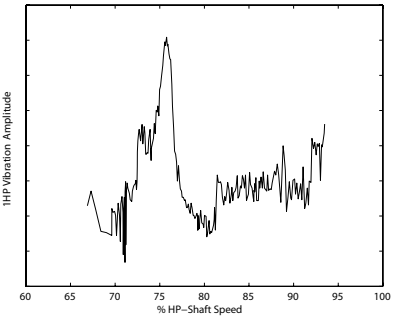

Fig. 1. Speed-based vibration signature $a(s)$ for an example engine (y-axis units anonymised)

low pressure ("LP") shaft for compression, and is passed to the *high pressure* ("HP") shaft, before ultimately being used within the fuel combustion chamber. A further *radial drive* shaft, mechanically connected to the HP shaft, is used to start the engine and provide power during take-off. Vibration energy at the frequency of rotation of each shaft is referred to as the *fundamental 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.

Engines are tested against a contractual limit H which must not be exceeded for the engine to be deliverable; i.e., $a(s) < H$.

An example vibration signature from the data set used in this investigation is shown in Figure 1. The vibration signature for the fundamental tracked order of an engine shaft is constructed with respect to the percentage speed *for that shaft*; thus, amplitude of vibration of the HP shaft is plotted against the percentage of maximum speed for the HP shaft.

The data set used in the investigation described in this article consists of 71 engine examples, initially divided into four sub-sets $D_1 \ldots D_4$, 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.[1] This unbalance is noted by engine

[1] Weights are deliberately applied to fan (and sometimes turbine) blades in order to correct any imbalances during rotation of those blades about the shaft to which they are connected. In order to determine the effect upon engine vibration of adding masses to fan blades, a small number of weights are applied, forcing the blades into a state of unbalance. Examples from class 4 are recorded from an engine undergoing this procedure.

Table 1. Data classified according to maximum vibration amplitude

Sub-set	$\lvert D_n \rvert$	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

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 HP shaft* 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* [2]. 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 $N = 10$ speed sub-ranges [7]. The n^{th} dimension of shape vector $\mathbf{x}^n$, for $n = 1 \ldots N$, is defined to be:

$$\mathbf{x}^n = \int_{s_{\min}}^{s_{\max}} a(s)\omega_n(s)ds \tag{1}$$

in which the vibration amplitude $a(s)$ is integrated over the speed range $s : [s_{\min}\ s_{\max}]$, using weighting functions $\omega_n(s)$, for $n : 1 \ldots N$.

In setting the interval of speeds $s : [s_{\min}\ s_{\max}]$, the operational range of each shaft must be considered independently. For example, the gas-turbine engines of the class considered within this investigation are idle at HP shaft speeds below 60% of maximum rotational speed, during which no measurements of vibration amplitude are made. Thus, the interval [60% 100%] of maximum shaft speed is used when quantising vibration signatures from the HP shaft.

Component-wise normalisation [1, 3] (in which the normalised vibration signature has zero mean and unit variance at each quantised speed) was found to emphasise the separation between "good" (D_1, D_2) and "bad" engines (D_3). Labelling the data by comparison with a contractual limit is intended only as initial guide to model construction.

4 Visualisation Using a NeuroScale RBF Network

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

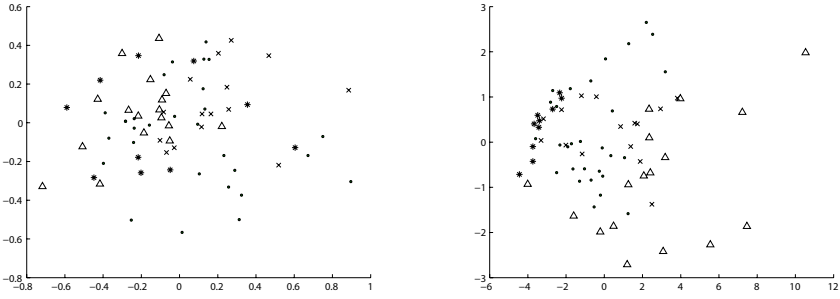

Fig. 2. Projections of HP shaft vibration signatures: (a) un-normalised; (b) component-wise normalisated. Data sub-sets $\{D_1 \ldots D_4\}$ are shown by $\{\bullet \times \triangle *\}$, respectively. Separation of "abnormal" D_3 patterns ($\triangle$) is clearly better in (b) than (a).

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*, $\mathbb{R}^d$). Typically $d > q$, $q = 2$ or 3. The *Sammon stress metric*[8] 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 $\mathbb{R}^q$:

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

in which the distance measure is typically Euclidean. The NeuroScale model [4, 5] 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).

Projection of all 10-dimensional shape vectors before and after component-wise normalisation is shown in Figure 2, with clear separation of "abnormal" patterns (from sub-set D_3) evident.

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:

$$z(\mathbf{x}) = \min_{k=1}^{K} \frac{d(\mathbf{x}, \mathbf{C}_k)}{\sigma_k} \tag{3}$$

where $d(\mathbf{x}, \mathbf{C}_k)$ is Euclidean distance. We propose a new definition of width σ_k:

$$\sigma_k = \sqrt{\frac{1}{I_k} \sum_{i=1}^{I_k} d(\mathbf{x}_i, \mathbf{C}_k)^2}. \tag{4}$$

for the I_k points which have closest cluster centre $\mathbf{C}_k$. This allows an intuitive interpretation of the magnitude of novelty scores: novelty scores $z(\mathbf{x})$ computed using (3) are the number of standard deviations that pattern $\mathbf{x}$ lies from its closest cluster centre, relative to the distribution of training data about $\mathbf{C}_k$. A threshold H_z is applied to $z(\mathbf{x})$ such that all patterns $z(\mathbf{x}) \geq H_z$ are classified "abnormal". H_z is set to best separate the "normal" and "abnormal" patterns in the data set.

Investigation of the placement of cluster centres is possible using the NeuroScale RBF technique. The position of cluster centres with respect to the data set may be determined by projection using the NeuroScale network previously trained using the patterns from the example data set. Selecting a model of normality for use can be assisted through use of the projections generated by the neural network, such that the cluster centres accurately represent the distribution of patterns from the data set.

6 Results of Application to Example Engine Data

The NeuroScale RBF neural network method is further used to examine the results of novelty detection. Selected results are presented within this section from the HP shaft, and the radial drive shaft.

6.1 HP Shaft Results

Figure 3 shows all HP shaft patterns projected into 2 dimensions.

Of the 17 "abnormal" patterns (from sub-set D_3), 12 were classified as "abnormal" by the novelty detection scheme. From these 12 patterns, 5 corresponded to engines with vibration signatures that, though above the simple contractual limit (thus placing them into sub-set D_3, which would thus not pass product testing), exhibited very similar signature shapes to engines that were deemed fit for service. This indicates that the comparison of an engine's maximum vibration level to a fixed contractual limit, as is the standard means of pass-off testing, may not adequately detect abnormal engines.

Of those engines which were deemed fit for service, based on comparison to the simple threshold, two were classified "abnormal" by the scheme investigated in this article. These corresponded to engines which were identified by engine developers as requiring rebuilding due to high vibration levels within the HP shaft.

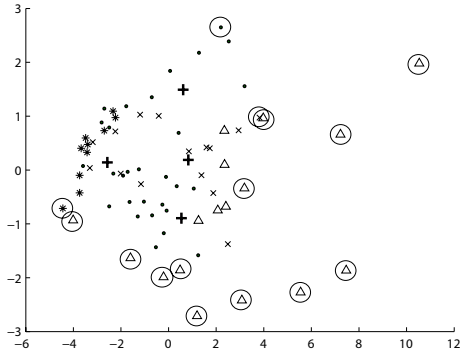

Fig. 3. Projection of all HP shaft patterns. Cluster centres C_K are shown as '+' symbols. Patterns classified as "abnormal" are circled. Data sub-sets $\{D_1 \ldots D_4\}$ are shown by $\{\bullet \times \triangle *\}$, respectively.

The NeuroScale projection also shows that engines that were deliberately unbalanced (from sub-set D_4) exhibit similar vibration signatures, which are classified "normal", agreeing with the observation that weights applied to the engine fan only affect vibration of the LP shaft.

6.2 Radial Drive Shaft Results

Due to no contractual vibration limit being specified for the radial drive shaft, patterns for this shaft could not be divided into the 4 sub-sets described previously. With no prior distinction made between "normal" and "abnormal" patterns, a different approach was taken, in which engines that were released into

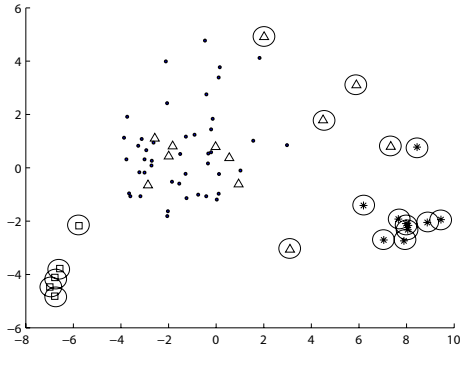

Fig. 4. Projection of all radial drive shaft patterns. Patterns classified as "abnormal" are circled. Engines that failed testing are shown as $\triangle$. Engines that were re-fitted during service are shown as $\square$. Engines with forced unbalance are shown as $*$.

service following pass-off testing formed the training set for the construction of the model of normality.

Figure 4 shows a NeuroScale projection of all patterns derived from radial drive shaft vibration signatures.

The NeuroScale projection shows that the older engines, re-fitted during service, form a distinct cluster in the lower-left of the plot, indicating that extended use of the engine in service results in a change in vibration characteristics. Furthermore, it can be seen that unbalanced patterns show significant difference to "normal" engines in the radial drive shaft, indicating that this condition is not limited only to the LP shaft.

It can also be seen that several engines which fail pass-off testing appear "normal" according to radial drive shaft vibration levels.

7 Conclusions

The method of vibration analysis presented in this paper has shown that the "shape analysis"process is effective for novelty detection in vibration data from new engine classes. Results from analysis of vibration signatures from each shaft of the example engine have shown that the majority of patterns derived from abnormal engine recordings are correctly classified, whilst false-positive detections associated with normal patterns are low.

The method presented provides a more accurate assessment of abnormality than comparison of vibration levels to a simple threshold, whilst visualisation by neural network allows changes in engine condition to be examined. Hence, increasing accuracy of engine assessment allows improved control of engine maintenance, and enhances the diagnostic process.

Acknowledgements

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

References

1. Bannister, P.R., Tarassenko, L., Shaylor, I.: SPANN-04-PB1 Shape Analysis Tools, Department of Engineering Science, University of Oxford, Technical Report (2004)
2. Bishop, C.: Neural Networks for Pattern Recognition. Oxford University Press, Oxford (1995)
3. Clifton, D.A.: Condition Monitoring of Gas-Turbine Engines. Department of Engineering Science, University of Oxford, Transfer Report (2005)
4. Lowe, D., Tipping, M.E.: Feed-Forward Neural Networks and Topographic Mappings for Exploratory Data Analysis. Neural Computing and Applications 4(2) (1996) 83-85
5. Nabney, I.T.: NETLAB, Algorithms for Pattern Recognition. Springer-Verlag, New York (2002)

6. Nairac, A., Corbett-Clark, T.A., Ripley, R.M., Townsend, N.W., Tarassenko, L.: Choosing an Appropriate Model for Novelty Detection. In: Proc. 5th IEE Int. Conf. on Artificial Neural Networks, Cambridge (1997) 117-122
7. 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 6(1) (1999) 53-56
8. Sammon, J.W.: A Non-Linear Mapping for Data Structure Analysis. IEEE Transactions on Computers C-18(5) (1969) 401-409