Constructing and Re-training Models for Condition Monitoring in Jet Engines

Peter R Bannister, David A Clifton and Lionel Tarassenko
Biomedical Signal Processing Research Group, Department of Engineering Science,
University of Oxford
Oxford, OX1 3PJ, UK
Tel: 01865 273121
Fax: 01865 273905
prb@robots.ox.ac.uk

Abstract

This paper describes a method for constructing a model of normal system state for the purpose of providing continuous condition monitoring of an aerospace gas-turbine engine.

We show how the model can be visualised in 2D to allow an intuitive identification of abnormal engine and how periods of abnormal engine condition can be identified by means of a *novelty score*. This approach is illustrated using vibration data from an Engine Health Monitoring (EHM) system.

This paper focuses on the specific case of modelling engine condition following a maintenance overhaul during which the characteristics of an engine can shift significantly, leading to false positive alerts based on the existing engine model.

1. Introduction

Techniques for prognostic fault detection have a high value amongst the manufacturers of high integrity systems, for example a gas aero turbine engine or power generation system. Recent work in the field of *novelty detection* has yielded techniques for monitoring such systems and alerting in the case of abnormal behaviour being detected(1). A key feature of these schemes is that no prior information regarding the nature of the fault is assumed: the novelty detection paradigm seeks only to identify areas of operation that differ to what it has been trained to recognise as normal. This distinction from fault-detection methods (which assume some *a priori* characterisation of the fault) is particularly important for the high integrity systems under consideration, as there may be a very large number of potential fault cases which cannot be exhaustively seeded and investigated during the development of the system.

This paper presents a method for detecting novel behaviour in multi-dimensional condition monitoring systems that is able to characterise the engine by monitoring combinations of standard EHM system outputs, making it possible to quickly and intuitively ascertain if an engine is behaving normally according to the previously learned model. This work demonstrates that it is possible to fuse existing EHM vibration data from multiple channels into *feature vectors* and to then determine an optimal mapping that can represent these high-dimensional features in 2D for visualisation.

It is shown that this low-dimensional representation of the score data can adequately capture differences between sets of flight data to track variations in the normal vibration profile for an engine in service.

In this paper we pay particular attention to the case of condition monitoring of an engine after overhaul, where the characterstics associated with normal condition may shift significantly following maintenance or repair.

Results are based on a 3-shaft Rolls-Royce aero gas turbine engine, demonstrating that the models may be used to analyse in-service flight data by graphical as well as numerical means. These approaches are complementary and lend themselves to high-level analysis of engine condition where the outputs of individual EHM modules are unable to detect a shift in engine condition.

The paper begins with a description of the steps that are taken to construct a data-driven model of engine condition before presenting both the visual and quantifiable results of flight data analysis. This is followed by a discussion of extensions of the scheme to cope with multiple datatypes and to enable automatic alerting.

2. Method

This section presents an overview of the data-types collected by the EHM system and details the steps that are necessary to render this data suitable for constructing a model of engine condition.

2.1 Input Data

Vibration data is captured using a QUICK^{TM} Engine Health Monitoring system(2) and processed using a number of software modules, or *Feature Detectors*. Amplitude and frequency information is captured at 5 Hz and each Feature Detector (FD) is designed to identify abnormal patterns in the vibration data that may be indicative of an engine failure. For the purpose of this investigation we focus on the following FD datatypes: the fundamental Tracked Order (TO), which is a measure of vibration amplitude at the frequency of engine shaft rotation; the fractional and multiple TOs associated with a set of common shaft harmonics; the residual energy that cannot be assigned to these tracked order peaks; and the broadband level up to 1kHz.

2.2 Pre-processing

The outputs of each of the FDs is median filtered to remove noise, data considered to lie below the idle speed of each engine shaft is rejected and the FD outputs (or *scores*) are then grouped together according to one of four feature vector designs.

2.2.1 Feature Vector Construction

The FDs provide both shaft-specific and engine-wide measures of particular parameters suggesting a number of models based on these EHM outputs. In this paper, results from the low-pressure shaft system of the engine are presented where the following scores are grouped together to form a feature vector:

[Fundamental TO; 2nd harmonic TO; 3rd harmonic TO; 4th harmonic TO; 5th harmonic TO; 1.5th harmonic TO; residual energy; broadband]

2.2.2 Speed Change Tracking

For medium to long-haul air travel, a typical flight will consist of a short period of high thrust at take-off and during landing with the remaining bulk of the flight (cruise) taking place under relatively steady conditions. As a result, the data that is collected by the EHM will comprise mainly the latter datatype and in order to prevent the engine model being unduly biased towards these regions of steady-state behaviour, data is filtered so that the model contains a balance of data from transient, high-power periods of operation and cruise. It is often (but not always) the case that engine condition will deteriorate more rapidly while the system is under the greatest stress so it is important the model is sensitive to changes in condition during these periods of operation.

In order to ensure this even representation of operating conditions, test data is filtered to reject any feature vectors whose associated shaft speed has changed by less than 0.5% since the previous sample. An added benefit to this approach is that the quantity of data being considered for the model is significantly reduced by rejecting this repeated data, making the problem of training the model itself significantly more tractable.

2.2.3 Speed partitioning

It is known that the vibration characteristics of an engine shift according to the current region of operation, whether cruise, climb or descent(3). To improve the sensitivity of the model, feature vectors associated with a given shaft are sub-divided against speed so that data gathered in the lowest quartile of the operating speed range is used to train a low-speed model with data from the remaining three quartiles also being modelled independently.

2.2.4 Normalisation

Once a set of feature vectors has been constructed using the methods described above, each element is component-wise normalised in order to equalise the multi-dimensional outputs of the FDs.

2.3 Model Training

Each model is described by the set of feature vectors used in its construction and these are used to train a projection function (used to display the high-dimensional feature vectors in 2D) and a data density estimator in parallel.

2.3.1 Visualisation

The mapping from feature-vector space into 2D is achieved through the use of a Neuroscale mapping(4). In this approach, a Radial Basis Function neural network is trained on the EHM data to give a mapping into 2D which aims to make the Euclidean distances between pairs of image patterns in the 2D visualisation space as close as possible to the Euclidean distances between the corresponding pair of patterns in the original space. Once trained, this mapping can be applied to new data, allowing easy comparison with the training data.

The Neuroscale algorithm is trained on a set of 500 prototype feature vectors that are generated by running the K-means clustering algorithm on the set of feature vectors gathered to train each model, generating 500 *centres* in the high-dimensional space.

This approach facilitates an intuitive presentation of data that typically covers multiple channels by means of a simple 2D plot. Examples are presented in section 3.

2.3.2 Density Estimation

To complement the visual presentation of the flight data (which in itself cannot be guaranteed to highlight all discrepancies between the training data and a subsequent flight), an overall measure of the engine condition with respect to the model, referred to as a *novelty score*, can be calculated using a probability density estimation technique, such as a Parzen window kernel density estimator(5).

By placing a Gaussian kernel over each feature vector in the training dataset (now reduced to 500 vectors), it is possible to construct a probably density estimate of the data, $p(x)$. Test data is compared to the Parzen window model and feature vectors a long distance away from the Gaussian kernels placed on the 500 training vectors (and therefore unrepresentative of the condition captured by the model) will yield a low probability density estimate for $p(x)$. The novelty score is defined as $-\log_e p(x)$ which will ensure that a high novelty score is generated for these areas of low probability.

Scores for test data are presented in section 5 and these demonstrate how an overall novelty score can be used to quickly and reliably identify periods of abnormal condition for an engine.

3. Test Data

The methods described in the previous section were applied to test flight data taken from a group of Rolls-Royce engines. The primary focus of this investigation was to determine if the change in engine characteristic resulting from an overhaul was sufficient to warrant re-training of the model. Anecdotal evidence has previously suggested that the characteristics of an engine can change dramatically after servicing (where major engine systems may be completely replaced) leading to false alerts when the original model is used to test post-overhaul flight data.

Models for an engine from this test programme are presented in the following sections. For this investigation, data from the low-pressure shaft system is used to construct a model of normal engine condition. During the course of flight-testing this engine was temporarily removed from the wing for overhaul, to be re-instated at a later date during continued flight-testing.

4. Visualisation Results

Figures 1 and 2 show the visual comparison between the model for this engine (plotted as black circles) and the test data (grey dots). Figure 1 presents a comparison between the original (pre-maintenance) model and post-maintenance data. It is possible to see from the scaling and spatial distribution of the data that there is a considerable discrepancy between the otherwise normal post-maintenance test data and the pre-maintenance model, which suggests that the model should no longer be relied on to monitor normality in the system from this point. Such changes are commonly seen after an engine reset and this observation seems to confirm that the engine condition models should be re-trained in the event of an engine being taken off the wing and subsequently returned.

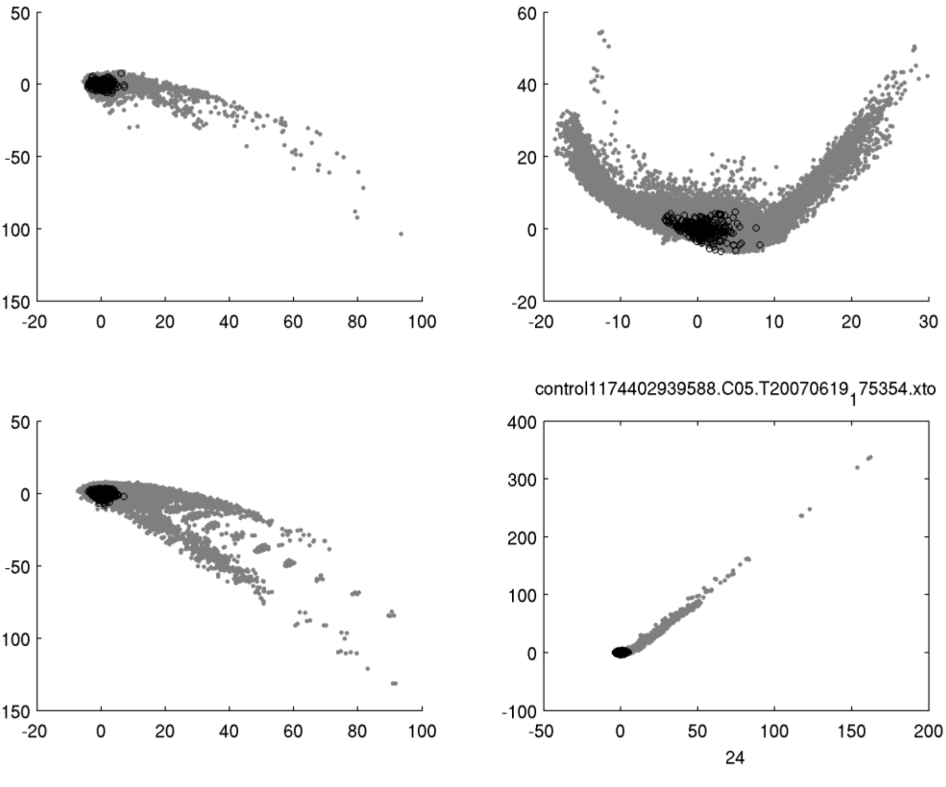

Figure 1. Use of a pre-maintenance model with post-maintenance data. Model prototype feature vectors (black) for this engine are shown after their projection into 2D along with test data from the same aircraft. The obvious disparity in clustering between the two sets of data can be attributed to the engine overhaul carried out between model training and test data acquisition. Each plot represents a different quartile of the engine speed range, beginning with the lowest speeds in the top left-hand corner and progressing to high-speed data in the bottom right-hand plot.

Figure 2 provides a similar comparison of the same post-maintenance flight data, this time compared to a model trained on (separate) post-maintenance data. The two sets of points are much more consistently distributed than those in Figure 1, although there is a shift in cluster centroids in three of the speed ranges, this offset is maintained through subsequent flights indicating that the engine's condition is stable.

These plots demonstrate the intuitive manner in which high-dimensional models can be compared in 2D, facilitating an immediate high-level assessment of an engine's condition based on existing EHM data.

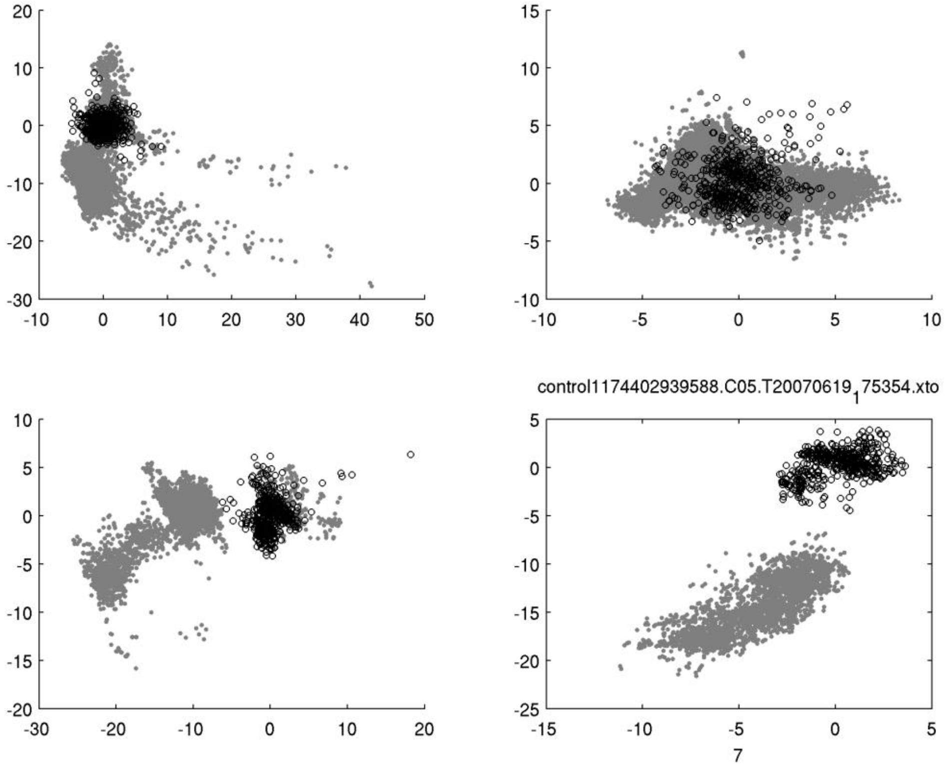

Figure 2. Use of a post-maintenance model with post-maintenance data. Model prototype feature vectors (black) for this engine are shown after their projection into 2D along with test data from the same aircraft. The obvious disparity in clustering between the two sets of data seen in Figure 1 has been significantly reduced and this close agreement between model and test data confirms that a re-trained post-maintenance model is required in the event of an engine overhaul. Each plot represents a different quartile of the engine speed range, beginning with the lowest speeds in the top left-hand corner and progressing to high-speed data in the bottom right-hand plot.

5. Density Estimation Results

Novelty scores (computed as $-\log_e p(x)$) are shown in Figure 3. This model corresponds to the data from the 3rd quartile speed range in Figures 1 and 2. The graphical analysis in the previous section is confirmed by these scores, which show a high degree of abnormality for the test data EHM data when compared to the pre-maintenance model (red) but a much closer correlation to the post-overhaul model (blue).

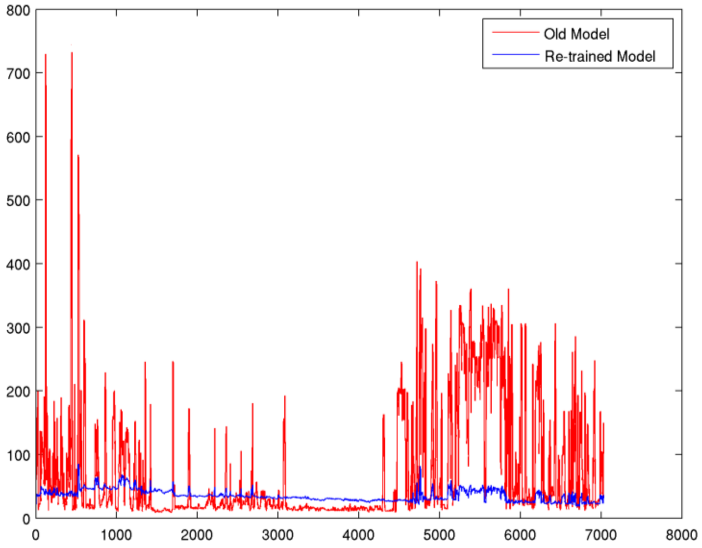

Figure 3. Density Estimation: Engine B. A comparison of the novelty scores generated using the pre-maintenance model (red) and the post-maintenance model (blue) confirms that the post-maintenance flight data necessitates the use of a new model for this engine that has been trained after overhaul has been carried out. Indices along the horizontal axis correspond to samples taken during the test flight.

6. Conclusions

This paper has demonstrated a method for combining and visualising data from a multi-channel EHM system so that periods of abnormal system condition may be identified using two complementary approaches: a graphical visualisation of the high-dimensional input data using a 2D projection and a quantitative measure of novelty, computed using a Parzen window density estimator.

More specifically, the results on in-flight data clearly demonstrate that there is a fundamental shift in the condition of an engine as characterised by these methods after a maintenance overhaul has been carried out. This result strongly suggests that engine-specific models should be reset after an overhaul to avoid false alerting, which may be very costly to the engine owner.

Future work will extend this modelling approach to data from multiple EHM systems with the aim of fusing vibration and performance data into a comprehensive multi-modal condition-monitoring tool. Principled methods for setting the novelty thresholds for these data-driven models have also been investigated(6) and are currently being validated for in-service use.

Acknowledgements

This work was carried out under the DTI HECToR grant. The authors would like to extend particular thanks to Oxford BioSignals Ltd. and Rolls-Royce Plc. for ongoing collaborations.

References

1. L Tarassenko, DA Clifton, PR Bannister, S King and D King. Novelty Detection. In Worden, K. *et al* (eds): Encyclopaedia of Structural Health Monitoring. Wiley, 2008.

2. P Hayton, S Utete, and L Tarassenko. QUOTE Project Final Report. Technical Report, Department of Engineering Science, University of Oxford, 2005.

3. The Jet Engine. Rolls-Royce Technical Publications, London, 2005.

4. I Nabney. NETLAB: Algorithms for Pattern Recognition. Springer-Verlag: Berlin, 2002.

5. E Parzen. On Estimation of a Probability Density Function and Mode. Annals of Mathematical Statistics, 33(3), 1962; 1065-1076.