SPECIFIC AND GENERIC MODELLING FOR JET ENGINE NOVELTY DETECTION

David A. Clifton
Dept. of Engineering Science, University of Oxford, Parks Road, Oxford, OX1 3PJ, UK
Oxford BioSignals Ltd. Brook House, 174, Milton Park, Abingdon, OX14 4SE, UK
davidc@robots.ox.ac.uk

Ben Haskins
Oxford BioSignals Ltd.
Brook House, 174, Milton Park,
Abingdon, OX14 4SE, UK
ben.haskins@oxford-biosignals.com

Peter R. Bannister
Dept. of Engineering Science
University of Oxford
Parks Road,
Oxford, OX1 3PJ, UK
prb@robots.ox.ac.uk

Lionel Tarassenko
Dept. of Engineering Science
University of Oxford
Parks Road,
Oxford, OX1 3PJ, UK
lionel@robots.ox.ac.uk

ABSTRACT

This paper describes a method for modelling the normal behaviour of a modern military gas-turbine engine, for which fault data are particularly scarce. Signatures of fundamental tracked order vibration amplitude are used to characterise the normal operation of the engine, with models formed in high-dimensional space trained on a fleet of engines of the same class. These models are shown to be able to identify fault behaviour in engine tests of the same type, using well-defined "PAT" engine tests. These generic class-wide models are then applied to a single engine undergoing repeated cyclic testing ("ASMET" cycles), to evaluate their performance with respect to the fleet-based models. Initially proven using a retrospective analysis, identification of potential engine events is shown to be possible using the visualisation models developed using the original fleet data, in which such "events" are shown as step-changes in signature visualisation. Models of normality formed from pass-off tests are shown to be inapplicable for endurance tests, due to the differing nature of each test type.

KEYWORDS
Condition Monitoring, Novelty Detection, Visualisation, Vibration, Gas-Turbine Engines

1. INTRODUCTION

1.1. The Need for Novelty Detection

Modern industrial companies manufacturing high-integrity systems (such as aircraft gas-turbine engines) are increasingly shifting their business model to provision of after-market services and system maintenance, in which a guarantee of system availability ("up-time") is given to customers [7]. Robust intelligent condition monitoring techniques are required for detection of potential failure in such systems. In addition to providing early warning of critical failure, thus allowing the avoidance of hazard, such techniques enable a preventative "needs-based" approach to system maintenance, in contrast to traditional dependence on preventative maintenance.

A *novelty detection* approach is taken to cope with the limited availability of failure data for complex systems such as the engine under consideration. Models of normal system operation are constructed from the large quantity of available "normal" data, with departures from the normal model identified as "abnormal" episodes. The degree of "disagreement" between model and observed data can provide an estimate of the severity of the event. This approach is also termed "one-class classification", in reference to the "one-sided" nature of the available data-sets, in which new samples are classified as being members or non-members of the single ("normal") class [1].

1.2. Generic vs. Specific Novelty Detection

High system complexity can result in large *inter-class* variability in observed data, in which each new system build can behave significantly differently to previous builds (such as comparing successive generations of gas-turbine engines). This makes transferring system knowledge from previous builds to new developments particularly difficult, and re-learning this knowledge for new classes can be expensive. Furthermore, individual instances within a class often exhibit high *intra-class* variability (such as comparing different engine units of the same generation).

This variability in system behaviour may make *generic* class-wide models of normality (such as models constructed for an entire type of gas-turbine engine) inappropriate for detailed modelling of individuals drawn from that class. In such cases, a *specific* model of normality may be required to be learned for each individual.

This paper compares the performance of current state-of-the-art novelty detection techniques for both generic and specific modelling of gas-turbine engine vibration data, using data from a new military gas-turbine engine currently in development. Section 2 introduces the domain of application, providing a brief introduction to the dynamics of gas-turbine engine operation, in order to then introduce the data set used throughout this study.

An extension to existing novelty detection techniques is described in Section 3 for application to a new class of military engine, in which a generic model of normality is constructed for the entire engine class. This method is shown to identify abnormal engine behaviour during "pass-off" ("PAT") testing, in which a well-defined manoeuvre is performed on a test bed. Visualisation techniques are introduced to graphically communicate novelty detections to a user in an intuitive manner.

Novelty detection using the generic model of normality is then applied, as described in Section 4, to a single engine undergoing cyclic ("ASMET") testing, in which test duration and contents are significantly different to pass-off testing. Visualisation techniques, trained using the generic data-set, are shown to allow identification of abnormal episodes in engine operation, with indication of a significant engine event shown to be possible several hours in advance of its occurrence, using a retrospective analysis of the data set.

Conclusions and recommendations are drawn in Section 5.

For reasons of commercial sensitivity, data and system descriptions have been anonymised where necessary by not divulging absolute values of operational parameters.

2. APPLICATION DOMAIN DESCRIPTION

2.1. Vibration Data from Gas-Turbine Engine Operation

Modern aerospace gas-turbine engines consist of several concentrically-mounted, rotating engine shafts. The modern military engine considered within this paper is a "two-shaft" model, in which atmospheric air is compressed in two successive increments by passing through the low pressure ("LP") and high pressure ("HP") compressors. The output pressure of the final shaft is sufficient for use within the combustion chamber, where aviation fuel is mixed with the pressurised air, and ignited to produce thrust [10]. Broadband measurements of vibration amplitude and phase are made using two non-invasive sensors, using the QUICK system [9].

2.2. Finding Tracked Orders in Vibration Data

Vibration energy is concentrated at the frequency of the fundamental harmonic of vibration associated with each shaft. Note that the fundamental harmonic of a shaft occurs at the frequency of rotation of that shaft. Thus, broadband vibration spectra for a two-shaft engine will exhibit peaks in energy at the rotational frequency of the LP and HP shafts. Tracking the amplitude and phase of vibration at these peak-energy frequencies gives the *tracked order* for each shaft [3]. These are sometimes elsewhere termed the "first engine order" of the LP shaft and the HP shaft.

2.3. Data Description

Two data-sets, described below, were used in the investigation described by this paper, corresponding to the generic and specific models of normality introduced in the previous section.

Data-Set I – Generic Class Data

This data-set consisted of tracked order time-series recorded from 127 engine tests, each test corresponding to a different engine build, as shown in Table 1. For each engine test, vibration amplitudes and phases were provided for the LP and HP fundamental tracked orders (hereafter referred to as "1LP" and "1HP", respectively), using the QUICK system [5,9].

Each engine in this data-set performed a two-minute acceleration from idle to maximum speed, followed by a two-minute deceleration back to idle.

Table 1 - Data Set I

Engine test number	Comment
1-116	Tests labelled as "Normal"
117-127	Tests labelled as "Abnormal"

It is important to note that the "abnormal" engines were labelled as such for *any* tests for which test engineers noticed an engine anomaly, whether or not it affected vibration measurements observed from the engine (such as, for example, an issue with oil temperature). Thus, the initial labels are treated only as a guide for model construction. Methods are later introduced for determining the veracity of these initial labels, as should be undertaken for any externally-labelled data prior to model construction.

This data-set is used in the construction of a generic class-wide model of normality, in Section 3.

Data-Set II – Specific Individual-Engine Data

This data-set contained a series of cyclic tests of the same engine, as shown in Table 2. Each test followed a complex schedule of acceleration and deceleration manoeuvres, in contrast to the simple acceleration/deceleration pair performed for Data-Set I.

Table 2 - Data Set II

Engine test number	Comment
1-918	Tests unlabelled

This data-set is used in the comparison of generic and specific models of normality, in Section 4.

3. GENERIC MODELS OF NORMALITY

3.1. Creating Signatures from Tracked Orders

We define a *vibration signature* to be the vibration amplitude and phase of a tracked order measured over a range of speeds of the corresponding shaft. For example, a signature may be constructed from vibration data associated with the 1HP tracked order measured as a function of the speed of the HP shaft. Amplitude and phase of tracked order vibration are measured across the speed range 0 to 100% maximum speed of the corresponding shaft. This speed range is subdivided into B equal bins, defining B sub-ranges ("bins") of the operating point (shaft speed). We note that typical phase variance is such that we should expect improved performance by considering phase values separately from amplitude [6].

In this investigation, we form signatures of vibration amplitude by integrating over a series of weighting functions, as described in [4], effectively making the B-dimensional signature a series of (scaled) fuzzy memberships over the B speed bins covering the speed range 0 to 100%. Due to the *curse of dimensionality*, the amount of data required to train a robust model increases exponentially with B [2]. Thus, the value of B must be selected to be large enough such that sufficient information is retained in order to perform successful novelty detection, and small enough such that the amount of data available allows effective training of a model of normality.

A separate signature is created from vibration amplitudes of each tracked order (i.e., 1LP, 1HP) per run. Thus, the data-set shown in Table 1 resulted in 127 signatures for 1LP, and 127 for 1HP. Each signature effectively summarises (in a vector of B dimensions) the vibration response of the engine during a test.

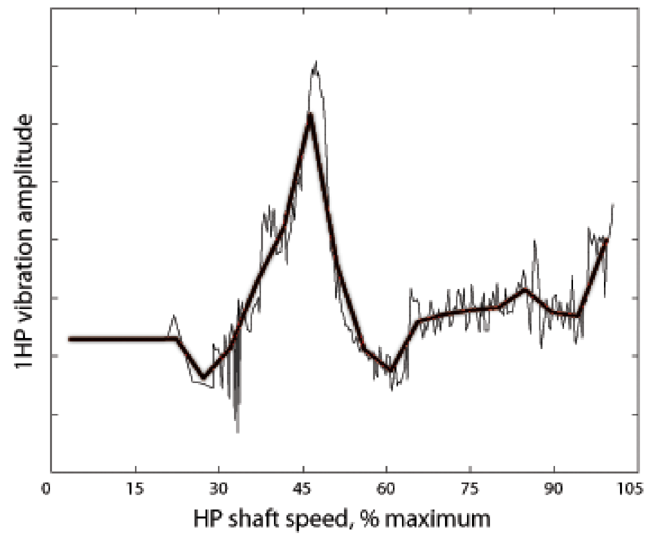

Figure 1 – Creating a Signature from Test Data

Figure 1 shows an example of comparing a 20-dimensional signature (shown in black) from a full-resolution signature formed with B = 400 bins (shown in grey). The overall *shape* of the underlying signature is characterised by the coarsely-quantised 20-D signature, which is used as the basis for detecting novelty.

3.2. Exploring the Data-Set with Visualisation

In order to examine the relationship between the sets of 127 signatures (for 1LP and 1HP tracked orders), a visualisation method is required in order to project these high-dimensional quantities into a smaller number of dimensions, suitable for human inspection.

This investigation uses the NeuroScale technique [8], in which a radial basis function neural network is used to implement Sammon's Mapping [11], which maps each 20-dimensional signature into 2-dimensions. NeuroScale implements a distance-preserving transformation, in which signatures that are "close" in 20-dimensional data space (defined using a Euclidean distance metric) appear close when projected into 2 dimensions. Similarly, signatures that are significantly distant from one another in 20-dimensional space appear appropriately distant when projected. Thus, it is anticipated that signatures constructed from "normal" engine data will cluster together in 2-dimensional projected space, allowing "abnormal" engines to be seen as significantly separate from the "normal" cluster.

The NeuroScale network was trained using the full data-set of generic signatures (Data-Set I), with network architecture parameters selected and validated empirically (details not shown here). A projection of the 127 signatures formed from 1HP data is shown in Figure 2. Note that the axes of the projection are unit-less, with the origin (0,0) located at the centroid of the training set.

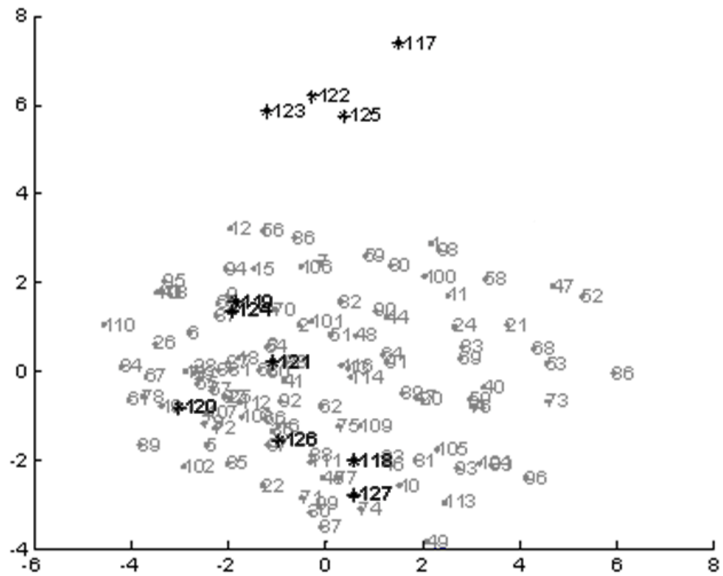

Figure 2 - NeuroScale Projection of 127 Signatures (1HP)

The figure shows signatures 1-116, initially labelled as "normal" (see Table 1), in grey. Signatures 117-127, initially labelled as "abnormal", are shown in black. The "normal" signatures form a cluster about the centroid (0,0) as anticipated, indicating that the engines from which those signatures were constructed exhibit similar vibration responses.

However, only four of the "abnormal" signatures are shown separated from the main cluster of "normal" signatures; the remaining six appear within the locus of "normality". On further examination of the time-series of vibration data from which these six "abnormal"-labelled signatures were derived, it was shown that the six events that resulted in the labelling had no observable consequence on vibration characteristics, thus it is expected that they should appear "normal" when considering only vibration data. This illustrates how labels provided by domain experts can be validated in a principled and rapid fashion, and such techniques can be powerful tools in deciding which signatures should be used in the training sets for constructing models of normality.

3.3. Automatic Novelty Detection Using a Generic Model of Normality

In order to perform automatic detection of abnormal signatures, a model of normality must be constructed. Novelty detection should take place in the high-dimensional data space of the signatures, rather than in the 2-dimensional projected space: much information is discarded in the projection from 20-D to 2-D, which should be retained for the purposes of novelty detection.

We form a generic model of normality using a training set of the 116 "normal" signatures for each tracked order, by placing a small number of prototype centres into 20-dimensional space using the batch *K*-means algorithm [11]. The number of centres *K* was selected using a detailed empirical method (not described here) such that the data space covered by "normal" signatures was adequately described without over-fitting, as occurs when too many centres are used to represent the training set (resulting in poor novelty detection ability when generalised to previously-unseen test signatures).

We here define a *novelty score*, z_i, for signature x_i:

$$z_i = \min\left\{ j:[1\ldots K] \bullet \frac{d\left(x_i, C_j\right)}{\sigma_j} \right\}$$

(1)

where d is the Euclidean distance norm between x_i and the j^{th} prototype centre C_j. We also define the width of the j^{th} centre, σ_j, to be based on the I_j signatures from the training set which have closest centre C_j:

$$\sigma_j = \sqrt{\frac{1}{I_j}\sum_{i=1}^{I_j} d\left(x_i, C_j\right)^2} \tag{2}$$

Signatures with novelty score $z > H$, for novelty threshold H, are classified "abnormal" with respect to the model. H is set to be the maximum novelty score from the training set; i.e., $H = \max\{i = 1 : 116 \mid z_i\}$.

Computing novelty scores z_i for all $i = 1 \ldots 127$ signatures, both 1LP and 1HP, resulted in correct automatic identification of the four "abnormal" signatures shown in Figure 2, with all other 123 signatures classified "normal" with respect to the model, as desired. This allows automatic detection of signatures from their original vibration data, with visualisation used to provide visual feedback for the system operator.

4. SPECIFIC MODELS OF NORMALITY

4.1. Application of a Generic Model to Repeated Cyclic Testing

Using Data-Set II, signatures were created for its 918 tests, using 1LP and 1HP tracked order vibration data as described previously. As noted in Section 2, this data-set was recorded from a single engine performing a fixed schedule of (ASMET) manoeuvres, which were longer in duration and more complex (involving more manoeuvres that moved through partial sub-ranges of the total speed range) than the pass-off tests used to create the initial models of normality.

The generic NeuroScale projection network trained as described in Section 3.2 was used to project all 918 signatures from 1LP and 1HP tracked order vibration data. The generic model was used to allow comparison between signatures from Data-Set I and II.

Figure 3 shows NeuroScale projections of 1HP signatures for the entire series of 918 tests, which build up in a cumulative manner (i.e., all previously-plotted points are retained in subsequent plots).

Figure 3a shows the initial 25 tests, in which test equipment was calibrated. Non-standard test-configuration tests (dashed line) are clearly distinguishable from the period of "normal" tests into which the series stabilises. Figure 3b shows stable behaviour of the engine over the next 76 tests, in which vibration characteristics change marginally from test to test, drifting in the direction indicated by the arrow.

Figure 3c shows a significant deviation from previously-stable behaviour. Retrospective investigation showed that this was due to a change in vibration characteristics occurring in an engine sub-system, affecting main shaft vibration. The engine returned to its previous locus of stable operation for 154 tests, as shown in Figure 3d, again slowly moving in the direction of the arrow.

Figure 3e shows the effect upon vibration characteristics due to a change in the data acquisition system (a gain-boosting circuit was removed from the acquisition loop, resulting in attenuated vibration amplitudes). All subsequent data were recorded with this new configuration, and the engine remained stable within this region of operation until test 911.

Figure 3f shows a significant change in vibration signatures from tests 911-918. An engine event resulting in halting of the test series occurred during test 918, showing that indication of possible changes in engine condition were produced by this investigation up to 7 tests in advance, in test 911. It should be noted that this analysis is retrospective, as discussed in Section 5.

4.2. Automatic Novelty Detection

The locus of "normal" behaviour for the repeated testing of the single engine in Data-Set II was noted to be significantly different to that observed for the generic model of normality, constructed using Data-Set I (comparison not shown here). Thus, the generic model would not automatically detect all the engine events observed from the visualisation exercise described above.

A specific model of normality, trained using a sliding window containing the last 100 tests of the engine (using the training procedure described in Section 3.3) resulted in sufficient engine-specific characterisation of normality to allow both the events shown in Figure 3c and Figure 3f to be detected. This is shown as step-changes in the visualisations, shown in Figure 3. The sliding window allowed identification of the two engine periods of interest, while not alerting on the gradual shifts in engine condition (shown in Figure 3b and Figure 3c).

The difference in models of normality for the generic and specific cases is believed to be due to the differing nature of the engine manoeuvres performed in Data-Sets I and II, and further experimental data from engines undergoing "pass-off" testing (as occurred in Data-Set I) shows that this generic model remains suitable for previously-unseen engines performing the same set of "pass-off" test manoeuvres.

5. CONCLUSIONS

A novelty detection scheme has been introduced for a new class of military engine, and shown to be suitable for detecting abnormality in well-defined "pass-off" (PAT) engine tests, verified in conjunction with engineering vibration specialists in the manufacturing company. This system provides greater discrimination than conventional techniques based on threshold exceedance, and has been adopted for use within the QUICK toolset [5,9] for analysis of the modern military gas-turbine engine considered in this investigation.

Application to a repeated series of different manoeuvres showed that visualisation techniques allowed advance warning of abnormal engine conditions to be provided, though engine-specific models of normality are required in order to account for differences in engine manoeuvre between the generic and specific data-sets considered.

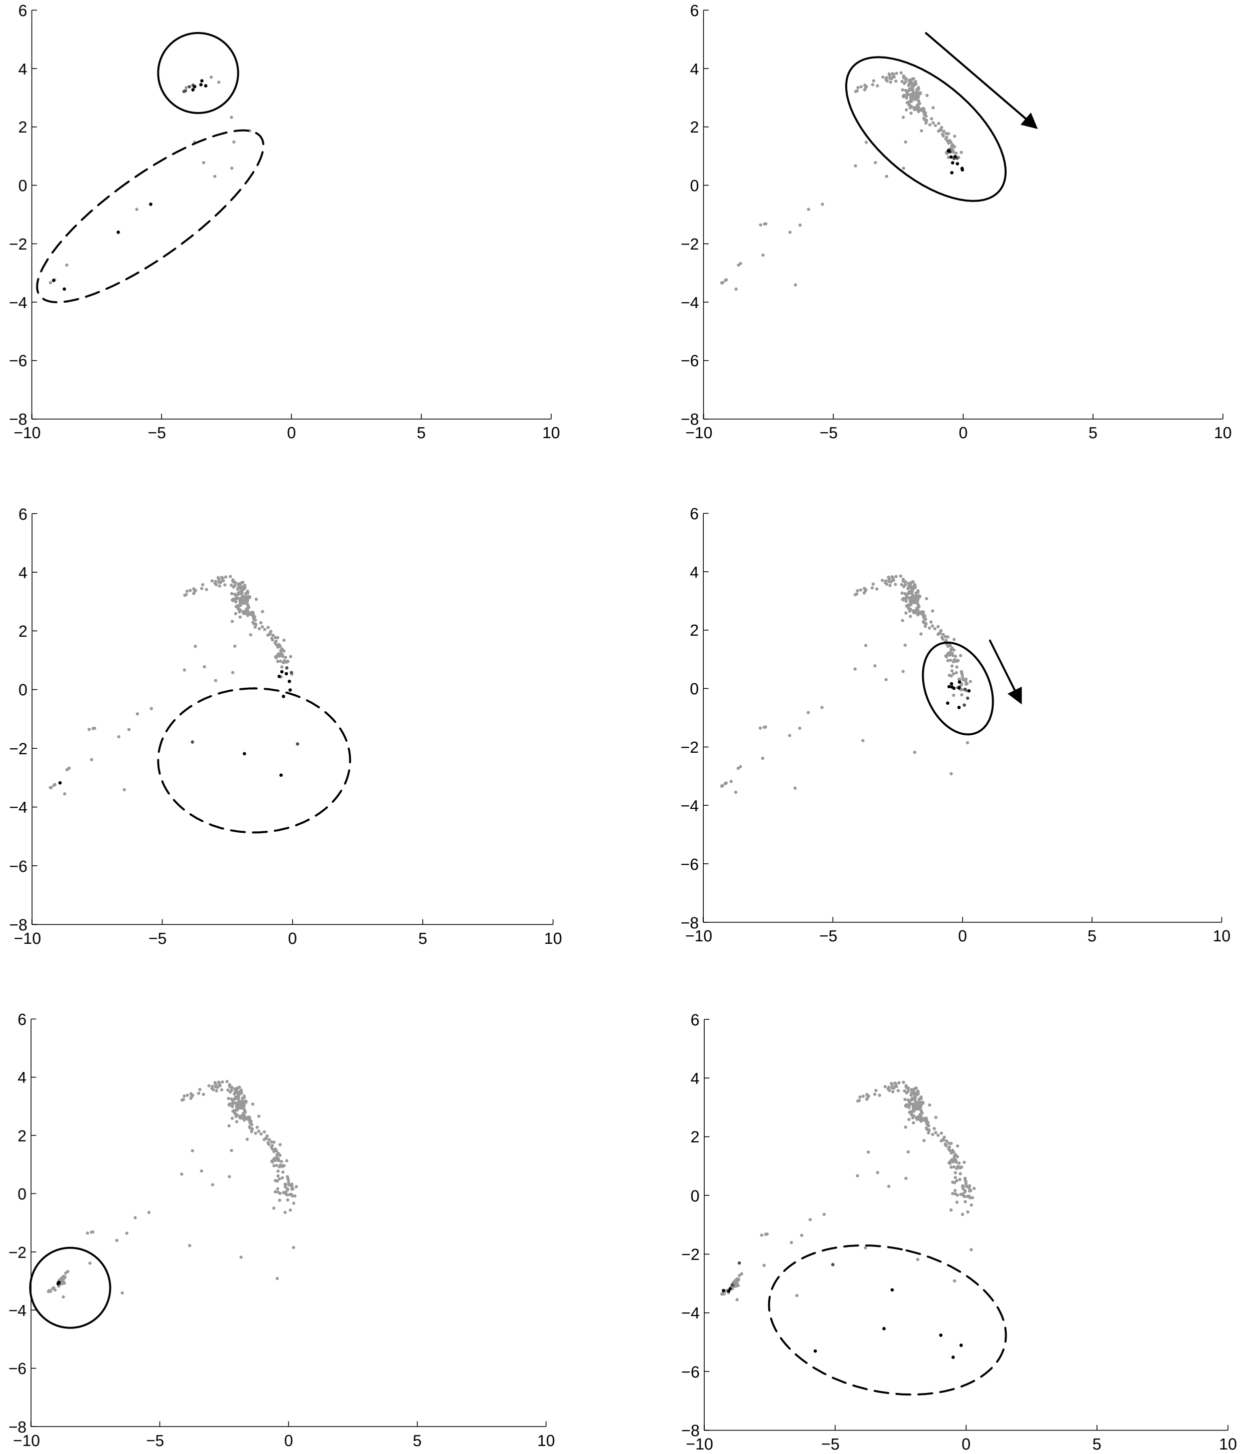

Figure 3 – (From top-left, in reading order) Visualisation of 1HP Signatures from Data-Set II:
a) Initial Configuring, b) Stable Operation, c) Component Event,
d) Return to Normality, e) Change in Acquisition Characteristics, f) Engine Event

The system described in this investigation is lightweight in terms of required computation and memory requirements, and is suitable for running in the real-time environment in which engine testing occurs.

Further work will concentrate on the use of this system during actual testing, as all analyses performed in this investigation were retrospective. More work in the use of this system to formulate actual engineering maintenance actions is also required.

Acknowledgements:

The authors gratefully acknowledge the support of Nicholas McGrogan and Simon Turner of Oxford BioSignals Ltd.; and Dave Morgan, Duncan Chase, Nigel Pitts, Paul Anuzis, Graham Hesketh, Dennis King, and Steve King of Rolls-Royce PLC.

The QUICK toolset provides implementations of intelligent vibration and performance analysis tools, further information of which can be obtained from Simon Turner, Industrial Business Manager, Oxford BioSignals Ltd., Brook House, 174, Milton Park, Abingdon, OX14 4SE.

References:

1. C.M. Bishop: *Pattern Recognition and Machine Learning*. Springer-Verlag, Berlin (2006)
2. C.M. Bishop: *Neural Networks for Pattern Recognition*. Oxford University Press, UK (1995)
3. D.A. Clifton, P.R. Bannister, and L. Tarassenko: *Learning Shape for Jet Engine Novelty Detection.* In: J. Wang et al. (eds.): Advances in Neural Networks III. Lecture Notes in Computer Science, Vol. 3973. Springer-Verlag, Berlin (2006) pp. 828-835
4. D.A. Clifton, P.R. Bannister, and L. Tarassenko: *Application of an Intuitive Novelty Metric for Jet Engine Condition Monitoring.* In: M. Ali (ed.): Advances in Applied Artificial Intelligence. Lecture Notes in Artificial Intelligence. Springer-Verlag, Berlin (2006) pp. 1149-1158
5. D.A. Clifton, P.R. Bannister, and L. Tarassenko: *Visualisation of Jet Engine Vibration Data for Novelty Detection*. Proc. of Natural Computing Applications Forum, London, UK (2007)
6. D.A. Clifton: *Condition Monitoring of Gas-Turbine Engines.* Transfer Report, Department of Engineering Science, University of Oxford (2005)
7. I.K. Jennions: *Cross-Platform Challenges in Engine Health Management*. Int. Conf. of Integrated Condition Monitoring, Anaheim, CA, USA (2006)
8. I. T. Nabney: *NETLAB: Algorithms for Pattern Recognition*. Springer-Verlag, Berlin (2002)
9. Oxford BioSignals Ltd.: *QUICK Technology*. White paper, Oxford BioSignals Ltd., UK (2006)
10. Rolls-Royce PLC.: *The Jet Engine*, 2nd Edition. Rolls-Royce PLC, UK (2005)
11. J.W. Sammon: *A Non-linear Mapping for Data Structure Analysis*. IEEE Transactions on Computers (1969)
12. L. Tarassenko: A Guide to Neural Computing Applications. Arnold, London (1998)