

Novelty Detection in Jet Engine Vibration Data

David Clifton, Peter Bannister, Lionel Tarassenko
Signal Processing & Neural Networks Research Group


1. Motivation

The identification of abnormal operating behaviour in gas-turbine jet engines is of critical importance in the avoidance of hazard. Recently, engine manufacturers are adopting a *condition monitoring* approach to maintenance, in which embedded intelligent data analysis systems process data from engine-mounted sensors to assess the “health” of engine components. These *health-monitoring systems* can provide early warning of engine failure by detecting abnormal operating behaviour.

This research project introduces a framework for the analysis of engine vibration data, providing on-line (in-flight) abnormality detection, and off-line (ground-based) engine monitoring, using data from the development of the Rolls-Royce Trent 900 engine (used by the Airbus A380, pictured right) and the Rolls-Royce EJ200 (used by the Typhoon Eurofighter, far below-right).


3. In-Flight Novelty Detection

A vibration signature (see right) constructed from many hours of normal flight data may be used as a model of normality *for that engine*, with new tracked order values observed during flight compared to the model on a sample-by-sample basis. Within each speed bin of a vibration signature, *decision boundaries* are set to define the range of values describing normality for that speed bin.



Actual flight data indicates that the distribution of vibration amplitudes within a speed bin is bimodal. Modelling these distributions with a 2-component Gaussian Mixture Model (left), decision boundaries may be set such that values are deemed novel if the probability of their observation is below some threshold (e.g., 10^{-9}).

Using three months of normal flight data to train a model of normality, a critical engine event was detected in subsequently tested flight data, with early warning of over 2 hours (right).




Other methods of on-line distribution modelling and threshold setting are under investigation, with adaptive and static models of normality being compared.

2. Constructing Signatures

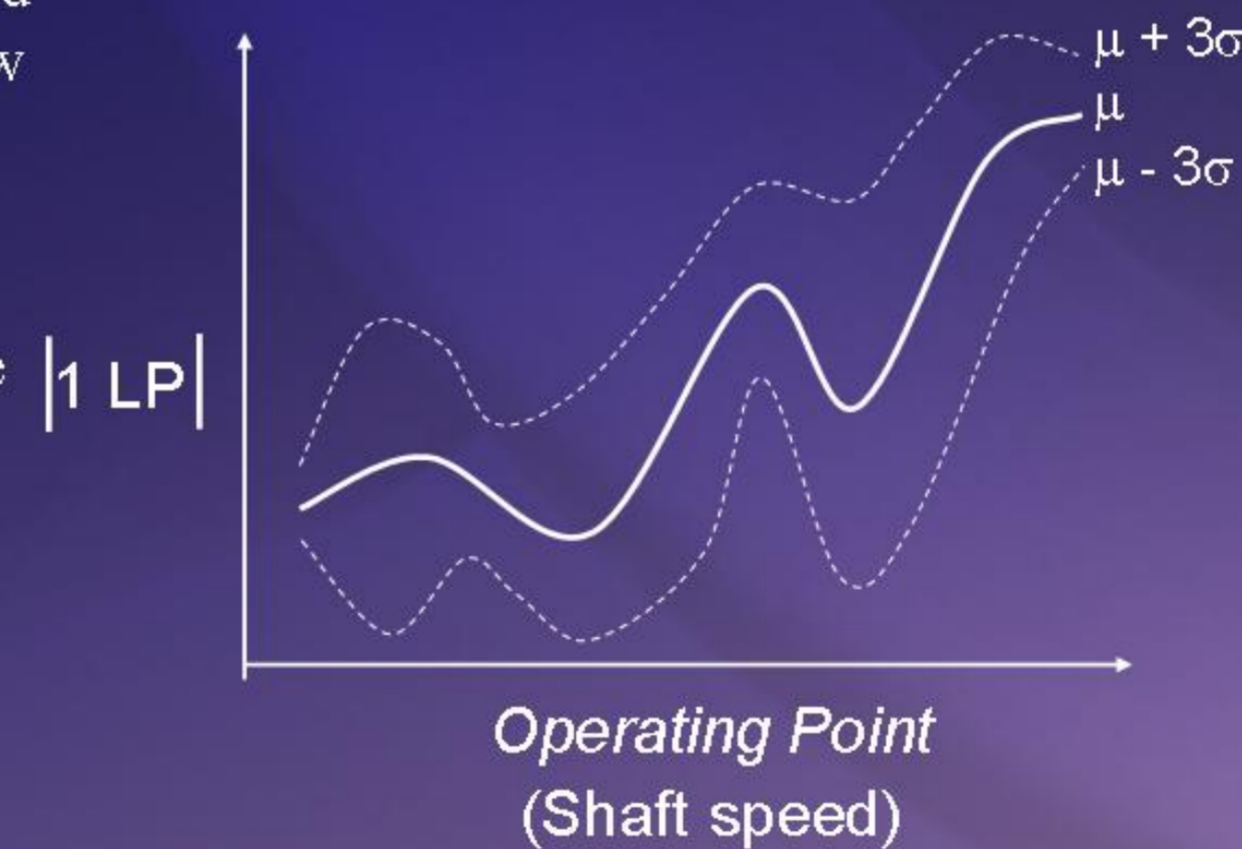


Acquire,
FFT



Case-mounted sensors measure engine vibration from which FFTs are computed. *Tracked orders* are extracted from spectral data (pictured right), defined to be vibration amplitude and phase within a narrow frequency band centred on the fundamental or harmonic of the rotational frequency of an engine shaft. Speed-based signatures of tracked order vibration and phase are constructed for each engine shaft, and used for event detection. Examples of abnormal engine behaviour are rare, and so a *novelty detection* approach is taken, in which departures from a model of normality (constructed from normal data) are identified.

Tracked Order Extraction



5. Conclusions, Future Research

Vibration signatures have been shown to provide early warning of engine failure, in both flight and ground-based analyses, providing protection against hazard during engine development programmes, and in service.

Future research will focus on the development of models of normality that integrate vibration data with performance data (engine temperatures, pressures, &c.). Time- and time/frequency-domain approaches to on-line novelty detection will also be investigated, providing complementary models to the frequency-domain analyses performed to date.



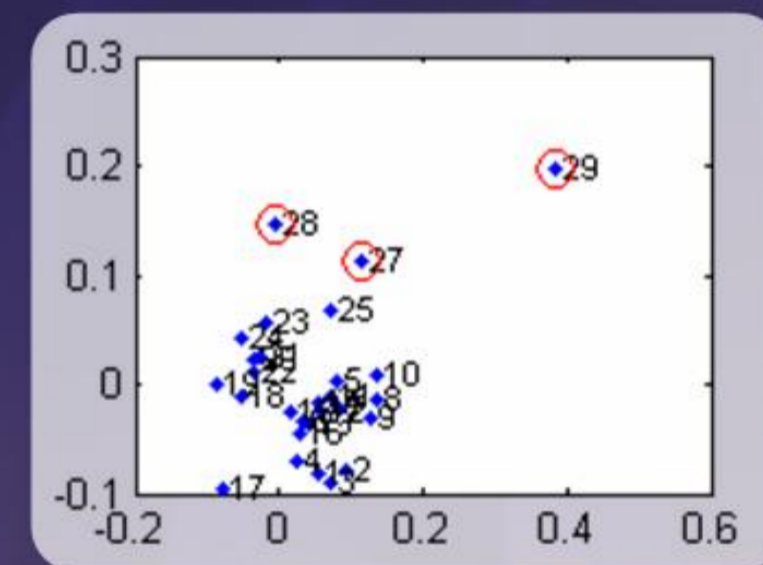
Acknowledgements

The author would like to thank Oxford BioSignals Ltd. and Rolls-Royce PLC for supporting this research, particularly the collaboration of Nicholas McGrogan, Ben Haskins, Dennis King, Paul Anuzis, Steve King, Mark Walters, and David Ault.


4. Ground-Based Novelty Detection

Intelligent algorithms for ground-based analysis systems are required to track engine performance during service life, across a fleet of aircraft. However, transfer of data from an engine to a ground-based analysis system is constrained by the extremely limited bandwidth available in conventional aircraft transmission systems - approximately 4 KB is available for an entire flight.

Vibration signatures are used to summarise engine performance over an entire flight, from which flight-by-flight novelty detection can be performed. Models of normality are formed in high-dimensional space from signatures constructed using $N = 20$ speed bins. Classification algorithms such as Support Vector Machines are used to identify signatures from “abnormal” flights.



Flight signatures from an example series of 29 tests is shown projected in 2-D using NeuroScale (above), the last 3 of which were identified as “abnormal” with respect to the majority. These corresponded to an engine fire only identified by conventional methods in test 29 (resulting in forced engine shutdown).



Vibration signatures from flights of a *fleet* of engines may also be compared by projection into 2-D (left). The example shows flights from an engine with abnormally high vibration (green) separated from flights of other engines. Flights from a pair of engines that suffered critical event (blue, red) are also visibly separate from the majority, at the bottom of the plot.

Features extracted from flight signatures may be plotted over time to provide a time-line of engine condition (below). The example shows feature values for an engine that suffered a critical event after six months of operation (flights shown in blue, average in black). Conventional methods observed the fault in the last test, whereas the example shows significant deviation from the previous norm three flights earlier.

