


Condition Monitoring of Gas-Turbine Engines



David Clifton
St. Cross College

Supervised by
Professor Lionel Tarassenko
Submitted: January, 2006

This Transfer Report is submitted to the
Department of Engineering Science, University of Oxford.

Condition Monitoring of Gas-Turbine Engines

David Clifton

St. Cross College, December, 2005

Summary

Condition monitoring assesses the operational health of gas-turbine engines, in order to provide early warning of potential failure such that preventative maintenance action may be taken. Gas-turbine engine manufacturers are increasingly offering a “service-based” approach to marketing their products, in which their customers are guaranteed certain availability of the engine after purchase. To achieve this, manufacturers take on the responsibilities of engine condition monitoring, by embedding health monitoring systems within each engine unit and prompting maintenance actions when necessary.

This report describes preliminary research into condition monitoring approaches for modern gas-turbine aircraft engines, and outlines plans for novel research to contribute to machine learning techniques in the condition monitoring of such systems, leading to the D.Phil. degree.

A framework for condition monitoring of aircraft engines is introduced, using signatures of engine vibration across a range of engine speeds to assess engine health. Inter- and intra-engine monitoring approaches are presented, in which a model of engine normality is constructed using vibration data from other engines of its class, or from the test engine itself, respectively.

Results of inter-engine analysis of final engine vibration tests prior to their release into service are presented, showing that the approach described within this report provides a more reliable estimate of engine condition than manufacturers’ conventional engine vibration tests, leading to better discrimination between “good” and “bad” engines.

Intra-engine analysis of an engine undergoing cyclic endurance testing, in which a set operational manoeuvre is performed repeatedly, shows that the method described in this report provides early warning of engine failure that eventually resulted in a hazardous engine fire, undetected by engine developers until it had to be shut down manually.

Future research is planned in application of this condition monitoring framework to an engine currently under development, improving upon existing methods and investigating new approaches, ultimately leading to the formulation of a general “black box” monitoring approach that can learn a model of system normality without prior knowledge of that system.

Contents

1	The Need for Condition Monitoring	1
1.1	Introduction	1
1.2	Aerospace Gas-turbine Engines	2
1.3	Existing Techniques for Condition Monitoring of Aerospace Gas-turbine Engines	3
1.4	Overview of this Report	6
2	Novelty Detection for Gas-turbine Engines	7
2.1	Introduction	7
2.2	Modelling Normality	7
2.3	Novelty Detection using TORCH	10
3	TORCH for Inter-engine Analysis	15
3.1	Introduction	15
3.2	Data Set of Vibration Signatures	16
3.3	Vectorisation	18
3.4	Normalisation	19
3.5	Visualisation	25
3.6	Clustering	29
3.7	Choosing a Model of Normality	30
3.8	Calculation of Novelty	31
3.9	Conclusion	33
4	TORCH for Intra-Engine Analysis	40
4.1	Introduction	40
4.2	Constructing TORCH Signatures of Vibration Phase	40
4.3	Intra-engine Analysis	42
5	Future Work	50
5.1	Flight Summary Compression	50
5.2	Improved Novelty Detection Methods	50
5.3	Time- and Time-Frequency Domain Analysis	51
5.4	Data Fusion, and Oil Debris Monitoring Systems (ODMS)	51
5.5	On-Line Learning	52
5.6	DPhil Project Plan	53
	Appendix A: Courses and Conferences Attended	54
	Bibliography	55

Chapter 1

The Need for Condition Monitoring

1.1 Introduction

Gas-turbine engines are critical to the operation of most industrial plants, aircraft, and heavy vehicles such as military armour and transport ships, and their associated maintenance costs can be high. Traditionally, operators of gas-turbine engines have attempted to reduce these costs by performing preventative maintenance actions at fixed intervals, in an attempt to avoid potential engine failure. More recently, gas-turbine engine manufacturers have been adopting a *condition monitoring* approach instead, in which intelligent data analysis systems are employed to assess the “health” of engine components. The engine’s maintenance needs are determined according to its operating condition, rather than maintenance being performed at fixed periods of time.

These *engine health monitoring* systems typically process data from engine-mounted sensors. Early warning of potentially hazardous engine conditions may result, the hope being that precursors of component failure may often be identified in advance of actual failure. This is a *prognostic* approach to condition monitoring, and is useful for types of faults that may be prevented if identified soon enough. Faults for which there are no such precursors (e.g., a “bird strike”, in which a bird impacts an engine fan) require a *diagnostic* approach. Such systems automatically identify engine faults that have occurred, and may recommend restorative maintenance actions appropriate to the type of fault.

Gas-turbine engine manufacturers are increasingly offering a “service-based” approach to marketing their products, in which their customers are guaranteed certain availability of the engine after purchase. To achieve this, manufacturers take on the responsibilities of engine condition monitoring, by embedding health monitoring systems within each engine unit and prompting maintenance actions when necessary.

Condition monitoring techniques are also applied to engines during the development process, and throughout product testing. During development, engines are typically re-built and tested many times. Assessment of component health during this process can avoid component damage and potential hazards.

1.2 Aerospace Gas-turbine Engines

This report is concerned with the condition monitoring of gas-turbine engines within the aerospace industry. These engines typically consist of a compressor section at the front of the engine that draws in air, pressurises it by up to forty times atmospheric pressure, and then delivers it to a combustion chamber at the rear. Fuel and compressed air are mixed and ignited within the combustion chamber, before being expelled from the rear of the engine, where the rapidly-expanding heated gas provides motive force. Rear-mounted turbines extract energy from these exhaust gases to drive the compressor section, to which they are connected via a shaft.

1.2.1 Stages of Compression

Modern aerospace gas-turbine engines divide the task of air compression from atmospheric pressure to that ultimately required within the combustion chamber into several stages. Many gas-turbine engines within the civil aerospace market involve three consecutive compression stages: the low pressure (*LP*), intermediate pressure (*IP*), and high pressure (*HP*) stages [58]. Air passes through each stage as it travels from the front of the engine to the rear, being further compressed by each, until it reaches the combustion chamber.

Each of the compressor stages is driven by its own turbine assembly, resulting in three corresponding turbine units situated within the exhaust stream at the rear of the engine. Each compressor is linked to its corresponding turbine by a separate shaft, which are mounted concentrically. In three-compressor engines, these are named the LP shaft, the IP shaft, and the HP shaft. The *operational point* of the engine is often defined in terms of the rotational speed of these shafts.

1.2.2 Engine Vibration Measurement

Transducers are mounted on various points of the engine assembly for the measurement of engine vibration. Vibration data used for investigations described in this report were acquired using a system [26] that computes Fast Fourier Transforms (*FFTs*) representative of engine vibration every 0.2 seconds, for each sensor output. Engine vibration is assumed to be pseudo-stationary over this measurement period such that the generated FFTs are assumed to be close approximations of actual engine vibration power spectra.

A *tracked order* is defined [24] to be the amplitude of engine vibration measured within a narrow frequency band centred on the fundamental or a harmonic of the rotational frequency of a shaft.

During normal engine operation, most vibration energy is present within tracked orders centred on the fundamental frequency of each rotating shaft; we define these to be *fundamental tracked orders*. Using the terms *LP*, *IP*, and *HP* to refer to engine shafts, we define fundamental tracked orders associated with those shafts to be *1LP*, *1IP*, and *1HP*, respectively.

Significant vibration energy may also be observed at harmonics of the rotational frequency of each shaft. These *harmonic tracked orders* may be expected to contain less vibration energy than corresponding fundamental tracked orders during normal engine operation. In the example of an LP shaft rotating with frequency 400 Hz, harmonic tracked orders may be observed at frequencies $400n$ Hz, for $n = 0.5, 2, 3, 4, \dots$. We define these harmonic tracked orders of the LP shaft to be *0.5LP, 2LP, 3LP, 4LP, \dots*

The system used to acquire data for the investigations described within this report automatically identifies peaks in vibration spectra corresponding to fundamental and harmonic tracked orders, using measurements of the rotational frequency of each shaft. From these peaks, a time series of vibration amplitude and phase for each tracked order is generated.

1.3 Existing Techniques for Condition Monitoring of Aerospace Gas-turbine Engines

1.3.1 Standard Methods

Guidelines for the production of condition monitoring systems applied to aircraft gas-turbine engines have been available for over twenty years [59–61]. Engine monitoring is performed using either “on-line” systems, mounted within the aircraft, that perform analysis of engine data during flight, or “off-line” ground-based systems, to which engine data is downloaded from the aircraft at the end of a flight.

On-line condition monitoring systems are subject to strict certification requirements from international aviation authorities. The cost of producing systems to these required standards is very high, primarily due to the technical challenge of suppressing the large numbers of false alarms generated by many condition monitoring algorithms [75]. Methods currently used in on-line condition monitoring are relatively simple due to these constraints.

Off-line condition monitoring systems are limited by the amount of engine data that may be transferred from an aircraft. Typically, data storage and transfer costs are high, which constrains the quality and quantity of data available to ground-based monitoring systems in commercial use.

Standard methods of engine condition monitoring involve comparison of recorded engine parameters with fixed operational limits set by engine manufacturers [61], or simple univariate analyses in which trends in engine data are monitored in order to detect sudden shifts in operational behaviour, often indicators of potential failure [27].

Current research on the condition monitoring of aircraft engines is focussed on providing more sophisticated methods of reliably determining engine condition and resulting maintenance recommendations, within the constraints of on-line and off-line monitoring systems. The following sections describe the general areas of current

research.

1.3.2 Fault Diagnosis using Expert Systems

Expert systems aim to capture the knowledge of domain experts, typically from engine manufacturers, in order to diagnose faults from engine data and recommend appropriate maintenance actions. A typical software package commercially available for this purpose is the TEXMAS system [11].

The accuracy of diagnosis may be improved through the construction of ever-larger databases of fault information for gas-turbine engines [7, 74]. *CASE-based Reasoning (CBR)* uses the intuition that solutions to new problems are likely to be similar to previously encountered problems of a similar nature. In this approach, a user-friendly software environment is utilised to capture expert knowledge and manage fault diagnoses. This can then be used to provide coordinated construction of expert-system databases for monitoring engines in civil aircraft between many different sites [48], using a framework for distributed Internet-based data analysis [19, 30]. CBR has also been applied to the management of the maintenance process for gas-turbine engines in military aircraft [1, 23].

However, the CBR approach is limited by the variation in operational behaviour between classes of aircraft engine, and even variation between engines of the same class, which can render the use of general fault databases ineffective when diagnosing engine-specific faults.

To avoid costs associated with reconstructing expert system databases for use with individual engines, the automated generation of the logical rules in the expert-system databases using “soft computing” techniques has been proposed.

“Fuzzy” expert systems consisting of fuzzy logic rules for characterising engine-specific faults have been combined with data-mining techniques [18] in which rules are formulated from trends in the value of measured engine parameters [81]. Performance parameter measurements are divided into segments modelled using linear regression, and described using clauses of fuzzy logic. Rules for diagnosing engine faults are formed using these fuzzy descriptions of each segment.

Within a conventional expert system, neural network methods have been used to match features detected in performance parameter signals with diagnosis rules [12, 14]. These methods have also been applied to fault diagnosis of space shuttle engines [54], and to military armour [22]. Here, diagnoses are created by combining sets of rules (comparing sensor output to fixed operational thresholds) and the outputs of a neural network classifying engine operation into one of several states (idle, full power, &c.).

The combination of expert system databases from multiple engines of the same class has been attempted, aiming to generate systems that might generalise to other engines of that class [50].

1.3.3 Thermodynamical and Frequency-based Modelling

Thermodynamical models of gas-turbine engines have been used for condition monitoring by comparison of a model derived from actual engine data (usually temperatures and pressures) with a model constructed from the theoretical relationship between those engine data [66]. The residual error between the two models may be used to identify possible engine failures [41].

Such methods combine a theoretical approach to thermodynamical model construction with knowledge derived from engine testing [53]. Models derived using finite-element analysis have been used in the condition monitoring of helicopter engines [21]. Linear flow models have been constructed from fuel-flow data and applied to condition monitoring of the engine combustion system [35].

Frequency-domain approaches to modelling engine parameters have been used in fuel-flow monitoring [17], exhaust gas monitoring [67], and more general diagnosis of engine faults [28, 80]. These models have been combined with auto-regressive, moving-average (*ARMA*) models for fuel-flow monitoring [9], and for analysis of engines used in industrial plants [6].

1.3.4 Statistical Pattern Recognition and “Soft Computing” Approaches

Determining engine condition by comparison of engine data with a model of normal behaviour may also be achieved using “soft computing” methods of model creation, such as neural networks, or using statistical pattern recognition techniques.

These methods offer the advantage that engine condition may be assessed without detailed theoretical system models. These *data-driven* approaches can be used to construct models of engine behaviour directly from engine data [16], in contrast to rule-based approaches that encode expert knowledge, and may be adapted for use with individual engine units, providing engine-specific condition monitoring. However, these methods can be sensitive to the quality and quantity of example data from which models of engine operation are constructed.

Neural networks have been used to identify abnormalities in the operation of gas-turbine engines within industrial plants [69, 70], in space shuttle engines [42], and in engines used by heavy ships [29].

The *support vector machine* (*SVM*) method, a statistical pattern recognition approach [77], has been used to identify examples of “abnormal” engine vibration behaviour from a “normal” training set in aircraft engine production tests [24, 65], and in water-pump gas-turbine engines [73]. These *one-class classifiers* learn a model of normality from a training set of normal examples, and identify patterns that fall outside that class as “abnormal”, in comparison with multi-class classifiers, which require approximately equal numbers of examples for each class.

Other methods of constructing models of normal engine behaviour within engine production tests using

data-driven techniques have been successfully demonstrated using data from military aircraft engines [45, 46].

Automated extraction of patterns describing normal and abnormal gas turbine engine operating conditions with engine vibration data has been investigated [2], set within a general fuzzy-neural analysis framework [3]. .

1.4 Overview of this Report

Analysis of engine tracked order vibration is believed by engine developers to provide useful indication of engine operational health. Many modes of engine failure affect observed tracked order vibration, and thus the assessment of engine condition through analysis of tracked order vibration can provide an indication of these failures.

Chapter 2 presents a framework for the condition monitoring of aerospace gas-turbine engines using tracked order vibration data. Several alternative, and complementary, general approaches to condition monitoring are described.

Chapter 3 presents an investigation of a condition monitoring method in which the operational health of an engine is determined with respect to the “normal” operation of a whole class of engines, within the framework presented in the previous chapter. The method is extended to the analysis of the Rolls-Royce *Trent 500* engine, showing that engines known to be operating abnormally are correctly identified.

Again within the framework described in Chapter 2, an investigation of an engine-specific method of condition monitoring is presented in Chapter 4, using engines of the Rolls-Royce *Trent 900* class. Early warning of hazardous engine conditions is shown to be possible using this method.

Conclusions are drawn in Chapter 5, with extensions to the described techniques considered. A plan for future work is discussed, with the ultimate goal of investigating condition monitoring methods for automatically learning a description of “normal” behaviour of a gas-turbine system.

Chapter 2

Novelty Detection for Gas-turbine Engines

2.1 Introduction

Examples of abnormal aircraft engine behaviour are few in comparison to the quantity of examples of normal behaviour, due to the rarity of engine failure. *Novelty detection*, in which departures from a model of normality are identified, is particularly suited to the condition monitoring of aircraft gas-turbine engines. A model of normality is constructed from normal engine data, departures from which are classified as “abnormal” events. This single-class approach to monitoring is appropriate due to the many possible modes of engine failure, of which the effects on engine state are not precisely defined, and which may vary from engine to engine.

This chapter introduces two complementary approaches to modelling normal engine behaviour. A framework is presented in which the resulting normal models may be used to identify the occurrence of abnormal engine behaviour.

2.2 Modelling Normality

2.2.1 Approaches to Constructing Models of Normality

Two approaches to constructing a model of engine normality are introduced:

- 1. An intra-engine approach:** a model of normality is constructed using data from a single engine. This engine-specific model may then be used to identify “abnormal behaviour”, defined as departures from normality, in subsequent data from that same engine. Thus, assessment of engine abnormality is performed by comparing it to its own previously normal behaviour; i.e. “the engine is its own control.”
- 2. An inter-engine approach:** a model of normality is created using data from more than one engine. Typically, the engines used to construct the model of normality are of the same class (e.g. the *Trent 900* class), and the model is assumed to represent the normal behaviour of that engine class. Thus, in this approach,

the abnormality of engine operating characteristics is assessed by comparison with a class-specific model of normality.

Advantages of Intra-engine Novelty Detection

The intra-engine approach requires sufficient normal data to construct a model of normality from a single engine. It is typically used when an engine is expected to be operated for long durations, in order to produce the required quantity of data.

An engine-specific model of normality may provide better characterisation of an engine's normal behaviour than a generic model, as illustrated in Figure 2.1. Behaviour which is abnormal for a single engine may be considered normal for the entire class of engines, due to the variability in operating characteristics between engines of the same class.




Figure 2.1: Normal behaviour for a single engine (region B) may cover smaller regions in data space than normal behaviour for the entire class to which that engine belongs (region A). Region C represents an area of the data space which is *abnormal* with respect to the single-engine normal region B, yet which is *normal* with respect to the class-wide normal region A.

The figure shows an example data space for two variables x_1 and x_2 , in which region A represents the area of data space corresponding to normal behaviour for the entire class of engines. A single example engine from that class has normal behaviour corresponding to region B in the data space. Region B is a subset of region A in this example, due to the normal behaviour of the single engine being less variant than (and a subset of) that of the entire engine class.

The advantages of the intra-engine approach are illustrated in the figure by considering the classification of region C, which corresponds to an area of the data space that would be considered abnormal for the single example engine (with respect to that engine's region of normality, region B). An intra-engine approach to novelty detection, in which data from region B may be used to generate a model of normality, would result in data from region C being correctly classified as abnormal for that engine.

If an inter-engine approach were adopted, a model of normality may be constructed from the whole class

of engines, corresponding to region A. Using that model for novelty detection would incorrectly classify those same data from region C as normal.

Advantages of Inter-engine Novelty Detection

The inter-engine approach to construction of a model of normality is appropriate in situations when there are insufficient data available to construct an engine-specific model of normality.

One such example situation in which an inter-engine approach is appropriate for determination of engine abnormality is final product testing prior to release of Rolls-Royce aircraft engines into service. A pre-defined test manoeuvre is performed by the engine, over a period of four minutes. The quantity of data recorded during this test may be insufficient to construct an engine-specific model of normality, due to the short duration of the test. Instead, a generic model may be generated using data recorded from many example engines previously performing the same test manoeuvre.

Engines sold to customers must operate within generic (class-specific) bounds of normality (corresponding to region A in the figure). Thus, even though each engine has slightly different characteristics, the abnormality of tests from engines of the same class performing the test manoeuvre may be determined with respect to the generic model of normality.

2.2.2 Supervised and Unsupervised Modelling of Normality

In both the intra- and inter-engine approaches to novelty detection, a model of normality is constructed using data which are deemed to represent normal engine operation in one of two ways:

1. Supervised model construction: normal data are labelled as such by an expert in the domain (typically an engine manufacturer or analyst, when the method is applied to gas-turbine engines).

In an intra-engine approach to condition monitoring, this could correspond to labelling normal periods of operation within data recorded from a single engine.

In an inter-engine approach, the engines from which example data are taken may be partitioned into “normal” and “abnormal” sets of engines.

2. Unsupervised model construction: no labelling of normal data is available, and so the determination of data normality must be made by other means.

In an intra-engine approach to novelty detection, it may often be assumed that an initial period of engine operation is normal, from which a model of normality is generated, whilst subsequent periods of operation are compared with that model.

In an inter-engine approach, the relative similarity of examples of engine behaviour derived from different engines may be assessed; those which differ significantly from the majority may be deemed abnormal, and *not* used to generate the model of normality.

During construction of a model of normality, available *a priori* domain expertise may be exploited to best characterise normal engine behaviour. It may be anticipated that a supervised approach to the construction of a normal model would be more likely to provide a more accurate characterisation of normality than an unsupervised approach.

This report presents examples of both intra- and inter-engine approaches using supervised learning. Extension to fully-unsupervised learning, in which models of normality are automatically learned from a system with limited or no *a priori* knowledge of that system, is considered as future work within Chapter 5.

2.3 Novelty Detection using TORCH

Constructing a model of normal engine behaviour may be performed by characterising tracked order vibration corresponding to normal engine operation. This section presents a framework for engine vibration analysis based upon “Tracked ORder CHaracteristics” (a Rolls-Royce acronym referred to as *TORCH*). The notion of constructing a speed-based representation of tracked order vibration is introduced. The use of TORCH for inter- and intra-engine monitoring, on both a sample-by-sample or episodic basis, is discussed, with the limitations of each technique also considered.

2.3.1 TORCH Vibration Signatures

We define a *TORCH vibration signature* to be the vibration amplitude and phase of a tracked order measured over a range of speeds of the corresponding shaft [24]. For example, a TORCH 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% 100% maximum speed] of the corresponding shaft. This speed range is subdivided into B equal bins. Typically $B = 400$, and each bin corresponds to a 0.25% sub-range of the maximum shaft speed. Within each bin $b = 1 \dots B$, the mean and variance (μ_b, σ_b^2) of vibration measurements occurring within that bin’s sub-range of shaft speeds are computed.

Phase and amplitude of vibration are considered separately, resulting in two TORCH vibration signatures for each tracked order. Vibration phase measurements are known to exhibit high variability within some sub-ranges of shaft speeds [40]. Separation of vibration phase and amplitude signatures allows analysis of vibration amplitude independent of vibration phase, which may otherwise dominate the analysis due to this

high variability.

An example TORCH vibration amplitude signature of a 1LP tracked order is shown in Figure 2.2. The figure shows mean vibration amplitudes μ_b , and $\mu_b \pm 3\sigma_b$ (i.e. \pm three standard deviations) for an example *Trent 900* engine. The signature was constructed from 200 hours of flight data recorded over a period of one month.




Figure 2.2: Vibration amplitude signature of the 1LP tracked order, for an example Trent 900 engine. Mean μ_b , and $\mu_b \pm 3\sigma_b$ vibration amplitudes within bins $b = 1 \dots 400$ are shown in blue and red respectively.

The figure shows that the engine from which the vibration signature was generated was operated in the range [0% 85%] of maximum shaft speed, and that the mean vibration amplitude for the 1LP tracked order of this engine (shown in blue) tends to increase with increasing speed of the LP shaft.

TORCH signatures may be constructed for any selection of engine tracked orders. Typically, fundamental tracked orders (i.e. 1LP, 1IP, and 1HP in an engine containing three shafts) are used to generate signatures, being those tracked orders that may be expected to contain most vibration energy during normal engine operation. The selection of the set of tracked orders from which to generate TORCH signatures is not discussed further in this chapter.

2.3.2 Using TORCH Signatures to Construct a Model of Normality

A model characterising normal tracked order vibration may be constructed from TORCH vibration signatures. These signatures are derived from a single engine in the case of intra-engine analysis, or from multiple engines of the same class in the case of inter-engine analysis.

Inter-engine analysis requires *normalisation* of signatures, to allow the comparison of TORCH signatures derived from different engines, and is considered in Section 2.3.4. Further TORCH vibration signatures (termed *test* signatures) may be compared to the model of normality in order to determine their relative abnormality. In an intra-engine approach, test signatures are derived from the same engine used to construct the model of

normality. In an inter-engine approach, test signatures are derived from an engine of the same class as that used to construct the model of normality. If a TORCH signature differs significantly from the model of normality, it is classified as “abnormal”.

Figure 2.3 illustrates the process of model construction and novelty detection. A model of normality is shown in this example to be constructed from TORCH vibration signatures 1 to n . Test TORCH signatures are then compared to that model. The novelty detection process categorises each tested TORCH signature as being either “normal” or “abnormal” (as shown in the figure), and *novelty scores*, which are defined below, as calculated for each signature.

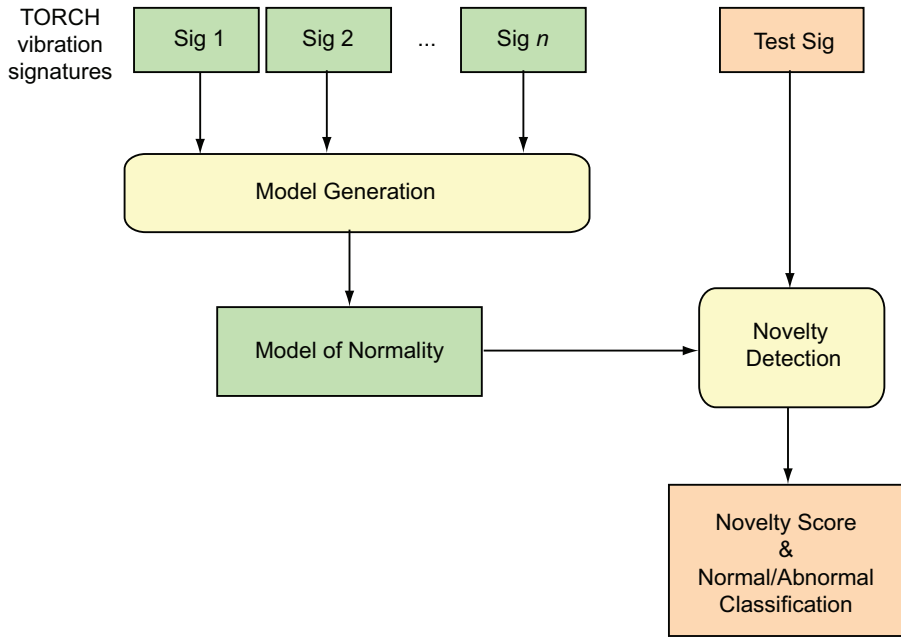


Figure 2.3: Model construction and novelty detection using TORCH signatures. Signatures 1 to n are used to generate a model of normality. Test signatures from the same engine are then compared to that model.

2.3.3 Calculating Novelty Scores from TORCH Signatures

The *novelty score* of a TORCH vibration signature with respect to a model of normality is defined as a measure of the “difference” or “distance” between that signature and the model. Definition of “distance” is deferred to the detailed descriptions of inter- and intra-engine methods in Chapters 3 and 4, respectively.

Thus, signatures which differ significantly from those used to construct the model of normality will be assigned high novelty scores, whilst signatures which are similar to those used to construct a normal model will be assigned low novelty scores. TORCH signatures with novelty scores higher than some threshold score H_η are classified “abnormal”.

Setting the Novelty Score Threshold

The value of H_η is selected to best separate TORCH signatures labelled “normal” from those labelled “abnormal”, where that labelling results from the use of either supervised or unsupervised means, as described in Section 2.2.2.

Each TORCH signature classified as “abnormal” is either a *true positive* (TP) or a *false positive* (FP). Similarly, “normal” classifications of TORCH signatures may be defined to be either *true negative* (TN) or *false negative* (FN).

The *sensitivity* S of a classifier of TORCH signatures is defined to be:

$$S = \frac{TP}{TP + FN}; \quad (2.1)$$

i.e. the proportion of all truly abnormal TORCH signatures that were correctly classified. The *positive predictive value* PPV of the classifier is defined to be:

$$PPV = \frac{TP}{TP + FP}; \quad (2.2)$$

i.e., the proportion of correct classifications of TORCH signatures as “abnormal”, with respect to all such classifications.

Selecting an optimal value for H_η corresponds to optimising S and PPV according to the costs associated with false negative and false positive classifications.

Applications in which a false negative classification carries a higher cost than a false positive classification require maximising S , while tolerating sub-optimal values of PPV . An example application of this type is a system used to identify cancerous growths in tissue samples, in which failure to detect a cancer (a false negative classification) is more unacceptable than a false positive classification.

Conversely, applications in which a false positive classification carries a higher cost than a false negative classification require maximising PPV while accepting sub-optimal values of S . Systems for the identification of abnormal operating conditions in gas-turbine engines are of this type: the cost associated with a false positive classification is high, often involving significant examination of equipment and maintenance processes.

Therefore the value of H_η in systems for the classification of TORCH signatures is selected such that PPV is maximised.

Generalising to Test Data

It is desirable that the performance of a classifier (defined in terms of S and PPV) is acceptable when presented with TORCH signatures derived from previously unseen data; i.e. the classifier must be able to *generalise* to test data.

Verifying the ability of the classifier to generalise requires a set of test data, distinct from the data used to construct the model of normality. Sufficient data must be provided both to construct the model of normality, and to test the generalisation performance. There is no accepted method of quantifying these data requirements, though heuristic methods of doing so are presented in the detailed descriptions of inter- and intra-engine TORCH signature analysis in Chapters 3 and 4, respectively.

2.3.4 Inter-engine Analysis using TORCH Signatures

Vibration signatures derived from different engines must be normalised before they can be compared. Results of applying various normalisation schemes to TORCH signatures for inter-engine comparison are presented within Chapter 3. The use of normalisation in the process of inter-engine model construction and novelty detection is illustrated in Figure 2.4.

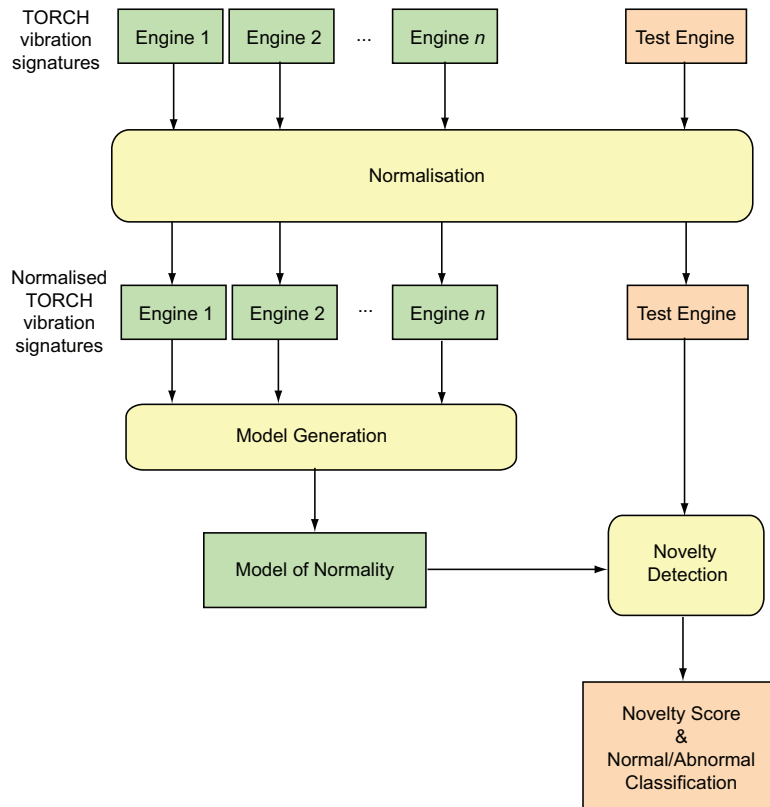


Figure 2.4: Inter-engine model construction and novelty detection including normalisation for the comparison of TORCH signatures from different engines.

A TORCH vibration signature consists of both mean and variance (μ_b, σ_b^2) of either vibration phase or amplitude, for $b = 1 \dots B$ shaft-speed bins, as defined previously. In normalising a TORCH signature, it must be ensured that appropriate normalisation is applied to μ_b and σ_b^2 . Chapter 3 presents a method for the inter-engine analysis of TORCH signatures which uses only the mean values μ_b . Inclusion of variance values σ_b^2 is considered as future work, within Chapter 5.

Chapter 3

TORCH for Inter-engine Analysis

3.1 Introduction

Section 2.3.4 introduced the notion of inter-engine analysis of engine condition using speed-based signatures of engine vibration. This approach may be applied to vibration data measured during engine testing prior to product release, termed *pass-off testing*, referring to testing in which an engine is “passed off” prior to entry into service. A method of performing inter-engine analysis using such data has been developed and results are presented in this chapter.

Previous work [46] has described a method of analysis using speed-based engine vibration signatures recorded from a two-shaft gas-turbine engine, the Rolls-Royce *Pegasus*. In this chapter, the method is extended to the pass-off testing of a civil three-shaft engine, the Rolls-Royce *Trent 500*.

3.1.1 The Need for Pass-Off Novelty Detection

Approximately 10% of Trent 500 engines exhibit unusually high vibration levels within the HP-shaft [4], which must be identified during pass-off testing before product release.

With present pass-off testing, Trent 500 engines are released for service if the maximum engine vibration level measured at any speed is below a pre-specified contractual limit. The current 10% rejection rate of Trent 500 engines occurs as a result of the maximum vibration level of the HP fundamental tracked order exceeding the contractual vibration limit.

Novelty detection, defined to be the identification of departures from normality, could be applied to the *shape* of vibration signatures from these engines. As described in Section 2.3.4, a model of normality can be constructed from TORCH vibration signatures from a large number of engines. Novelty detection can then be applied to new vibration signatures to determine how close they are to the model of normality.

This could identify engines which pass the simple contractual test but whose vibration signature is significantly different from the signatures found in a “normal” reference set of engines. This technique is termed *Shape Analysis*, and is the subject of this chapter.

3.2 Data Set of Vibration Signatures

3.2.1 Signatures for Acceleration and Deceleration

During pass-off testing, a Trent 500 engine performs a controlled two-minute acceleration from idle to maximum speed, followed by a two-minute deceleration back to idle. The vibration amplitude levels of the fundamental tracked orders are recorded during these tests, from which the speed-based vibration signature is constructed for each engine shaft.

An example vibration signature from the data set used in this investigation, plotted against shaft speed, is shown in Figure 3.1. Note that, as described in Section 2.3.1 the signature for each fundamental tracked order is constructed with respect to the percentage speed *for that shaft*; the amplitude of vibration of the HP shaft is plotted against the percentage of maximum speed for that shaft. Figure 3.1 illustrates the lack of vibration measurements below 65% HP shaft speed: the HP shaft is operating at sub-idle speeds below this point, and no vibration data are collected for those speeds.




Figure 3.1: Vibration signature for HP shaft acceleration, engine 71100.

3.2.2 Illustration: HP Shaft Acceleration

The data set used in the investigation described in this chapter consists of 71 engine examples recorded during pass-off tests, as shown in Table 3.1. Each example contains an acceleration and a deceleration signature of shaft vibration, for the LP, IP, HP, and *radial drive* shafts. The latter is a shaft connected to the HP shaft by a geared mechanism, used to provide initial drive to the HP shaft during engine start-up. The results of the analysis of the HP shaft acceleration vibration signatures are described in detail in this chapter, with results from the other shafts reviewed at the end of the chapter.

The four sub-sets of engine examples shown in Table 3.1 were each recorded from a different source of pass-off tests:

Sub-set	Dates	Number of Examples	Engine Numbers
A	24/02/2003 - 27/08/2003	45	71100 - 71160
B	09/04/2002 - 30/09/2003	11	71045 - 71170
C	19/03/2003 - 21/06/2004	5	71040 - 71051
D	16/05/2000 - 18/05/2000	10	70005

Table 3.1: Data sets used in inter-engine analysis of Trent 500 tracked order vibration.

Sub-set A: examples recorded from engines that were released into service following standard pass-off testing at Rolls-Royce Plc., Derby, UK;

Sub-set B: examples recorded from engines that were rejected after failing pass-off testing at Rolls-Royce Plc., Derby, UK;

Sub-set C: examples recorded from engines overhauled during their service-lives, undergoing further pass-off testing after an extended period of commercial use, at Hong Kong Aero Engine Services Ltd.;

Sub-set D: examples recorded from engines to which weights were deliberately attached to the fan during testing, forcing it into a state of unbalance, at Rolls-Royce Plc., Derby, UK.¹

Examples from sub-set D (unbalanced engines) were found to display approximately normal behaviour as far as HP shaft vibration is concerned; the unbalance only affects the LP shaft (with very large changes in vibration amplitude). Novelty detection *applied to the HP shaft* should therefore not identify examples from sub-set D as being novel.

The initial classification of examples by source cannot be used to determine which engines represent “normal” behaviour. Some engines that were released into service (sub-set A) exhibited vibration levels that exceeded contractual limits², whilst some of the overhauled engines (sub-set C) also failed their pass-off test after re-build.

Thus, the 71 examples were reclassified into four classes according to their maximum vibration amplitude a_{MAX} and the contractual vibration limit H_{MAX} , as shown in Table 3.2.

Class	Number of Examples	Description
1	29	Examples from sub-sets {A, B, C} for which $a_{MAX} < 0.9H_{MAX}$
2	15	Examples from sub-sets {A, B, C} for which $0.9H_{MAX} \leq a_{MAX} \leq H_{MAX}$
3	17	Examples from sub-sets {A, B, C} for which $a_{MAX} > H_{MAX}$
4	10	Examples from sub-set D

Table 3.2: Data sets for inter-engine analysis, after reclassification.

¹Weights 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. 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 Source D are recorded from an engine undergoing this procedure.

²During pass-off testing, engine vibration levels at specified shaft speeds are compared to contractual limits. Vibration data used in the investigation described in this report were recorded across the full range of shaft speeds at which each engine was tested, and thus the maximum vibration levels may exceed the contractual limit.

3.3 Vectorisation

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* [5]. A TORCH vibration signature, as defined in Section 2.3.1 consists of vibration values across a range of shaft speeds divided into 400 bins.

In order to avoid the curse of dimensionality, each signature is summarised by a *shape vector* \mathbf{x} . This is conducted by computing a weighted average of the vibration amplitude values $a(s)$ over N speed sub-ranges [46]. The n^{th} dimension of shape vector \mathbf{x}^n , for $n = 1 \dots N$, is defined to be:

$$\mathbf{x}^n = \int_{s_{\min}}^{s_{\max}} a(s)\omega_n(s)ds \quad (3.1)$$

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 $n : 1 \dots N$.

The value of N must be sufficiently large to avoid the unnecessary discarding of useful signal information, whilst being sufficiently small to avoid problems associated with the curse of dimensionality. This chapter compares quantisation using $N = 10$ and $N = 20$ (used previously in [4]).operated during a

A histogram of the speeds at which the vibration amplitudes were measured for the entire HP shaft data set is shown in Figure 3.2. It can be seen that a speed sub-range of [60% 100%] maximum speed is appropriate to cover the speeds contained within the example data set. The weighting functions used are overlapping trapezoids, as shown in Figure 3.3.

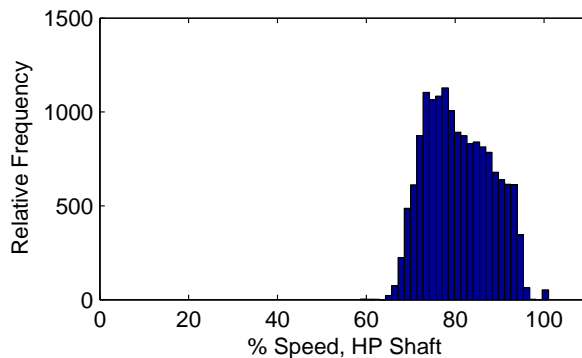


Figure 3.2: Histogram of HP shaft speeds expressed as a percentage of maximum speed across all tests.

Examples of shape vector construction from vibration signatures is shown in Figure 3.4. Each signature has its corresponding shape vector superimposed upon the original signature, for $N = 10$ and $N = 20$.

The effect of quantisation can be seen in the figure: peaks in vibration amplitude, at approximately 75% maximum shaft speed in (a), at 72% and 78% in (b), and at 82% in (c), are not closely represented for $N = 10$ quantisation. These peaks are represented more accurately with $N = 20$ quantisation, though their peak amplitude is not retained.




Figure 3.3: Example of weighting functions ($n : 1 \dots 10$), used to create shape vectors.

Though detail of signal variation is discarded in the quantisation process, this investigation aims to determine if novelty detection is successful when using the information retained in TORCH vibration signatures. The production of novelty detection methods that are effective when applied to summarised or quantised data is particularly important in aerospace gas-turbine analysis, in which the cost of communicating full engine data from an engine to a ground-based system can be high. Many ground-based “off-line” systems have only a summary of engine flight data available for analysis, due to restrictions on data transfer.

3.4 Normalisation

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

One aim of normalisation is to transform vibration signatures such that “normal” and “abnormal” signatures are more easily separable. This section describes several normalisation methods, shows results of their application to example vibration signatures (and to 10- and 20-dimensional shape vectors generated from those signatures), and concludes by identifying the method most suitable for separating “normal” and “abnormal” examples.

3.4.1 Normalisation Methods

Five normalisation methods were investigated for scaling vibration signatures \mathbf{x}_i (and the 10- and 20-D shape vectors constructed from them) prior to modelling. Results of applying each method to three example vi-




Figure 3.4: Example of shape vector generation: vibration signatures (shown in black) are used to construct corresponding 10-D and 20-D shape vectors (shown in red and blue, respectively). (a) Trent 500 engine 71100; (b) Trent 500 engine 71101; (c) Trent 500 engine 71104.

vibration signatures and their corresponding shape vectors are presented. Of the example signatures, two are “normal” ($\mathbf{x}_a, \mathbf{x}_b \in \text{class } \mathbf{1}$), and one is “abnormal” ($\mathbf{x}_c \in \text{class } \mathbf{3}$). These signatures are shown unnormalised in Figure 3.5(a), alongside their corresponding 20- and 10-dimensional shape vectors in Figure 3.5(b) and (c), respectively.

Signature \mathbf{x}_a shows vibration levels typical of most “normal” signatures in the data set, across all measured shaft speeds. A further example of normal HP shaft vibration is shown by \mathbf{x}_b , in which vibration amplitude is relatively low below 80% HP shaft speed, but which significantly increases in magnitude above 80% shaft speed. Signature \mathbf{x}_c shows abnormally high vibration amplitude below 80% shaft speed.

The term *pattern* is used throughout this section to refer to either a vibration signature or a shape vector to which normalisation is applied. A pattern is denoted \mathbf{x}_i^n , for the $i = 1 \dots I$ patterns in the data set, and $n = 1 \dots N$ elements of that i^{th} pattern. Thus, $I = 71$ for the data set described in Section 3.2. When referring to a typical TORCH vibration signature, $N = 400$ (as specified in Section 2.3.1); when referring to a




Figure 3.5: Example vibration signatures: two normal (\mathbf{x}_a , \mathbf{x}_b shown in blue and black), and one abnormal (\mathbf{x}_c shown in red). (a) Original vibration signatures; (b)-(c) 20-D and 10-D shape vector representations of the example signatures, respectively

10-dimensional shape vector, $N = 10$.

Unit Normalisation

Unit normalisation scales all n elements of a pattern \mathbf{x}_i such that

$$\max(\mathbf{x}_i^n) = 1. \quad (3.2)$$

This attempts to preserve the shapes of patterns by scaling such that all vibration amplitudes vary within the same range, $[0 \ 1]$.

Unit-normalisation applied to the three example vibration signatures and their corresponding shape vectors is illustrated in Figure 3.6.




Figure 3.6: Example of unit normalisation applied to three vibration signatures: two normal (\mathbf{x}_a , \mathbf{x}_b shown in blue and black), and one abnormal (\mathbf{x}_c shown in red). (a) Unit-normalised vibration signatures; (b)-(c) unit-normalised 20- and 10-D shape vectors, respectively.

This method causes the range of variation in vibration amplitude from “normal” engines (which may consist of generally low vibration amplitudes, such as \mathbf{x}_a) to be scaled up to the same range as for “abnormal” engines. Pattern \mathbf{x}_a has been scaled in all three plots such that it appears to have higher normalised vibration levels than the “abnormal” pattern \mathbf{x}_c below 80% maximum shaft speed. The abnormally high vibration levels below

80% maximum speed in \mathbf{x}_c no longer appear abnormal. This increasing overlap between the two classes opposes the aim of improving class separation, thus unit normalisation is unsuitable for this application.

Amplitude Normalisation

Amplitude normalisation scales all n elements of a vibration signature \mathbf{x}_i such that

$$\sum_{n=1}^N \mathbf{x}_i^n = 1. \quad (3.3)$$

This normalises the area under the curve of pattern \mathbf{x}_i . Application of this method to the three example vibration signatures and their corresponding shape vectors is shown in Figure 3.7.

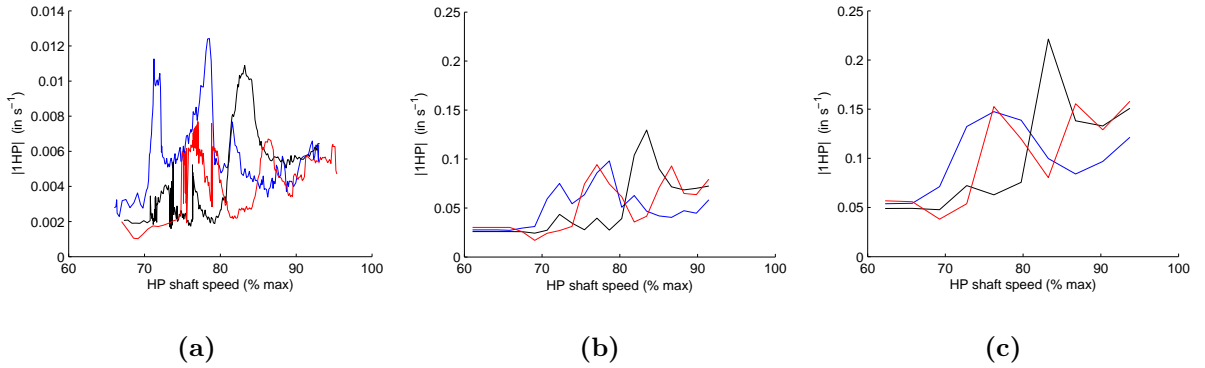


Figure 3.7: Example of amplitude normalisation applied to three vibration signatures: two normal (\mathbf{x}_a , \mathbf{x}_b shown in blue and black), and one abnormal (\mathbf{x}_c shown in red). (a) Amplitude-normalised vibration signatures; (b)-(c) amplitude-normalised 20- and 10-D shape vectors, respectively.

From the figure, it can be seen that patterns \mathbf{x}_a and \mathbf{x}_c have been scaled to overlap with \mathbf{x}_b . The abnormally high vibration amplitude in \mathbf{x}_c below 80% maximum shaft speed has been scaled to be generally lower than pattern \mathbf{x}_a . As with unit normalisation, the increase in overlap between the classes caused by amplitude normalisation makes it unsuitable for this application.

Energy-based Normalisation

Energy-based normalisation transforms the n elements of pattern \mathbf{x}_i such that

$$\sum_{n=1}^N (\mathbf{x}_i^n)^2 = 1. \quad (3.4)$$

This normalisation scheme aims to retain information associated with the distribution of energy throughout each pattern. Results of applying this method to the example vibration signatures are shown in Figure 3.8.

The relationship between “normal” and “abnormal” classes after energy-based normalisation is approximately similar to that seen after application of amplitude normalisation. Again, pattern \mathbf{x}_a has been scaled to overlap significantly with pattern \mathbf{x}_c . This increase in overlap between the two classes, when applied to vibration signatures and to both 10- and 20-dimensional shape vectors, makes this method of normalisation unsuitable for separation of “normal” from “abnormal” patterns.




Figure 3.8: Example of energy-based normalisation applied to three vibration signatures: two normal (\mathbf{x}_a , \mathbf{x}_b shown in blue and black), and one abnormal (\mathbf{x}_c shown in red). (a) Energy-normalised vibration signatures; (b)-(c) energy-normalised 20- and 10-D shape vectors, respectively.

Zero-mean, Unit-variance Normalisation

We define the zero-mean, unit-variance normalisation function $N(\mathbf{x}_i)$ to be

$$N(\mathbf{x}_i) = (\mathbf{x}_i - \mu_i) / \sigma_i \quad (3.5)$$

where (μ_i, σ_i) are calculated from the n elements of \mathbf{x}_i^n :

$$\mu_i = \frac{1}{N} \sum_{n=1}^N \mathbf{x}_i^n \quad (3.6)$$

$$\sigma_i = \left(\frac{1}{N-1} \sum_{n=1}^N (\mathbf{x}_i^n - \mu_i)^2 \right)^{\frac{1}{2}} \quad (3.7)$$

Figure 3.9 shows the result of normalising the three example patterns using this method.




Figure 3.9: Example of zero-mean, unit-variance (ZMUV) normalisation applied to three vibration signatures: two normal (\mathbf{x}_a , \mathbf{x}_b shown in blue and black), and one abnormal (\mathbf{x}_c shown in red). (a) ZMUV-normalised vibration signatures; (b)-(c) ZMUV-normalised 20- and 10-D shape vectors, respectively.

Because this method preserves variation in vibration amplitudes with respect to the mean vibration amplitude of a pattern, the degree of overlap between patterns may increase. A small peak in vibration amplitude within a pattern that has a low overall variance in vibration amplitude will be scaled such that it appears similar to extremely large peaks in amplitude occurring within patterns with much higher overall amplitude variance.

Discrimination between “normal” and “abnormal” patterns becomes poor, as the degree of overlap between the two classes increases. This is shown within the figure, in which pattern \mathbf{x}_a is scaled such that it contains generally higher vibration levels than pattern \mathbf{x}_c , in all three plots.

Component-wise Normalisation

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

$$N(\mathbf{x}_i) = \{n : [1 \dots N] \bullet (\mathbf{x}_i^n - \boldsymbol{\mu}^n) / \boldsymbol{\sigma}^n\} \quad (3.8)$$

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

$$\boldsymbol{\mu}^n = \frac{1}{I} \sum_{i=1}^I \mathbf{x}_i^n \quad (3.9)$$

$$\boldsymbol{\sigma}^n = \left(\frac{1}{I-1} \sum_{i=1}^I (\mathbf{x}_i^n - \boldsymbol{\mu}^n)^2 \right)^{\frac{1}{2}} \quad (3.10)$$

The results of applying this method of normalisation to three example shape vectors is shown in Figure 3.10.




Figure 3.10: Example of component-wise normalisation, applied to three vibration signatures: two normal (\mathbf{x}_a , \mathbf{x}_b shown in blue and black), and one abnormal (\mathbf{x}_c shown in red). (a) Component-normalised vibration signatures; (b)-(c) component-normalised 20- and 10-D shape vectors, respectively.

Figure 3.10 shows that “abnormal” pattern \mathbf{x}_c has been scaled such that its high vibration levels within the speed range [70% 80%] maximum shaft speed are greater than those of “normal” patterns \mathbf{x}_a and \mathbf{x}_b . Furthermore, vibration levels of pattern \mathbf{x}_c in the speed range [60% 65%] are normalised to approximately 4.0. By comparison with the original signatures in Figure 3.5, it may be seen that these high normalised values at the lower shaft speeds of \mathbf{x}_c are caused by relatively small increases in vibration amplitude over \mathbf{x}_a and \mathbf{x}_b . However, because “normal” engines exhibit very little variation in vibration amplitude at lower shaft speeds, such normalised values are considered appropriate for \mathbf{x}_c .

Pattern \mathbf{x}_b , prior to normalisation, showed a peak in vibration amplitude in the range [80% 85%] maximum shaft speed, as seen in Figure 3.5. After component-wise normalisation, shown in Figure 3.10, this peak is

significantly reduced (to below 3.0 standard deviations from the mean), indicating that it is not unusual for vibration signatures in the data set to exhibit high vibration amplitudes above 80% shaft speed, as was described previously.

Thus, component-wise normalisation appears to provide suitable separation between “normal” and “abnormal” example patterns in terms of their overall shape, whilst removing dependence upon absolute values of vibration amplitude.

3.5 Visualisation

In order to confirm the results of the investigation of normalisation, the data set was visualised by projecting the set of 10- and 20-dimensional shape vectors into 2 dimensions. By inspection of the distribution of normalised data in 2 dimensions, the extent to which “normal” and “abnormal” patterns are separated by each normalisation method may be assessed.

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* [62] is based upon the distance d_{ij} between pairs of points (x_i, x_j) in \mathbb{R}^d , and the distance d_{ij}^* between the corresponding pair of points (y_i, y_j) in \mathbb{R}^q :

$$E_{\text{sam}} = \sum_{i=1}^N \sum_{j>i}^N (d_{ij} - d_{ij}^*)^2 \quad (3.11)$$

in which the distance measure is typically Euclidean.

3.5.1 Sammon’s Mapping

Sammon’s mapping [62] attempts to minimise (3.11) as follows:

1. Initialise the projected positions (y_i, y_j) in \mathbb{R}^2 of all P patterns (x_i, x_j) in \mathbb{R}^{20} , randomly.
2. Calculate the pairwise error E_{sam} for all $P(P-1)/2$ inter-pattern distances, as defined in (3.11).
3. Minimise E_{sam} using gradient descent [13]:

$$\forall \mathbf{y}(t) : 1 \dots P, \quad \mathbf{y}(t+1) = \mathbf{y}(t) - \alpha \frac{\partial E_{\text{sam}}}{\partial \mathbf{y}} + \mu \Delta \mathbf{y}(t-1) \quad (3.12)$$

which updates the positions of point \mathbf{y} in the direction of the steepest error gradient. The size of the update is controlled by the learning rate α (typically small). The final term is a momentum term, which reduces the possibility of the descent becoming trapped in local minima of the error function. μ is in the range $0 \leq \mu \leq 1$, and determines the contribution of the momentum in the direction of the previous step $\Delta \mathbf{y}(t-1)$.

4. Repeat 2-3 until E_{sam} is below some threshold, or a maximum number of iterations is reached.

This method of visualisation is used in Chapter 4, and is suited to projection of small numbers of patterns.

3.5.2 NeuroScale

With the NeuroScale model [38], a radial basis function (*RBF*) neural network is used to parameterise the mapping from \mathbb{R}^d to \mathbb{R}^q , in which E_{sam} is minimised. This method allows new test patterns to be projected in \mathbb{R}^q without generating a new mapping, whereas Sammon’s mapping is defined only for the patterns in its training set.

NeuroScale was used to project shape vectors derived from the Trent 500 data set, with $d = 10$ and $d = 20$ inputs and $q = 2$ outputs (for 2-dimensional projection).

The projection of un-normalised 10-dimensional shape vectors derived for HP shaft acceleration manoeuvres is shown in Figure 3.11(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 network architecture used consisted of a single hidden layer of 30 nodes for both the 10-D and 20-D input patterns, with a training set consisting of all 71 patterns in the data set.

Each node in the hidden layer corresponds to a radial basis function centre in 10-D or 20-D space, the initial positions of which were set to be those of 30 patterns randomly selected from the training set.

The output weights of the RBF network (i.e., those from hidden layer to output layer) were initialised using Principal Component Analysis (PCA). The $N = 2$ eigenvectors of the data covariance matrix with highest corresponding eigenvalues are found, and these are the N *principal components* of the training set. They represent the projection of the data in N -dimensional space that retains maximum variance information from high-dimensional space. They provide an initial projection of the data, giving targets for the output layer from which initial values of the output weights are found.

Training the RBF network is a two-stage process. In the first stage, the parameters of the radial basis functions are set so that they model the unconditional data density. In the second stage, the output weights are set by optimising an error function using methods from linear algebra [44].

The resulting projections from each NeuroScale network are shown in Figure 3.11(b)-(f). Note that the axes of NeuroScale projections have no units of measurement: they are axes that allow minimisation of E_{sam} . Equivalent projections of the 20-dimensional set of shape vectors are shown in Figure 3.12. Similar results were obtained when the RBF network contained 20, 40, and 50 nodes in its hidden layer.

From the projection of the un-normalised data in Figure 3.11(a), it can be seen that several of the engines from “abnormal” class 3 are intermixed with points from “normal” class 1, and class 2. The projected points

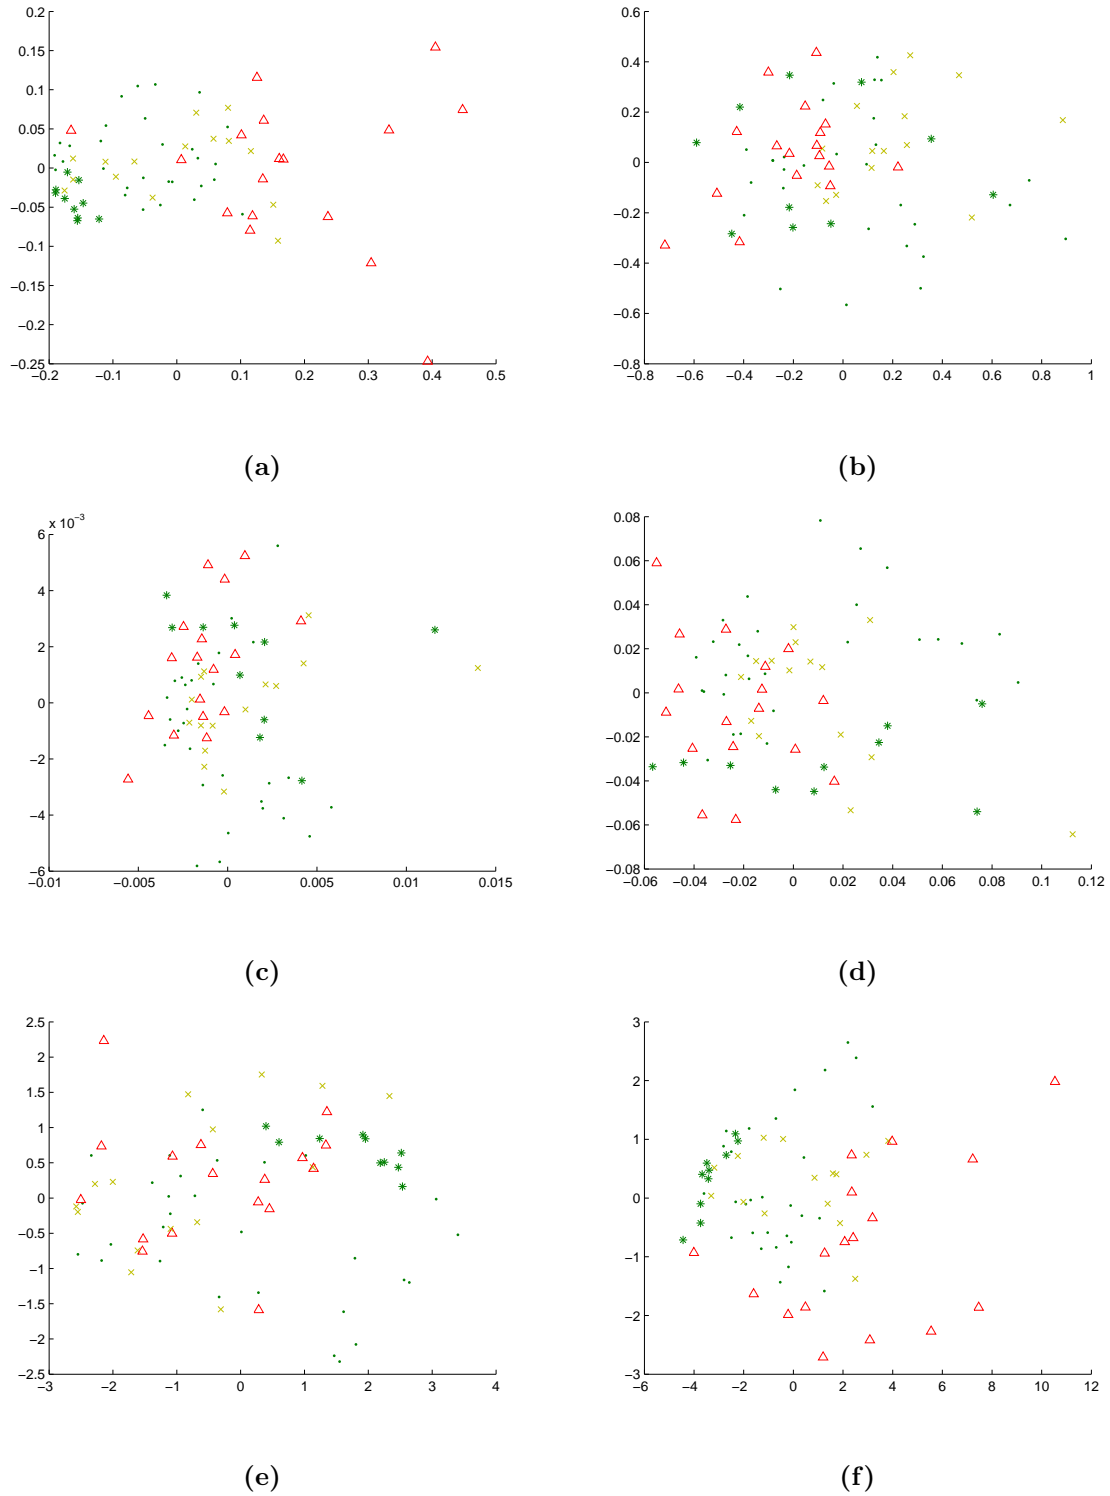


Figure 3.11: Projections of 10-D shape vectors for HP shaft accelerations: (a) un-normalised; (b) unit normalisation; (c) amplitude normalisation; (d) energy-based normalisation; (e) zero-mean, unit-variance normalisation; (f) component-wise normalisation. Labels are shown in Table 3.3.

Class	Symbol	Description
1	Green ●	Normal - far below contractual limit
2	Yellow ×	Normal - close to contractual limit
3	Red △	Abnormal
4	Green *	LP shaft imbalanced

Table 3.3: Class Labels within Visualisations. Classes correspond to those described in Section 3.2

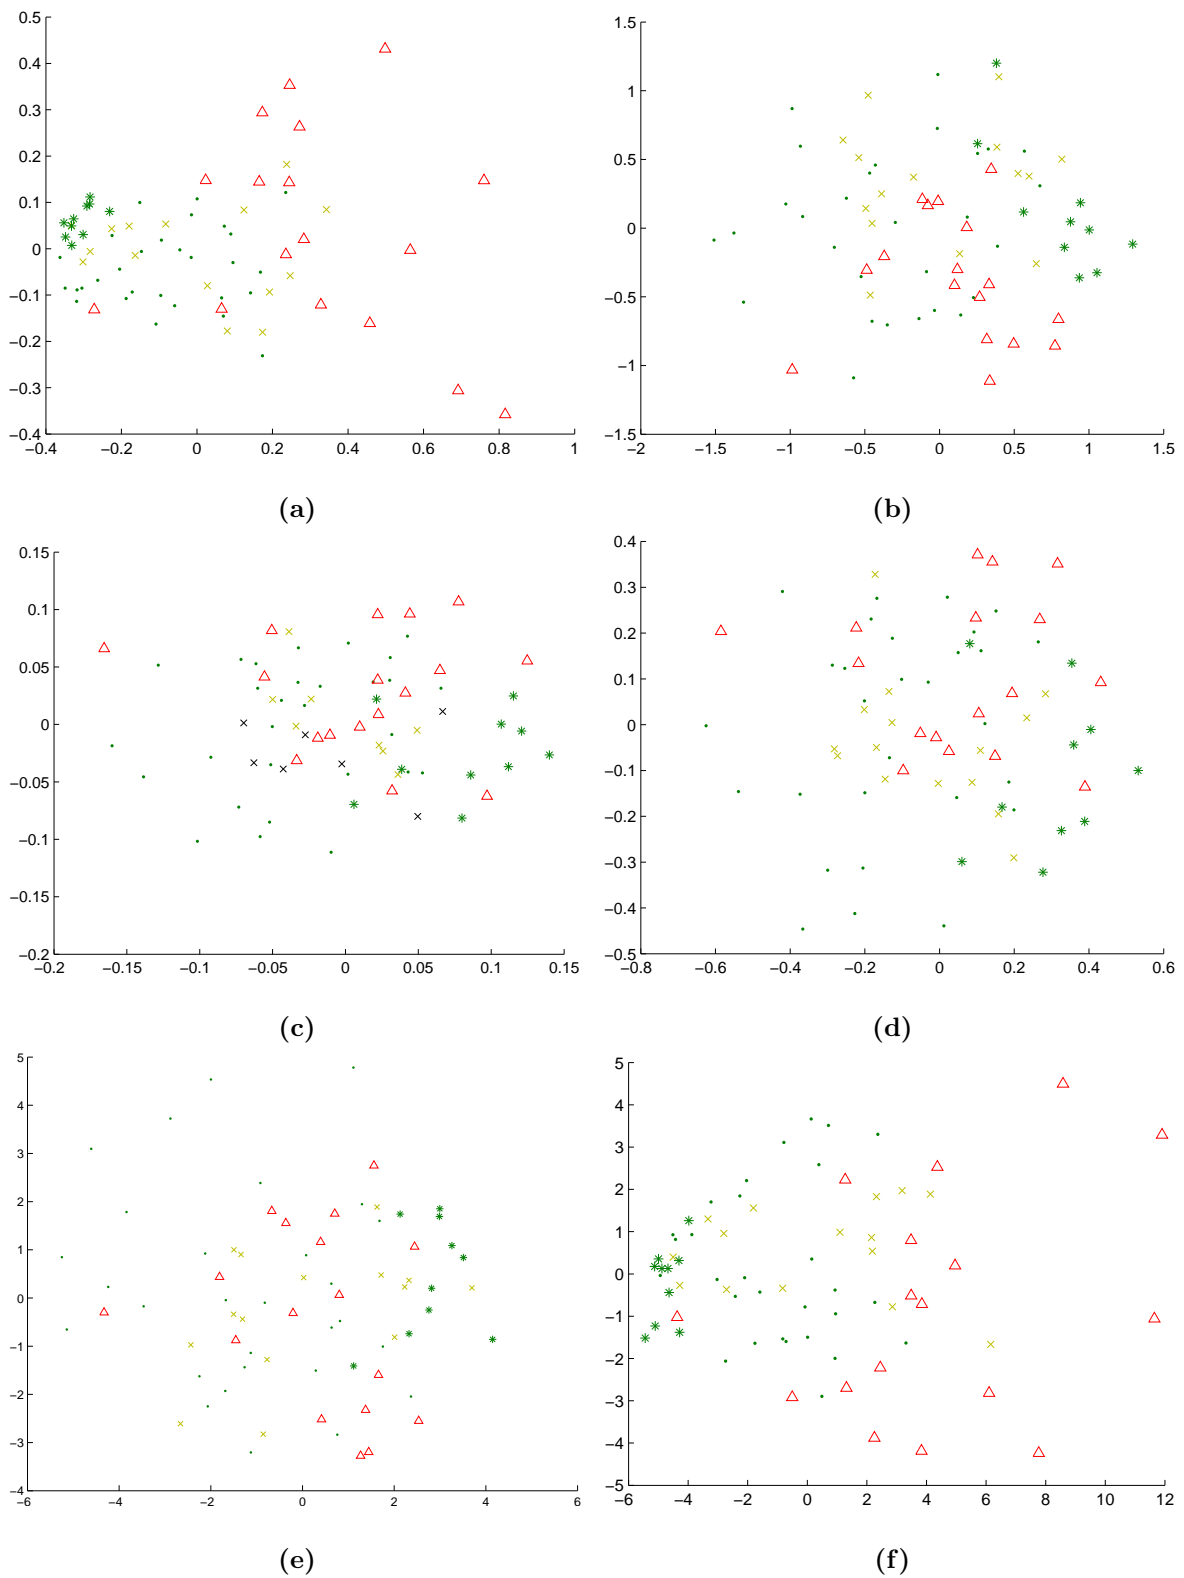


Figure 3.12: Projections of 20-D shape vectors for HP shaft accelerations: (a) un-normalised; (b) unit normalisation; (c) amplitude normalisation; (d) energy-based normalisation; (e) zero-mean, unit-variance normalisation; (f) component-wise normalisation. Labels are shown in Table 3.3.

from the deliberately-unbalanced engines of class 4 form a tight cluster, indicating that their 10-dimensional shape vectors are very similar.

Figure 3.11(b) shows patterns projected after unit normalisation. The overlap between “normal” and “abnormal” examples, illustrated previously in Figure 3.6, is seen across the whole data set, in which the “abnormal” class 3 patterns are projected throughout the distribution of “normal” class 1 and 2 patterns. Even the unbalanced engines, which were tightly clustered in the projection of the unnormalised shape vectors, are widely distributed across the space occupied by other classes. Thus, the projection of the entire data set after unit normalisation confirms the unsuitability of unit normalisation for this application.

Similarly, the overlap between “normal” and “abnormal” patterns seen previously in Figures 3.7-3.9 is also confirmed by projection of the data after those normalisation schemes have been applied, as shown in Figure 3.11(c)-(e). In each case, the “abnormal” class 3 patterns are distributed throughout the projected space occupied by the “normal” class 1, class 2, and class 4 patterns. Thus, amplitude-, energy- and ZMUV-normalisations are also shown to be unsuitable for this application.

Figure 3.11(f) shows the data set projected after component-wise normalisation. The separation seen previously between “normal” and “abnormal” example patterns in Figure 3.10 is confirmed by the projection: the “abnormal” pattern \mathbf{x}_c used previously is the pattern from class 3 shown in the upper-right of Figure 3.11(f), while the “normal” patterns \mathbf{x}_a and \mathbf{x}_b are projected within the centre of the cluster of normal patterns. The figure shows that this separation generally extends to a distinct separation between almost all “normal” and “abnormal” shape vectors.

These trends appear similarly in Figure 3.12, and it can be seen that the projections of 10- and 20-dimensional shape vectors are very similar. Though the 20-dimensional shape vectors more closely represent vibration signatures, the separation between “abnormal” (class 3) patterns and other classes has not increased. This may be caused by the effect that distances between patterns becomes more uniform with increasing pattern dimensionality. As the number of dimensions in a pattern increases, the contribution to the Euclidean distance between patterns caused by small numbers of highly abnormal dimensions in a pattern is reduced, relative to the increasing number of other (normal) dimensions. The small contributions to the Euclidean distance between patterns from the increasing number of other, normal dimensions eventually dominate.

3.6 Clustering

After normalisation of all shape vectors, a model of normality may be constructed to characterise “normal” shape vectors. Patterns which give rise to deviations from this model, if significant in magnitude, may then be classed as novel. This section describes the application of a clustering algorithm to a training set of normal

patterns, in order to construct a model of normality with a relatively small number of prototype patterns. Use of this model as a basis for novelty detection is discussed in Section 3.8.

3.6.1 K -means Clustering

Batch K -means clustering is an iterative method of producing a set of prototype patterns $\boldsymbol{\mu}_j$ (for $j = 1 \dots K$) that represent the distribution of a (typically much larger) set of patterns \mathbf{x}_i . The K -means algorithm is often used to compress large training sets into computationally-tractable smaller sets.

Each pattern \mathbf{x}_i is assigned to its nearest cluster \mathbf{C}_j , centred at $\boldsymbol{\mu}_j$. Thus, each cluster \mathbf{C}_j has a population of patterns associated with it. An error function may be defined to be the total sum-of-squared distances within each cluster, summed for all K clusters:

$$E = \sum_{j=1}^K \sum_{i \in \mathbf{C}_j} \|\mathbf{x}_i - \boldsymbol{\mu}_j\|^2 \quad (3.13)$$

where cluster \mathbf{C}_j has centre $\boldsymbol{\mu}_j$, the mean of its population of N_j patterns:

$$\boldsymbol{\mu}_j = \frac{1}{N_j} \sum_{i \in \mathbf{C}_j} \mathbf{x}_i \quad (3.14)$$

A K -means model is trained by initially setting the centres $\boldsymbol{\mu}_j$ to be a set of randomly-selected patterns \mathbf{x}_j . The algorithm then proceeds by iterative application of the following two steps, until there is no further decrease in the error function from (3.13):

1. $\forall j : 1 \dots K$, compute $\boldsymbol{\mu}_j$ from (3.14);
2. $\forall i : 1 \dots N$, assign pattern \mathbf{x}_i to the population of its nearest cluster centre \mathbf{C}_j .

This is the *batch* K -means algorithm, in which cluster centres are assigned new positions after all patterns have been assigned to a cluster (in the second step, above). This is used in preference to the *adaptive* K -means algorithm, in which cluster centre positions are updated after selection of a pattern at random from the training set. This method is susceptible to becoming trapped in local minima, though it has advantages for “on-line” learning, in which data are acquired sequentially [68].

Due to the random initialisation of cluster centres, the positions of those centres at the conclusion of the training process may differ between different models trained upon the same data set. Typically, the algorithm is run several times (each producing different cluster centre locations $\boldsymbol{\mu}_j$), and one of the candidate models is selected based on some metric of fitness in characterising the data \mathbf{x}_i . This is discussed in a later section.

3.7 Choosing a Model of Normality

The training set was used to generate a series of models of normality using the batch K -means algorithm. 5 models were constructed for $K = 3 \dots 20$. Thus, $5 \times 18 = 90$ models of normality were generated, for each of

the two sets of shape vectors (10-D and 20-D), from which one must be selected for use in novelty detection with each set of shape vectors. This section describes the model selection process.

Figure 3.13 shows how two metrics were used for selecting a model of normality. The upper plot shows the value of the error function E defined in (3.13) for each of the 90 models, with increasing K . Note that there are 5 models for each value of K ; i.e. models 1 to 5 have $K = 3$ cluster centres, models 6 to 10 have $K = 4$ cluster centres, etc.

The upper plot of the figure shows that the value of the E generally decreases as the value of K increases. As the number of cluster centres in a model is increased, the fit of those centres to the data becomes better, eventually reaching the ultimate minimum (perfect fit) when the number of cluster centres equals the number of patterns. However, as the model-fit increases, “over-fitting” occurs in which case the model will not be sufficiently general to apply to new, test patterns. Traditionally, a validation set of patterns is applied to decide on the optimal value of K .

A good model is one in which no cluster centre has few patterns associated with it. If the latter occurs, however, it indicates an inappropriate choice of cluster centres. As described previously, we define the *population* of a cluster \mathbf{C}_j to be those patterns \mathbf{x}_i for which \mathbf{C}_j is the closest cluster.

The lower plot of Figure 3.13 shows the size of cluster populations for one of the 90 models. The illustrated model is a poor choice: it shows that two cluster centres ($j = 2, 6$) have large populations (14 and 13 patterns, respectively), whilst two other cluster centres ($j = 1, 9$) have populations containing a single pattern.

As K increases, models will have a significantly higher likelihood of containing cluster centres with small populations, as the training set is divided into increasing numbers of partitions. If all of the clusters within a model have populations greater than a single pattern, that model is marked with a green dot on the plot of error function E , shown in the upper plot of Figure 3.13. It can be confirmed from the figure that as K increases, models tend to have no green dot associated with their location in the plot of the error function E .

A model was selected with $K = 4$ cluster centres, corresponding to model number 10 within Figure 3.13. The 4 cluster centres within this model had approximately the same sized populations (9, 10, 11, and 13 patterns), fulfilling the requirement that no cluster centre within the model should have few patterns associated with it.

3.8 Calculation of Novelty

The novelty value z_i of pattern \mathbf{x}_i is computed with respect to the K cluster centres within the selected model:

$$z_i = \min_{j=1}^K \frac{d(\mathbf{x}_i, \mathbf{C}_j)}{\sigma_j} \quad (3.15)$$




Figure 3.13: Model evaluation metrics: error (upper plot), and cluster populations for an example model (lower plot). The exemplar is indicated in the upper plot by a red cross. Green dots in the upper plot indicate that all cluster populations within that model contain more than a single pattern.

where $d(\mathbf{x}_i, \mathbf{C}_j)$ is Euclidean distance, and σ_j is the width of cluster j :

$$\sigma_j = \sqrt{\frac{1}{I_j} \sum_{i=1}^{I_j} d(\mathbf{x}_i, \mathbf{C}_j)^2}. \quad (3.16)$$

for the I_j points which have closest cluster centre \mathbf{C}_j

Figure 3.14 shows a cumulative plot of the novelty values of all patterns in both normal classes (1 and 2) and abnormal classes (3). The normal class of patterns is indicated by a blue line, and shows the number of patterns (on the y -axis) that have the corresponding novelty value or lower (on the x -axis). If we consider Figure 3.14(a) (10-D shape vectors), approximately 25 patterns from the normal classes have a novelty value of 1.0 or lower; increasing the novelty threshold to 2.0 encloses all of the normal patterns, but also begins to include some of the abnormal class (indicated by the red line).

A novelty threshold H_η above which patterns are classified as abnormal may then be set, corresponding to a vertical line on the cumulative distribution plot, as shown in Figure 3.14, in which the threshold is set to include all normal patterns. For the case of the 10-D shape vectors, 5 abnormal patterns have a novelty value below the novelty threshold, and so would be classified as “normal”; for the 20-D shape vectors, the corresponding number of patterns is 3. Figure 3.15 shows the 2 of these abnormal patterns which had lowest novelty scores for both 10- and 20-D shape vectors. Vibration amplitude at the upper end of the range of shaft speeds, exceeding the contractual limit of $H_{\text{MAX}} = 0.4 \text{ ins}^{-1}$, caused these signatures to be placed within the abnormal (class 3) data sub-set. These high vibration levels are not retained after quantisation, as may be seen by the red and blue lines in the figure, corresponding to the 10- and 20-D shape vectors, respectively.

There is a high cost associated with false-positive novelty detections in engine pass-off testing, so the positive predictive value of the novelty detection process should be maximised (as defined in Section 2.3.3). Thus, the




Figure 3.14: Cumulative number of patterns by novelty score: (a) 10-D shape vectors; (b) 20-D shape vectors. Normal patterns are shown in blue, abnormal patterns in red. An example novelty threshold is shown as a vertical dashed line that gives maximum model specificity.

novelty threshold is set such that all patterns in the training set of normal examples are classified “normal”.

If the sensitivity of the process is increased such that all abnormal examples are correctly identified as such (using a novelty threshold $H_\eta = 1.1$ for the case of 10-D vectors, as shown in Figure 3.16(a)), 11 patterns from the normal training set would be classified as “abnormal”, which would not be acceptable in this application.

3.9 Conclusion

Conclusions drawn from the analysis of HP shaft data are presented below, followed by summaries of the application of the process to other shafts. For each case, results from the 10- and 20-dimensional sets of shape vectors are compared.

3.9.1 HP Shaft

Figure 3.17 shows all patterns projected into 2-D, using class symbols defined in Table 3.3. Those patterns classified “abnormal” are circled in blue.

Classification of “Abnormal” Patterns

Using a novelty threshold value $H_\eta = 2.1$, 12 of the 17 abnormal patterns (class 3) were classified as “abnormal”, when $N = 10$ dimensions were used in quantisation. This increases to 14 of 17 abnormal patterns when $N = 20$ dimensions are used in quantisation ($H_\eta = 1.6$).

Class 3 contains patterns from sources A, B, and C (defined in Section 3.2.2) as shown in Table 3.4. The table also indicates the number of these patterns within class 3 that were classified as “abnormal” by the novelty detection process.




Figure 3.15: Two abnormal (class 3) vibration signatures (shown in black) whose corresponding 10-D and 20-D shape vectors (shown in red and blue, respectively) had lowest novelty scores of all abnormal patterns. The contractual limit H_{MAX} is shown by the dashed line. (a) Trent 500 engine 71123; (b) Trent 500 engine 71119.

Source	Patterns in Class 3	Classified “abnormal” w.r.t. 10-D model	Classified “abnormal” w.r.t. 20-D model
A	3	0	1
B	7	6	7
C	7	6	6

Table 3.4: Class 3 patterns, divided by source, with indication of “abnormal” classification after novelty detection.

6 of 7 patterns from engines which failed pass-off testing (source B) were correctly classified for $N = 10$, as were engines which were overhauled following an extended period of commercial use (source C). However, all patterns derived from engines which were released into service (source A) were classified “normal” for $N = 10$, within class 3.

The vibration level of these 3 engines from source A exceeded the contractual vibration limit for the HP shaft, thus placing them into class 3, using the criterion defined in Section 3.2.2. However, their corresponding vibration signatures were “normal” overall, with respect to the training set. 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.

Using 20-D shape vectors produces better classification performance than 10-D shape vectors. All 7 patterns from engines which failed pass-off testing (source B) were correctly classified for $N = 20$. 1 of the 3 engines from source A that exceeded the contractual vibration limit was also classified as “abnormal”.

Classification of “Normal” and “Unbalanced” Patterns

The “normal” and “unbalanced” classes (1, 2, and 4), in which maximum engine vibration levels were below the contractual limit in the HP shaft, were all classified “normal” with regard to the model, for $N = 10$ with




Figure 3.16: Cumulative number of patterns by novelty score: (a) 10-D shape vectors; (b) 20-D shape vectors. Normal patterns are shown in blue, abnormal patterns in red. An example novelty threshold is shown as a vertical dashed line that gives maximum model sensitivity.




Figure 3.17: Projection of all HP shaft patterns for (a) 10-D shape vectors; (b) 20-D shape vectors. Patterns classified as “abnormal” are circled in blue. Labels are as in Table 3.3.

two exceptions:

- One pattern with maximum vibration levels close to the contractual limit for the HP shaft (class 2) was classified as “abnormal”. This was an engine identified by engine developers as requiring rebuilding due to high vibration levels within the HP shaft. With $N = 20$, a second engine from class 2 was also identified as “abnormal”.
- One pattern taken from an engine that had been deliberately unbalanced (class 4) was classified as “abnormal”. Section 3.2.2 noted that the 10 patterns taken from unbalanced engines should be classified as normal with respect to the HP shaft model. With $N = 20$, this pattern was correctly classified as “normal”.

When training a model of normality, training patterns which appear significantly separated from the majority of other training patterns may be deemed unrepresentative of normality (being “outliers”), and excluded from the training set. This “pruning” of the training set may result in these outlying patterns being classified as “abnormal”, which is the cause of the detections described above.

3.9.2 LP, IP, and Radial Drive Shaft Conclusions

The same process of shape vector construction, model training, and selection of a novelty score threshold was performed using data sets for the LP, IP, and radial drive shafts. The results are reviewed below.

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 3.2.2. Figure 3.18 shows a NeuroScale projection of all patterns generated from LP shaft vibration signatures. Class labels are as defined in Table 3.3.




Figure 3.18: Projection of all LP shaft patterns for (a) 10-D shape vectors; (b) 20-D shape vectors. Patterns classified as “abnormal” are circled in blue. Labels are as in Table 3.3.

Results are almost identical for both the $N = 10$ and $N = 20$ sets of shape vectors.

6 patterns derived from unbalanced engines 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 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 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 class 3 (as defined in Section 3.2.2). This provides support for the claim that

comparison of maximum vibration amplitude to a fixed contractual limit does not provide a reliable indication of abnormality.

IP Shaft Results

All abnormal patterns (class 3) derived from IP shaft vibration signatures are correctly classified by the novelty detection process, as shown in Figure 3.19, for both $N = 10$ and $N = 20$ sets of shape vectors. 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.




Figure 3.19: Projection of all IP shaft patterns for (a) 10-D shape vectors; (b) 20-D shape vectors. Patterns classified as “abnormal” are circled in blue. Labels are as in Table 3.3.

Radial Drive Shaft Results

As there is no contractual vibration limit specified for the radial drive shaft, patterns for this shaft could not be reclassified into the 4 classes described in Section 3.2.2. With no prior distinction made between “normal” and “abnormal” patterns, a less-informed approach was taken, in which data from Source A (engines that were released into service following pass-off testing) formed the training set for the construction of the model of normality.

Figure 3.20 shows a NeuroScale projection of all patterns derived from radial drive shaft vibration signatures. Symbols used to plot patterns from each source are shown in table 3.5.

Source	Symbol	Description
A	Blue ●	Released into service
B	Red △	Rejected during pass-off testing
C	Red □	Re-fitted during service life
D	Green *	LP shaft imbalanced

Table 3.5: Source labels within radial drive shaft visualisation

As with the LP and IP shafts, results from novelty detection of radial drive data are similar for both $N = 10$ and $N = 20$ sets of shape vectors. Note that NeuroScale projections may be arbitrarily rotated or inverted.




Figure 3.20: Projection of all radial drive shaft patterns for (a) 10-D shape vectors; (b) 20-D shape vectors. Patterns classified as “abnormal” are circled in blue. Labels are as in Table 3.5.

The figure shows that the projection of patterns derived from the Source A engines (the training set) form a cluster in the centre of the plot in Figure 3.20(a). Distributed throughout this cluster are the projections of 7 of the 12 patterns from Source B (engines that were rejected following pass-off testing). The remaining 5 patterns from Source B are projected outside this main cluster, and classified as “abnormal” with respect to the model of normality.

This indicates that several engines from Source B had vibration amplitudes (measured in the frequency range corresponding to the fundamental tracked order of the radial drive shaft) that were similar to those engines released into service (Source A). This motivates the need to reclassify patterns into “normal” and “abnormal” classes, rather than using the original classification of data by their source, should a contractual limit be specified for the radial drive shaft.

The figure shows that projection of the patterns derived from deliberately unbalanced engines (Source D) form a cluster in the lower-right of the plot in Figure 3.20(a), separated from the main cluster of projections of the training data. These engines are classified as “abnormal”, indicating that the condition of deliberate unbalance results in significant change of vibration amplitudes in the radial drive shaft. Prior to this analysis, it was assumed that unbalance only affected the vibration amplitudes of the LP shaft tracked orders.

Figure 3.20(a) also shows that patterns derived from Source C engines (those re-fitted following a period of use in commercial service) form a cluster in the lower-left of the plot, separated from the main cluster of training set patterns. These engines are classified as “abnormal” with respect to the model of normality.

This result indicates that the vibration signatures of engines may change significantly during their operational lives. Engines from Source C are older than those from other sources, having been used in service, and their significantly different vibration signatures are correctly classified as “abnormal”. When selecting a training set of data, it is necessary to choose example “normal” patterns recorded under similar conditions to those

anticipated to be used in the recording of test patterns. Better sensitivity may be achieved if two different models are used: (i) for engines which have just been built (“new engines”); (ii) for engines which have been in service for a specified number of hours (e.g., 10,000 hours).

3.9.3 Overall Conclusion

The method of vibration analysis presented in this chapter has shown that the “Shape Analysis” inter-engine analysis process, using TORCH signatures, is effective for novelty detection in Trent 500 engine data. Results from analysis of vibration signatures from each shaft of the Trent 500 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 quantisation of vibration signatures with shape vectors of $N = 20$ elements allows closer representation of engine vibration levels (particularly peaks in the vibration signature), as shown in Section 3.3. IP, LP and radial drive-shaft show similar results for both $N = 10$ and $N = 20$. However, results for $N = 20$ with HP shaft data show that more abnormal patterns are correctly classified, and that an unbalanced engine (from class 4), which should appear normal in HP vibration levels, is no longer incorrectly classified. Two engines with maximum vibration levels close to the contractual limit for the HP shaft are classified as “abnormal” when $N = 20$, one of which was determined by Rolls-Royce engineers to require rebuilding. Thus, it is suggested that 20-dimensional shape vectors are used in further analyses using this method.

One anticipated benefit of the application of this method to Trent 500 engines is the identification of engines that have maximum vibration amplitude below the contractual limit, during “pass-off” testing (which would thus be released into service), but which go on to suffer vibration-related faults.

The method presented in this chapter provides a more accurate assessment of abnormality than comparison of vibration levels to a simple threshold. Increasing accuracy of engine assessment allows better control of engine maintenance (as introduced in Chapter 1.1) and, if used during pass-off testing, better achieves the goal of determining if an engine is fit for release into service.

Chapter 4

TORCH for Intra-Engine Analysis

4.1 Introduction

Chapter 2 introduced an intra-engine approach to novelty detection using TORCH vibration signatures. This chapter presents a method of using TORCH information for intra-engine analysis, identifies problems encountered in the construction of a TORCH signature of vibration phase, and proposes a solution to avoid such problems.

Results of applying this intra-engine method of novelty detection to fundamental tracked order vibration data are presented, showing that the method can be used for condition monitoring of aircraft gas-turbines engines. The investigation shows that abnormal engine operational behaviour can be identified in advance of engine failure using this technique.

4.2 Constructing TORCH Signatures of Vibration Phase

As described in Section 2.3.1, vibration phase and amplitude are treated independently when constructing TORCH signatures. Measurements of vibration phase associated with the LP shaft of Trent-series engines are made using a transducer that samples engine vibration phase once per shaft revolution (a so-called *once-per-rev* device), in which phase measurements are aligned with respect to a constant reference phase value. Any two vibration phase signatures Φ_1 and Φ_2 constructed from vibration phase measurements using this type of transducer may be compared, even if derived from different runs of the engine, due to the constant reference value.

However, the phase information associated with IP and HP shafts is typically derived from a *pseudo-once-per-rev* transducer, in which the initial phase offset of the sensor may *not* be constant between runs of the engine. In order to compare a phase signature Φ_1 with a phase signature Φ_2 derived from a different engine run, the constant phase offset ϕ_o between Φ_1 and Φ_2 must be estimated. Φ_2 has to be *realigned* by subtracting ϕ_o from all phase values in Φ_2 , to allow comparison with Φ_1 .

4.2.1 Realigning Vibration Phase Measurements

The constant phase offset ϕ_o between two TORCH phase signatures Φ_1, Φ_2 , measured against shaft speed divided into B bins of equal width, may be estimated by minimising a cost function $E(\phi_o)$:

$$E(\phi_o) = \sum_{b=1}^B \frac{(\mu_2^b - \mu_1^b - \phi_o)^2}{(\sigma_2^b)^2} \quad (4.1)$$

where μ_n^b and $(\sigma_n^b)^2$ are the means and variances of phase signature Φ_n , respectively, for shaft-speed bins $b = 1 \dots B$. Minimising $E(\phi_o)$ corresponds to finding a constant phase offset ϕ_o that minimises the squared difference between the phase signature means. An example of realigning a phase signature is shown in Figure 4.1, in which $\phi_o = 0.4$ radians.




Figure 4.1: Phase offset example. A signature for the phase of vibration of the 1LP tracked order Φ is shown in blue. Realigning $\Phi - \phi_o$, with $\phi_o = 0.4$ radians, is shown in red. Note that phase is plotted in the range $[0 \ 2\pi]$.

The squared differences in phase mean for speed bin b are weighted by the variance of the phase signature to be aligned $(\sigma_2^b)^{-2}$; i.e. speed bins with low phase variance contribute more to the squared difference than speed bins with high phase variance.

4.2.2 Minimising the Phase Offset Error Function

Minimisation of the one-dimensional function $E(\phi_o)$ across the domain of phase offsets $\phi_o : [0 \ 2\pi]$ may be performed using the *Golden Section Search* process [55]. This method converges upon the minimum of a one-dimensional function by iterative bisection of that function’s domain (here $[0 \ 2\pi]$). It can be shown that if each bisection is made at a point $(3 - \sqrt{5})/2$ through the remaining interval (hence the Pythagorean “golden ratio” of approximately 0.38 : 0.62), each iteration reduces the search interval at an optimal rate.

The minimum value of $E(\phi_o)$ corresponds to the value of ϕ_o that best aligns phase signature Φ_2 to phase signature Φ_1 . This is assumed to be a close approximation of the actual phase offset between initial states of the phase-measurement transducer used in the two different engine runs.

4.2.3 Illustration of Phase Realignment

This use of this process to find ϕ_o is illustrated with an example phase signature, shown in Figure 4.2. A signature of mean vibration phase of the 1LP tracked order μ_1^b is plotted in blue, for $b = 1 \dots 400$ speed bins, covering the range [0% 100%] maximum shaft speed. A second phase signature μ_2^b is generated by adding normally-distributed noise with distribution $N(1, 1)$ to the original signature μ_1^b , plotted in green.




Figure 4.2: Phase realignment example. A signature for the phase of vibration of the 1LP tracked order is shown in blue. A second signature generated by adding normally-distributed noise to the first is shown in green.

Using the Golden Section Search method to minimise (4.1) for these two example phase signatures results in an estimated phase offset $\phi_o = 0.96$ radians, a good estimate of the actual phase offset $\phi_o = 1.0$ radians.

When the superimposed noise used to generate μ_2 was drawn from the normal distribution $N(2, 1)$ instead, the phase offset was estimated to be $\phi_o = 1.96$ radians, again a close estimate of the true phase offset.

4.3 Intra-engine Analysis

This section presents a method of intra-engine novelty detection using TORCH signatures constructed from vibration data recorded from a series of runs of a single *Trent 900* engine. It is shown that, using these techniques, engine behaviour can be determined to be abnormal in advance of engine failure.

Vibration signatures constructed by engine-mounted monitoring systems from flight data are transferred from the aircraft at the end of a flight. These vibration signatures must be quantised due to transmission bandwidth limitations. The effect of quantisation on flight signatures, and their subsequent use for novelty detection, is also investigated as part of the analysis presented in this chapter.

4.3.1 Data Used in this Investigation

A data set consisting of tracked order vibration data (amplitude and phase) and shaft speeds was used in this investigation, recorded from a Trent 900 engine, as shown in Table 4.1. Data were recorded during ground-based “cyclic” testing of the engine, in which a defined manoeuvre was repeated 29 times in succession. Figure 4.3 shows LP shaft speed throughout four test cycles. A period of constant operation at 83% max. speed (which varies in duration between tests) is followed by a defined series of accelerations and decelerations over a period of 60 seconds, after which the engine is decelerated to “idle” speed (approximately 16% max. speed for the LP shaft).

Name	Date	Notes
29 Cyclic Tests	08/05/2004	Development Engine 90006-1

Table 4.1: Data set used in intra-engine analysis of Trent 900 tracked order vibration

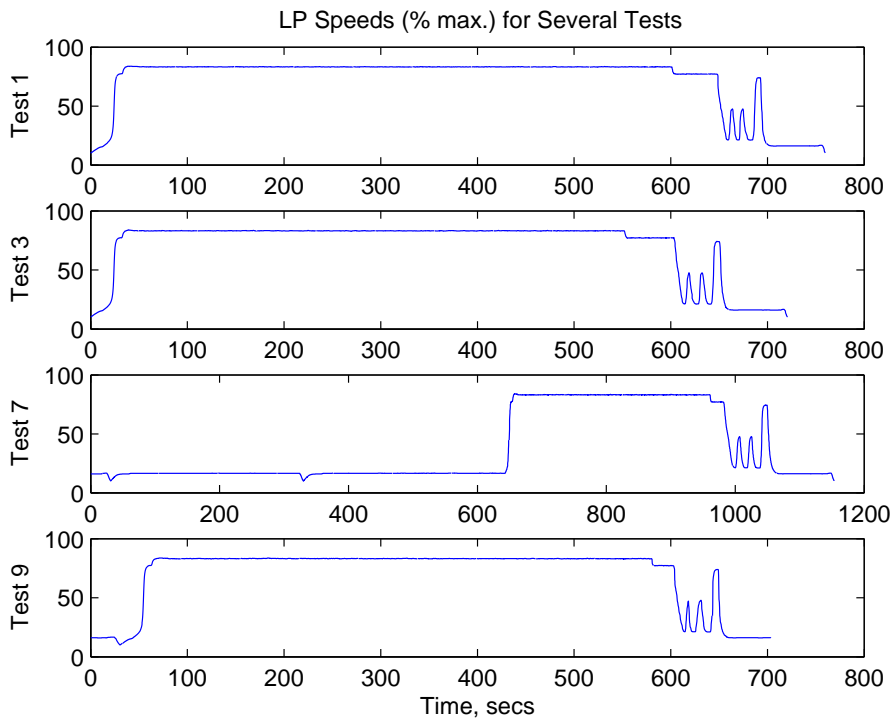


Figure 4.3: LP shaft speeds for four cyclic tests, showing a period of sustained engine operation at high speed, followed by a defined series of engine accelerations and decelerations

A fire occurred in the engine that was not identified during the testing process until test 29, when the engine required manual shutdown.

4.3.2 Constructing TORCH Signatures

Vibration signatures constructed during flight are quantised prior to their transmission to ground-based data analysis systems at the end of a flight:

- i. Vibration amplitude is linearly quantised using 94 quanta covering the range $[0 \ 3] \text{ ins}^{-1}$;

- ii. Vibration phase is linearly quantised using 94 quanta covering the range $[0 \ 2\pi]$ radians;
- iii. Shaft speed is linearly quantised using $B = 20$ speed bins covering the range $[15\% \ 100\%]$ max. speed.

To investigate the effect of such quantisation on TORCH signatures, and their use for novelty detection, three signatures are created for both vibration amplitude and phase, for each of the 29 tests, as shown in Table 4.2.

Signature number	Amplitude quantisation	Phase quantisation	Speed quantisation (B)
1	unquantised	unquantised	400 speed bins
2	as (i)	as (ii)	400 speed bins
3	as (i)	as (ii)	as (iii)

Table 4.2: Quantisation used in creating TORCH signatures

4.3.3 Examination of Quantisation Effects

Figure 4.4 shows vibration amplitude signatures for tests 1, 5, and 10. It can be seen that the 400-D amplitude signatures for each test appear generally similar, with vibration amplitude increasing with shaft speed similarly between tests. Signatures constructed from quantised amplitude values, as described in Section 4.3.2, appear very similar to the corresponding unquantised signatures. The 20-D signatures can be seen to follow the general shape of the unquantised signatures, while not retaining high-frequency “noise-like” variations in vibration amplitude. This investigation aims to determine if adequate information is retained in the 20-D signatures to perform reliable novelty detection.

Figure 4.5 shows phase amplitude signatures for tests 1, 5, and 10. Again, signatures appear similar between tests. Phase decreases with increasing shaft speed. Approximately one-third of the way through the speed range, phase values decrease below zero, and are mapped to 2π radians, due to the phase measurement range $[0 \ 2\pi]$ radians. This transition appears to occur at slightly higher speeds in test 1 than in tests 5 and 10. Several sharp decreases in phase are evident for tests 5 and 10, at approximately speed bin 180. Linear quantisation of phase values in these signatures appears to discard little information. The 20-D signatures again follow the general shape of the 400-D signatures, while discarding high-frequency variation. The sharp decreases in phase are represented as small decreases in speed bin 9.

Vibration amplitude signatures for the final three tests of the series are shown in Figure 4.6. Comparison with the amplitude signatures from the earlier three tests (in Figure 4.4) shows that there is little variation in amplitude signatures throughout the series. Test 29 shows large vibration amplitude at the top of its speed range, corresponding to the fire that resulted in manual shutdown of the engine.

Phase signatures for the final three tests are shown in Figure 4.7. These signatures show significant differences to phase signatures from earlier tests. Large decreases in phase can be seen in test 27 at approximately speed

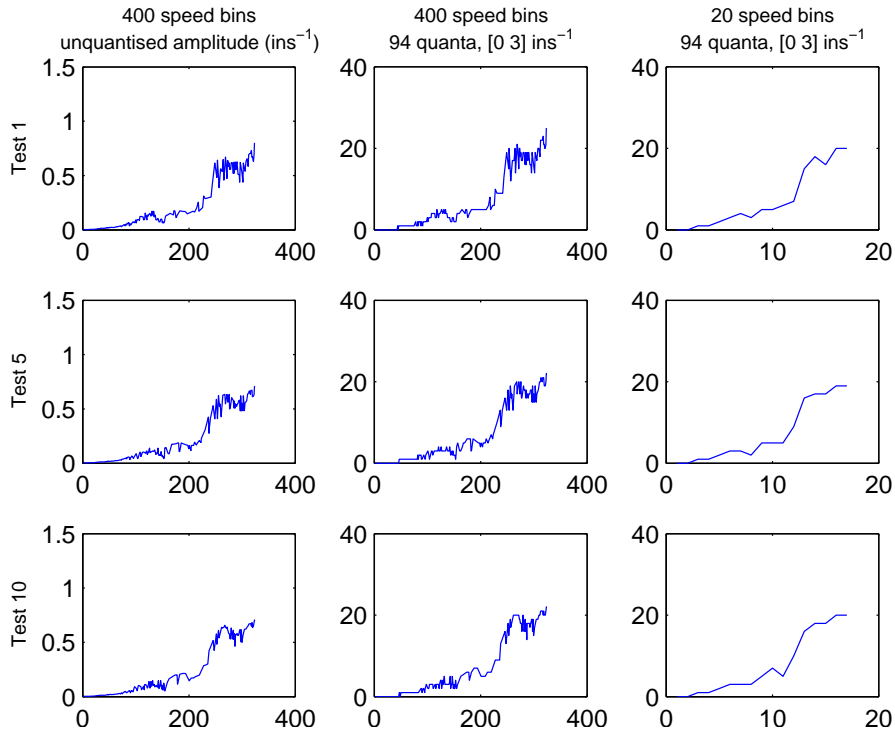


Figure 4.4: Amplitude signatures for cyclic tests 1, 5, and 10. Unquantised signatures are shown in the left column; quantised amplitude values are used in the middle column; quantised amplitude values and $B = 20$ speed bins are used in the right column.

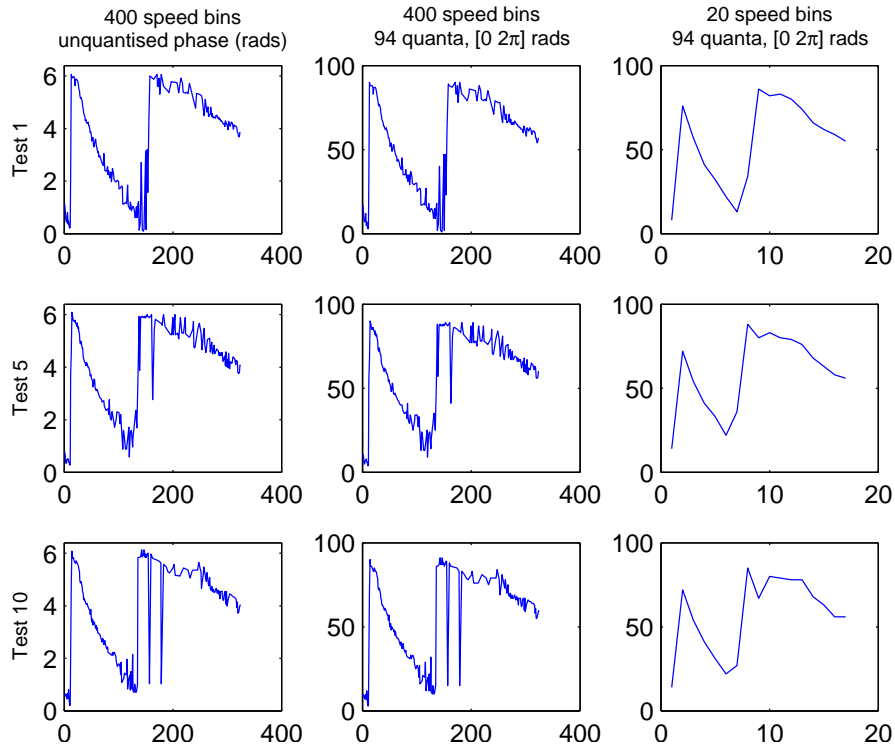


Figure 4.5: Phase signatures for cyclic tests 1, 5, and 10. Unquantised signatures are shown in the left column; quantised phase values are used in the middle column; quantised phase values and $B = 20$ speed bins are used in the right column.

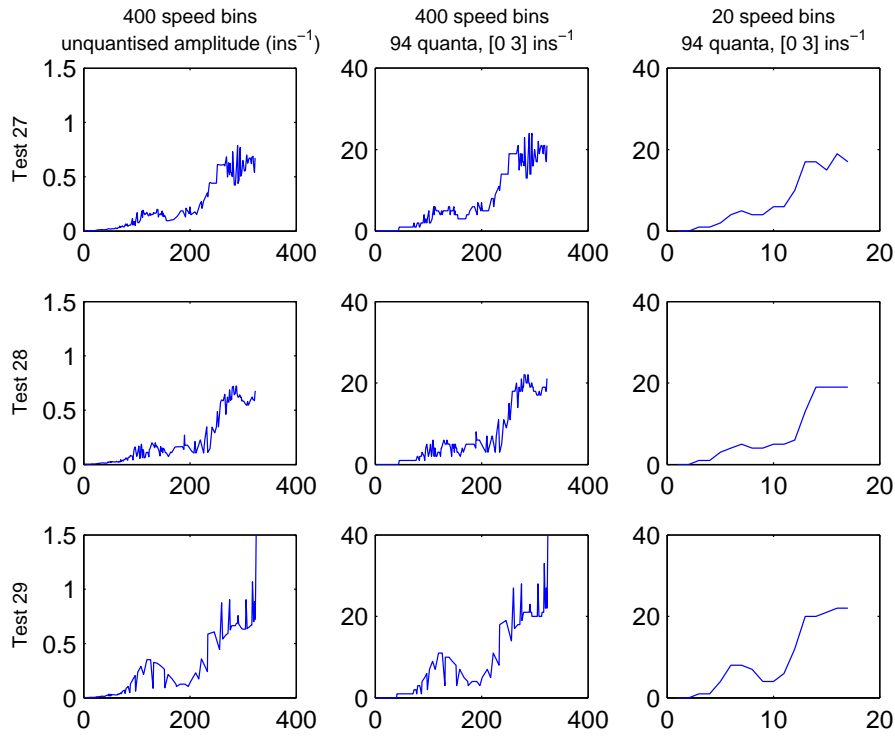


Figure 4.6: Amplitude signatures for cyclic tests 27, 28, and 29 (one test per row). Unquantised signatures are shown in the left column; quantised amplitude values are used in the middle column; quantised amplitude values and $B = 20$ speed bins are used in the right column.

bin 210, and in speed bins 180 to 200 in test 29. These phase changes were linked by Rolls-Royce engineers to the appearance of small fires in the LP shaft of the engine, which were not identified during testing, and which culminated in the large fire during test 29.

4.3.4 Visualisation

To explore the differences between vibration signatures from each of the 29 tests, visualisation is performed. Due to the small number of patterns available for training of a projection mapping, a NeuroScale network approach (as described in Section 3.5) is not appropriate. Using the rule of thumb that states the number of RBF kernels in the network (i.e., nodes in the hidden layer) should be at least as great as the number of input nodes (and preferably much greater) [68], this would require at least 20 kernels. The rule of thumb that states the size of the training set should ideally be an order of magnitude greater than the number of RBF kernels [68] leads to a requirement of at least 200 training patterns.

A more appropriate method for projection of small numbers of patterns is Sammon’s mapping, described in Section 3.5.1.

Projections of the 20-D vibration signatures for the 29 tests are shown in Figure 4.8, allowing comparison of signatures constructed from unquantised vibration data with signatures constructed from data quantised using (i) and (ii) in Section 4.3.2.

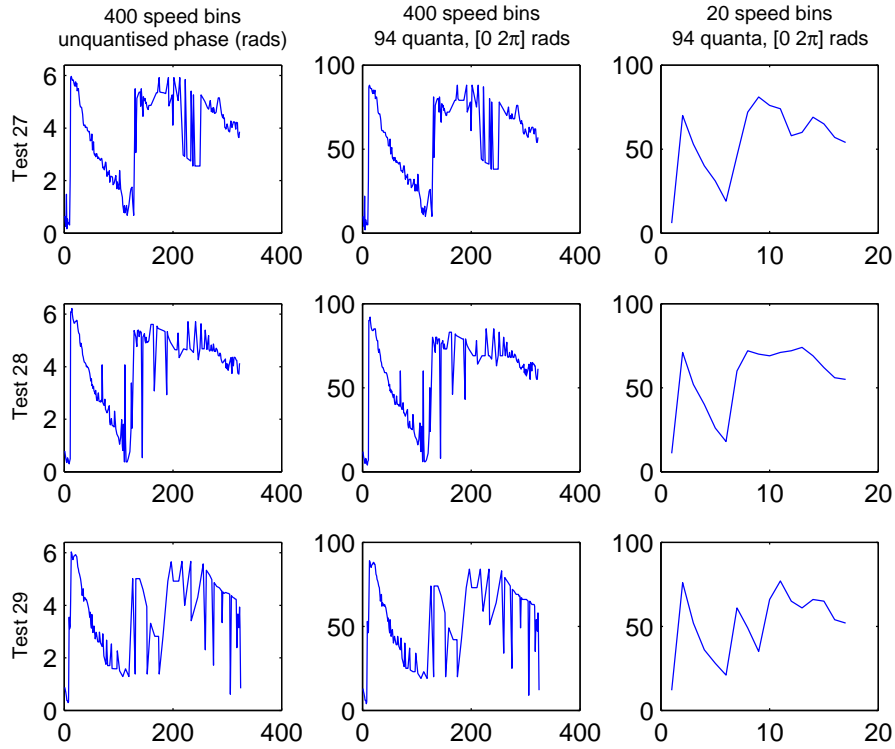


Figure 4.7: Phase signatures for cyclic tests 27, 28, and 29 (one test per row). Unquantised signatures are shown in the left column; quantised phase values are used in the middle column; quantised phase values and $B = 20$ speed bins are used in the right column.

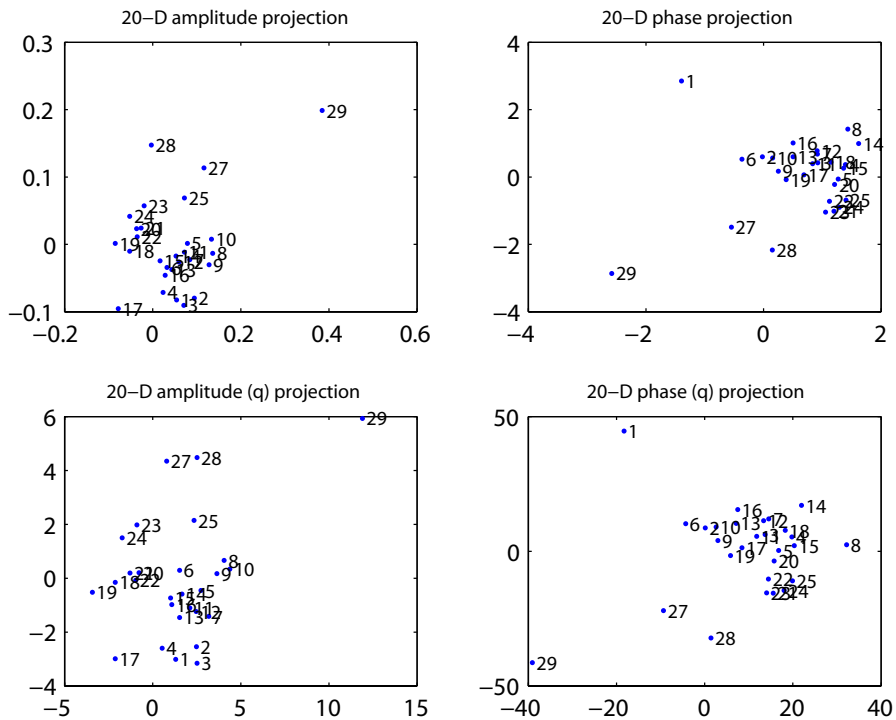


Figure 4.8: Visualisation of 20-D signatures for 29 cyclic tests. Unquantised and quantised amplitude signatures are shown in the upper left and lower left plots, respectively; unquantised and quantised phase signatures are shown in the upper right and lower right plots, respectively.

It can be seen from the visualisations of amplitude signatures that tests 27, 28, and 29 are separated from the main cluster formed by all other signatures. This is evident for signatures constructed using both quantised and unquantised amplitude values. These three tests appear further separated from the main cluster in visualisations of phase signatures. Again, this separation is evident for both quantised and unquantised phase data.

Test 1 appears to be separated from the main cluster of patterns in the visualisations of phase signatures. This is confirmed in Figure 4.9, in which test 1 appears significantly different in speed bins 7, 8, and 9 to those of six other tests, which represent “normal” patterns that lie within the main cluster of points in Figure 4.8. It can be seen that the phase increase to 2π in the phase signature of test 1 is shifted one speed bin to the right when compared with those of the other tests. This confirms the observation made in Section 4.3.3.

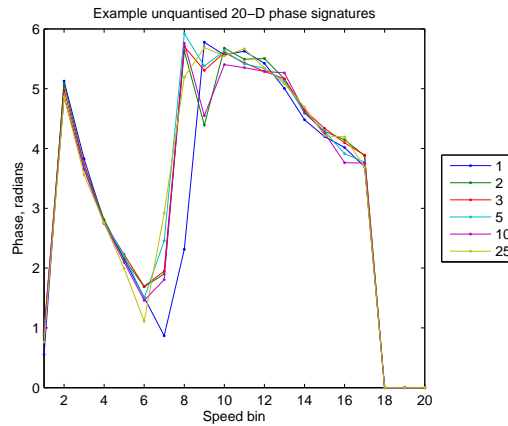


Figure 4.9: 20-D phase signatures for six cyclic tests.

4.3.5 Novelty Detection

As described in Section 3.6.1, k -means clustering was used to construct a separate model of normality (with $k = 2$) for each of vibration amplitude and phase (both quantised and unquantised). These used the 20-D signatures from the first 6 tests of the series, assuming that these first tests represent normal operation of the engine, noting that test 1 contains a phase difference to other signatures as observed previously. Novelty scores of patterns were again calculated using

$$z_i = \min_{j=1}^K \frac{d(\mathbf{x}_i, \mathbf{C}_j)}{\sigma_j}, \quad (4.2)$$

and a novelty threshold H_η set as described in Section 3.8. Values of H_η for the four sets of signatures are shown in Table 4.3. Results of novelty detection using these models are shown in Figure 4.10.

Data set	Quantisation	H_η
20-D amplitude	unquantised	4.0
20-D amplitude	quantised	3.8
20-D phase	unquantised	4.5
20-D phase	quantised	3.0

Table 4.3: Novelty threshold H_η for each set of TORCH signatures

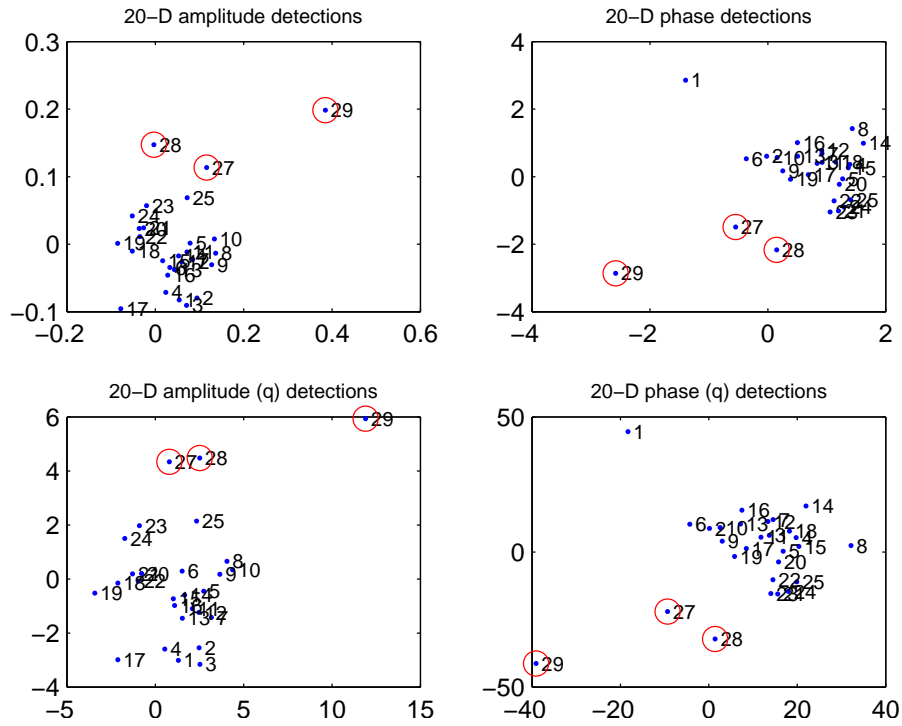


Figure 4.10: Visualisation of 29 cyclic tests, with patterns classified as “abnormal” with respect to corresponding models of normality circled in red.

The final three tests of the series were detected as being abnormal with respect to the models of normality for both amplitude and phase signatures.

4.3.6 Conclusion

This investigation has shown that abnormality in the tracked order vibration of a single engine is evident in TORCH signatures constructed using both amplitude and phase data. Abnormality corresponding to a hazardous condition (undetected during testing) is evident in signatures constructed from data up to 36 minutes prior to manual engine shutdown. Increased vibration amplitude and unusual variation in vibration phase during the final three tests were identified by Rolls-Royce engineers as pre-cursors of the hazard. These are represented within the corresponding TORCH signatures.

A method of novelty detection similar to that used for inter-engine analysis in Chapter 3 was shown to retrospectively identify the abnormal TORCH signatures. However, the amount of training data available in this case suggests that further work and much more example data would be required in order to provide robust novelty detection.

The quantisation of vibration amplitude and phase as will be used for the transmission of vibration signatures from Trent 900 engines after a flight was shown to retain sufficient information to perform novelty detection for the test data used here.

Chapter 5

Future Work

5.1 Flight Summary Compression

The downloading of aircraft flight data is often subject to tight restrictions on available transmission bandwidth. In order to maximise information throughput, measurements of engine vibration and shaft speeds are typically quantised. A number of data compression schemes will be investigated to ensure that sufficient information is retained to allow periods of abnormal engine behaviour to be identified using novelty detection techniques. The compression of flight information using difference encoding techniques as used in speech compression will be explored, such as Differential Pulse Code Modulation (or Delta Modulation), in which differences between sample values and predicted values are encoded, rather than actual sample values. Incorporating minimum and maximum vibration information at all shaft speeds into TORCH vibration signatures may further increase the robustness of novelty detection.

5.2 Improved Novelty Detection Methods

The identification of tracked orders at unexpected frequencies (i.e., tracked orders that occur at neither fundamental nor harmonic frequencies of shaft rotation) has been shown to provide indication of abnormal engine behaviour in Trent 500 data [24, 25]. It is proposed that a method be investigated for identification of the most significant of such unknown tracked orders in Trent 1000 engine vibration data, and how such information might be conveyed by a flight summary.

A model of the distribution of energy between tracked orders in an engine shaft has previously been used for novelty detection in Trent 500 data, and it is proposed that this be extended for use with Trent 1000 data.

As the Trent 1000 engine development programme proceeds, it is proposed that new methods of novelty detection be investigated to identify any engine-specific operational abnormalities that may arise. This is considered continuous work, involving the analysis of all Trent 1000 data as it becomes available (as described previously).

We propose applying methods of intra-engine analysis using TORCH signatures to Trent 900 and 1000 engine

data, initially with ground-based data from development test-beds, then on data acquired during aircraft flights. This consists of both full vibration recordings and flight summary data.

5.3 Time- and Time-Frequency Domain Analysis

Research to date has focused on the analysis of frequency domain spectral data. However, time domain data also contains information that may be exploited for the diagnosis of engine abnormality. A model of normal signal behaviour may be constructed using *auto-regressive moving-average* techniques [6, 9], deviations from which could identify abnormal conditions.

Time-frequency analysis of engine data combines both domains for a more complete representation of engine state, describing the evolution of the spectral content of a signal through time. Wavelet analysis [51] is suggested for constructing a model of normality in both frequency and time. The wavelet transform represents a signal in terms of a family of wavelet basis functions, shifted in time and scaled in frequency. These basis functions are typically irregular in shape, making them suitable for the analysis of non-stationary signals with sharp changes, and compactly supported in time, enabling localisation of a signal's features in time [36, 39, 78].

5.4 Data Fusion, and Oil Debris Monitoring Systems (ODMS)

Indicators of abnormal engine behaviour are often evident in many different types of data, such as vibration data, performance data, and ODMS information (in which engine oil content is analysed). This may be exploited by construction of models of normality that use multiple data types.

It is proposed that *Parzen window* [16, 57] methods be investigated for application to engine data, in which the unconditional probability distributions of normal data is approximated, allowing the probability of observed test values to be estimated. This method has been successfully applied to fusing data from patient physiological parameters in order to identify abnormality [71, 72].

The use of neural network based approaches to data fusion are also suggested, with a focus on *support vector machines* [77], which have been shown to be suitable for data fusion and novelty detection in many industrial applications [10, 24, 65, 73].

Robust novelty detection is also possible by combining classifications of engine data generated from models of normality of single types of data; e.g., combining classifications from vibration-based models, and performance-based models. Techniques proposed for application to engine monitoring include:

- Mixture-of-experts techniques, in which different “expert” networks are trained to classify patterns from different regions of input space [5, 8, 31, 32]. A “gating network” then combines the experts’ classifications of an input pattern, favouring those trained to localise on the region of input space that contains the input

pattern.

- Neural network ensembles [5, 34, 37, 49], in which *committees* of neural networks (usually trained independently of one another) are used to classify input patterns.
- AdaBoosting methods, in which large numbers of classifications from a set of classifiers (termed “weak” classifiers) are combined into a single “strong” classification [63]. These methods typically modify the training set during the training process by emphasising repeated training on patterns which are misclassified [56, 64].

5.5 On-Line Learning

Learning a model of normality from a “black box” system of which little or no prior information is available is proposed for further research, allowing system-specific models of normality to be constructed during operation, without supervision. Application of on-line learning systems to engines such as the Trent 900 and 1000 is suggested, allowing comparison of results with those obtained using conventional novelty detection methods, with the aim of producing a general novelty detection system that may be applied *ab initio* to any monitoring system.

Unlike sequential learning, training patterns in on-line learning are typically used to update the model when they become available, after which time the pattern is never seen again [68]. For applications in which off-line (or “batch”) learning is not applicable, new training techniques are required. Methods exist for the on-line training of neural networks [43, 47, 52], and of support vector machines [33]. The technique of “boosting” has also been shown to be effective in on-line training of classifiers [20]. On-line learning has been successfully applied to ensembles of neural networks [15].

It is proposed that methods for automatically selecting parameters from all those available for modelling be investigated, possibly using a sensitivity analysis approach, in which those signals most indicative of system condition are identified. On-line learning requires investigation of methods for automatically determining when model training should start and end, and how existing models of system normality may be updated as new data become available.

5.6 DPhil Project Plan

Work Package	Duration	2006				2007		
		Q1	Q2	Q3	Q4	Q1	Q2	Q3
Data Compression	26		■	■				
Improved Novelty Detection	26	■	■					
Time Domain Analysis	13			■				
Time-Frequency Domain Analysis	13			■				
Data Fusion	26				■	■		
On-Line Learning	39			■	■	■		
Thesis Writing	26						■	■

Figure 5.1: DPhil project plan, completion deadline October, 2007

Appendix A: Activities

Courses

- Managing Your DPhil** (Oxford, November 2004) Providing advice on the planning of DPhil research over three years and time-budgeting, this course was run by the Department of Engineering Science.
- Managing Team Conflicts** (Oxford, June 2005) This workshop in the Department of Engineering Science focused on management of an academic team, the resolution of conflicts between team members, and the structuring of a team's working relationships for maximum productivity.
- Poster Presentation Training** (Oxford, August 2005) This course gave training for the design and presentation of academic posters at conferences, and was organised by the Department of Engineering Science.
- Shape Analysis Training** (Derby, 2005) Organised and presented a day-long training course for Rolls-Royce engine analysts and developers on the use of Shape Analysis techniques for novelty detection in jet engines.
- Novelty Detection in Military Gas-Turbine Engines** (Bristol, October 2005) Presented results of novelty detection methods applied to the Eurofighter/Typhoon engine project to the engine development team.
- Professional Registration Training** (London, October 2005) An all-day IEE course providing support for the application of Chartered Engineer status.
- Engineering in the U.K.** (Beijing, China, October 2005) Held a seminar for graduate students in the Beijing Institute of Technology introducing the process of graduate research in U.K. universities, and describing research from current and previous positions.

Conferences and Seminars

- Rolls-Royce Annual Research Conference** (Derby, June 2005) Attended presentations included topics such as novelty detection in gas-turbine engines, identification of material failure in components, and thermodynamical modelling of gas-turbine engines.
- Industrial Applications of Artificial Intelligence and Expert Systems** (Bari, Italy, June 2005) A three-day residential conference with streams on novelty detection, statistical pattern recognition, machine learning, and visualisation.
- HECToR Quarterly Research Review** (Bath, December 2005) A review of research completed by the various universities and companies involved with the DTI-funded HECToR project for intelligent sensor systems, fault analysis, and novelty detection.
- Robotics Research Group Weekly Seminars** (Oxford, June 2004-present) Weekly presentations from researchers in fields relevant to the Robotics Research Group, attended seminars including: "*Object and Scene Recognition in Large Datasets*", "*Receiver Operating Characteristic Analysis*", "*Posterior Sampling Techniques*", "*Variational Bayesian Multinomial Probit Regression*", "*Probabilistic Non-linear PCA*", and "*Generative Probabilistic Approach to Visualizing High-Dimensional Data*".

Academic Papers

- Learning Shape for Jet Engine Novelty Detection** (Chengdu, China, May 2006) Submitted to the International Symposium for Neural Networks 2006.
- Application of an Intuitive Novelty Metric for Jet Engine Condition Monitoring** (Annecy, France, June 2006) Submitted to Industrial Applications of Artificial Intelligence and Expert Systems 2006.

Bibliography

- [1] Abu-Hakima S, Davidson P, Halasz M, Phan S. Jet Engine Technical Advisor (JETA). In *Proc. of the 2nd Int. Conf. on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems*, 1989.
- [2] Austin J, Davis R, Fletcher M, Jackson T, Jessop M. DAME: Searching Large Data Sets within a Grid-enabled Engineering Application. *Proc. of the IEEE*, 2005.
- [3] Austin J, Kennedy J, Lees K. The Advanced Uncertain Reasoning Architecture, AURA. In *Proc. of the Weightless Neural Network Workshop*, 1995.
- [4] Bannister P, Tarassenko L, Shaylor I. RR shape analysis tools. Technical report, Oxford University and Oxford BioSignals Ltd., Feb 2004.
- [5] Bishop C. *Neural Networks for Pattern Recognition*. Oxford, 1995.
- [6] Booth C, McDonald JR, Donald P, Lines N, Cooke N, Smith C. Model-based Condition Monitoring of Gas Turbine Engines for Power Generation Duty. *IEEE Power Engineering Review*, 2001.
- [7] Brown GV, Lawrence C, Manchala DW, Palazzolo AB, Klusman S, Kascak AF, Montague GT. In Operation Detection and Correction of Rotor Imbalance in Jet Engines Using Active Vibration Control. *SAE Journal of Aerospace*, 103(1), 1994.
- [8] Chen K, Xu L, Chi H. Improved Learning Algorithms for Mixture of Experts in Multiclass Classification. *Neural Networks*, 1999.
- [9] Chiras N, Evans C, Rees D. Nonlinear Gas Turbine Modeling Using NARMAX Structures. *IEEE Transactions on Instrumentation and Measurement*, 2001.
- [10] Clifton LA, Yin H, Zhang Y. Support Vector Machine in Novelty Detection for Multi-channel Combustion Data. In *Proc. for the 3rd Int. Symposium on Neural Networks*, 2006.
- [11] Collinge K, Schoff K. TEXMAS - An Expert System for Gas Turbine Engine Diagnosis and More. *SAE Aircraft Gas Turbine Engine Monitoring Systems*, 1988.
- [12] DePold HR, Gass FD. The Application of Expert Systems and Neural Networks to Gas Turbine Prognostics and Diagnostics. *ASME Journal of Engineering for Gas Turbines and Power*, 116, 1994.
- [13] DeRidder D, Duin RPW. Sammon's Mapping Using Neural Networks: A Comparison. *Pattern Recognition Letters*, 1997.
- [14] Dietz WE, Kiech EL, Ali M. Pattern-based Fault Diagnosis Using Neural Networks. In *Int. Conf. on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems, Tullahoma, TN, USA*, volume 1, 1988.
- [15] Dimitrakakis C, Bengio S. Online Policy Adaption for Ensemble Classifiers. In *Proc. of the European Symposium on Artificial Neural Networks*, 2004.
- [16] Duda RO, Hart PE, Stork DG. *Pattern Classification*. Wiley, 2 edition, 2001.
- [17] Evans C, Rees D, Hill D. Frequency Domain Identification of Gas Turbine Dynamics. *IEEE Transactions on Control Systems Technology*, 1998.
- [18] Fayyad U, Piatetsky-Shapiro G, Smyth P, Uthurusamy R. *Advances in Knowledge Discovery and Data Mining*. AAAI Press / MIT Press, 1996.
- [19] Foster I, Kesselman C, Nick J, Tuecke S. Grid Services for Distributed System Integration. *Computer*, 35(6), 2002.
- [20] Freund Y, Schapire RE. A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting. *Journal of Computer and System Sciences*, 1997.
- [21] Ganguli R, Chopra I, Haas DJ. Helicopter Rotor System Health Monitoring Using Numerical Simulations and Neural Networks. In *AIAA Modeling and Simulation Technologies Conference, New Orleans, USA*, 1997.
- [22] Greitzer FL, Kangas LJ, Terrones KM, Maynard MA, Wilson BA, Pawlowski RA, Sisk DR, Brown NT. Gas Turbine Engine Health Monitoring and Prognostics. In *International Society of Logistics Symposium*, 1999.
- [23] Halasz M, Davidson P, Abu-Hakima S, Phan S. *Applied Intelligence*. Springer, 1991.
- [24] Hayton P, Scholkopf B, Tarassenko L, Anuzis P. Support Vector Novelty Detection Applied to Jet Engine Vibration Spectra. In *NIPS Proceedings*, 2000.
- [25] Hayton P, Tarassenko L. Condition Monitoring of Engine Data. *Proc. of Royal Society*, In Press.
- [26] Hayton P, Utete S, Tarassenko L. QUOTE Project Final Report. Technical report, University of Oxford, 2003.
- [27] HercWatch. AeroSupport HercWatch Product Specification, 2003.
- [28] Ho N, Lozano P, Mangoubi R, Martinez-Sanchez M. A Model-based Vehicle Health Monitoring System for the Space Shuttle Main Engine. In *34th AIAA/ASME/SAE/ASEE Joint Propulsion Conference, Cleveland, Ohio, USA*, 1998.

- [29] Illi OJ, Greitzer FL, Kangas LJ, Reeve TJ. An Artificial Neural Network System for Diagnosing Gas Turbine Engine Fuel Faults. In *48th Meeting of Mechanical Failures Prevention Group*, 1994.
- [30] Jackson T, Austin J, Fletcher M, Jessop M. Delivering a Grid Enabled Distributed Aircraft Maintenance Environment (DAME). In *UK e-Science Proceedings*, 2003.
- [31] Jacobs RA, Jordan MI, Nowlan SJ, Hinton GE. Adaptive Mixtures of Local Experts. *Neural Computation*, 1991.
- [32] Jordan MI, Jacobs RA. Hierarchical Mixtures of Experts and the EM Algorithm. *Neural Computation*, 1994.
- [33] Kivinen J, Smola AJ, Williamson RC. Online Learning with Kernels. *IEEE Trans. on Signal Processing*, 2004.
- [34] Krogh A, Vedelsby J. Neural Network Ensembles, Cross Validation, and Active Learning. *Advances in Neural Information Processing Systems*, 1995.
- [35] Lieuwen TC. *Investigation of Combustion Instability Mechanisms in Premixed Gas Turbines*. PhD thesis, Georgia Institute of Technology, 1999.
- [36] Liu S, Du R, Yang S. An Improved Algorithm for Wavelet Packets and its Applications to Vibration Diagnosis in Diesel Engines. *Trans. of Chinese Soc. for Internal Combustion Engines*, 2000.
- [37] Liu Y, Yao X. Ensemble Learning Via Negative Correlation. *Neural Networks*, 1999.
- [38] Lowe D, Tipping ME. Feed-forward neural networks and topographic mappings for exploratory data analysis. *Neural Computing and Applications*, 1996.
- [39] Luo GY, Osypiw D, Irle M. On-Line Vibration Analysis with Fast Continuous Wavelet Algorithm for Condition Monitoring of Bearing. *Journal of Vibration and Control*, 2003.
- [40] McGrogan N. Prototype Feature Detection. Technical report, University of Oxford, Oxford BioSignals Ltd., Rolls-Royce Plc., 2005.
- [41] Merrington GL. Fault Diagnosis in Gas Turbines Using a Model-based Technique. *ASME Journal of Engineering for Gas Turbines and Power*, 116, 1994.
- [42] Meyer CM, Maul WA, Dhawan AP. SSME Parameter Estimation and Model Validity Using Radial Basis Function Neural Networks. In *Proc. of the 1994 Conf. on Advanced Earth to Orbit Propulsion Technology*, 1994.
- [43] Musavi MT, Ahmed W, Chan KH, Faris KB, Hummels DM. On the Training of Radial Basis Function Classifiers. *Neural Networks*, 1992.
- [44] Nabney I. *Netlab: Algorithms for Pattern Recognition*. Springer, 2002.
- [45] Nairac A, Corbett-Clark T, Ripley R, Townsend N, Tarassenko L. Choosing an Appropriate Model for Novelty Detection. In *IEE 5th International Conference on Artificial Neural Networks*, 1997.
- [46] 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.
- [47] Napolitano MR, Kincheloe M. On-Line Learning Neural Network Controllers for Autopilot Systems. *Journal of Guidance, Control and Dynamics*, 1995.
- [48] Ong M, Ren X, Allan G, Kadiramanathan V, Thompson HA, Fleming PJ. Decision Support System on the Grid. In *International Conf. on Knowledge-based Intelligent Information and Engineering Systems, New Zealand*, 2004.
- [49] Opitz D, Maclin R. Popular Ensemble Methods: An Empirical Study. *Journal of Artificial Intelligence Research*, 1999.
- [50] Pena JM, Famili F, Letourneau S. Data Mining to Detect Abnormal Behavior in Aerospace Data. In *KDD*, 2000.
- [51] Percival DB. An Introduction to Spectral Analysis and Wavelets. In *Proc. of Int. Workshop on Advanced Mathematical Tools in Metrology*, 1993.
- [52] Platt J. A Resource-Allocating Network for Function Interpolation. *Neural Computation*, 1991.
- [53] Polycarpou MM. An On-line Approximation Approach to Fault Monitoring, Diagnosis, and Accommodation. *SAE Journal of Aerospace*, 103(1), 1994.
- [54] Powers WT, Cooper AE, Wallace TL, Buntine WL, Whitaker K. Space Shuttle Main Engine Plume Diagnostics: OPAD Approach to Vehicle Health Monitoring. *SAE Journal of Aerospace*, 102(1), 1993.
- [55] Press WH, Vetterling WT, Teukolsky SA. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge, 1997.
- [56] Ripley BD. *Complements to 'Pattern Recognition and Neural Networks'*. Cambridge University Press, 1996.
- [57] Ripley BD. *Pattern Recognition and Neural Networks*, chapter 6. Cambridge University Press, 1996.
- [58] Rolls-Royce Technical Publications Dept. *The Jet Engine*. Renault Printing, 1996.
- [59] SAE. Aircraft Gas Turbine Engine Monitoring System Guide. Technical report, Aerospace Recommended Practice, 1981.
- [60] SAE. A Guide to Aircraft Turbine Engine Vibration Monitoring Systems. Technical report, Aerospace Information Report, 1986.
- [61] SAE. A Guide to the Development of a Ground Station for Engine Condition Monitoring. Technical report, Aerospace Information Report, 1994.
- [62] Sammon JW. A non-linear mapping for data structure analysis. *IEEE Transactions on Computers*, 1969.
- [63] Schapire RE. A Brief Introduction to Boosting. In *Proc. of the 16th Int. Conf. on Artificial Intelligence*, 1999.
- [64] Schapire RE. The Boosting Approach to Machine Learning - An Overview. In *MSRI Workshop on Nonlinear Estimation and Classification*, 2002.

- [65] Scholkopf B, Williamson RC, Smola AJ, Shawe-Taylor J, Platt JC. Support Vector Method for Novelty Detection. *Advances in Neural Information Processing Systems*, 2000.
- [66] Simani S. Identification and Fault Diagnosis of a Simulated Model of an Industrial Gas Turbine. *IEEE Transactions on Industrial Informatics*, 2005.
- [67] Subbu R, Goebel K, Frederick DK. Evolutionary Design and Optimization in Aircraft Controllers. *IEEE Transactions on Systems, Man, and Cybernetics*, 2005.
- [68] Tarassenko L. *A Guide to Neural Computing Applications*. Arnold, 1998.
- [69] Tarassenko L, Nairac A, Townsend N, Buxton I, Cowley P. Novelty Detection for the Identification of Abnormalities. *International Journal of Systems Science*, 2000.
- [70] Tarassenko L, Nairac A, Townsend N, Cowley P. Novelty Detection in Jet Engines. In *IEE Colloquium on Condition Monitoring, Images, External Structures and Health*, 1999.
- [71] Tarassenko L, Townsend N, Clifford G, Burton J, Price J, Cove S, Braithwaite E. Multi-parameter Monitoring for Early Warning of Patient Deterioration. In *Proc. for the IEE Conf. on Medical Applications of Signal Processing*, 2002.
- [72] Tarassenko L, Townsend N, Clifford G, Mason L, Burton J, Price J. Neural Networks and Data Fusion for the Next Generation of Patient Monitors. In *Natural Computing Applications Forum*, 2001.
- [73] Tax D, Ypma A, Duin R. Support Vector Data Description Applied to Machine Vibration Analysis. In *Proc. of 5th Annual Conf. of the Advanced School for Computing and Imaging, Heijen, NL*, 1999.
- [74] Trammel C, Vossler G, Feldmann M. UK Ministry of Defence Generic Health and Usage Monitoring System. *American Helicopter Society Annual Forum*, 3, 1997.
- [75] Tumer IY, Bajwa A, Los Angeles C. A Survey of Aircraft Engine Health Monitoring Systems. In *AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, 1999.
- [76] Turney P, Halasz M. Contextual Normalization Applied to Aircraft Gas Turbine Engine Diagnosis. *Journal of Applied Intelligence*, 1993.
- [77] Vapnik VN. *The Nature of Statistical Learning Theory*. Wiley, 1998.
- [78] Wang L, Yin H. Wavelet Analysis in Novelty Detection for Combustion Image Data. In *Proc. of the 10th CACSC, Liverpool*, 2004.
- [79] Webb A. *Statistical Pattern Recognition*. Wiley, 2002.
- [80] Winston H, Clark R, Buchina G. An Application of Model-based Reasoning to Gas Turbine Diagnostics. *AI Magazine*, 1995.
- [81] Yu LJ, Cleary DJ, Cuddihy PE. A Novel Approach to Aircraft Engine Anomaly Detection and Diagnostics. In *19th Conf. of Uncertainty in Artificial Intelligence*, 2003.