

Machine learning techniques for the health monitoring of rotating machinery in nuclear power plants

Jason J. A. Costello

A thesis submitted for the degree of Doctor of Engineering to
Institute for Energy and Environment
University of Strathclyde

2019

The copyright of this thesis belongs to the author under the terms of the United Kingdom Copyright Acts as qualified by University of Strathclyde Regulation 3.50. Due acknowledgement must always be made of the use of any material contained in, or derived from, this thesis.

Signed:

Date:

Abstract

This thesis explores the development of data-driven and machine learning methods in application to the health monitoring of rotating plant items being used in the primary and secondary cycles of the Advanced Gas-cooled Reactor (AGR) nuclear power plants in the UK. The methods fall broadly into two categories: the statistical augmentation of a pre-existing knowledge-based system for turbine generator vibration alarm analysis, and the development of a machine learning model for the exploration of long-term predictive measures of asset health for AGR gas circulator units. Both of these topics are unified in their engineering context, and the overall aim of the approaches employed: to provide improved decision support using data to reliability staff tasked with monitoring key nuclear assets.

A self-tuning methodology for knowledge-based system parameterisation and data selection in rotomachinery vibration monitoring is introduced, providing a comparative study of numerous methods and case studies for features of interest in both steady-state and step change conditions. These approaches were developed using a historical dataset taken from a turbine generator in use at an AGR, with time series streams from multiple component channels.

An event-driven approach to asset health is presented, utilising a support vector machine & logistic regression hybrid model to estimate particular states of interest associated with the gas circulator duty cycle. This approach to health monitoring (examining responses during semi-regular refuelling events) is shown to correlate highly with the remaining useful life of a circulator unit which eventually underwent an unexpected failure, and provides a potential quantitative metric for preventing repeat instances.

Acknowledgements

Firstly I'd like to express my gratitude to Prof. Stephen McArthur and Dr. Graeme West. Their supervision, guidance and patience(!) has been central to this work. Prior to joining Steve's team, I had no concept of how data could be used to solve problems: this influence has shaped much of my career since. Graeme's thoughtful discussions and continued encouragement long after my disappearance from the department are a huge part of what ensured this thesis eventually came into existence. Thank you both.

Thanks are also due to EDF Energy, who provided invaluable data and funding for this research. Thanks to my industrial supervisors Graeme Campbell and Chris Gilroy, who both gave useful insights into the work of reliability engineers and vibration monitoring.

My family have always been a limitless source of encouragement for everything I do, so thank you Mum, Dad & Hammy (and Tully, Dexter and Padraig), and to my Taylor family Poppy & Mary.

I think all other halves of doctoral candidates have to endure some level of interrupted service from their partner, but my wife Claire deserves special credit for putting up with my 'process' for so long. We've had thesis writing sessions in Italian hotel rooms, London studio flats, and even at 70mph on the motorway. You've always been there for me, so thank you Topo.

Final thanks go to our Scottish Terrier Macallan, who joined the family just in time to see this document completed. Good timing, wee man.

Fir' a moose

Contents

List of figures	xiv
List of tables	xv
List of acronyms	xv
1 Introduction	1
1.1 Introduction to the research	1
1.2 Contributions from the research	3
1.3 Overview of the thesis	4
1.4 Associated publications	5
2 Condition monitoring of nuclear-context rotating plant	7
2.1 Nuclear energy	7
2.1.1 Background	7
2.1.2 Nuclear fission	10
2.1.3 UK nuclear programme	12
2.2 Advanced Gas-cooled Reactor (AGR)	14
2.2.1 Reactor core	15
2.2.2 Primary cycle	16
2.2.3 Secondary cycle	17
2.2.4 Associated rotating machinery	18
2.3 Rotating machinery condition monitoring	19

2.3.1	Historical perspective	20
2.3.1 a)	Drivers for monitoring	20
2.3.1 b)	Knowledge, experience and data	21
2.3.1 c)	Maintenance regimes	22
2.3.2	Vibration monitoring	23
2.3.2 a)	Vibration	23
2.3.2 b)	Measurements	25
2.3.2 c)	Transducers	26
2.3.3	Steam turbines	27
2.3.3 a)	Turbine stages	28
2.3.3 b)	Generator	30
2.3.4	Gas circulators	30
2.4	Modern condition monitoring	31
2.4.1	Hardware, storage and visualisation	31
2.4.2	Monitoring strategies	33
2.4.2 a)	Components, systems and complexity	33
2.4.2 b)	Alarms and notifications	35
2.4.2 c)	From diagnosis to prognosis	37
2.4.3	Intelligent condition monitoring	38
2.4.3 a)	Definition	38
2.4.3 b)	Types of intelligent system	39
2.4.3 c)	Applied to generation-based rotating machinery	40
2.4.3 d)	Applied to rotating machinery in other domains	41
2.4.3 e)	Applied to nuclear generation	41
2.4.3 f)	Applied in power and electrical scenarios . . .	42
2.4.3 g)	Other applications	42
2.5	Challenges and opportunities	43
2.5.1	Challenges	43

2.5.1 a)	Energy security	43
2.5.1 b)	An ageing nuclear fleet	44
2.5.1 c)	An ageing nuclear workforce	45
2.5.2	Opportunities	45
2.5.2 a)	Increased data availability	45
2.5.2 b)	Increased technique sophistication	46
3	Statistical methods and machine learning	47
3.1	Statistical methods	47
3.1.1	The Gaussian distribution	48
3.1.2	Probability density function	48
3.1.3	Cumulative distribution function	49
3.1.4	Empirical statistical analyses	50
3.1.4 a)	Kernel density estimation	51
3.1.4 b)	Empirical cumulative distribution function	56
3.1.5	Parametric vs. non-parametric	57
3.1.5 a)	Parametric models	57
3.1.5 b)	Non-parametric models	58
3.2	Machine learning fundamentals	58
3.2.1	Introduction	58
3.2.2	Key concepts	59
3.2.2 a)	Supervised learning	61
3.2.2 b)	Unsupervised learning	62
3.2.2 c)	Classification	63
3.2.2 d)	Regression	64
3.2.2 e)	Clustering	65
3.2.3	Deep learning	67
3.2.4	Overfitting	67
3.3	Machine learning approaches	68

3.3.1	Linear classification	68
3.3.2	Perceptron learning	70
3.3.3	Logistic model	71
3.3.4	Kernel methods	73
3.3.4 a)	Support vector machines	74
3.3.5	Ensemble methods	75
3.3.5 a)	Random forests	75
3.3.5 b)	Gradient boosting machines	76
3.3.6	Clustering	76
3.3.6 a)	k-means clustering	76
3.3.6 b)	Gaussian mixture models	77
3.3.6 c)	Hierarchical clustering	77
3.4	Machine learning in reliability	77
3.4.1	Kernel methods in CM	78
3.4.2	Other techniques	80
3.4.3	Potential	80
4	Self-tuning diagnostics in rotating machinery	82
4.1	Routine alarms	83
4.2	Knowledge-based system	85
4.2.1	Rule base	86
4.2.2	Signal-to-symbol transformation	89
4.2.2 a)	Extracting impulses	90
4.2.2 b)	Extracting trends	91
4.2.2 c)	Extracting steps	92
4.2.3	Channel and machine profiles	95
4.2.4	Proof-of-concept channel profiles	95
4.2.5	Retrospective	97
4.3	Learning the machine profile	98

4.3.1	Learning parameters	99
4.3.1 a)	Envelope-based	101
4.3.1 b)	Event-based	102
4.3.2	Data selection - envelope-based	102
4.3.2 a)	Unimodality	102
4.3.2 b)	Rolling KS-statistic	106
4.3.2 c)	Machine state phase space	108
4.3.2 d)	Extracting the dominant distribution	109
4.3.3	Data selection - event-based	110
4.3.3 a)	Standard deviation-based	111
4.3.3 b)	Density-based	113
4.4	Case study: Tuning channels	115
4.4.1	Channel profile vs. data	115
4.4.2	Tuning process and results	117
4.4.2 a)	Empirical tuning distributions	117
4.4.2 b)	Selecting parameters	118
4.5	Case study: typical step changes	123
4.5.1	Technique selection	123
4.5.2	Tuning process and results	125
4.5.2 a)	Empirically selected steps	125
4.5.2 b)	Selecting parameters	126
4.6	Discussion	129
5	A data-driven degradation model for circulators during refuelling	130
5.1	Problem definition	130
5.1.1	Low power refuelling	131
5.1.2	Transient monitoring	133
5.2	Strategy	134
5.2.1	Importance of the LPR	134

5.2.2	Data-driven model	135
5.3	Model selection - theory	135
5.3.1	<i>k</i> -fold cross validation	136
5.3.2	Hyperparameters	137
5.3.3	Classifiers	138
5.3.3 a)	Linear classifiers	138
5.3.3 b)	Higher-order classifiers	139
5.3.4	Dataset overview	139
5.3.5	Labelling LPR behaviour	142
5.3.6	Training data feature vector and dimensionality	143
5.4	Model selection - results	144
5.4.1	Grid search results	144
5.4.1 a)	Linear classifiers	144
5.4.1 b)	Higher-order classifiers	145
5.4.2	Discussion	148
5.5	Decision support	149
5.5.1	Identifying LPR state data	150
5.5.2	Phase space view	152
5.6	Predictive capabilities	156
5.6.1	Model output evolution	157
5.6.2	State estimation	157
5.6.2 a)	State classification	158
5.6.2 b)	RUL estimation	159
5.7	Discussion	163
6	Conclusions & future work	165
6.1	Conclusions	166
6.1.1	Turbine generator monitoring	166
6.1.2	Gas circulator monitoring	168

6.2	Future work	170
6.2.1	Routine alarm analysis	170
6.2.2	Gas circulator monitoring	172
6.2.3	Machine learning and reliability	172

List of Figures

2.1	Illustration of a nuclear fission reaction.	10
2.2	Civil nuclear sites in the UK, showing the AGR and PWR stations presently operational	13
2.3	Core photographs of Hinkley Point B, showing the nature of the in-core graphite brick structure	15
2.4	Overview of the primary cycle of the AGR core, with annotation of CO ₂ gas flow [Non96]	17
2.5	Primary and secondary loop of an AGR power plant	18
2.6	Single frequency signal component with annotated measurements	25
2.7	Vibration measurement example, showing a vibrating rotating element oscillating in the horizontal axis	27
2.8	Schematic of a steam turbine generator	28
2.9	Photograph of steam TG [NRC]	29
2.10	Gas circulator unit [Non96]	32
3.1	Three example PDFs, drawn from the Gaussian family	50
3.2	Example cumulative distribution function of Gaussian PDF	51
3.3	Four kernel functions commonly used in non-parametric statistics: uniform, triangular, Gaussian & Epanechnikov	52
3.4	Variation of the kernel bandwidth, with $h = \{0.1, 0.5, 1, 10\}$	55
3.5	Example machine learning process, where a program is constructed and learns to play board game Go	60

3.6	Two bivariate data relationships with a high likelihood of an underlying relationship, which would typically be quantified through regression	65
3.7	An example clustering problem, visualised in D^3 phase space	66
3.8	Supervised learning example (classification), with comparison of ideal and over-fitted decision functions	68
3.9	Typical linearly separable binary classification training domain in two dimensions	69
3.10	Comparison of standard discrete and logistic threshold functions	72
4.1	Example coupling of response and operational observable	84
4.2	Example feature in the response observable without a corresponding operational change	85
4.3	Inference diagram of typical rule chain used in the knowledge-based system	87
4.4	Example of impulse extraction, w.r.t. bounds of operation	90
4.5	Example of features and calculation used in extracting trends	91
4.6	Example of features and calculation used in extracting steps	93
4.7	Overview of the parameter tuning steps to build a machine profile	100
4.8	Example unimodal and multimodal distributions	103
4.9	Comparative illustration of ECDFs for unimodal and bimodal PDFs	104
4.10	Comparative illustration of ECDFs for unimodal and bimodal PDFs, with KS-statistic metric D_n annotated for each	105
4.11	Example channel time series, with rolling KS-statistic series	107
4.12	Distribution of KS-statistic values	107
4.13	Tuned distribution example	108
4.14	Vibration-load phase space with D_n filter, with crisp load boundary	109

4.15	Vibration-load phase space without D_n filter, with crisp load boundary	109
4.16	Overview of standard deviation-based identification of change-points in vibration time series data	112
4.17	Evolution of PDE at various stages of an example time series with step conditions	114
4.18	Rolling KS-statistic applied to example load profile	115
4.19	Time series and distribution of example overall amplitude subset, with annotated channel profile parameters	116
4.20	Time series and distribution of example overall amplitude subset, with annotated channel profile parameters (incorrect channel)	118
4.21	Empirical distributions for vibration observables	119
4.22	Empirical distributions for operational observables	120
4.23	Empirical kernel density estimates for operational observables, with annotated existing channel profile limits	120
4.24	Empirical kernel density estimates for vibration observables, with annotated existing channel profile limits	121
4.25	Density-based envelope setting of revised channel profile values	121
4.26	Illustration of FWHM taken from a density estimate	122
4.27	Standard deviation- & density-based changepoints on channel 9 time series	123
4.28	Standard deviation- & density-based mean step profiles	124
4.29	Standard deviation- & density-based hierarchical clustering . .	125
4.30	Upstep profiles for Turbine A vibration channels, extracted using the density-based methodf	127
4.31	Downstep profiles for Turbine A vibration channels, extracted using the density-based method	128

5.1	Typical generator load profile during AGR refuelling events with LPR highlighted	132
5.2	Typical vibration response to changing conditions during AGR refuelling events	133
5.3	k -fold cross validation procedure	137
5.4	Full (2006-2010) duty cycle of the GC used to construct the model, with zoomed LPR instance	140
5.5	Duty cycle of the selected LPRs used to train the model	141
5.6	Drive and non-drive end vibration phase space of the refuelling events	141
5.7	Kernel density estimate across the range of load values for the entire dataset	142
5.8	Mixture of Gaussians ($N = 3$) model of load data, identifying the bounds of the three load-based behaviours	143
5.9	CV grid search scores for each candidate linear model	146
5.10	Classification accuracy of RBF SVM with varying C and γ	147
5.11	Class vs. predicted output, with deltas shown in red	151
5.12	Load values of predicted class results	151
5.13	Load values of predicted class results for drive end vibration	152
5.14	Load values of predicted class results for non-drive end vibration	153
5.15	Confusion matrix of model output from time series values taken from periods of refuelling	153
5.16	Phase space of drive end vibration, with predicted labels	154
5.17	Phase space of drive end vibration, with predicted labels	155
5.18	Multiple segmentations of the operational data in drive end space, providing an abstraction of the vibration data for examination	156
5.19	Misclassification of the <code>online</code> behaviour with operation	157

5.20	Most probable state estimates for each of the refuelling behaviours in chronological order	159
5.21	Remaining useful life compared with probabilistic output of LPR model for late class	161
5.22	Remaining useful life compared with probabilistic output of LPR model for mid2 class	162

List of Tables

2.1	Nuclear generation from selected nations [WNA19]	8
2.2	Estimated emissions (tonnes of CO ₂ per GWh) from generation [MHA18]	9
3.1	Selection of kernel functions	53
4.1	Definitions of time series primitives	89
4.2	Operational observable profile	96
4.3	Channel 5 (front LP-stage A bearing) machine profile	96
4.4	Channel 6 (rear LP-stage A bearing) machine profile	96
4.5	Channel 9 (front LP-stage C bearing) machine profile	96
4.6	Identified steps per method	123
4.7	Turbine A overall level steps	128
4.8	Turbine A first order magnitude steps	129
4.9	Turbine A second order magnitude steps	129
5.1	CV and test scores for linear classifiers	145
5.2	CV and test scores for RBF SVM diagonal	148
5.3	Correlation between late temporal class & true RUL	163

List of acronyms

AGR	Advanced gas-cooled reactor
AI	Artificial intelligence
CDF	Cumulative distribution function
CM	Condition monitoring
ECDF	Empirical cumulative distribution function
FFT	Fast Fourier transform
GC	Gas circulator
HMM	Hidden Markov model
KBS	Knowledge-based system
KDE	Kernel density estimation
ML	Machine learning
PDF	Probability density function
PHM	Prognostics and health monitoring
PWR	Pressurised water reactor
RVM	Relevance vector machine
ROMAAN	Rotating machinery alarm analyst
SCADA	Supervisory control and data acquisition
SVM	Support vector machine
TG	Turbine generator
XML	eXtensible markup language

Chapter 1

Introduction

1.1 Introduction to the research

Condition monitoring (CM) has a long and distinguished history as a scientific and engineering discipline. Since the industrial revolution in the late 18th century, the world has become increasingly reliant on machinery and automated processes. Industries such as transportation, manufacturing, construction and energy production all require almost innumerable numbers of mechanical, chemical and electrical processes and machines. Accordingly, the proper maintenance and analysis of these assets represents an important area for the organisations and practitioners in these fields, helping ensure that systems remain operational without the risk of unexpected failures and outages.

The supply of reliable and secure energy in today's technologically-centred economy is a commodity more important [Bra10, YM13, CI17] than ever before. With the level of consumption steadily rising [Ama14, IEA17], keeping asset availability high is critical in meeting the market requirements from both industry and consumers. Coupled with the ever-increasing threats from climate change, a common approach for nation states to consider is to build a diverse generation mix [PK15], of which nuclear energy can play a significant

role.

Nuclear generation in the UK has a long legacy and potentially expansive future as a part of the energy strategy of the country. Irrespective of the next stages for prospective new builds and developments in next-generation nuclear technologies, the pre-existing fleet in the UK continues to provide 18% [WNA18a] of the energy cross-section, and requires continued attention in the areas of operational efficacy and safety. A premium is placed on reliability and sound quantitative decision making by the operator on all of the primary, secondary and auxiliary processes involved in the plants, given the importance of maintaining output economically and safely.

Among the primary and secondary cycle plant items for nuclear operators, turbine generators (TGs) and gas circulators (GCs) are instrumented to survey a variety of data streams for the purposes of monitoring behaviour and performance empirically. These records are well integrated into the reliability engineering workflow: regulatory requirements exist [iso09] dictating acceptable levels of vibration for observables such as amplitude time series and frequency components. Data-based resources for the reliability engineer are also increasing in availability and sophistication, corresponding analogously with the 'big data' [Loh12] revolution seen elsewhere in technology trends.

While there has been historical successes [Ran04, Tav08] in applying quantitative and automated techniques to common CM problems within rotating machinery CM, there remains analyses still reliant on manual legacy approaches to extract diagnostics from machine data. There is a wealth of useful information to be derived from the steadily improving data view of key rotomachinery systems in contexts such as nuclear generation, and the increased adoption of data-driven techniques can take advantage of this - resulting in improvements to machine availability and generation reliability. However, the necessity of clear decision-making and accountability when considering safety critical as-

sets (such as primary cycle nuclear machines like GCs) means that any data-driven inference needs to be provided in a concise and explicable manner to engineers who may not be domain experts in the fields of data science & intelligent systems. This has been a recognised issue [PCLH15] for the application of so-called ‘black box’ techniques to problems in various domains. Any methods employed should provide clear explanations or employ intuitive visualisation techniques when presenting decision support to the end-user.

This thesis demonstrates several techniques which exploit the increased data resource for nuclear rotomachinery reliability engineering; in particular approaches from the discipline of *machine learning* (ML); while considering the importance of the presentation of clear reasoning. Two suites of approaches are demonstrated; a self-tuning augmentation of an existing knowledge-based intelligent system for alarm analysis on TGs, and a data-based model of GC re-fuelling for predictive health monitoring purposes. From these examples, the value of robust intelligent system methodologies in data-rich but explicability-focused scenarios is presented.

1.2 Contributions from the research

The key novel contributions from this work can be summarised as:

- Augmentation of an existing knowledge-based intelligent system with ML and statistical inference techniques, providing an improved hybrid intelligent system tackling the engineering problem of routine alarm analysis in TGs,
- A self-tuning framework for vibration diagnostics, allowing for the application of a routine alarm knowledge base across an entire asset family under a single maintenance regime,
- The use of techniques in statistical inference to automatically define pe-

riods of system normality and transient behaviour in rotomachinery vibration data,

- Construction of a data-driven classification model with the ability to accurately label historical periods of refuelling events from the vibration response data,
- Development of an empirical model mapping event state to remaining useful life, providing predictive metrics to anticipate future failures in GC units,
- Presentation of a new GC phase space view for repeated asset events, which visually provides feedback on typical vibration characteristics across the machine,

Alongside these research outcomes, a number of industrial deliverables make up the contributions from the project. Much of the work outlined has been developed into a prototype vibration analysis toolkit (the Rotating Machinery Alarm Analyst, ROMAAN, system), which was demonstrated on data taken from machines at two UK nuclear power plant sites.

1.3 Overview of the thesis

This thesis opens with two chapters covering important areas of background and contextual information. Chapter 2 discusses the history and present day setup of nuclear generation in UK, before introducing the rotating asset class within this context. The chapter goes on to provide an in-depth overview of the state-of-the-art in rotating machinery CM, before concluding with the opportunities where improvements can be made using intelligent systems techniques.

Chapter 3 introduces the discipline of machine learning, providing a broad overview of the field, relevant techniques and its history of application in sim-

ilar engineering areas. It then concludes by presenting the argument for ML solutions to the engineering problems introduced in the previous chapter.

Chapter 4 introduces the first novel contribution from the research: a data-driven augmentation of an existing knowledge-based system for use in the vibration analysis of rotating machinery. The ‘self-tuning’ framework underpinning this is discussed in detail, providing insight into the development and reasoning behind the selected approaches. Several empirical and statistical methods for time series analysis were developed as part of this section, which are all also introduced and demonstrated.

Chapter 5 presents a novel data-driven model for examining refuelling events in GC units. This section illustrates the process of constructing and testing this model utilising a variety of ML candidate techniques, before discussing the potential for its application in long-term GC health monitoring.

Finally, Chapter 6 concludes the thesis by summarising the two strands of work, discussing their overlaps and recommending on potential next stages for this research.

1.4 Associated publications

- J. J. A. Costello, G. M. West and S. D. J. McArthur, “Machine learning model for event-based prognostics in gas circulator condition monitoring”, *IEEE Transactions on Reliability*, 66(4), pp 1048-1057 (2017)
- J. J. A. Costello, G. M. West, S. D. J. McArthur and G. Campbell, “Self-tuning routine alarm analysis of vibration signals in steam turbine generators”, *IEEE Transactions on Reliability*, 61(3), pp 731-740 (2012)
- J. J. A. Costello, G. M. West, S. D. J. McArthur and G. Campbell, “Investigation of gas circulator response to load transients in nuclear power plant operation”, *The Eighth American Nuclear Society International Topical Meet-*

ing on Nuclear Plant Instrumentation, Control and Human-Machine Interface Technologies, San Diego CA., USA (2012)

- J. J. A. Costello, G. M. West, S. D. J. McArthur and G. Campbell, "Self-tuning diagnosis of routine alarms in rotating plant items", *The Eighth International Conference on Condition Monitoring and Machinery Failure Prevention Technologies*, Cardiff UK, (2011)
- V. M Catterson, J. J. A. Costello, G. M. West, S. D. J. McArthur, C. W. Wallace, "Increasing the adoption of prognostic systems for health management in the power industry", *Chemical Engineering Transactions*, 33(3), pp. 271-276 (2013)
- C. J Wallace, J. J. A. Costello, G. M. West, S. D. J. McArthur, M. Coghlan, "Integrated condition monitoring for plant-wide prognostics", *Chemical Engineering Transactions*, 33(3), pp. 859-864 (2013)

Chapter 2

Condition monitoring of nuclear-context rotating plant

This chapter introduces the main setting of the research in this thesis: namely the Advanced Gas-cooled Reactor (AGR) and its associated rotating assets that ensure the effective and safe operation of this reactor design. Alongside this, the concept of condition and health monitoring for engineering assets is presented.

2.1 Nuclear energy

2.1.1 Background

Nuclear fission as a source of exploitable energy was proposed in the first third of the 20th century, with physicist Enrico Fermi heading the development of the first fission-based reactor - the Chicago Pile 1, or CP-1 - in 1942 [Fer46]. This was shortly followed by the Experimental Breeder Reactor 1, or EBR-1, which marked the first instance nuclear fission was used to produce electricity, providing power to a number of everyday lightbulbs in December 1951 [Mic01].

Since then nuclear power has evolved into an important modern generation

Table 2.1: Nuclear generation from selected nations [WNA19]

<i>Country</i>	<i>Current</i>			<i>Planned</i>	
	<i>No. reactors (NPPs)</i>	<i>Net MWe</i>	<i>%</i>	<i>No. reactors</i>	<i>Gross MWe</i>
Canada	19 (7)	13553	14.6	2	1500
China	45 (23)	42976	3.9	43	50900
France	58 (19)	63130	71.6	1	1720
Germany	7 (3)	9444	11.6	0	-
India	22 (7)	6219	3.2	14	10500
United Kingdom	15 (8)	8883	19.3	3	5060
United States	98 (60)	99376	20.0	14	3100

means, with around 450 civil nuclear reactors around the globe supplying an estimated 11% [WNA18b] of the world’s electricity as of 2018. Thirty countries [WNA18b] have active civil nuclear programmes. A selection of some of the major countries utilising nuclear power plants (NPPs) is provided in Table 2.1, providing details of their operable nuclear capacity in each state along with the confirmed plans for future reactors . Note that the percentages given indicate the proportion of the total generation mix in that nation.

Even Germany, where the governing authority have elected to rescind plans for future atomic energy development, has an existing NPP-based function that sees nuclear-based generation continuing for the next twenty years as the current capacity is phased out (Neckarwestheim 2 plant is scheduled for decommissioning in 2036 [WNA13]). As shown by these figures, the ongoing operation and management of atomic energy is important (at least in the short-to medium-term) for a large proportion of the major economic powers.

This international interest in nuclear can be attributed to a variety of factors, but one of the most overriding considerations for future generation capacity (of any format) is its associated ‘*carbon footprint*’; the estimated rate at which the construction, operation and decommissioning of a generation facility expels CO₂ and other greenhouse gases into the Earth’s atmosphere. The reduction of carbon emissions was identified by the United Nations in the seminal Kyoto Protocol [dB08] as of great importance in tackling climate change, with thirty-

Table 2.2: Estimated emissions (tonnes of CO₂ per GWh) from generation [MHA18]

<i>Generation type</i>	<i>Emissions</i>		
	2015	2016	2017
Coal	909	931	918
Gas	382	378	357
All fossil fuels	625	497	460
All fuels (incl. nuclear & renewables)	335	265	225

seven parties (including the European Union and its 28 states) agreeing to meet binding targets in reducing their output by 2020.

This has rendered management of emissions a massive issue for industry as a whole; often with the tender of projects being decided on the expected greenhouse gas expulsion rate [UKC08] of the proposal. Within energy, this is made particularly pertinent as much of the existing conventional capacity (coal-fired, natural gas) have an ongoing associated emission rate when producing power. Table 2.2 shows the marked difference in emission ratios between conventional approaches in coal and gas when compared against a mixed including nuclear and renewable solutions. The inclusion of nuclear and renewables markedly reduces the CO₂ output per GWh unit generated. Coupled with the existing base-load penetration of nuclear across many developed nations, maintaining and developing the sector is a key point for many national administrations. From a more quantitative perspective, studies over the previous decade [POS11, PAH⁺17] have consistently placed nuclear among the best performing low-carbon technologies available to the UK moving forward.

Each of these points contribute to the fact that existing and new build nuclear capacity has an important role to play in addressing global energy requirements.

2.1.2 Nuclear fission

While numerous NPP designs exist, the general principle surrounding their operation can be considered fairly uniform. The process of energy generation comes from the fissile¹ material (commonly Uranium-235 or Plutonium-239) contained in the reactor core undergoing radioactive decay into fission products. From a first principles perspective, the process occurs when a large atomic nucleus is stimulated and decomposes into multiple smaller component atomic nuclei, an amount of free neutrons and energy in the form of gamma radiation [SF07]. The difference between the binding energies of the larger parent nucleus and the resultant daughter nuclei accounts for emitted energy. Fig. 2.1 provides a schematic example of a fission process.




Figure 2.1: Illustration of a nuclear fission reaction.

The free neutrons previously contained in the larger nucleus go on to interact with neighbouring atomic nuclei; stimulating them to decay in a similar manner and in turn create their own fissile products, thus propagating the reaction on. This is commonly referred to as a chain reaction in the nuclear literature; with the major aim being to maintain a self-sustaining but control-

¹The definition of 'fissile' refers to the ability of a material to maintain a fission reaction, whereas 'fissionable' materials can undergo nuclear fission but their low probability of neutron emission makes them unsuitable for sustaining a reaction.

lable chain reaction, or *criticality*. The day-to-day operation of a nuclear facility is largely concerned with the maintenance and management of the conditions favourable to the ongoing criticality.

In order to optimise the absorption rate of produced free neutrons in fission reactions by nearby atomic nuclei, the kinetic energy of the neutrons is reduced through means of a material known as the *moderator*. This slows high energy neutrons down to a value corresponding to a greater efficiency of reaction likelihood, as faster neutrons often have a smaller neutron absorption coefficient in fissile materials. The slowed neutrons, or *thermal neutrons*, are brought to lower energies by moderators including deuterium oxide ($^2\text{H}_2\text{O}$), a heavy isotope of water, or materials such as graphite. The moderator is normally the major enclosing structure or fluid of the reactor core. For example, the Pressurised Water Reactor (PWR) design submerses the fuel assemblies in a deuterium oxide fluid moderator. Each moderator has individual characteristics and requirements that directly affect the nature of the NPP, making many designs of reactor quite different due to their selection of this material.

As the chain reaction continues, the fissile material eventually depletes in atomic structure to an isotope unable to maintain a continued fission process, moving from a majority of fissile to fissionable. This creates radioactive isotopes of the fuel that are no longer useful for generation, or 'waste products'. The spent fuel requires to be stored safely for a sufficient length of time in order to allow the long term radioactivity of the material to fall to a safe environmental level. Management of spent fuel and waste products in the nuclear fuel cycle represents a point of contention for the nuclear industry, with a number of solutions being brought forward in order to address such issues.

2.1.3 UK nuclear programme

The UK began its development in civil nuclear energy with the formation of the United Kingdom Atomic Energy Authority (UKAEA) in 1954 [BBC54], which paved the way for the opening of Calder Hall A NPP in 1956. This represented the world's first reactor generating at rate consistent with (then) commercial requirements [JE56]. Calder Hall, along with numerous sites across the UK, was of the MAGNOX design: a British endeavour investigating the use of CO₂ gas and graphite material as the reactor coolant and moderator respectively. MAGNOX NPPs utilised enriched Uranium as their primary fuel source, which also allowed for the creation of weapons-grade materials as a by-product of the nuclear fuel cycle.

Towards the conclusion of the MAGNOX site licenses, the UK industry began work on the design of its successor. The Advanced Gas-cooled Reactor (AGR) is an evolution from the first generation MAGNOX, taking many of the principles of its predecessor and introducing technological improvements. Both the moderator and coolant remain the same, however the AGR's reactor temperature is higher than that of the MAGNOX. This allows for greater thermal efficiency [SF07] comparatively.

The first AGR-type facility was constructed for experimental and testing purposes at Windscale in 1963 [Nat13]. Development of the first of the commercial AGR (note that there are ad-hoc features unique to each reactor site) began in 1965, at Dungeness B. The first operational AGRs began generation in 1976 at the twinned design sites of Hunterston B and Hinkley Point B, followed by sites at Heysham (two NPPs), Hartlepool and Torness. Alongside this, a single PWR was commissioned to complement the majority gas-cooled capacity, with a single site opened at Sizewell in 1995. The prominence of the AGR as the main UK reactor design is the reason for the focus on this reactor type in this research.

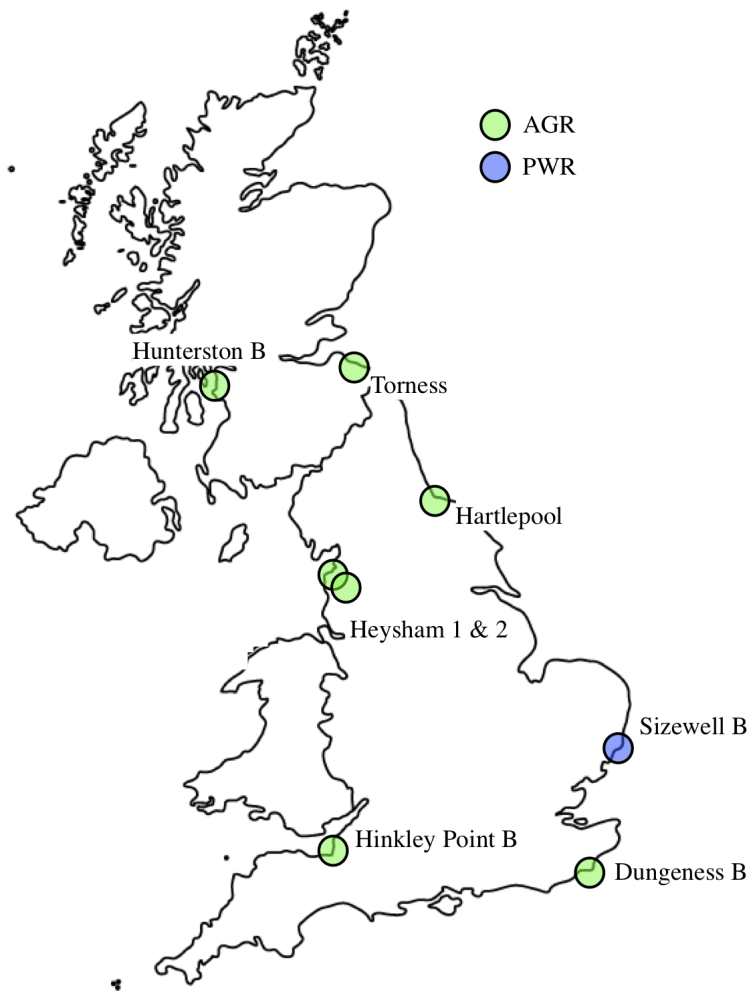


Figure 2.2: Civil nuclear sites in the UK, showing the AGR and PWR stations presently operational

An illustrative overview of the currently operational (as of 2019) UK civil nuclear sites is provided in Fig. 2.2, with two reactors operating on each AGR site and a single reactor on each of the PWR/MAGNOX sites. In total, the UK operates 16 nuclear units for civil energy purposes, with a contribution of 9,231MW at around 18% of the generation used nationwide [NEI13]. With this substantial share to the energy mix, the existing nuclear builds across the British Isles represent a vital base-load asset. The lifetime of the AGR design has been extended at various stages [Bri05] [NEI12], in line with continual assessment and improvements to the AGR.

For the reasons outlined in the Section 2.1.1, this existing capacity looks

set to continue operating well beyond its original life expectancy and be augmented through development and investment in future NPPs. The operational life of the AGR design has been extended on regulatory review and research interest [WWJ⁺10, BRTL11] beyond its initial sixty year tenure. Furthermore, as shown in Table 2.1, at least four new build reactors were confirmed at the time of writing. Further to this, the UK government has outlined detailed plans for around 16GW more nuclear-driven capacity by 2030 [HG08, HG13, WNA18a] approximately corresponding to twelve reactors in total. Most prominently in 2018, this includes a third nuclear site at Hinkley Point.

2.2 Advanced Gas-cooled Reactor (AGR)

Classified within the Generation-II reactor group, the AGR was constructed with an initial design life of 30-40 years, a figure which has since been extended on regulatory review to over 60. The AGR is the reactor design studied in this research for a number of reasons:

- The AGR is the major reactor model utilised in the current generation of NPPs in the UK, and is responsible for a considerable proportion of the UK's energy mix (all but 2 reactors in the 19.6% outlined in Table 2.1)
- There a number of unique features related to the AGR (explored later in this section) and its operation, along with its repeated life-extension cycles, that make deep understanding of the design important for the UK nuclear operator

This short section will discuss some of the key features of the design.




Figure 2.3: Core photographs of Hinkley Point B, showing the nature of the in-core graphite brick structure

2.2.1 Reactor core

The reactor itself is centred around large core fashioned from graphite, which acts as both an integral structure and nuclear moderator. This is constructed from a combination hollowed cylindrical and squared bricks which interlock to form the whole body, with core-length channels facilitating the entry of fuel and movement of control rods about the core. Fig. 2.3 provides photographs taken during the construction stages of one of the Hinkley Point B units, giving a useful illustration of the inter-brick formation of the graphite core and the differences between fuel and control rod channels.

The fuel channels throughout the core house the nuclear material, which is enriched UO_2 , in contrast with the U^{237} used in the previous MAGNOX generation. These are assembled into long processions of fuel and associated instrumentation known as 'fuel stringers'. Through a process of conversion and manufacture from uranium ore, UO_2 is extracted from mined materials and fabricated into pellets. These are stacked into fuel rods that make up the fuel stringer, with each rod consisting of around 64 fuel pellets [Wes11]. Across the core, there are over 300 fuel channels containing various levels of rods in their

corresponding stringers, depending on the operational history of the plant.

In order to exercise control over the fission reaction in the core, the control rod channels allow for the rods of boronated steel to be mechanically inserted as and when required. Spaced interstitially among the fuel channels, the insertion of control rods brings the reaction rate down due to the high absorption properties of the rod material. AGR rods consist of two types: bulk rods, for long-term low power periods which are generally used during specific events, and regulating rods, which are involved in the fine adjustment of maintaining a steady power output [WWJM10]. For unplanned shutdown scenarios, nitrogen is also available to be injected into the coolant to arrest any criticality unable to be brought under control by conventional means.

2.2.2 Primary cycle

The primary cycle comprises the nuclear-side coolant flow through the reactor into the AGR boilers (in order to generate steam), and is driven by high pressure CO₂. This gas is held at approximately 4,000kPa, and propagates through the fuel assemblies housed within the fuel channels from the bottom of the core, through the graphite core itself and then directed to the boilers by a structure known as the gas baffle. On its pass of the criticality in the graphite core, energy is transported via heat from the core to the boiler units by the procession of this gas. The selection of CO₂ is largely attributed to its chemically inert properties in radioactive environments.

In order to propel the CO₂ through the primary cycle, each AGR uses eight induction motor-based assets known as *gas circulators* (GCs). Two of these rotating machines correspond to each quadrant of the core; being responsible for the successful propagation of coolant to each of the channels in that area. The GCs represent one of the major assets under consideration in this thesis, and will be revisited in greater detail in later sections. An illustrative overview of

the core and primary cycle is shown in Fig. 2.4, providing explanatory examples of the core with its channels, boilers, control rods and GCs. Note that a single GC unit is annotated in this diagram - in total, eight of these will be positioned around the core.




Figure 2.4: Overview of the primary cycle of the AGR core, with annotation of CO₂ gas flow [Non96]

2.2.3 Secondary cycle

In the design stages of the gas-cooled UK reactors, one of the earliest defined aims for the project was the outlet temperature and pressure of the primary cycle be conducive to the use of conventional steam-driven plant already utilised in coal-fired and other conventional generation types [Non96]. This decision allowed for the wealth of expertise and support already existent in these areas to be used, with analysis and experience from over a century of use informing their everyday operation.

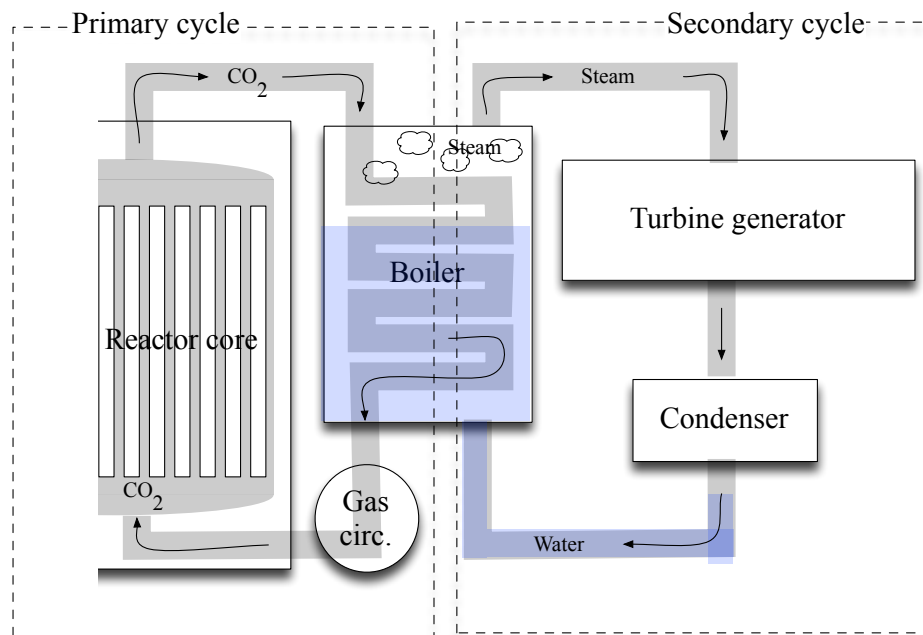


Figure 2.5: Primary and secondary loop of an AGR power plant

From each of the boiler units, superheated steam is output to inlets for the turbine generators (TGs). The TGs are the mechanism by which the heat energy generated in the reactor core is converted into mechanical work. On exiting each of the turbine stages, the steam is then condensed back to water in order to be reused in the boiler for future conversion. A schematic of the secondary cycle, along with a corresponding primary cycle, is provided in Fig. 2.5.

2.2.4 Associated rotating machinery

Though referred to within the context of their cycle loops in the reactor throughout previous sections, it is of use to list the rotating machinery utilised within the AGR and its support systems in order to gain perspective of the extent and importance the class has. Within a single AGR system, the assets falling into the rotating machine class can be listed as follows:

- Gas circulators (primary cycle, eight units),

- Main boiler feed pumps (primary cycle, one unit),
- Steam turbine generators (secondary cycle, one unit),
- Various auxiliary rotating machines.

The ongoing analysis and upkeep of these is central to the safe and effective operation of the reactor, with each playing a notable role in the day-to-day considerations of the operator. Accordingly, there exists regulatory and licensing requirements directly in reference to the correct operation, upkeep and reporting for these machines. Much of the existent knowledge and best-practice has been derived from years of rotating machinery research and development.

The study of the health and behaviour of these items falls under the general umbrella of *machinery condition monitoring*; a discipline rooted in mechanical and electrical engineering, physics and statistical analysis. Reliability engineering staff are employed by the operator to provide support analysis to station-side operational staff by studying both the medium- to long-term condition of these systems. This topic and a more general overview of the rotating machines under consideration in this thesis are discussed in the following section.

2.3 Rotating machinery condition monitoring

Not limited to nuclear generation, rotating machinery is a major asset family used in a wide array of engineering systems including transportation and manufacturing processes. Fundamental machine components including engines, lathes, propellers, rotors and even the wheel itself all provide useful work by function of their rotational characteristics. Accordingly, the ongoing study and maximisation of availability for such assets and mechanisms is central to a wide spectrum of engineering analysis; rendering their ongoing condi-

tion monitoring an important consideration for safety and economic-conscious operators.

This section provides a top-down overview of the typical rotating machines examined throughout this research, detailing their general theory of operation and the characteristics pertinent to their successful health monitoring. Parallel to these discussions, an overview of the common condition monitoring strategies is provided, outlining the challenges, drivers and opportunities facing the modern-day asset manager and monitoring engineer.

2.3.1 Historical perspective

The umbrella term ‘condition monitoring’ encompasses a wide variety of approaches [Moh18], technologies [GP16] and techniques rooted in the analysis of machine asset health and reliability. In order to successfully manage components and assets within an engineering system, as much information regarding the characteristics of the machine in question needs to be brought together from a myriad of available data sources. Monitoring approaches are developed in order to support maintenance decision making, and to allow for counteractive measures to be taken against emergent faults and failure conditions.

2.3.1 a) Drivers for monitoring

The successful maintenance of engineering machinery has a variety of underlying drivers and considerations. Most prominently, safety is a high priority for many disciplines: catastrophic failures are simply not an option for any assets where a fault could endanger staff, the general public or the environment. In order to operate in safety-critical environments such as nuclear, aerospace and nautical applications, there are often associated regulatory requirements for adequate maintenance to safeguard against potentially dangerous fault scenarios. These measures complement design considerations such as redun-

dancy, contributing to the overall ‘defence in depth’ strategy [Moh11, YJH16] seen in the nuclear industry to combat failure and failure repercussions.

Beyond the obvious safety considerations, the economic effects of an unplanned asset outage can be very damaging for any organisation dependent on a particular machine class. Any process centred around the ongoing operation of an asset class has an associated monetary value per hour. For example, if a coal-fired plant loses the capacity to generate from a combined-cycle gas turbine due to an unplanned machine failure, the operator loses revenue for every potential operating hour the system remains unavailable.

2.3.1 b) Knowledge, experience and data

The rotating machinery health monitoring community, like many maintenance-focused sectors, has historically built on the available knowledge through a combination of first principles study [BH02] and lessons learned [Nea06, EMAQAG16] from hands-on experience. This paradigm is common across a wide range of mechanical and electrical analysis problems, where the inherent complexity of the system in question often rules out an entirely physics-based model of faults and anomalies.

Accordingly, the acquisition, archiving and scrutiny of machine data is particularly important in the ongoing analysis of machines under condition monitoring regimes where a complete knowledge domain isn’t available. Operators in a wide array of disciplines have identified this concept, and the previous decade has seen a noted increase in the volume, variety and velocity of data associated with rotating machines.

The machinery observables² of importance are often asset-dependent, with particular machines have differing measures corresponding to useful metrics for behaviour. For the purposes of this research, we define the collection of

²The term ‘observable’ is used in lieu of the more common engineering term ‘parameter’ as portions of this thesis refer to machine learning parameters - a different concept entirely.

machine observables used to reason about a particular system as that system's '*machine view*'. For example, an off-shore wind turbine condition monitoring system will likely utilise sensing technology to survey gearbox conditions such as temperature. However, a monitoring system developed to collect information regarding a circuit breaker unit will focus more on electrical observables including load and voltage. While there exists a difference in the machine views, there are often strong parallels in monitoring practices across assets.

2.3.1 c) Maintenance regimes

Maintenance regimes can be broadly categorised into the following approaches:

- **Failure-based maintenance;** where assets are maintained or replaced as and when they reach a failure criteria (e.g. a light-bulb being replaced when blown),
- **Time-based maintenance;** where assets are maintained or replaced at a given time interval, selected in order to minimise unexpected failures (e.g. an automobile's regular MOT assessment to determine road-worthiness),
- **Condition-based maintenance;** where assets are examined at a suitable fidelity to ascertain their ongoing condition and health in order to effectively take preventative measures against incipient faults or problems,
- **Predictive maintenance;** also referred to as prognostics, where assets are modelled and condition is extrapolated into the future in order to make predictions about future states and potential failures.

The selection of the maintenance regime to be used is dependent on the machinery under scrutiny itself and the resources available to the asset manager. A succinct argument for a particular regime needs to be made: for example, it would be considered uneconomical to spend a large amount of resources

on the condition-based maintenance of an easily replaceable system or component. Similarly, mission-critical assets need to be given greater consideration of their ongoing health as a function of use and planned usage, in order to avoid potentially disastrous results.

Generally, a combination of time-based and condition-based regimes are employed for rotating machinery (there are not many large scale engines that can be run to destruction in an economical business model, for example).

2.3.2 Vibration monitoring

Vibration has been a mainstay in the analysis of rotating machinery since the 1950s [Fos67]: being used as a direct or indirect indicator of state, condition and asset health. While it is certainly not the only measure condition monitoring engineers use in analysis of this kind, with temperature and oil-based techniques also employed, it has seen the most success and adoption in benchmarking normality and outlining machine anomalies [Ran04] for rotating machinery assets.

2.3.2 a) Vibration

From a fundamental perspective, vibration is a quantitative measurement of motion of a body as it oscillates to and from a point of equilibrium. This periodic motion is set off by an initial displacement, from which the body then oscillates back to the initial state. Within classical mechanics, vibration is a well-studied phenomenon with an extensive literature [Ton02, Tho96] associated with its first principles. All of the research herein is concerned with the study of *forced vibration*; or vibration where the displacement driving the process is time dependent. In rotating machines, this displacement is often in the form of the rotary motion itself, undesirable component contact or the application of external load. In contrast, the study of free vibration is centred around

examples where an initial displacement is the only driver experienced - for example, in the case of the harmonic motion exhibited by a simple pendulum. Forced vibration response remains at a constant amplitude when the driving force is constant, whereas free vibration response decays with time after the original impulse [Sch04].

There are a huge number of potential sources of vibration-based behaviours in a moving body utilised in engineering systems, not least in the mechanisms encountered in rotomachinery. Naturally, bodies undergoing any form of motion will undergo some degree of vibration, which is to be expected as these are the result of unavoidable mechanical and thermal processes such as friction existent in any real-world system. For a machine dynamicist, a distinction is often made between nominal and excessive vibration in order to classify undesirable changes in asset conditions and behaviour. Large values are generally considered as unwanted for two reasons. Firstly, extreme vibrating conditions subject the rotating machine to undue internal stresses which have the potential to lead to mechanical damage itself. There is also the problem of unwanted contact between otherwise separated parts and components in the system. A simple example of a conditions like this would be in driving an engine at a rate greater than its components are designed to cope with. Conversely, excessive vibration can itself be symptomatic of an existing or emergent problem in the asset, presenting a powerful feature of vibration as an analysis tool. Signatures consistent with faults such as shaft cracks, rotor imbalance or gearbox wear have been extensively studied and documented for diagnostic purposes [Mit81, PRB03]. In fact, experienced vibration experts have the ability to diagnose each of these faults, make judgement calls on their severity and advise on future maintenance action - all based the scrutiny of the vibration data.




Figure 2.6: Single frequency signal component with annotated measurements

2.3.2 b) Measurements

The two major vibration measurands are *amplitude* and *frequency*. Amplitude corresponds to the incident signal level on the transducer, predominantly denoted in either a distance or distance/time unit base (though a number of legacy systems also use a voltage measurement proportional to the actual movement experienced). Depending on the particular engineering discipline and aims of the monitoring system, the amplitude can be measured using a number of different conventions. *Peak to peak* (pk-pk) amplitude measurements take into account the procession of the entire waveform from maximum to minimum values. In contrast, *peak* and root mean squared amplitude consider the maximum or scaled-maximum value against the zero or equilibrium state of the signal. Each of these values are annotated in Fig. 2.6.

Peak to peak is the among the most commonly used measurements in the analysis of rotomachinery due to the nature of multiple frequency components existent in real-world vibration signals. Considering both the positive and negative components of the overall displacement allows for the analysis of both

of these to be made in the post-processing in creating a frequency-based view of the system. Conversely, there are numerous examples where the fidelity of peak to peak is not required. For example, in the measurement of time series velocity components. In such instances, peak or RMS measurements are often utilised [BH02].

Frequency (defined as the inverse of the signal period, or $\frac{1}{T}$, denoted in Fig. 2.6) is the periodic, repetitive rate of any oscillation or recurrent event. With particular regards to vibration, the frequency measures the cycles per second of the vibratory oscillation incident on the physical system. This is particularly useful when compared with the frequency of operation of the rotating machinery itself, allowing for the nominal characteristics of the machine and secondary behaviours potentially coinciding with damaging behaviour to be separated and identified. Many types of machinery fault are manifested primarily in this frequency space [CF04], so it is important to examine these components.

2.3.2 c) Transducers

The sensing technology used in vibration monitoring is predominantly based on the measurements of position and displacement between the component's equilibrium and excited mechanical states. *Accelerometers* are mounted at strategic positions on asset casings and components in order to record the experienced vibration levels. This short section provides a high-level overview of transducer principles in vibration monitoring. It should be noted that this is intended as an illustrative summary to allow the reader to better understand vibration as a data source: further details on vibration instrumentation are available from a large existing literature on the subject [NK03, dS12].

Fig. 2.7 gives an illustrative overview of a typical instrumentation setup. The transducer would be mounted on an encasing or nearby structure of the

component, and measures the relative distance between itself and the excited rotor. Measuring this difference in separation at a regular interval allows for a time series of the vibration level to be created.

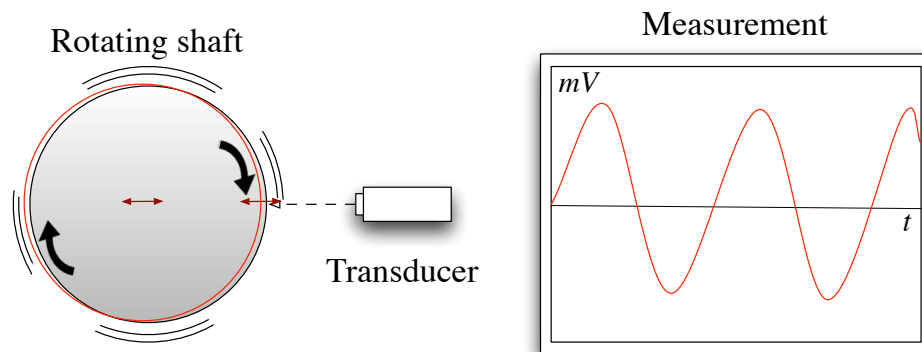


Figure 2.7: Vibration measurement example, showing a vibrating rotating element oscillating in the horizontal axis

2.3.3 Steam turbines

Steam turbines are a member of the turbomachinery family, reliant on the flow of high-temperature pressurised steam to provide useful mechanical work. A noted invention of the industrial revolution, the conception and construction of the first steam turbine has largely been credited to Sir Charles Parsons, who built the first unit in 1887 [Par11]. Among the early applications for these machines was propulsion and transportation; with designs being applied to both marine and locomotive vehicles.

Within the scenario of power generation, these machines (also commonly referred to as 'turbo-generators' (TGs)) make use of input thermal energy in driving a rotational mechanism which is in turn exploited by a partnered generator unit to produce electricity. Steam monopolises the power industry as the major driver of TGs, with as much as 90% [Wis00] of the generation means around the world today being at least in-part facilitated by it. The ubiquity




Figure 2.8: Schematic of a steam turbine generator

of steam-based generation makes the understanding and upkeep of these machines an important topic for the reliability industry as a whole.

Fig. 2.8 provides an illustration of a three stage steam turbine generator unit, depicting the input and outputs of the system along with each of the turbine components. Each of these turbine stages differ in their pressure characteristics: from high pressure (HP) stages, to intermediate pressure (IP) to low pressure (LP). From a mechanical perspective, we can distinguish the operational characteristics of a steam turbine into two functional areas for clarity of discussion: the *turbine stages* and the *generator*. The following sections expand on both of these areas.

2.3.3 a) Turbine stages

The first functional priority of the turbine unit is to create a rotary procession from the supplied inlet steam to the machine. This is achieved through interaction with the turbine blades of the machine's rotor, which are oriented to alter the momentum of the incident steam flow. A variety of blade setups exist (which can be largely grouped as *impulse* or *reaction*-based) which interact with the steam in differing manners. However the main outcome remains the creation of a force acting on the machine rotor to drive movement [BH02]. This principle governs the operation of all steam turbines.

A typical generation-context unit is comprised of a number of turbine blade stages at varying pressures, a characteristic dependent on the rating and in-

tended application of the machine. Multiple blade stages allows for greater efficiency of movement to be drawn from the steam supply, providing as optimal a generation rate as possible. Steam leaving directly from the boiler goes through a blade stage before being recycled through a re-heating process for use at a lower pressure blade stage, increasing the total net energy output from the system. In application to gas-cooled nuclear stations and conventional fossil fuel-fired plant, a steam turbine would tend to consist of a single high pressure (HP) turbine stage, a single intermediate pressure (IP) turbine stage and two or more low pressure (LP) turbine stages [JLD⁺91]. The stage components of TGs used for energy generation means represent a large body of equipment; requiring adequate housing in often volumes power plant premises. A photographic example of a TG is shown in Fig. 2.9 for a representation of scale, particularly for the turbine stages.




Figure 2.9: Photograph of steam TG [NRC]

2.3.3 b) Generator

The generator component consists of a driven rotor and a stator, each with a complement of windings forming a particular electromagnetic polar orientation when provided with a generated or supplied potential. There are a number of means by which electrical energy can be created between the windings of the rotor to the stator, but generally this is facilitated by exploitation of Faraday's law of induction [Ben96]; stating that a moving (or time-variant) magnetic field will induce an electrical potential across a conductor. An electromagnet is created within the generator (normally through energising the rotor windings) and the motion of the rotor creates the required moving magnetic field, therefore inducing an electromotive force (EMF) in the stator windings to the output of the generator. This propulsion requirement of the rotor is the main function of the steam-driven components in the steam turbine, effectively converting thermal energy (steam) into useful mechanical work (movement of the rotor) and then into electrical energy (EMF generated in the stator windings). This basic principle governs the operation of all generator units [WWS13], irrespective of the origin of the mechanical work.

2.3.4 Gas circulators

Purpose-designed and -built for use in the AGR, the gas circulators (GC) function as primary cycle CO₂ propagators. As discussed in 2.2.2 there are eight circulator units per reactor, each corresponding to a particular area of the core. Mounted under the main reactor structure, these machines maintain the desired temperature conditions in-core, and transfer the fissile energy production to the boiler units onto the secondary cycle. This dual purpose nature of the GC function renders them an important asset in the AGR system.

The circulators are variable speed machines; they have the ability to operate

at a variety of duty cycles dependent on the requirements of the operator. For example, the stepped load regimes immediately following a planned outage require the circulator units to be brought up in discrete levels to a maximum value, maintaining a controlled rise in coolant flow in line with the return to criticality. Along with this regulation, fine tuning of the circulation rate at steady-state operation can be altered through manipulation of the inlet guide vane (IGV) component of the GC machines. A photograph of a typical AGR GC design is provided in 2.10.

In common with the TGs under the auspices of nuclear maintenance professionals, the standard means of monitoring GCs is rooted on vibration-based monitoring. The instrumentation setup can be considered analogous to the systems deployed on the TGs, meaning large overlaps of the analysis used on each asset class. However, the role of the GC in the AGR's operation creates unique duty cycle and response features when compared to other turbomachinery. The bespoke nature of the circulator units in relation to the reactor design makes them an important source of insight for AGR health, but with comparatively less study associated with their own characteristics when contrasted with a generic steam TG.

2.4 Modern condition monitoring

2.4.1 Hardware, storage and visualisation

As referred to previously in 2.3.2 c), the interrogation of a system undergoing vibration is achieved utilising mounted transducers on key strategical components and areas of the machine in question. The physical mounting of an accelerometer or eddy current-based instrument on these components allows for the vibration characteristics to be noted and examined. Once acquired, these signals need to be presented and managed in an extensible and effective




Figure 2.10: Gas circulator unit [Non96]

manner, in order to extract the greatest value from their contents. A number of post-processing and visualisation considerations need to be addressed when presenting information back to the engineer; with numerous time series and frequency-space abstractions being available to inform decision making.

A variety of vibration monitoring hardware and software vendors trade in both the regional and global markets; including organisations such as Bentley Nevada. In the UK, Beran Instruments provide end-to-end vibroacoustic solutions to the needs of rotating plant reliability experts. Their PlantProTech transducer-interrogator-software package is a good example of the industry standard of systems presently operated by the CM professional.

With the continued advancement of computer technology, the reliability industry has generally followed the technology sector at large in embracing improvements in storage and hardware capabilities. Ever-larger machine datasets are becoming become available for real-time and retrospective analyses [JLB06, MH14, XSW⁺17, AHMR18].

A number of health monitoring areas have also seen the application of wireless sensor networks in their data collection processes; an agile technology allowing for quick and reliable interrogation of engineering systems [PFKC08] by taking advantage of techniques such as energy harvesting [SZ16] to make deployment onto assets straightforward. The development of these related sensor networks is notable as they provide access to larger still volumes and fidelities of streaming data.

2.4.2 Monitoring strategies

2.4.2 a) Components, systems and complexity

Zio [Zio09] provides a still relevant historical overview of reliability engineering, documenting the advance from rudimentary testing towards the scientifi-

cally mature approaches rooted in probability and statistics utilised today. Of particular interest is the identified shift in the 1960s and 1970s from component-level analysis (often focused on consumable device components such as transistor valves [Cop84]) to introducing system and subsystem-wide reliability measures. Such techniques infer root cause and probabilistic measures of historical and potential failures at the interaction level between components in complex systems, including the rotating machines considered in this research. While the failure rate of a single element within an engineering asset was previously the major unit of study [Vee80, BZS76, HW58], the 'scaling up' of such measures has become necessary to provide meaningful decision support to reliability staff tasked with the maintenance schedule of an entire asset or class of asset. Reliability also underwent an incentivisation process through this period; with the inherent positives of increased availability being championed alongside the more traditional production output successes.

This evolution reflects the changing requirements of the modern asset manager. As monitored plant items continue to increase in complexity, the modelling approaches utilised must follow similarly in their sophistication. For example, a historical precedent was generally made in reliability studies in considering the state of a component; it would either occupy a *functional* or *faulty* classification [Zio09]. However, today's reality much more closely represents a multi-state or continuous functional state when considering dynamic assets including rotary machinery, engines and other energy-based plant items. The performance output of a particular machine will often not vary in a binary fashion between 100% and 0%, and making this distinction is worthwhile when considering the maintenance strategy for the asset in question. This change in perspective has opened the potential for quantifying health as a continuous quantity, in contrast with the discreteness of previous philosophies, leading naturally to the concepts of benchmarking, extrapolation and

prediction of state.

The term '*soft computing*' is widespread in the reliability research literature; especially when coupled with established maintenance arenas such as energy, transportation and production³. For the purpose of this discussion, this refers to the application of nonlinear modelling techniques to reason about the inherent complexity, uncertainty and 'fuzziness' [Zad93] of engineering systems, taking into consideration both quantitative measurements from data and qualitative heuristics from previous experience. Techniques including neural networks [KWM92, Alt09] (with recent developments in deep neural nets [FLL⁺16, LSG⁺16, SJZW17]), fuzzy logic [FMJ⁺11, CNGT97] and genetic algorithms [JN00, CPL96] have all been exercised in maintenance-type scenarios with valuable results. ,

2.4.2 b) Alarms and notifications

In view of the varied and large data streams that health monitoring professionals are required to survey, an intuitive policy for keeping abreast of system changes as and when they occur is via use of alarm-based strategies. Such approaches provide a particularly useful machine view from a maintenance perspective; a large proportion of the data-based features of interest can be attributed to points of change or anomalous behaviour. This allows for multivariate observable spaces for assets, or entire families of asset, to be surveyed effectively and target analysis to operational periods of interest.

The rotomachinery industry commonly uses limits of acceptable behaviour for each of the monitored observables in the corresponding machine view. Breaches of these boundaries will trigger latched notifications for examination at a later time, and engineering staff will use these alerts at both a single- and

³The term is less mature in fields owing a direct heritage to computer science and modern technology, for example - datacentre or technical computing reliability, where the distinction between soft computing and general artificial intelligence tends not to be made

multi-level in making decisions regarding the ongoing operation of the plant item in question. In fact, a number of industry agreed standards exist [iso09] in defining these thresholds of tolerable behaviour, some of which are central in the regulation efforts of nuclear and other safety-premium scenarios. A key strength of this practice is in the explicability of particular values crossing a pre-determined standard boundary, without the requirement for specific analysis on the machine example or scenario.

Despite the dimensionality- and labour-reductive benefits to taking an alarm-centric approach when devising a monitoring strategy, there remains a number of disadvantages associated with the current best practice associated with the procedure:

- **Alarms can be numerous;** While it is true that setting allowable bands of behaviour allows reliability staff to focus only on operational features that deviate from some nominal baseline, the amount of notifications requiring post-analysis can still remain high. This is especially true when considering the multiplicity of assets often under the charge of small CM engineering teams,
- **Alarms can be brittle;** The majority of vibration-based alarms define allowable limits of operation at set values, without considering the dynamism of the monitored piece of plant. For example, increased or decreased duty cycle within an operational regime may breach an alarm boundary but be non-indicative of a relevant change to the reliability engineer,
- **Alarms can be inflexible;** Through the life cycle of an item of machinery, the inevitable degradation processes that occur during normal use will have an effect on the item response to steady state conditions. Hard-coded boundaries will not take this evolution of state into consideration,

and may in fact prompt unrequired analysis at the single alarm level when coupled with the strong regulatory requirements surrounding operational standards.

2.4.2 c) From diagnosis to prognosis

The previous decade has seen a notable rise in interest [KHV06, HZTM09, SZ15, RSWZ18, WGC⁺18] for predictive reliability measures: robust metrics for the quantitative (and, by necessity, probabilistic) measurement of future states, data features and potential failures. The trend can be attributed, along with the system-level reliability engineering paradigm discussed in 2.4.2 a), to the move towards more condition-based maintenance regimes (as defined in 2.3.1 c)) as the health monitoring profession looks to move with improved technology and expectation. *Machinery prognostics*, as a scientific sub-discipline of engineering CM, has experienced a flurry of academic interest and some industrial developments [SHM11, KMF⁺16], due to its promise of identifying problematic behaviours before they occur; avoiding costly outages and scheduling preventative maintenance.

Numerous distinguished engineering organisations such as NASA [NAS] & General Electric [Bon06] have focused resources of developing prognostic techniques. A large proportion of the literature discusses the requirement for predictive ability [HACF05, JSB04, JLB06, TBT15] in modern CM, suggesting roadmaps and measures of success. When benchmarked against other analysis functions in reliability engineering such as fault detection and diagnosis, prognostics is a comparatively young discipline. However, some developed and deployed systems have already brought initial success in industrial applications [LWZ⁺14].

This study does not seek to provide a definitive prognostics representation for the rotating plant items in question, but it is useful nonetheless to formu-

late a definition (at least in the interim) for clarity and the sake of argument. Among the most succinct and well cited definitions is provided by Sheppard *et al.* [SKW09]; where prognostics is the estimation of the *remaining useful life* (RUL) of a particular asset or process. The RUL is dependent on the asset class, data availability and requirements of the developed system but the unifying characteristic is RUL estimates have either an explicit or latent connection to the operational hours available before a minimum functional condition boundary is breached. Selection of this RUL parameter is an area of study in itself [Cob10].

2.4.3 Intelligent condition monitoring

Through this section, computational techniques that fall into the broad category of ‘intelligent’ are discussed; with a strong focus on those used in reasoning about engineering asset condition. The state-of-the-art is framed alongside a variety of important application examples across a wide variety of health monitoring scenarios.

2.4.3 a) Definition

Krishnakumar [Kri03] describes the techniques associated with the term ‘intelligent system’ as:

“...nature-inspired, mathematically sound, computationally intensive problem solving tools and methodologies... [which are] ideally suited for tasks such as search and optimization, pattern recognition and matching, planning, uncertainty management, control, and adaptation”

While the exact definition of the term is likely to prompt lively debate among practitioners from a variety of fields (in both reliability engineering

and beyond), for the purposes of this discussion we will consider this description as satisfactory. Two important additions (or clarification of existing conditions) are the abilities of *diagnosis* (inference of state or features of state, more broadly falling into the category of data interpretation [Hop11]) and *prognosis* (the extrapolation of state to make predictive reasoning on future states and conditions).

For the types of condition monitoring problems encountered through this research, therefore, the goal of an intelligent system is to provide some form of automated, repeatable analysis to staff which aids them in making good decisions regarding asset health. This output can be a wide variety of different results: from a quantitative measure with no real abstraction of the modelled item at hand, to more detailed decision support advising on a particular course of preventative action regarding particular components or processes.

2.4.3 b) Types of intelligent system

We can consider intelligent systems techniques to fall into three broad categories, as follows:

knowledge-based; systems built upon heuristic and expert-derived knowledge, often relying on sets of rules to reason,

data-based; systems which take advantage of asset data in order to reason statistically,

model-based; systems comprising of a modelled representation of the asset (physics-based or otherwise), which use outputs of this model to reason.

These headings are not exclusive; in fact many 'hybrid' systems have proved successful in fusing features from each of these categories to solve particular engineering problems.

2.4.3 c) Applied to generation-based rotating machinery

Rotomachinery used throughout the energy industry has seen success in intelligent system development [JLB06, HZTM09] applied to their monitoring processes for fault detection [LYZC18], classification [GCGP16, CZC⁺16] and predictive analyses. As discussed through 2.3.2, decisions are largely based upon vibroacoustic analysis [Ran04] and the identification of particular diagnostic conditions.

Knowledge-based techniques have historically dominated rotating plant decision support systems, with a number of major knowledge-based systems (KBS) breaking new ground in recent years for both rotary asset-specific and general engineering intelligent system development. From a pure KBS perspective, a few rule-based inference systems elicited through formal knowledge engineering [SWdH⁺83, SAA⁺00] processes have been created and applied to steam turbine generators [YZX⁺11, MAK14]. Further augmentations to such techniques have also proved successful [EP08], notably resulting the VIBEX (VIBration EXPert) system [YLT05]. Gas turbine units have also seen KBS development, with a rule base being used in conjunction with model-based and qualitative techniques [TM92, ZLCM14] to positive effect.

Data-driven approaches [CWT⁺11, TMMZT12, BMZ12] have grown in stature over the past decade, attributable to the increased rotomachinery data availability coupled with the rising sophistication in required decision support. The maturity of such techniques (discussed in greater depth through Chapter 3) has allowed for previously unattainable levels of quantitative and statistical reasoning to be applied to generation asset reliability problems where plentiful historical information is available. Approaches including neural networks [ACS⁺08], hidden Markov models (HMMs) [MM07, KCMT13], ensemble methods [BMZ12] and kernel machines [ZZXC10, ZD12] have all been explored with success. Concurrent to these technique-based solutions are the

continual improvements made to the signal pre-processing, with feature engineering and signal representation [GPJW12, HYHY10] also being investigated.

2.4.3 d) Applied to rotating machinery in other domains

Adjacent to the intelligent CM progress being made in conventional generation plant, the wind energy industry has also embraced AI-based technologies [CLEP12], which is likely due to the distributed and multiplicative nature of the average wind turbine fleet. Wind-specific monitoring approaches have been constructed with support vector machines (SVMs) [LSoO11], neural networks [SF11, KL11] and predictive analytic techniques such as particle filters [BCFR12] and Markov chains [BD10]. There is also a significant interest in weather modelling [SBH⁺05] for predictive and scheduling purposes using a similar family of techniques, illustrating the broad applicability of such methods.

The aerospace sector has historically represented one of the more data-rich monitoring arenas, naturally making it one of the trailblazers in terms of intelligent technique experimentation and adoption. This is likely driven by the 'high integrity' [CBT07] nature of rotary assets in such scenarios; a property mirrored in nuclear-context machines of the same family. Statistical [CTM⁺08] and AI-based [HSTA00] approaches have all seen application to aero engines, with particular focus on the prognostics [BRG02, HWB06, ZMHF16] leg of intelligent CM research.

2.4.3 e) Applied to nuclear generation

Considering the monitoring of an NPP as a whole, the improvement of reliability engineering processes has continued to grow in importance for many reactor designs. This is in part due to the policy of life extension associated with many of the recent generation of nuclear plants [VR11, IOU⁺12], with de-

ployments of the AGR, CANDU [Tap08] and BWR [BAG⁺14] types each seeing consideration for use beyond their original commissioning period. While more traditional disciplines such as materials science [AAK17] and innovations in inspection techniques [MWMM16] have historically made most of the contributions, data-driven approaches have seen success [BMZ15, ZD10, WMT09, LWDG17] in recent years as well.

2.4.3 f) Applied in power and electrical scenarios

After the generation process, power needs to be provided through the distribution and transmission networks on the grid to homes and centres of industry - a complex engineering problem with its own reliability requirements. Consequently, CM is also mature in this industry; with a plethora of intelligent systems being deployed across a wide spectrum of power system monitoring tasks including circuit breakers, transmission lines and power transformers. The complexity of the power system in the infrastructure of developed countries has necessitated a blend of knowledge- [SRM⁺08], model- [DMM03] and data-based [NM11, RCMJ11] decision support systems and studies. Numerous protocol standards, such as *supervisory control and data acquisition* (SCADA), exist which allow the interrogation of remote assets; with a number of deployed systems being designed specifically [GMW13, DMM⁺06] for such data types.

2.4.3 g) Other applications

Reliability and operational considerations have become a central issue for the distributed high performance and cloud computing industries, with Google recently publishing [Gao14] insights into a neural network approach for data centre procedures. The data centre represents a new frontier in reliability science [ZCB10], with the hugely parallel physical and virtual machines presenting a unique scenario where engineers need to understand system health at

scale from the outset.

2.5 Challenges and opportunities

This short section condenses the topics touched on throughout this chapter to a variety of key points of importance. The principle aim of this discussion is to extract the perceived shortcomings in the present approach to health monitoring in the rotomachinery discipline (in particular, the practice surrounding rotating plant in the nuclear environment) and to present the opportunities for research and innovation to meet these challenges.

2.5.1 Challenges

2.5.1 a) Energy security

Numerous studies [Bra10, YM13] have identified that the rising trend in energy consumption shows no short-term signs of abating and is expected to continue well into the future. The reliance on affordable energy at the global scale has prompted intense debate in the topic of *energy security* [ACN15] with relation to engineering, national policy and macro-economical thinking across nearly all of the nations around the world today. Some uncertainty surrounds the term (with it often being used in both economic and political [LMR10] arguments), however Winzer [Win12] provides a useful quantification as:

“... [energy security] is the absence of, protection from or adaptability to threats that are caused by or have an impact on the energy supply chain.”

With this in focus, a number of academic policy and canvassing studies [CVS⁺11] illustrate the perceived importance nuclear has to play in the energy

security of the UK; especially when coupled with the low carbon future concerns of modern society. Elsewhere, the clear ramifications of a catastrophe such as that experienced in Fukushima Daiichi [Hat12], Japan) have subtler long-term effects on the decision made on nuclear [HH13].

2.5.1 b) An ageing nuclear fleet

The AGR core has itself an entire field of researchers working on understanding its condition and health, from both a materials science and empirical [WJM⁺06, MRS10] perspective. While the rotating assets discussed in this study do not undergo degradation processes related to the nuclear environment experienced in-core, the implicit dependence and dependencies between the TGs and (perhaps especially, being a primary cycle unit class) the GCs with regards to the whole reactor condition are of high importance. This is especially true when considering the regulatory-centric nature of the UK nuclear industry, where the examination of rotating assets is made with the context of the AGR clearly in focus.

The UK government and regulatory bodies have extended the operational lifetime of the AGR at several points in the last couple of decades [NEI12], mirroring the trend in NPP lifetime extension around the globe [BRTL11]. Maintaining and optimising the throughput of nuclear-context assets, irrespective of their relative proximity to the criticality, is vitally important for operators. This viewpoint is confirmed by the importance placed on lessons learned in historical asset failures [Kal72] on UK NPP sites. As the core itself begins to deteriorate, the effectiveness of the assets surrounding it will ensure the functionality, efficiency and safety is kept sustainable. Keeping this effectiveness at a satisfactory level presents another challenge.

2.5.1 c) An ageing nuclear workforce

The human resource issue in the UK for nuclear, and energy engineering as a whole, has been a topic of concern in recent years. Public sector forecasting [COG09] estimates that 70% of the workforce surveyed in 2009 are due to retire by 2025, decimating the numbers of skilled professionals in the discipline. Reductions of expertise at this scale compounded with the plans for new build is a major hurdle for the UK; especially with the heightened variety and complexity, as discussed in 2.4.2, of analysis being asked of reliability engineers today and in the years to come.

While emergent technology does not seek to entirely replace the workforce, the codification of tacit knowledge in fields of expertise like those of reliability in the nuclear industry can benefit from knowledge-based and other AI-driven systems. Rule bases [TMMS07] intuitively fit this role, but it is argued that the volume and variety of advanced analytics required for future energy systems calls for the increased adoption of statistical, data-driven analyses to complement crisp inference.

2.5.2 Opportunities

2.5.2 a) Increased data availability

Section 2.4.1 made reference to the notable volumes of operational information in the form of historical data now available to reliability professionals. It is difficult to overstate the value of such data as a resource when reasoning about rotomachinery and other such engineering assets; as discussed, the machines in question are often complex and any source of empirical insight into their behaviour is of value. The growing data volumes lend themselves to approaches in AI and machine learning, where the potential performance is often directly proportional to the amount of data available for a given problem.

This data explosion seen in engineering maintenance is by no means limited to this discipline alone. At the most extreme, disparate fields such as particle physics [Bru11], biostatistics [ABO10] and econometrics have all seen direct benefits in adopting the so-called '*big data*' philosophy in their development efforts to tackle domain-specific problems. While the volume, variety, veracity, value and velocity of information directly sourced from engineering assets are not quite at the same levels, the industry has begun to take note of these advances from further afield [Rus14, LLBK13].

2.5.2 b) Increased technique sophistication

The previous ten years have seen the fields of applied AI and machine learning explode into the public consciousness, with numerous high-profile and everyday products now powered by these technologies. Where in the past, systems reliant on statistical and automated inference were limited to fringe prototypes and demonstrations, more and more engineering problems are being solved by AI-based solutions. Nowhere is this more prevalent than the internet industry, where numerous billion dollar companies have built their successes on innovation [JM15] in systems which provide useful inference from data. Areas rich in data all benefit: including social media [LK12], commercial speech recognition [PP08] and more recently healthcare [JM15]. Open source development libraries have become more and more available [PA13] and easier to build reliable software with [Aea16], bringing intelligent technique construction towards the realm of engineering as opposed to remaining in research environments.

The successes reaped from other data abundant fields provide an indication of the potential rewards awaiting the nuclear and reliability industries as a whole when they adopt such techniques as central to their monitoring strategies.

Chapter 3

Statistical methods and machine learning

This chapter discusses a number of relevant statistical methods, before introducing the topic of machine learning: a computer science sub-discipline focused on the extraction of patterns and information from data using a combination of statistical and computational techniques. Key methods and approaches are introduced, and then contextualised against the domain area of application.

3.1 Statistical methods

In this section, a number of key concepts from statistics are examined in order to better contextualise and explain methods introduced later. The theory discussed herein is not intended as a rigorous mathematical foundation to ML and statistical techniques in general; merely enough detail to allow for the succinct demonstration of the approaches used throughout the research.

3.1.1 The Gaussian distribution

Encountered throughout engineering and the physical sciences, the *Gaussian* is a continuous probability distribution which describes random variables distributed symmetrically about a mean μ with a variance σ^2 . It is expressed as:

$$f(x|\mu, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} \quad (3.1)$$

The prevalence of the Gaussian distribution in real processes can be attributed to a few key reasons [Lyo13]: its importance to many exponential family distributions defined by the central limit theorem [Haz01], and its accuracy in describing many physical phenomena.

3.1.2 Probability density function

Some of the most important advances in technology have the advent of probabilistic techniques and reasoning in science to thank for driving their progress.

In the most general sense, probability is a measure of likelihood; how likely it is that a particular event will occur. This is an immensely useful property when dealing with random variables; observables which have some component of stochasticity associated with their measurement. In reasoning about real world data, uncertainty can be derived from a variety of sources: the finite nature of observed data points, measurement noise and complex losses to the environment or other obscured factors.

Random variables (in the strictest sense) are idealised mathematical objects. However, we can consider many observables encountered empirically in engineering (for example, vibration [Lal14]) to follow similar principles.

There is a large amount of mathematical discussion available regarding the fundamentals of probability theory (Breiman [Bre68] provides an excellent

overview on the subject).

A common means of examining data-based behaviours and their associated probabilities is in the use of a *probability density function* (PDF). PDFs provide an associated probabilistic measure of likelihood for a particular observable value in any number of dimensions. PDFs are analogous to the familiar histogram, differing in that the associated $p(x)$ is a real-valued continuous function as opposed to a binned discrete count. $p(x)$, as a function describing probability, is subject to the constraints

$$p(x) \geq 0 \tag{3.2}$$

$$\int_{-\infty}^{\infty} p(x) dx = 1. \tag{3.3}$$

Note that property (3.3) dictates that the integral of the function needs to equal unity. This allows for probability density values of > 1 , so long as the area bounded under the function integrates to one. Fig. 3.1 provides three example PDFs, each corresponding to a Gaussian distribution of differing parameters ($[\mu, \sigma^2] = [0, 0.5], [-0.3, 0.8] \ \& \ [1, 0.4]$).

The major conclusion to be derived from these figures is: in the event of a random example being taken from a process described by an associated $p(x)$, what is the probability of the example taking value x .

3.1.3 Cumulative distribution function

The cumulative distribution function (CDF) is a continuous real-valued distribution that examines the probability of a value of X , with an associated probability density $p(x)$, existing under or equal to the value x . More formally, a CDF examines the relationship:

$$F_{CDF}(x) = Pr(X \leq x). \tag{3.4}$$




Figure 3.1: Three example PDFs, drawn from the Gaussian family

CDFs are monotonic increasing functions, representing the continuous analogue to discrete cumulative frequency analysis techniques. Fig. 3.2 provides an example CDF taken from a Gaussian distribution. The examination and comparison of CDF's characteristic in relation to commonly utilised distributions is a standard technique in hypothesis testing and inference. For example, ascertaining the feasibility of a data population being taken from a particular distribution could be achieved by examining the similarity of the example with respect to an idealised CDF.

3.1.4 Empirical statistical analyses

In working with observables from real data where the parent distributions aren't defined *a priori*, there needs to be an intermediate between the functions that describe textbook distributions and the collected data itself. A histogram can be used to help visually examine the data distribution, but the strength of




Figure 3.2: Example cumulative distribution function of Gaussian PDF

tools like the PDF and CDF lies in their ability to create a continuous representation of probability density. Therefore, we look to empirically create density estimates of these distributions.

To approach this, the associated probability distributions can be empirically estimated from the available data through a number of techniques inspired by mathematical empiricism. This section discusses these approaches, under the general heading of *empirical techniques*.

3.1.4 a) Kernel density estimation

Kernel density estimation (KDE) [Cam02] utilises a finite set of observations, our input dataset \mathbf{x} , drawn from a behaviour of interest in order to create an empirical approximation of the associated PDF. The KDE approach achieves this through the assignment of a *kernel*, a standardised weighting function selected from a variety of common distribution families, to each instance in the dataset.

The global estimation of the PDF is then created through a summation over each the assigned kernel functions, creating a smoothed aggregation of the associated probability density. The frequency-based nature of this smoothing process often leads KDE to be referred to as a ‘smoothed histogram’. There are a variety of different kernel functions that can be utilised; some of the most commonly used kernels are provided in Fig. 3.3 and Table 3.1.

We can express the previous definition mathematically. Let $\mathbf{x} = \{x_1, \dots, x_n\}$ be set n observable inputs drawn from dataset \mathcal{D} with associated idealised probability density $p(x)$; we are concerned with the estimation of $p(x)$ through the KDE technique. $\hat{f}(x|h)$ represents our created empirical best-estimate, the

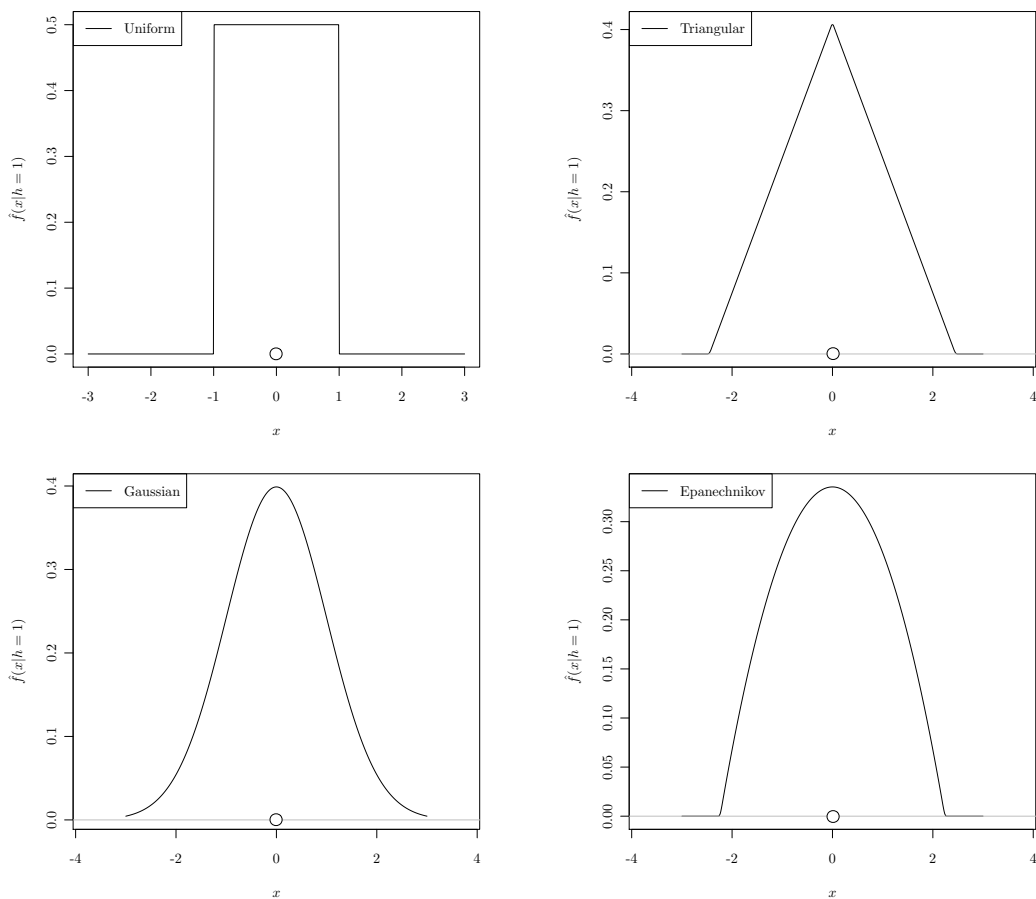


Figure 3.3: Four kernel functions commonly used in non-parametric statistics: uniform, triangular, Gaussian & Epanechnikov

Table 3.1: Selection of kernel functions

Kernel type	Function
Uniform	$K(u) = \begin{cases} \frac{1}{2}, & u \leq 1 \\ 0, & \textit{else} \end{cases}$
Triangular	$K(u) = \begin{cases} 1 - u , & u \leq 1 \\ 0, & \textit{else} \end{cases}$
Gaussian	$K(u) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{u^2}{2}\right)$
Epanechnikov	$K(u) = \begin{cases} \frac{3}{4}(1 - u^2), & u \leq 1 \\ 0, & \textit{else} \end{cases}$

density estimation, of the PDF $p(x)$ associated with the dataset. Our KDE, $\hat{f}(x|h)$, is expressed as:

$$\hat{f}(x|h) = \frac{1}{nh} \sum_{i=1}^n K\left(\frac{x - x_i}{h}\right), \quad (3.5)$$

where $K(\dots)$ is the utilised kernel function and h is a free parameter referred to as the *bandwidth* [Sil87]. The bandwidth parameter corresponds to the variance of each kernel, allowing for the smoothing component of the KDE

to be optimised as required. The selected kernel function must satisfy the condition

$$\int_{-\infty}^{\infty} K(x)dx = 1, K \geq 0 \quad (3.6)$$

ensuring that the probability mass of the global estimate function does not exceed unity, as required by (3.3). Once again, note that this property allows for the KDE function to have density values of greater than one.

By far the most commonly used kernel in empirical data analysis is the Gaussian, owing to its flexible and accurate representation of measured processes encountered in real world applications. However, there exists ample discussion regarding the merits of each of the commonly used $K_h(u)$ functions, with evidence existing that the selection of the kernel function is much less important than the selection of the associated bandwidth h . Despite this, the Gaussian kernel is used exclusively throughout this research due to its prevalence in modelling physical systems [Lyo13].

Fig. 3.4 shows the effect of h selection on the smoothing properties of the KDE, with the bandwidth of a Gaussian kernel altered on a generated dataset; 100 randomly drawn instances from a Gaussian with properties mean $\mu = 50$, standard deviation $\sigma^2 = 3$. There are a few methods [JMS96] of selection for a suitable h value of varying complexity. However, as a general heuristic; the KDE should be smoothed to an extent that removes unnecessary kernel-based features without undue loss of behaviour information from the associated $p(x)$.

For the example illustrated in Fig. 3.4, bandwidth values $h = 0.1$ and to a lesser extent $h = 0.5$ under-smooth the density estimate, retaining features characteristic of the individual data instances. It can be argued that while $h = 10$ provides the most Gaussian-like representation, the low instance count of the dataset should retain some non-Gaussian characteristics in an ideal KDE.

With this in mind, it is argued that $h = 10$ over-smooths the data, and that in this case the *almost*-Gaussian KDE renders by $h = 1$ is the preferable bandwidth value.

KDE-based techniques are useful in providing inference to empirical engineering problems for a number of reasons.

- They are entirely non-parametric: all that is required for a KDE to be created is the data itself. This negates any requirement for *a priori* modelling or reasoning about the behavioural process in the first instance. In fact, this property of KDEs make them particularly useful in data explo-




Figure 3.4: Variation of the kernel bandwidth, with $h = \{0.1, 0.5, 1, 10\}$

ration, where they are often used to determine general characteristics of data-based processes.

- The feedback provided by a KDE is visually intuitive. Peaks in the density function correspond to more probable modes of behaviour, making it quick to ascertain likely modes of behaviour. In providing an output to an engineer about some physical process, for example, the KDE function clearly presents any dominant tendencies and a measure of how often each mode has been observed historically. This allows for lengthy periods of historical operation to be condensed and presented in a single diagram.

3.1.4 b) Empirical cumulative distribution function

Similar to the KDE technique, empirical approaches exist in the approximation of the CDF from real-world data. *Empirical cumulative distribution functions* (ECDFs) are created to estimate the CDF associated with a sample population, with the function is determined by:

$$\hat{f}_{ECDF}(x) = \left(\frac{1}{n}\right) \sum_{i=1}^n \begin{cases} 1, & \text{if } \{x_i \leq x\} \\ 0, & \text{else} \end{cases} \quad (3.7)$$

As touched upon in Section 3.1.3, an accurate representation of a CDF can be useful in a variety of decision making processes associated with statistical inference. In particular, the underlying process underlying observed data can be attributed to a particular family of distribution from the application of ECDF-based inference. As means of an example, an ECDF of an observed vibration measurement from a piece of engineering plant can be compared with an ideal Gaussian CDF in order to determine how Gaussian-like the process is.

3.1.5 Parametric vs. non-parametric

It is useful to examine the concepts of *parametric* and *non-parametric* models. These concepts cover both statistical methods and ML techniques, and understanding their approach to inference problems is important when considering potential solutions.

The differences between these are rooted in how each utilises the data or training domain \mathcal{D} during the inferential stages. Both paradigms have associated advantages, disadvantages and design considerations which impact on their selection for particular problems. It remains important to have a functional grasp of their differences in order to select an appropriate model for the task at hand.

3.1.5 a) Parametric models

Parametric machine learning techniques, as the name suggests, model hypothesis functions through use of a finite number of parameters, sometimes referred to as weights. Common notation defines a parameter vector consisting of i parameters as $\underline{\theta} = \{\theta_1, \dots, \theta_i\}$ where $i \ll n$ for n training examples, and the corresponding hypothesis function defined parametrically as $h_{\underline{\theta}}(\mathbf{x})$ or $h(\mathbf{x}; \underline{\theta})$. An accurately weighted $\underline{\theta}$ encompasses all the required features of a training domain \mathcal{D} , allowing for the training set to be discarded as $\underline{\theta}$ can be used to make accurate hypotheses henceforth. The memory advantages of parametric approaches are obvious; there are no implications on training set size with regard to storage of the trained model, as future predictions made by the trained model only require the learned parameters.

The number of parameters explicitly used in the model is dependent on a number of factors, but it can be asserted that a higher parameter count tends to be required when learning more complex behaviours.

3.1.5 b) Non-parametric models

In contrast, non-parametric models are dependent only on the data or training examples (or a subset of these) at each stage in order to generate hypotheses. Commonly referred to as ‘instance’-based techniques, each data point of \mathcal{D} is mapped to some feature space representation, from which the model inference is then made. This means that non-parametric models require storage of at least a subset of the data used in training to provide future hypotheses on new data.

These methods and algorithms address an important issue associated with parametric models; where the selected parametric view does not represent the characteristics of the data sufficiently to provide accurate enough hypotheses. A common issue in modelling hypotheses parametrically is selecting a suitable model family without detailed knowledge of the problem at hand. Non-parametric models are said to be ‘led by the data’; allowing for the internal structure and characteristics of the training data to exert a more explicit influence on the resulting model outputs. Persisting a subset of the data can also bring about issues when scaling non-parametric models to tackle large-scale problems.

Non-parametric models can be subject to performance and over-fitting issues (a concept explored later in Section 3.2.4), as their complexity grows linearly with the volume of data used in their construction.

3.2 Machine learning fundamentals

3.2.1 Introduction

Among the range of problems computer science and artificial intelligence (AI) has sought to address, engineering systems that learn from experience has

remained one of the most prominent. *Machine learning* (ML) shares theoretical and semantic roots with many of the other sub-fields in AI; most notably AI planning, robotics and the formal modelling of first principles knowledge. Rooted in statistics and statistical inference, the term ‘statistical learning theory’ is often been used in the past in lieu of the term ‘machine learning’, at least to describe a core of the computational concepts exercised by ML techniques and approaches.

ML is an enormous growth field in technology¹, with the applications increasing in multiplicity as more areas become rich in available data. The potential problem solving and commercial impact of learning technologies have not went unnoticed by the cutting edge of the hi-tech sector [Rag18], with large engineering firms like Google [Aea16], Apple [ST17], Amazon [SGS⁺18] and [KSB⁺18] all researching and investing heavily in ML. In light of increased sophistication and data availability, it is only logical for ML to see deployment to new arenas of application; including the engineering reliability industry.

The section will introduce the field of ML, with some focus on the context and requirement for useful ML in a variety of today’s condition monitoring problems. In particular, the techniques used throughout this study will be introduced and discussed before their application to the domain is examined in later sections.

3.2.2 Key concepts

To provide a complete, well-defined problem or set of problems that encompasses all potential learning scenarios is beyond the scope of this thesis (and perhaps even current ML research in general), but it will be of some use to define a loose format to the learning theory and tasks encountered. Mitchell

¹For a perspective on scale, Google Scholar returns around 4.2 million article results in 2018 for the query ‘machine learning’ - which is almost a million more than the term ‘climate change’ which has around 3.4 million entries

[Mit97] provides a succinct definition in:

“A process (computer program or otherwise) is said to learn from experience E with respect to some class of tasks T and performance metric P , if its performance at tasks in T , measured by P , improves with experience E .”

As means of an example, a process might be in developing a program to play the strategy game Go competitively against a skilled human opponent. Learning the fundamentals of such a game may represent a complete set of logical rules (allowed moves, end-game requirements, perhaps some rudimentary heuristics), however this does not necessarily equate to proficiency in playing the game². The nuances of Go strategy and ‘game playing’ - in the game-theoretic sense of term, with regards to the actions of an opponent - are likely only to be achieved through sufficient study of previous games and experiences. Figure 3.5 broadly illustrates such a system.

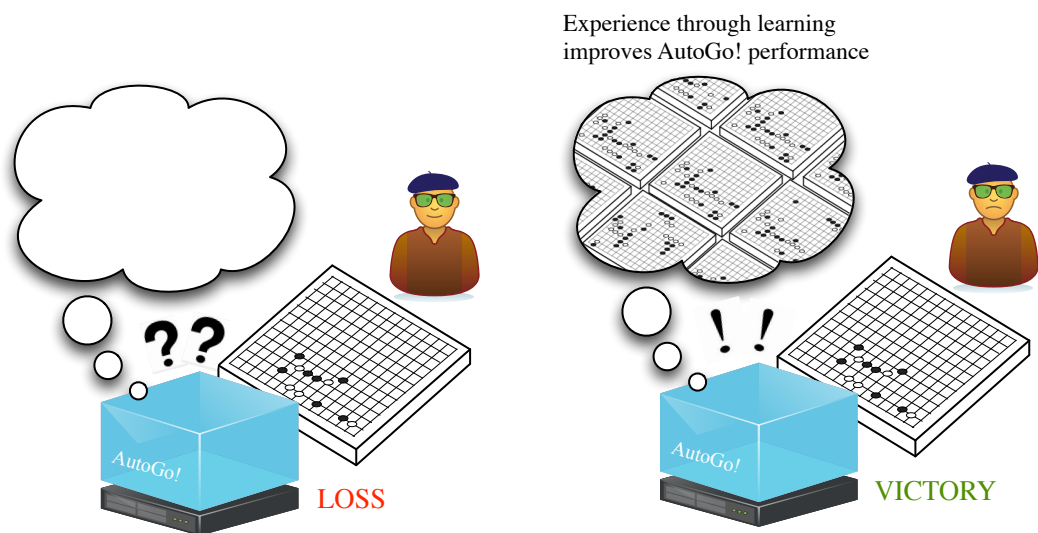


Figure 3.5: Example machine learning process, where a program is constructed and learns to play board game Go

²In fact, Go is a notable example of a game with a simplistic ruleset that remains unsolved from a computational perspective. In contrast, draughts/checkers is an example of a solved game.

In this example, the task T is ability to play Go competitively, the performance metric P will be the result of matches against a human opponent or group of human opponents and the experience E will be historical games played during a training phase or continued use of the program. Prior to sufficient training, 'AutoGo!' will perform poorly against an opponent as its strategy does not contain enough modelled knowledge³ regarding good play. There are a number of means by which a program such as 'AutoGo!' might seek to learn and model competence at the game beyond simply the rules of play during a hypothetical training phase. This is a topic in itself, with Google-backed organisation DeepMind successfully innovating the area of automatic Go play [Gib16] to defeat top-level human opponents. At this stage of the discussion, it is merely important to note that there is an improvement in the performance metric through a learning process of the task at hand. This concept of learning through experience or more generally, learning representation from data, is a key principle of all ML systems and algorithms.

ML can be broadly considered in two categories of approach: *supervised* and *unsupervised* learning. These differ in the availability of target examples (in E) during the training process.

3.2.2 a) Supervised learning

Supervised learning encompasses a large set of inference problems, where the creation of an ML model is undertaken to map a number of *inputs* to corresponding *outputs* or *targets*. The targets can be of discrete or continuous nature ; the most important note is that there is an *a priori* mapping of input to example targets by which the model looks to create a function with.

Formally, we take a set of n variables as an input pattern $\mathbf{x} = \{x_1, \dots, x_n\}$

³The knowledge and reasoning in this thesis will predominantly refer to inductive means. While deductive reasoning encompasses a large component of AI practice, this research is concerned with empiricism and learning from experience.

and map these to a target y . In order to achieve this, supervised learning uses a training domain of m mapped pairs of \mathbf{x} and y (or *labelled data*) $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1, \dots, m}$ to create a hypothesis function $h(\mathbf{x}) = y$. This is constructed such that it provides output, or hypotheses, for all possible input values provided to the function after training. That is, given \mathbf{x}' , a new input vector that may not have been included in the training domain \mathcal{D} , $h(\mathbf{x}')$ provides a suitable corresponding output of y' consistent with the intended target behaviour. This is generally achieved through minimisation of an loss function $\mathcal{L}(h; \mathbf{x}, y)$ during the training phase, which provides a measure of the difference between the actual output hypothesis of the model and the desired target function. Loss functions come in a variety of different formats, such as Euclidean distance, accuracy metrics or mutual information. The selection of these is often dependent on the problem nature and requirements from the model.

3.2.2 b) Unsupervised learning

Unsupervised learning techniques seek to identify patterns and structure from data without any *a priori* labelling information, as was the case in supervised scenarios. Formally, only the input pattern $\mathbf{x} = \{x_1, \dots, x_n\}$ is available: there is no corresponding set of y targets. Unsupervised techniques therefore seek to discover structure, patterns, and regularities in the data empirically. While this can be more challenging than in cases where fully labelled training data is available, the successful application of unsupervised approaches has the potential to tackle many more problems than in supervised, as labelled data is often difficult to acquire.

Pure unsupervised, or unsupervised augmenting supervised elements ('semi-supervised') techniques are being applied to speech recognition, natural language processing [CW14] and many other data-rich environments. Techniques range from statistical methods such as principal component analysis [Jol11] &

independent component analysis [HKO04], to deep neural network-enabled approaches like generative adversarial networks [RMS15].

3.2.2 c) Classification

In examples of *classification*, assignment is made of an input pattern \mathbf{x} to one of K discrete class labels. The learned $h(\mathbf{x})$ provides the model for assigning the corresponding class identified by the model for \mathbf{x}' , rendering $y \in \{c_1, \dots, c_K\}$ - a discrete output provided from a finite set. There exists a number of sub-problems under the umbrella of classification, of which it will be useful to examine three. Note that the following terms are provided with particular focus on application to engineering reliability problems, and these may vary in other application areas.

Unary classification ($K = 1$) creates a single membership class from training data \mathcal{D} . The output of $h(\mathbf{x})$ is therefore a decision on the membership of test data to this label. \mathcal{D} is often a nominal, routine or steady-state behaviour that requires continued monitoring or surveillance. Anomaly detection techniques such as this are therefore utilised to automatically flag changes in behaviour from the established normal.

Binary classification ($K = 2$) examines two behaviours within \mathcal{D} , often creating a decision boundary between them separating the extent of each. This particular family of classification is associated with some of the most rudimentary machine learning approaches used as benchmarks in the field, with particular approaches often being examined on their ability to perform in binary classification scenarios.

Finally, *multi-class classification* ($K > 2$) seeks to provide automated labelling of any number from a family of classes. This introduces a number of issues in terms of label discrimination between classes; the decision strategy between hypotheses now must examine likelihood of membership to a label

against all other potential labels for test data.

Typical classification tasks exist in:

information security; where an example task might be to label a test input \mathbf{x}' representing credit card use as '*Normal*', with failure to do so implying suspicious behaviour (unary classification or anomaly detection),

e-mail spam filtering; where an example task might be to label a test input \mathbf{x}' representing an e-mail as '*Spam*' or '*Not Spam*' (binary classification),

handwriting recognition; where an example task might be to label a test input \mathbf{x}' representing a 128×128 bitmapped image of a written character as a letter from the English alphabet (multi-class classification).

3.2.2 d) Regression

Regression techniques provide an estimate or prediction of a continuous value dependent on the properties of a given input pattern $\underline{\mathbf{x}}$, rendering $y \in \mathbb{R}$ or $y \in \mathbb{R}^D$ for D -dimensional problems. In this task, the learned function $h(\mathbf{x})$ is used to determine the relationship between a dependent variable and a number of independent variables, and uses this to provide output when examined with new test data \mathbf{x}' .

For example, it might be useful to learn the functional relationship between two variables in a sample of historical data as this may provide some information about the population the sample was taken from. Parameterising common functions such as linear, power or exponential relationships between observables is often the first stage of analysis for many data problems. Fig. 3.6 shows two bivariate data examples, taken from linear and exponential-type populations, along with some example regression curves overlaid.




Figure 3.6: Two bivariate data relationships with a high likelihood of an underlying relationship, which would typically be quantified through regression

A successful $h(\mathbf{x})$ effectively ‘fits the curve’ in the bivariate-type examples illustrated; however, it should be noted that regression is not limited in dimensionality to examples where a ‘curve’ can be easily visualised. Regression-based inference can be carried out in any number of dimensions, though this is not always a preferred approach due to the ‘curse of dimensionality’ [Tru79] problem.

Generic regression analysis precedes ML and even modern computation by a notable period, with a rich history existing within mathematical and statistical research surrounding the topic. Legendre’s paper [Leg05] on the trajectory of comets provided an appendix introducing the technique of *least squares*, a now ubiquitous method in the empirical fitting of continuous functions which has become synonymous with regression.

3.2.2 e) Clustering

Clustering can be considered the among the most important subset of commonly encountered unsupervised problems. A well formulated cluster-based procedure can be utilised to direct analysts towards features of interest in

very large datasets; potentially extracting regularity from previously considered overly large data. This form of dimensionality is particularly popular in building scalable ML systems for applications with little domain knowledge available.

At the most generic level, a clustering technique will seek to group together examples in order to determine any macro-organisation to the data, commonly using a distance or statistical metric in defining potential group membership.




Figure 3.7: An example clustering problem, visualised in D^3 phase space

Fig. 3.7 provides an example of a typical clustering-type phase space. Three distinct groupings of data exist in the 3-dimensional space, and without labels the empirical properties of the dataset need to be examined in order to enable classification of future examples into these groups. The phase space itself can be pre-processed in order to maximise the likelihood of a suitable clustering result being found through dimensionality reduction techniques such as principal component analysis [Jol11] and t-SNE embeddings [MH08].

3.2.3 Deep learning

While not explicitly examined in this work, it is useful to acknowledge the volume of work being undertaken under the heading of *deep learning*. This refers to a family of neural network-based techniques which rely on large volumes of training data to construct networks with large numbers of nodes and weights. Deep networks have the ability to create highly complex, nonlinear representations of almost any learning problem (given enough training data and compute resources), as they allow for features themselves to be learned [GBC16] as opposed to engineered. Their training can be an expensive and time-consuming process, and interpreted their outputs in an explicability-focused domain can be a challenge.

3.2.4 Overfitting

It is important during the process of training models to avoid fitting the classifier or regressor to reflect every single training example used in the training stage, as this would likely be training the model on noise as opposed to the useful signal in the data. The aim of a well-trained ML model is to provide useful and accurate predictions on unseen or future inputs: not to predict the already seen, training set inputs.

Fig. 3.8 provides a clear comparison of a suitable and an over-fitted classification model. The data \mathcal{D} used for training this example comprises two distinct classes, separable by a non-linear decision function. In learning the decision boundary to provide a classification on test data, the diagram (a) uses a broad ellipse as the decision boundary. Conversely, the approach from (b) creates a complex high order function modelling the behaviour interfaces with great detail. While it is true that (b) successfully classifies each of the training instances existing in \mathcal{D} , with (a)'s decision boundary encompassing a misclas-

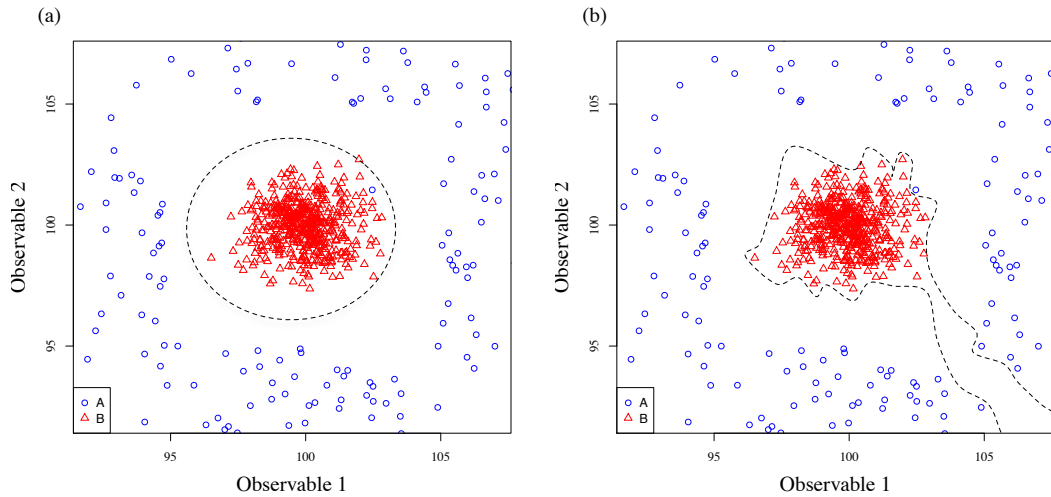


Figure 3.8: Supervised learning example (classification), with comparison of ideal and over-fitted decision functions

sified instance, the computational complexity associated with generating (b)'s hypothesis makes the approach of (a) preferable. The intricacy of (b)'s decision function does not well represent the relatively simplistic task in this example, instead introducing unnecessary features in the decision making procedure. Applying these learned models to similar examples of ellipse-like binary classification test sets would illustrate the lack of generalisability of (b).

3.3 Machine learning approaches

In this section, the ML methods employed in the research are introduced to provide context on how the data-driven models were constructed.

3.3.1 Linear classification

Consider a linearly separable binary classification problem, defined by a training set $\{(\mathbf{x}_i, y_i)\}_{i=1, \dots, m}$ where m is the number of training tuples and $y \in \{-1, 1\}$. y_i is the label for the i -th multidimensional input pattern \mathbf{x}_i . A visual representation of a two dimensional example of such a training domain is

provided in Fig. 3.9 for clarity.

From the illustration in Fig. 3.9, it can be seen that an infinite number of linear separating hyperplanes (known as *decision boundaries*) between the two classes prospectively exist, with each being a straight line bisecting the training domain and acting as a discriminator.

Linear decision boundaries in two dimensional examples take the generic parametric form $\theta_1 \mathbf{x} + \theta_0$; the familiar 'equation of a straight line'. Classification from a successfully learned separating hyperplane on a test input pattern \mathbf{x}' is achieved by examining the decision function:

$$y(\mathbf{x}'|\underline{\theta}) = \text{sgn}(\underline{\theta}^T \mathbf{x}' + \theta_0) \quad (3.8)$$




Figure 3.9: Typical linearly separable binary classification training domain in two dimensions

where $sgn(\cdot)$ denotes a *sign* or *threshold function*, defined as:

$$sgn(x) = \begin{cases} 1, & x \geq 0 \\ -1, & x < 0 \end{cases} \quad (3.9)$$

This should be intuitive as the points along any given hyperplane satisfy $\underline{\theta}^T \mathbf{x} + \theta_0 = 0$, so any hyperplane that satisfies a linearly separable classification problem can say that points corresponding to $\underline{\theta}^T \mathbf{x} + \theta_0 \geq 0$ belong to one class, and the opposite by extension.

It should be noted that the parametric nature of the linear discriminant function as described do not necessitate parametric-style learning approaches; as will become evident when examining kernel methods in application to linear classification.

3.3.2 Perceptron learning

One of the earliest techniques⁴ investigated in classification problems [Ros58], the *perceptron* is an example of an *artificial neural network* (ANN) - a learning algorithm directly inspired by the cognitive processes undergone in the brain [Heb49]. While used less often in application today (due in part to its limited abilities in non-linear, non-linearly separable scenarios), it's useful to understand the origins of linear classification approaches.

A great number of ML problems have been tackled using perceptron-descendant processes: multi-layer ANNs used in autonomous vehicle research [Pom89], networks with recurrent properties in examining temporal data [Xin07] and cascaded networks used to investigate highly non-linear features within features [WRM12].

⁴Note that the perceptron is more formally an optimisation algorithm that can be used in a variety of different ML techniques. The perceptron discussed herein is the simplest example of such a technique, and might be alternatively labelled a 'linear discriminant model'

The learning of a suitable hypothesis function is achieved through the algorithm:

- A primary state is initialised (typically $\underline{\theta}, \theta_0 = 0$),
- With values for $\underline{\theta}, \theta_0$, the function is examined for each input example - comparing the hypothesised output with each target label. In the instance $\underline{\theta}, \theta_0 = 0$, every example is labelled $y = +1$ due to (3.9),
- When a disparity between an output and target label exists (when $y = +1$ is hypothesis against an example labelled $y = -1$, for example), the perceptron amends the values of $\underline{\theta}, \theta_0$ to better reflect the training examples.

$\underline{\theta}$'s values are updated at each iteration using the *perceptron rule*. For the i^{th} parameter:

$$\theta_i \leftarrow \theta_i + \alpha(y - h_{\theta}(\mathbf{x})) \times x_i \quad (3.10)$$

where α is a selected parameter known as the *learning rate*; the selection of which impacts the magnitude with which updates are made to the values of $\underline{\theta}$.

It has been proven that the perceptron will converge successfully in linearly separable scenarios using (3.12). Issues begin to arise when the training domain *is not* linearly separable; the XOR function was famously proposed as the perceptron's downfall [Min88]; a limitation addressed by the development of multilayer perceptrons in the 1980s. Thus in order to use perceptron learning-based methods, the phase space for training needs to have a potential linear classifier which solves the decision problem in that data space.

3.3.3 Logistic model

Linear discriminant functions reliant on solely a discrete sign function as described in 3.3.1 are characterised by a 'hard' threshold; training examples close




Figure 3.10: Comparison of standard discrete and logistic threshold functions

to the hypothesis hyperplane are as equally representative of the class as examples further away. However, it can be useful to introduce a graduation of the importance assigned to examples at the boundary using some function dependent on their distance from any candidate hyperplane.

As opposed to applying a stepwise thresholding function, the *logistic* or *sigmoid function* can be used to output class membership hypotheses on candidate test data. The threshold function recast to take logistic properties is defined as:

$$y(\mathbf{x}'|\underline{\theta}) = \frac{1}{1 + e^{\underline{\theta}^T \mathbf{x}'}} \quad (3.11)$$

The logistic function is compared to the discrete sign function of (3.9) in Fig. 3.10, illustrating its continuous, smoothed nature.

In contrast to the perceptron algorithm of the previous section, the pro-

cess of using the logistic function for linear modelling⁵ seeks to minimise an L_2 loss function. This is due to the continuous nature of the logistic function as $h(y|x_i) \in [0, 1]$, as opposed to the discrete $h(y|x_i) = \{0, 1\}$ as with the perceptron algorithm. Thus, minimising the loss function is achieved through updating the i^{th} parameter:

$$\theta_i \leftarrow \theta_i + \alpha(y - h_\theta(\mathbf{x})) \times h_\theta(\mathbf{x})(1 - h_\theta(\mathbf{x})) \times x_i \quad (3.12)$$

The continuous nature of the thesholding function is particularly useful when dealing with noisy datasets where absolute linear separation is not possible due to instance overlap at the boundary. Such properties have made the approach popular in application to a wide variety of areas.

3.3.4 Kernel methods

Rather than considering a family of parametrised discriminant functions and iterating through error space to reach some loss function minima, kernel methods recast the points of the training domain as a given kernel function mapping in a new feature space [Bis06] in order to perform inference. Utilising the training domain in a feature mapping in this fashion makes kernel methods a non-parametric family of approaches.

A generic mapping form can be expressed as:

$$k(\mathbf{x}_1, \mathbf{x}_2) = \phi(\mathbf{x}_1) \cdot \phi(\mathbf{x}_2) \quad (3.13)$$

where $\phi(\cdot)$ is a defined kernel mapping function. Kernel space views of \mathcal{D} can be made utilising a variety of kernel mapping selections, which is paral-

⁵Much of the ML literature use the term 'logistic regression' to encompass linear modelling using the logistic function. To simplify the terminology in this research and avoid unnecessary confusion between classification and regression, 'logistic regression' will not be used to describe the technique

led to the kernel selection process outlined in 3.1.4 a). The most rudimentary kernel function is the linear kernel $\phi(\mathbf{x}) = \mathbf{x}$.

3.3.4 a) Support vector machines

The support vector machine (SVM), largely credited to Vapnik [CV95], has applications to both regression and classification problems. For classification, SVMs tackle decision function selection through the creation of a maximum-margin between data classes, their associated kernels and the function itself.

Applied to linearly separable cases, it should be obvious that any perceptron-based classifier can provide multiple hypothesised decision functions which successfully separate the training classes and fulfil the model requirements. Use of the logistic function (as examined in 3.3.3) to threshold generated hypotheses on data instances further informs the selection of a plane among these candidates. However, this takes into consideration each of the examples provided in the training domain, and therefore can be susceptible to outliers.

Using the feature space concept introduced in 3.3.4, the threshold expression (3.9) can be recast as:

$$h(\mathbf{x}'|y) = \text{sgn}\left(\sum_i \theta_i y_i (\mathbf{x}' \cdot \mathbf{x}_i) - \theta_0\right) \quad (3.14)$$

The SVM inherently seeks the case which generalises best through its inferential process of maximising the class discriminant margin; providing the most effective hyperplane by how distant it is from the class fringes. This is achieved by placing greater emphasis on the instances closer to the candidate separating function (the support vectors), and seeking the best hypothesis using this information. The optimisation goal for the SVM can take a few forms, with the most common variants being $L1$ - and $L2$ -SVMs. For $L1$ -SVMs, this is:

$$\text{minimise } \frac{1}{2} \|\theta\|^2 + C \sum_{i=1}^m \xi_i \quad (3.15)$$

whereas for $L2 - SVMs$ this is changed to:

$$\text{minimise } \frac{1}{2} \|\theta\|^2 + \frac{C}{2} \sum_{i=1}^m \xi_i^2 \quad (3.16)$$

where C is the margin hyperparameter and ξ is the slack variable - both hyperparameters which dictate the specificity and generalisability of the SVM approach to dealing with outliers and errors [KS03].

SVMs have proven successful across a very wide range [MSR⁺97, OFG97, OMS⁺15] of data scenarios, often ranking among the most accurate in predictive measures across numerous domains.

3.3.5 Ensemble methods

ML techniques can be scaled to incorporate multiple discrete models informing a final decision, by averaging or aggregating their individual outputs. These are known as *ensemble methods*, and this section will introduce two of the most successful families of approaches in this area.

3.3.5 a) Random forests

A decision trees is a graph-based conditional inference modelling technique which branches data points into discrete decisions based on a learned set of rules. *Random forests* represent the ensembling of numerous decision trees [Ho95], and taking an aggregate of their output in order to come to a final decision. Each tree is trained concurrently on a subset of the training data in order

to prevent overfitting and ensure the ensemble is approximately even in terms of individual model contributions. Random forests often perform well in high dimensional spaces [CKY08] due to this ensembling and subsetting property, which introduces stability and robustness when generalising from the training data.

3.3.5 b) Gradient boosting machines

A related tree-based approach is the *gradient boosting machine*, which differ from random forests in the manner in which decisions trees in the ensemble are constructed [Fri01]. Each tree is trained one at a time, and subsequent models are added to the ensemble when adding branches to the tree has diminishing returns in terms of model performance. Gradient boosting machines have exhibited similar, or in some cases improved [OP11], benchmarking over random forests.

3.3.6 Clustering

This section outlines the clustering methods used in this research.

3.3.6 a) k-means clustering

k-means clustering [Bis06] classifies a dataset into *k* groupings, where every data instance belongs to the cluster with the nearest mean in Euclidean space. The value *k* is defined *a priori*, which makes it suitable when there is a likely number of clusters that required to be identified. It can be used for cluster identification, or as a dimensionality reduction method to potentially improve the results of any vector-based ML [Mur12] approach.

3.3.6 b) Gaussian mixture models

Mixture models are a family of statistical methods used for calculating the likelihood of multiple statistical distributions contributing to an overall observed dataset. Gaussian mixture models (GMMs) specifically seek out multiple Gaussian distributions (as introduced in Section 3.1.1) in a population, and aim to provide the model parameters of the hypothesised contributory distributions.

There are a number of potential ways to implement GMMs: however the methods used in this study are reliant on the expectation-maximisation [Moo96] methodology. This consists of initialising N Gaussians with random parameters, and the probability that every data point in the population came from each of the proposed distributions is calculated and then aggregated. The parameter space of the distributions is then explored in order to maximise the total expectation of the mixture hypothesis.

3.3.6 c) Hierarchical clustering

The aforementioned clustering methods do not consider inter-instance or inter-cluster relationships when defining a suitable grouping solution. In practice, clusters might share properties or sub-populations of the data: something that could be useful to know when considering large, multi-dimensional datasets. *Hierarchical clustering* [Lia05] addresses this by constructing inter-cluster relationships based on similarity to build a holistic ranking and aggregative view of the proposed groupings.

3.4 Machine learning in reliability

The exploitation of ML techniques in commercial settings is typically associated with areas of application where large datasets are the norm. Industries be-

yond the research disciplines of theoretical computer science and statistics, including information retrieval, finance, e-commerce and bio-informatics, have seen adoption of many *de facto* standard ML algorithms [MLY17, Gao14] and approaches to common data problems. Under the general umbrella of ‘data science’, machine learning tools are also being used in central roles in astrophysics, digital advertising and recommendation systems.

Engineering, and more specifically reliability engineering, has embraced a handful of key techniques on a largely requirements-driven basis. Typically, a problem has been encountered and a solution using a particular class of ML technique was applied to tackle the requirements of that example. There has been comparatively little study in generalising such systems across wide categories of reliability type scenarios; resulting in the various sub-fields of the discipline adopting small numbers of widely used techniques within their research spheres.

There has been notable application of ML (and intelligent systems in general) in the so-called ‘mission-critical’ or ‘function-critical’ areas of defence and aerospace. These areas are characterised by a wealth of archived time series data, which lends itself to a statistically-driven approach to tasks in surveillance, control and automation. Unmanned aerial vehicle navigation [KLO⁺12], impact damage monitoring and online jet engine behaviour inference [CTM⁺08] each have had recent developments.

3.4.1 Kernel methods in CM

The SVM has proved to be a useful tool for many machinery CM scenarios, with the potential of the technique being heralded in publications as early as 2007 [WY07]. Both classification (fault identification) and regression (prediction and prognostics) of the algorithm have seen experimentation and use in the field with success.

Bearing element [YS02] and induction motor faults, magnetic measurements in rotating electric machines [PNA⁺02] and compressor valve failures [CX10] each have example systems which use a standard (or simple variation, through data pre- or post-processing) classification SVM implementation. Multi-class [SSR08] systems have also been investigated, including using the SVM as part of an ensemble of methods [NBTT12].

SVM regression naturally corresponds to the the empirical learning process associated with latent degradation functions that may be existent in medium and long-term time series data. Future state change indication metrics, for example in shaft misalignment prediction [OJBH06] applied to rotomachinery, map well to extrapolating a learned function to some time horizon, with the data-led properties of the SVM accounting for any potential non-linearities. A combination of both SVM classification and regression methods was used in electronic prognostic [SP07] systems; combining state estimation and extrapolation to predict the incidence of future failure states.

The crisp output of a learned SVM function can be further augmented with application of the relevance vector machine (RVM) algorithm [Tip00]; a technique which has been especially popular among PHM research circles as it augments the standard SVM approach with probabilistic output. Crack growth through fatigue [ZD12], bearing condition [CWT⁺11] and battery lifetime [SGPC09] have each used the Bayesian enhancements of the RVM to provide parameter trajectory estimations with a corresponding confidence interval for RUL outputs, a feature difficult to present with a simple SVM procedure. This is especially useful in delivering engineering decision support, where communicating the degree of certainty associated with a prediction is often as important as the hypothesis itself.

3.4.2 Other techniques

Both ‘shallow’ and deep neural network (NN) ML approaches have also been historically popular in reliability; with numerous nuclear-specific [SAT03, EB04, VBKR03, SJZW17] systems developed over the previous decade and a half. This rate of uptake has the potential to be even higher as more is invested in data-driven solutions to problems in heavily regulated sectors; as historically NNs have been met with some criticism for their ‘black-box’ nature.

Random forests have seen adoption [CSS⁺15, Yan06] due to their impressive results coupled with their clear explicability when providing reasoning behind decision-making. The flexibility of Gaussian processes [IAT13] in data-driven problem spaces has also made them an area of interest for empirical condition monitoring studies

3.4.3 Potential

The already existent developments in reliability-based ML aside, the CM discipline as a whole has a great opportunity to benefit from the further development and application of learning systems to the process of maintaining engineering machinery and assets. Outlined in 2.5.2 a), the volume of operational information now available to the reliability professional is enormous and growing notably each year. Meeker and Yilli [MH14] discuss the direct implications of the ‘big data’ revolution on the reliability field, road-mapping a number of the key growth areas alongside points which need further development in order to provide meaningful feedback from the myriad data streams becoming available. Extracting meaningful inference from data is vital to the success of future reliability systems; and bridging any perceived gap between purely data-driven systems and the existent tacit knowledge of the domain experts should be high on the list of priorities.

Considered alone, simply the increased volume of data presents two important issues to the CM professional:

- How is the current reliability practice undertaken by CM engineers scaled up to meet the multiplicity of data streams and data archives?
- How can the current reliability practice be improved upon with this previous unavailable quantity and variety of data about machines?

Many of the processes undertaken by the reliability professional rely on the analysis of data-based features; of which a large number are well-defined pattern recognition tasks. This examination of low-level data streams for familiar 'signatures' can be addressed by the creation of well-designed ML systems with the ability to interpret streams of data and information. Where perhaps the greatest potential for success lies is in the combination of knowledge and automatic inference; allowing existing domain expertise to inform the empirical approaches of ML-type systems. The research presented in this thesis deals with this concept in particular.

At the other extreme of the scale, many hypothesised analysis methods can rely on often complex aggregations and 'scaling-up' of data-based features across large volumes of information; a process which is often inconceivable for even a team of engineers to undertake manually. Examples might be in the automated detection of a specific time series pattern from terabytes of historical operation records; or a post-processing metric of a complex system change which required numerous levels of computation to provide a meaningful decision support output. Investigating and developing systems which learn from data can address these scenarios.

Chapter 4

Self-tuning diagnostics in rotating machinery

This chapter describes the first of the major contributions from this research: the development of statistical techniques in application to automated alarm inference for rotating machinery in the nuclear generation environment. The engineering practice of alarm analysis is introduced, before a number of inferential approaches are presented which improve on the existing procedure undertaken by professionals in the domain.

The novel contributions from this chapter can be listed as:

- Augmentation of an existing knowledge-based intelligent system with ML and statistical inference techniques, providing an improved hybrid intelligent system tackling the engineering problem of routine alarm analysis in TGs,
- A self-tuning framework for vibration diagnostics, allowing for the application a routine alarm knowledge base across an entire asset family under a single maintenance regime,
- The use of techniques in statistical inference to automatically define periods of system normality and transient behaviour in rotomachinery vi-

bration data.

4.1 Routine alarms

With regards to the common health monitoring approaches introduced throughout Chapter 2, the industry standard in vibration-based condition monitoring continues to rely on examinations of steady-state limits and boundaries of operation in the identification of undesirable machine behaviours. These measures do correspond to industry-agreed standards and metrics [iso09]. However, the dynamic and non-stationary nature of power plant operation is often not considered in greater depth beyond these high-level features.

The day-to-day schedule of a NPP like the AGR is made up of numerous distinct events including online maintenance, ad-hoc adjustments to system settings and operator interventions such as channel refuelling. Each of these carry the potential to impart change to the conditions experienced by the plant items, including those involved both explicitly and implicitly with the main primary cycle nuclear system. A majority of the change behaviours covering these scenarios are not to be considered anomalous, as there is as there is a clear causality between a planned change in system state and any alteration in the experienced machine state. This clear link between signals such as vibration response and operational observables forms the basis for identifying periods of change non-indicative of system damage. A hypothetical example of this sort of data feature is provided in Fig. 4.1, where there is a clear temporal correlation between the vibration response change and the underlying operational observable. A typical example of this might be an adjustment made to the reactor load output, which results in a corresponding change in state for the rotating assets in the primary and secondary cycle.

Events of this type can exceed the aforementioned alarm and alerting lim-




Figure 4.1: Example coupling of response and operational observable

its, resulting in a call to action and analysis for the engineering team. With the large number of data streams monitored by the CM professional and the strong regulation required in the nuclear industry, these incidents will often require analysis to verify their benign nature. When scaled to the numerous machines reliability teams are tasked with monitoring, a significant proportion of engineering effort is dedicated to the examination of these events. For example, a typical UK-based NPP will house two turbine generators and up to eight gas circulator units, each with upwards of twenty monitored vibration channels instrumented with the potential to produce time series measurements multiple times per second. This is before considering other non-primary and secondary auxiliary rotating plant items which may be used in other areas of the plant.

Instances lacking a corresponding operational tie (represented by the first response feature shown in Fig. 4.2) are the main focus of the majority of vibration monitoring and surveillance efforts, where this might be due to system damage or unexpected behaviours.

Alarms with this operational change-response link that are non-indicative of damage (as in Fig. 4.1) are referred to as *routine alarms*. The associated engineering analysis required from these events can be rudimentary: often as simple as noting correlation between an operational change and the vibration




Figure 4.2: Example feature in the response observable without a corresponding operational change

response. Despite this, the volumes of alerts and machines requiring sign-off can become time-consuming and open the potential for human error. This has prompted the investigation of intelligent approaches to augment the current processes for dealing with routine alarms, with the aim of automating simplistic analyses and allowing reliability teams more time for examining complex machine health indicators.

4.2 Knowledge-based system

The techniques outlined in this chapter build on existing investigations into knowledge-based inference [Tod09, TMMS07] for steam turbine diagnostics. This research focused on the initial stages of automating the routine alarm process with a knowledge-based system (KBS) approach. The typical engineering workflow undertaken by vibration diagnostics professionals was modelled in detail, using the principles of knowledge engineering [SWdH⁺83] to document both the explicit and tacit understanding of the diagnosis of a routine alarm directly from the expert. Augmenting this existing rule-based approach was selected due to the explanatory power inherent with knowledge-based systems: crisp rules make it clear why a particular automated decision was taken by any

intelligent system.

A rule base for turbine generator routine alarm diagnostics was created, which takes bearing pedestal-mounted vibration and, operational signals as input and delivers decision support on potential alarming instances. This system provides a rule-based conclusion on common alarm causes, following the same inference process an industry professional would. In order to use the knowledge codified in the KBS, the process extracts time series features from turbine data which are subsequently used to build an overview of the behaviour observed by the machinery. This section provides a technical overview of the system, before outlining the areas for improvement addressed in this chapter.

4.2.1 Rule base

The rule base comprises around 180 crisp IF-THEN clauses, with each of the instances corresponding to either an intermediate or final step in reaching a decision regarding the cause of an alarm notification. The structure of a typical final rule chain in making an alarm conclusion is shown in Fig. 4.3. Each of these steps outline the logical process taken by an engineer in examining a vibration observable time series to determine if a behavioural change has occurred due to a step in an operational observable.

The inference process accompanying Fig. 4.3 can be outlined as:

The **alarm fired** node needs to be asserted TRUE.

Three **vibration observable** nodes of four available need to be asserted TRUE: any trend in vibration needs to be within a set period of the alarm ($\text{VIB trend period} \leq \text{alarm period} + \text{threshold}$) and must last at least as long as the alarm period ($\text{VIB trend period} > \text{alarm period}$). The vibration needs to be high (VIB

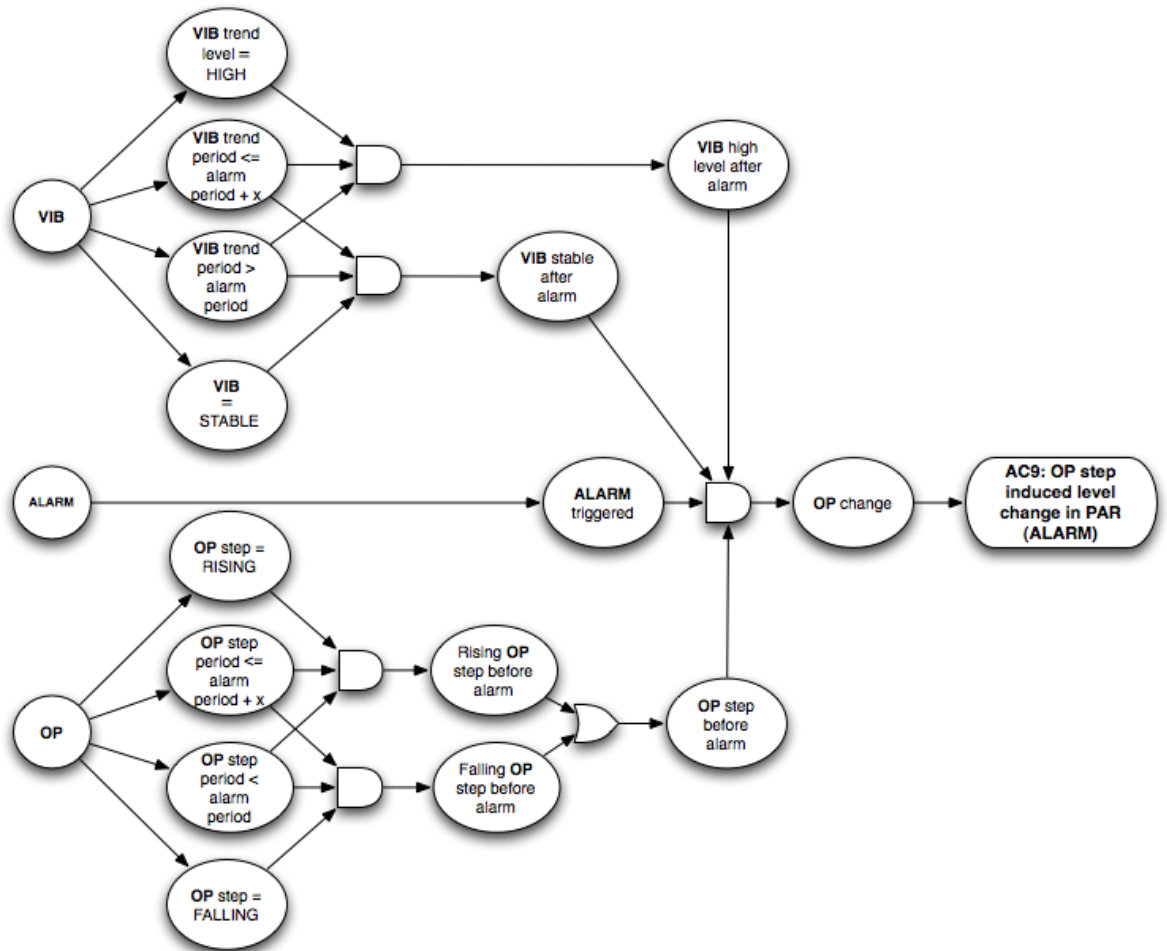


Figure 4.3: Inference diagram of typical rule chain used in the knowledge-based system

trend level = HIGH) and stable (VIB level = STABLE).

Three **operational observable** nodes of four available need to be asserted TRUE: any step in operational observable needs to be within a set period of the alarm ($OP\ step\ period \leq alarm\ period + threshold$) and must last at least as long as the alarm period ($OP\ step\ period > alarm\ period$). The operational observable step can either be rising ($OP\ step = RISING$) or falling ($OP\ step = FALLING$).

This discrete, modular approach to decision-making was selected as it resembles the logical process undertaken by the domain experts themselves. For example, a vibration engineer working on routine alarm investigations will identify a step change feature as a key underlying root cause when looking for clues to explain a machine state change, as opposed to examining each individual data point in the time series discretely. The similarity of this symbolic approach to that utilised by KBS-enabled inference systems is very powerful in providing clear, explicable reasoning behind automatic inference and decision support.

The particular rule illustrated in Fig. 4.3 falls neatly into a correlative-style problem: there is an observed effect and the objective is to ascertain the most likely root cause from a selection of potential stressors. This is the common format for each of the rules in the knowledge base. While it is also useful to take a more quantitative approach (calculating, for example, the pairwise Pearson's coefficient [RN88] for each of the observables in question) the associated explicability of a detected and discrete feature that aligns with the evidence an expert would normally look for is powerful in providing evidence in the post-diagnosis context. This is especially true considering the required transparency and conciseness required when building engineering systems for highly regulated industries like nuclear generation.

4.2.2 Signal-to-symbol transformation

The features used by the rule system are generated through a process of *signal-to-symbol transformation* (STST). This is a classical AI technique which takes continuous variables (i.e. time series streams) and maps these to a discrete representation (e.g. trend and step events) which can be reasoned about through propositional logic. A powerful aspect of this approach is clarity: showing evidence for extracted features and how they relate to straightforward rule-based inference is a concise approach for reasoning about problems and delivering effective decision support.

Three time series features are defined by the system: trend, step and level/impulse. The level parameters help define the spot artefacts which correspond to impulse changes (short deviations from outside the expected distribution of values). Table 4.1 provides a definition of each of the symbol types in relation to time series data.

Table 4.1: Definitions of time series primitives

Primitive	Definition
<i>Level/Impulse</i>	Corresponds to thresholds of expected behaviour; defined by parameters of upper & lower limits $\{\Lambda_{Upper}, \Lambda_{Lower}\}$
<i>Trend</i>	Used in the identification of rising and falling trends in the data. Defined by parameters of trend tolerances $\{\tau_{Upper}, \tau_{Lower}\}$, and trend period T_{Trend}
<i>Step</i>	Used in the classification of rising and falling step changes in the data. Defined by parameters of lead & tail tolerances $\{\tau_{Lead}, \tau_{Tail}\}$, minimum step magnitude $ S $ and step period T_{Step}

These are defined for each observable covered by the KBS, meaning that the parameter values for a wide variety of observables and potential machine states need to be initialised correctly in order for the system to provide inference correctly on TG data streams.

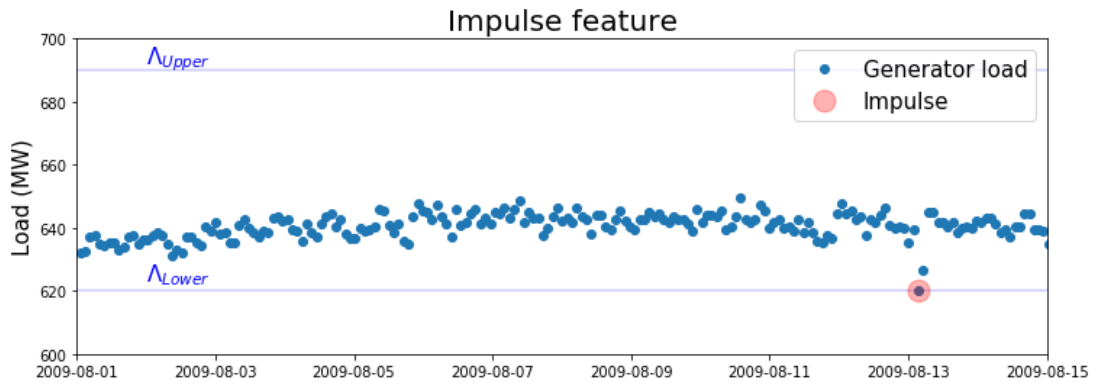


Figure 4.4: Example of impulse extraction, w.r.t. bounds of operation

4.2.2 a) Extracting impulses

Identifying anomalous instantaneous measurements from each time series stream is achieved through setting bounds of operation relative to the expected value of the observable, and flagging those data points which fall outside such envelopes. The limit parameters $\{\Lambda_{Upper}, \Lambda_{Lower}\}$ define these boundary values, and are set independently for each of the observables on the machine.

Any instantaneous value which exceeds or falls below the set boundary values of Λ_{Upper} and Λ_{Lower} will generate an impulse symbol in the STST module.

The impulse extraction procedure can be expressed as:

$$M_{Impulse} = [0, 0, \dots]$$

for x_t **in** x **do**

if $x_t > \Lambda_{Upper}$ **or** $x_t < \Lambda_{Lower}$ **then**

$$M_{Impulse}(x_t) = 1$$

end if

end for

where x is the time series of any given monitored observable, x_t is a single time series value at time t , and $M_{Impulse}$ is the index mask of x denoting the indexes of impulses.

4.2.2 b) Extracting trends

Trends are identified by slicing the data into regular periods, and examining rudimentary features of the time series values within these. *Trend periods* are flagged by comparing the start mean (μ_S , first N measurements) with the end mean (μ_E , final N measurements) and evaluating the following:

$$\Delta_{Trend} = \frac{(\mu_E - \mu_S)}{\mu_S} \quad (4.1)$$

where Δ_{Trend} is the calculated trend difference. If this value exceeds τ_{Upper} , or is lower than τ_{Lower} , then that qualifies the period as a trend symbol. 4.1 is normalised by μ_S to provide the change as a factor of the original value.

The initial construction of the KBS settled on the period $t_t = 24hrs$ (to remain consistent with the daily cadence of manual examination) and $N = 5$

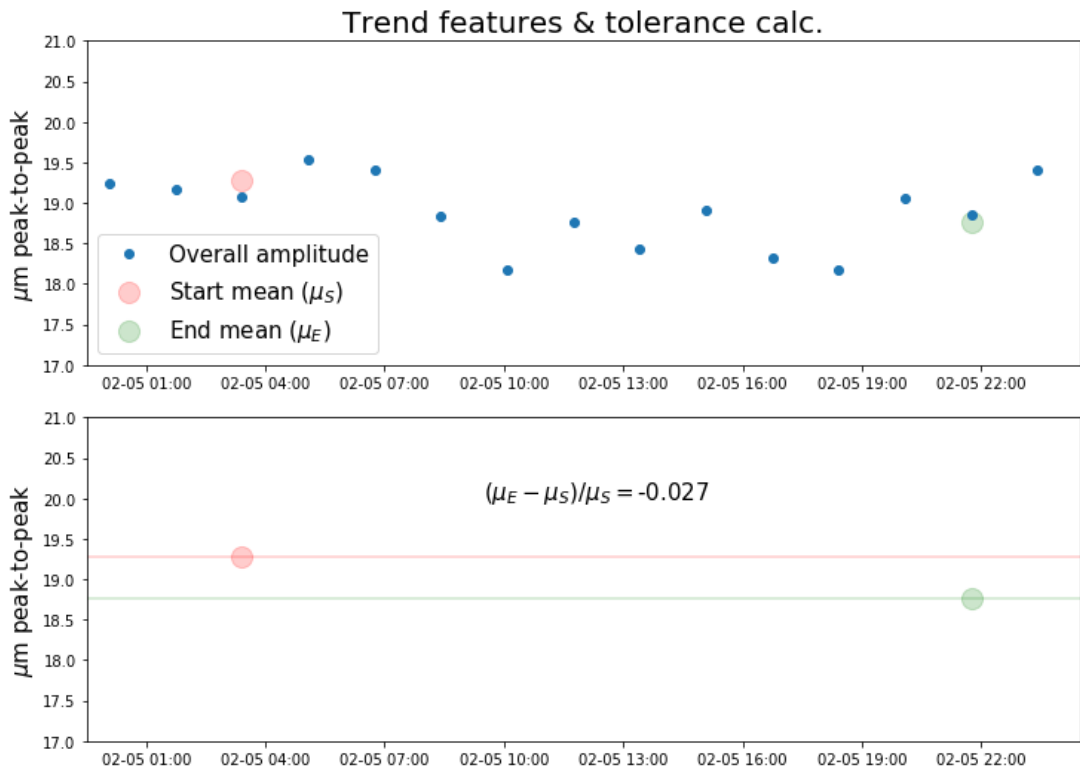


Figure 4.5: Example of features and calculation used in extracting trends

values for comparison at the start and end of the periods. The difference in the dimensions between these (hours for trend window, number of measurements for pre- and post-averages) can be attributed to the potential for non-regular measurements to be taken at varying fidelities of data acquisition (e.g. measurements may be available at sub-minute and minute period intervals at different parts of the assets operational history). Values for 4.1 for each of the periods are compared with the STST time series primitive parameters $\{\tau_{Upper}, \tau_{Lower}\}$ (depending on the directionality of the change) in order to identify trending periods of operation for a given observable. An example set of measurements with the extracted Δ_{Trend} is illustrated in Fig. 4.5. The normalised difference between the start and end behaviours (represented by the values of the green and red scatter points) is extracted, to give a quantitative metric for the trend.

The trend labelling procedure can be expressed as:

$$\mu_S = \mu(\mathbf{x}[: N])$$

$$\mu_E = \mu(\mathbf{x}[length(\mathbf{x}) - N :])$$

$$\Delta_{Trend} = \frac{(\mu_E - \mu_S)}{\mu_S}$$

if $\Delta_{Trend} > \tau_{Upper}$ **or** $\Delta_{Trend} < \tau_{Lower}$ **then**

$$trend = True$$

end if

4.2.2 c) Extracting steps

In contrast to the discretised approach taken for trend feature extraction, *step changes* are labelled by evaluating the time series context for each data point with respect to the parameters defined in Table 4.1. This allows for instances to be captured at an increased fidelity than the higher-level trend symbols.

The extraction process follows steps:

- The forward first-order finite difference $(x_{t+1} - x_t)$ is calculated for the full time series under analysis,

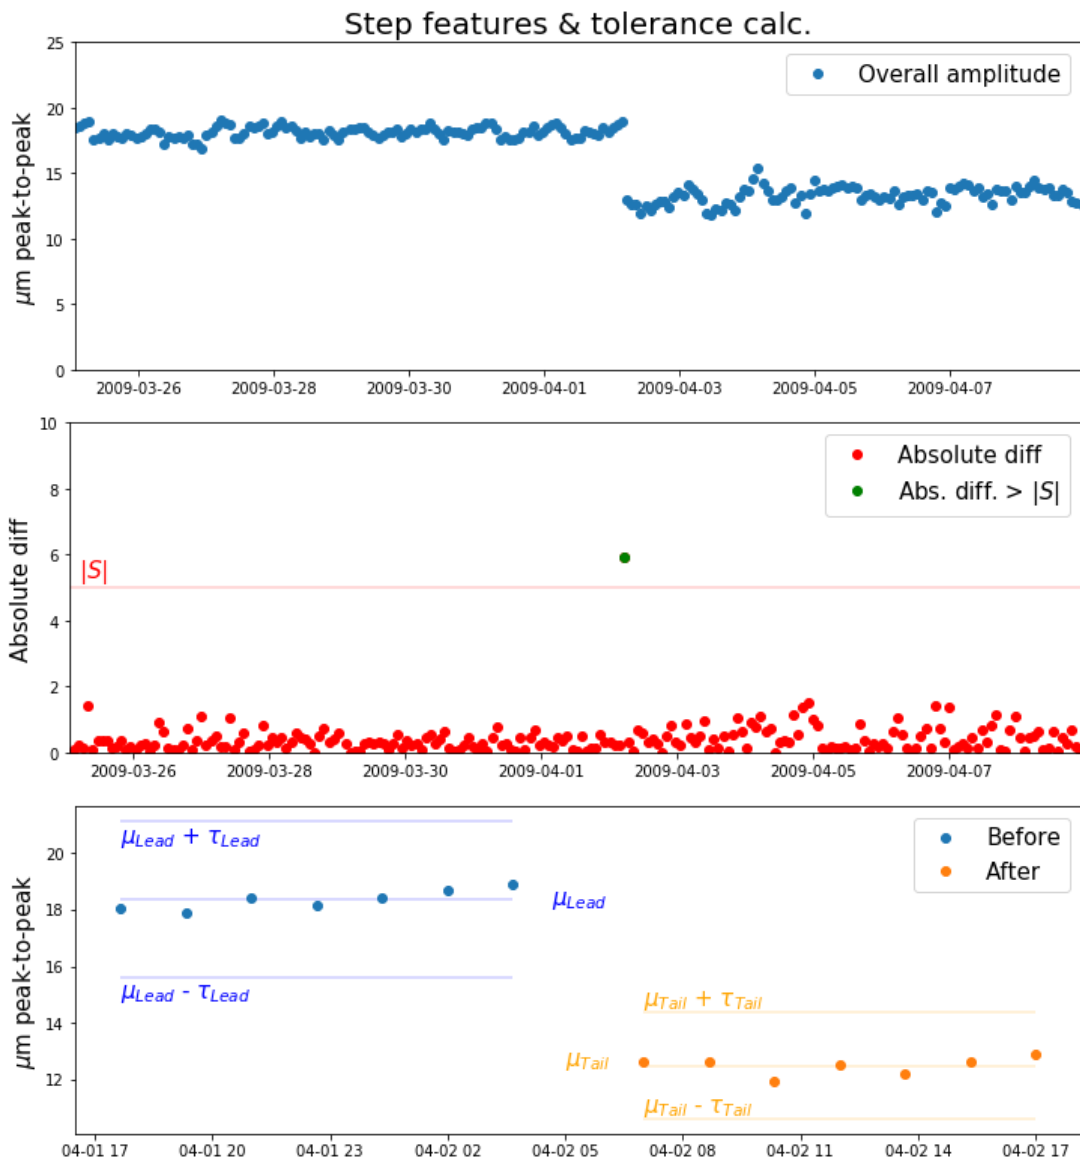


Figure 4.6: Example of features and calculation used in extracting steps

- The indexes of point-to-point deviations greater than the minimum step change magnitude $|S|$ are identified in the time series (*naive* step changes, as there is no consideration of the behaviour pre- and post-change),
- Each of these naive changes have their lead and tail periods analysed in order to ascertain the suitability for a legitimate step labelling (*stable* step changes).

Considering only stable step changes is important to safeguard against in-

stantaneous changes being flagged erroneously as steps (in cases where noise or signal artefacts might be responsible for quickly resolving deviations or point artefacts). The periods encapsulated within the time T_{Step} before and after the point change need to exhibit a level of relative stability and have a distinct enough difference between their distributions to be considered. This is captured by the lead and tail tolerance parameters $\{\tau_{Lead}, \tau_{Tail}\}$, which provide upper and lower bounds for which the lead and tail periods should fall within. These are additive and subtractive factors which dictate what bounds the before and after behaviours should be fully encompassed by to qualify as stable step changes.

The mean values μ_{Lead} and μ_{Tail} should also maintain a difference $> |S|$ over and above the initial naive change delta. This specifically removes single deviations from otherwise stable periods, which would otherwise be seen as potential step changes. Each of these phases in parametric step change extraction are illustrated in the example provided in Fig. 4.6. It should be noted that step change features will also be flagged as the previous defined impulse symbols, representing a special case of level-type changes.

The step tolerances $\{\tau_{Lead}, \tau_{Tail}\}$ and step period T_{Step} were set to $\{0.2, 0.2\}$ and 12 hours uniformly across each of the observables. Specific $|S|$ values were set on an individual basis for each observable through consultation with engineering staff and their expectations of step deviations qualifying as anomalous.

The step extraction procedure can be expressed as:

$$\mathbf{x}' = x_{t+1} - x_t$$

for x'_t **in** \mathbf{x}' **do**

if $x'_t \geq |S|$ **then**

$$\mathbf{x}_{Lead} = x[t - T_{Step} : t]$$

$$\mathbf{x}_{Tail} = x[t : t + T_{Step}]$$

if $\max(\mathbf{x}_{Lead}) < \mu(\mathbf{x}_{Lead}) + \tau_{Lead}$ **and** $\min(\mathbf{x}_{Lead}) > \mu(\mathbf{x}_{Lead}) - \tau_{Lead}$ **and**

```

 $max(\mathbf{x}_{Tail}) > \mu(\mathbf{x}_{Tail}) + \tau_{Tail}$  and  $min(\mathbf{x}_{Tail}) > \mu(\mathbf{x}_{Tail}) - \tau_{Tail}$  then
    step = True
end if
end if
end for

```

4.2.3 Channel and machine profiles

The term ‘channel’ is often used in the vibration monitoring literature to refer to a given accelerometer (or otherwise) data feed from a specific component of a rotating asset. In practice, many TG health monitoring hardware and software packages refer to individual observable time series as numbered channels. For example, the front bearing accelerometer instrument located in the HP stage of a TG is referred to as ‘channel 1’ throughout.

The parametric approaches described in the previous subsections need to be initialised on a per-channel, per-machine basis. Thus, a given channel requires a full set of level, trend and step parameters to be assigned: referred to in the KBS as a *channel profile*. At the level above this, all the channel profiles collected together for a given TG can be considered the *machine profile*. These terms are used throughout the remainder of this chapter when discussing the orientation of KBS parameters.

4.2.4 Proof-of-concept channel profiles

As a result of the research and KBS developed in [TMMS07], initialised channel profiles exist for 3 channels for a single TG instance. Tables 4.2-4.5 provide the parameters that were used as a result of this study, including the operational observable values in Table 4.2 which are applied to each channel uniformly. Each of these were created as generic starting points in the development of the

Table 4.2: Operational observable profile

Operational observ.	Λ_{Upper}	Λ_{Lower}	τ_{Upper}	τ_{Lower}	T_{Trend}	τ_{Lead}	τ_{Tail}	$ S $	T_{Step}
Generator load	690	620	0.1	0.1	24	0.2	0.2	50	12
Rotor current	2600	1900	3	3	24	0.1	0.1	750	12
Generator MVARs	140	10	1	1	24	0.1	0.1	100	12

Table 4.3: Channel 5 (front LP-stage A bearing) machine profile

Vibration observ.	Λ_{Upper}	Λ_{Lower}	τ_{Upper}	τ_{Lower}	T_{Trend}	τ_{Lead}	τ_{Tail}	$ S $	T_{Step}
Overall amplitude	12	5	0.4	0.4	24	0.2	0.2	5	12
Order 1 magnitude	10	3	0.4	0.4	24	0.2	0.2	3	12
Order 2 magnitude	3	0.5	1.0	1.0	24	0.2	0.2	2	12
Sub-sync magnitude	12	2	1.5	1.5	24	0.2	0.2	5	12
Sub-sync frequency	30	5	1.5	1.5	24	0.2	0.2	5	12

Table 4.4: Channel 6 (rear LP-stage A bearing) machine profile

Vibration observ.	Λ_{Upper}	Λ_{Lower}	τ_{Upper}	τ_{Lower}	T_{Trend}	τ_{Lead}	τ_{Tail}	$ S $	T_{Step}
Overall amplitude	15	8	0.4	0.4	24	0.2	0.2	5	12
Order 1 magnitude	12	5	0.4	0.4	24	0.2	0.2	3	12
Order 2 magnitude	9	3	1.0	1.0	24	0.2	0.2	2	12
Sub-sync magnitude	12	2	1.5	1.5	24	0.2	0.2	5	12
Sub-sync frequency	30	5	1.5	1.5	24	0.2	0.2	5	12

Table 4.5: Channel 9 (front LP-stage C bearing) machine profile

Vibration observ.	Λ_{Upper}	Λ_{Lower}	τ_{Upper}	τ_{Lower}	T_{Trend}	τ_{Lead}	τ_{Tail}	$ S $	T_{Step}
Overall amplitude	25	15	2.0	2.0	24	0.2	0.2	5	12
Order 1 magnitude	35	15	2.0	2.0	24	0.2	0.2	3	12
Order 2 magnitude	40	20	2.0	2.0	24	0.2	0.2	2	12
Sub-sync magnitude	12	2	1.5	1.5	24	0.2	0.2	5	12
Sub-sync frequency	30	5	1.5	1.5	24	0.2	0.2	5	12

expert system.

These parameters were initialised as part of the knowledge elicitation process used to build the original KBS, taking from domain experts in the field of vibration monitoring for rotating plant assets.

4.2.5 Retrospective

The KBS has a number of central advantages. As with many systems enabled by STST, providing a strong reasoning chain with any presented decision support is very useful in communicating to end-users. This is especially important in the monitoring of nuclear assets, where clear explicability can be a regulatory requirement. Each fired alarm has a set of time series symbols and a specific KBS rule, which makes reporting on the cause of particular alarms straightforward.

Furthermore, codifying large bodies of tacit domain knowledge can be used in standardising the approaches taken by individuals in repeatable diagnostic procedures. This can help inform best-practice guidelines, and utilised in the training of future domain experts and engineers.

The main shortcoming of the approach is the simplicity of the channel profile representation from a data perspective. In demonstrating the functionality of the rule base, channel-specific case studies were used and currently the system generalises to assume that there is no channel-, machine- or asset-specific differences to be considered. As shown in Section 4.2.4, initialised channel profiles only exist for two of the turbine-level components: bearings on the LP-stage A & C bearings).

These profiles are also static: there currently exists no functionality for determining behavioural characteristics of specific examples of rotating machinery at both the asset class level (e.g. turbine *A* vs. turbine *B*), or for extensibility to new asset classes (e.g. gas circulators, boiler feed pumps). This introduces the potential likelihood of errors in time series feature classification during the STST stage when deploying the KBS to previously unconsidered machines. The initialisation of the symbolic parameters would need to be repeated for each new instance, which represents a time-consuming effort if the process outlined in [Tod09] is repeated.

Furthermore, the system also disregards the case where the symbolic representation of time series primitives for a given machine might alter and evolve with continued operation. The potential for empirical changes to manifest in the vibration data (and therefore the machine profile representation) with machine degradation introduces scope for predictive indicators and metrics to be extracted from repeated use of the KBS if it had adaptive capabilities. Data-driven techniques for 'updating' the machine profile do not currently exist with the system.

With these advantages taken into consideration, it was decided that augmenting this KBS approach with further statistical methods that take advantage of the available data on turbine generators was a worthwhile approach to tackling routine alarms in reliability engineering.

The remainder of this chapter explores these areas: utilising data-driven methods in to augment the existing KBS approach to improve its generalisability, and investigate the evolving machine profile with extended machine operation. This represents a hybrid method for tackling automatic inference in routine alarm analysis by leveraging codified knowledge with empirical and generalisable statistical analyses.

4.3 Learning the machine profile

In contrast to the static values assigned to the machine profile of a new TG, a number of automated approaches were investigated in order to improve this process for use in the engineering environment. The problem of extracting a representation of state from labelled or semi-labelled data is presented as a typical ML-type problem for the following reasons:

- **The volume of available data:** Both nuclear-specific and general vibration monitoring disciplines are ingesting large data volumes, with greater

instrumentation sophistication, higher bandwidth of data output from transducers and lower cost of archiving the data.

- **The inherent complexity of health monitoring:** Beyond the rudimentary monitoring practices discussed in 4.1, it is of high importance to both operators and equipment manufacturers to best understand any degradation process associated with their machinery. This is not a presently well-understood process from a first principles standpoint, and it is presented that data-driven techniques can provide insight into any existent degradation mechanisms.
- **The evolution of machine state:** In evaluating the potential for any PHM or predictive metrics, one of the underlying assumptions is that there is a demonstrable data-based change in behaviour away from normality [LWZ⁺14] with a degradation process.

4.3.1 Learning parameters

From a previously unseen machine instance, our primary aim is to infer the best representation of system normality *at that time in the life cycle of the machine*. For a machine with k observables, we can define a machine profile α with respect to the STST symbolic parameters (introduced in Section 4.2.2) as follows:

$$\alpha_{[1,\dots,k]} = \begin{bmatrix} \Lambda_{Upper}(x_1) & \cdots & \Lambda_{Upper}(x_k) \\ \Lambda_{Lower}(x_1) & \cdots & \Lambda_{Lower}(x_k) \\ \tau_{Upper}(x_1) & \cdots & \tau_{Upper}(x_k) \\ \tau_{Lower}(x_1) & \cdots & \tau_{Lower}(x_k) \\ T_{Trend}(x_1) & \cdots & T_{Trend}(x_k) \\ \tau_{Lead}(x_1) & \cdots & \tau_{Lead}(x_k) \\ \tau_{Tail}(x_1) & \cdots & \tau_{Tail}(x_k) \\ |S|(x_1) & \cdots & |S|(x_k) \\ T_{Step}(x_1) & \cdots & T_{Step}(x_k) \end{bmatrix} \quad (4.2)$$

Note that each column in α is a channel profile. The parameters encapsulated by the machine profile can be broadly grouped into two types: *envelope-* (those involved with typical behaviour - level/impulse primitives: $\Lambda_{Upper}, \Lambda_{Lower}$) and *event-based* (those related to specific instances of machine scenarios - trend/step primitives).

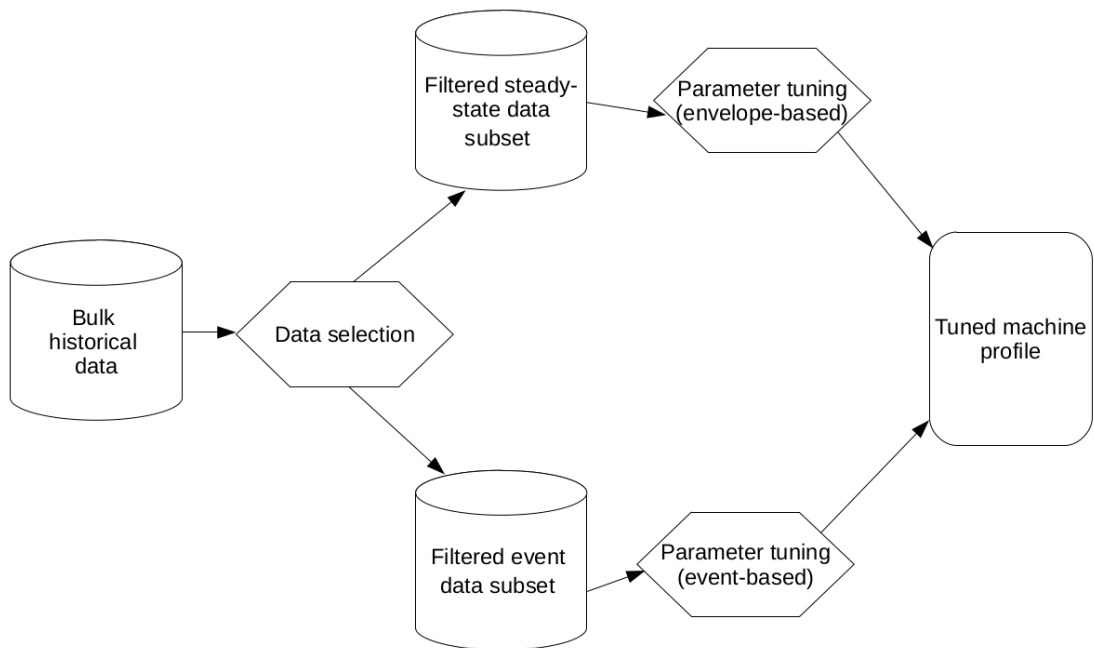


Figure 4.7: Overview of the parameter tuning steps to build a machine profile

The steps taken to create a suitable machine profile can be enumerated as:

- Select representative periods of data from the available bulk historical time series data,
- Extract relevant statistical moments and parameters from steady-state periods of the data in order to set parameters associated with envelope-based features,
- Extract relevant statistical moments and parameters from examples of step & reactor transient changes in the data in order to set parameters associated with event-based features.

A flowchart of the parameter learning process is provided in Fig. 4.7, illustrating the steps taken from bulk historical data towards a suitably initialised machine profile ready for KBS use.

4.3.1 a) Envelope-based

Inferring the allowable boundaries of operation for a previously unexamined observable or set of observables can be done by statistically profiling data which corresponds to normal behaviour. Selecting applicable periods of operation representative of these conditions can be approached in a variety of ways depending on the use-case: for this study, data is selected from an operational perspective i.e. the observed values recorded by the CM surveillance systems at times of regular, steady-state conditions. For example, an ideal candidate for this sort of period would be taken a sufficient length of time after a system event or transient, allowing for conditions to recover back to a less perturbed state. This assumption allows for envelopes of allowed behaviour (bounded by the channel profile values Λ_{Upper} and Λ_{Lower}) to be set with confidence that the learned parameters represent a single state and not multiple contributing or overlapping machine states.

4.3.1 b) Event-based

In contrast to steady state patterns that characterise envelope-based features in the previous subsection, event-based reasoning in the inference process require parameters defining the nature of the vibration event (specifically, the step change). In order to tune these parameters accurately, a training set of example step changes corresponding to normal behaviour is required to define each of the step parameters (as introduced in 4.1). The process for selecting this seed data and extracting the relevant parameter settings are explored in the following section.

4.3.2 Data selection - envelope-based

Example data from periods of steady state routine operation are most suitable for defining the bounds $[\Lambda_{Upper}, \Lambda_{Lower}]$. Hand-selecting subsets of the time series for these purposes would be an onerous task: this section discusses some of the central properties in defining suitable envelope-based tuning data, and how this can be identified automatically through use of repeatable statistical methods.

4.3.2 a) Unimodality

A major consideration for suitable steady-state data is how it compares to an idealised *unimodal* distribution. Unimodality [CC04] applies to data consisting of a single distinct mode value. Strictly, this corresponds to a single maximum probability or frequency value: but distributions with multiple local maxima are often considered to be non-unimodal for data analysis purposes. Fig. 4.8 shows comparative examples of unimodal and multimodal behaviours (note that the multimodal example given has a local maxima second mode) with annotated KDEs for each of the sampled distributions. KDEs are useful for con-

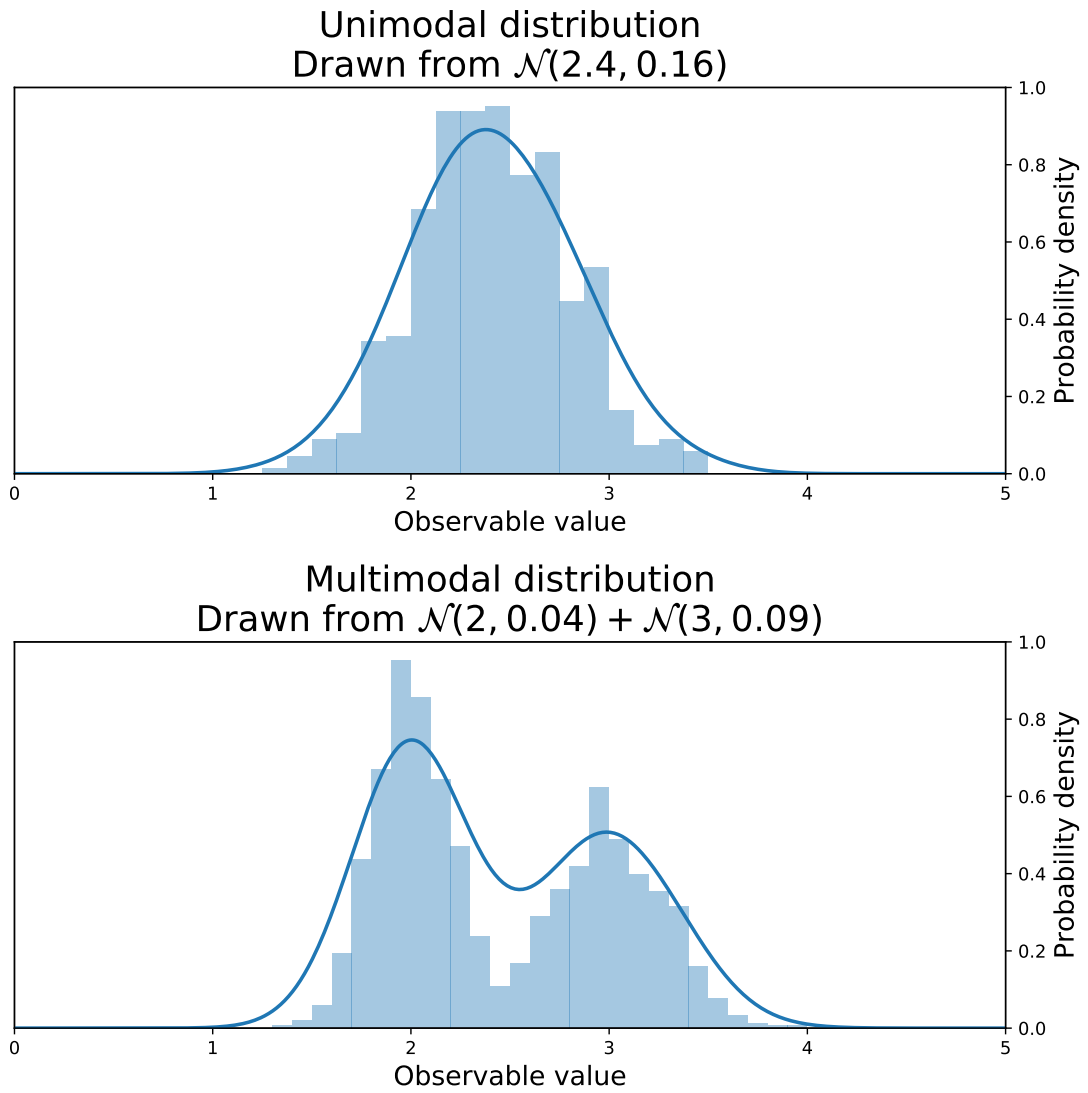


Figure 4.8: Example unimodal and multimodal distributions

sidering distribution modality: both in visualisation and statistically [Sil81].

Λ_{Upper} and Λ_{Lower} are dependent on a single expected value \bar{x} with random variation, so discriminating against obviously multimodal data in the setting of these parameters is important. As with many scenarios with data selection and pre-processing, guaranteeing strict unimodality can remain difficult. Cases where the PDF might have ambiguity regarding modality (i.e. skewed unimodal or a single dominant mode amongst multiple contributing behaviours) can be mitigated against by examining the probability density values themselves: utilising features such as the full width half maximum

(FWHM) [WH13] of the function. This provides a non-parametric measure to aid this setting factor-based boundaries around a detected mode in the data distribution.

Testing for unimodality can be accomplished by examining the CDF (introduced in Section 3.1.3) of candidate data subsets. Distributions where $CDF(x)$ is convex $< \bar{x}$ and concave $> \bar{x}$ meet the criteria of unimodality [HH85]. Mapping to finite populations of time series data, this heuristic is used on the ECDF measure compared directly to the CDF of an idealised Gaussian distribution with $\sigma_{Ideal} = 1$ and $\bar{x}_{Ideal} = \bar{x}$ (as described in Section 3.1.4 b). Examples of both unimodal and multimodal ECDFs in direct comparison to an ideal unimodal continuous function are provided in Fig. 4.9

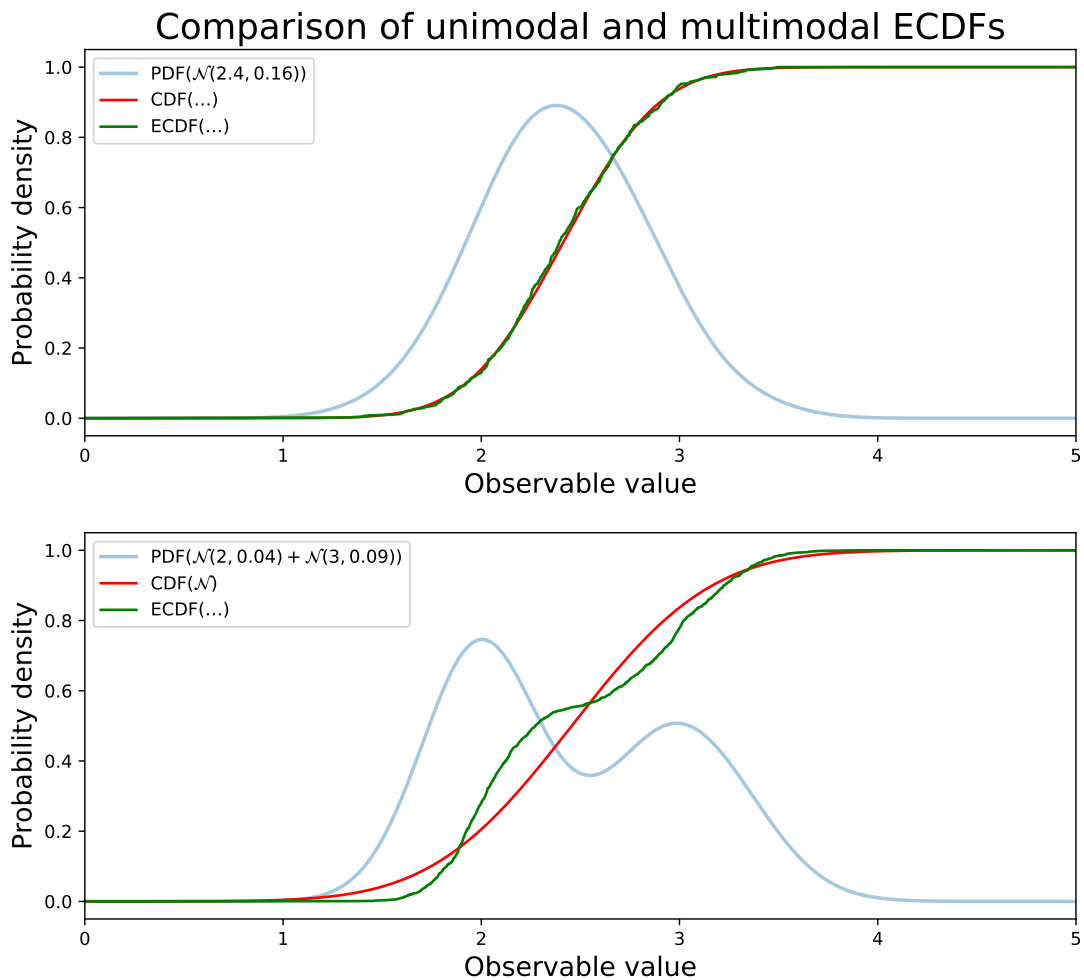


Figure 4.9: Comparative illustration of ECDFs for unimodal and bimodal PDFs

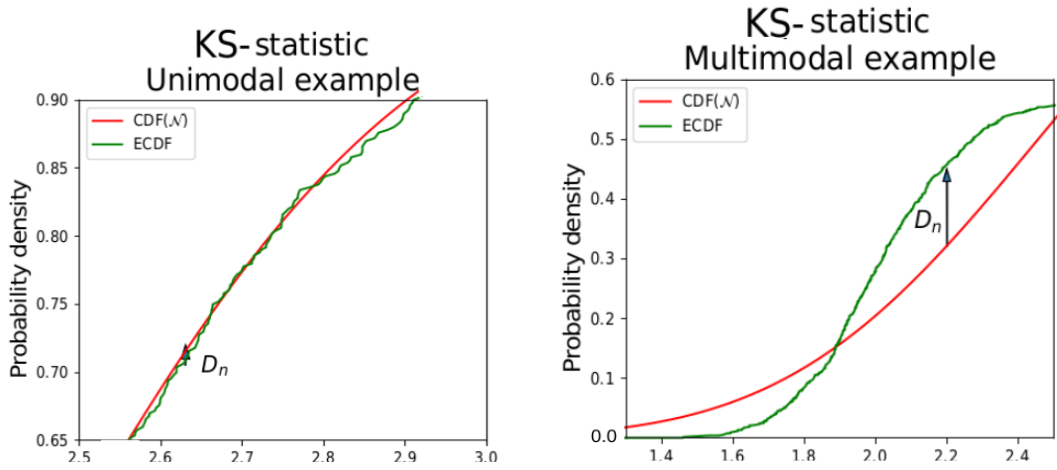


Figure 4.10: Comparative illustration of ECDFs for unimodal and bimodal PDFs, with KS-statistic metric D_n annotated for each

Numerous techniques exist for hypothesis testing of distribution modality types [PSJ⁺13]. Kolmogorov and Smirnov introduced the KS-statistic [Mas51] to quantitatively evaluate differences between target and empirical distributions, and this can be defined as:

$$D_n = \sup \|f_{ECDF} - f_{\mathcal{N}}\| \quad (4.3)$$

where f_{ECDF} is the ECDF function calculated from the data population and the function is the target Gaussian with (μ, σ^2) equivalent to those from the data population.

Minimising the difference between the empirical and ideal is the objective in selecting data using these techniques, with candidate data periods with smaller differences being closer to the ideal distribution. A quantitative metric which defines this is the quantity D_n : the maximum deviation between the empirical and idealised CDFs as shown in Fig. 4.9, or $\sup \|f_{ECDF} - f_{\mathcal{N}}\|$. When selecting from a number of candidates, the data period with the smallest value of this will correspond to the most suitable. KS-statistic values for each of the examples illustrated in Fig. 4.9 are provided in Fig. 4.10.

4.3.2 b) Rolling KS-statistic

The dynamic nature of vibration signals from TGs means periods of unimodal and multimodal behaviour are to be expected when considering historical operation at the macro level. Taking the full time series history of any given turbine channel is highly likely to contain numerous modes of behaviour from transient conditions and outliers. Identifying suitable distribution subsets to build channel profiles with is dependent on extracting dominant unimodal periods throughout the long-term time series', with confidence that the filtered data is representative of the machine's normal behaviour.

To achieve this, a technique for such periods of operation is presented which is based on the principle of calculating the KS-statistic on a rolling basis with a moving window. The process can be outlined as follows:

- The time series is batched into a rolling window of N_{D_N} observations,
- D_n is calculated for each window step (window set to 50 values),
- The distribution of windowed D_n values is filtered to include only values from the 10th percentile (selected as this represents the lowest rolling D_n values) - those corresponding to the windows showing the strongest unimodality properties. These windows have the lowest D_n values, so are relatively speaking the most likely to be unimodal.

Filtering the data corresponding to windowed KS-statistic values in the 10th percentile is dependent on the operational period being largely stationary with transient condition changes.

Figs 4.11 and 4.12 illustrate this procedure with an example overall amplitude TG series. The resulting empirically selected representation distribution is shown in Fig. 4.13, from which the relevant machine profile parameters can be derived.

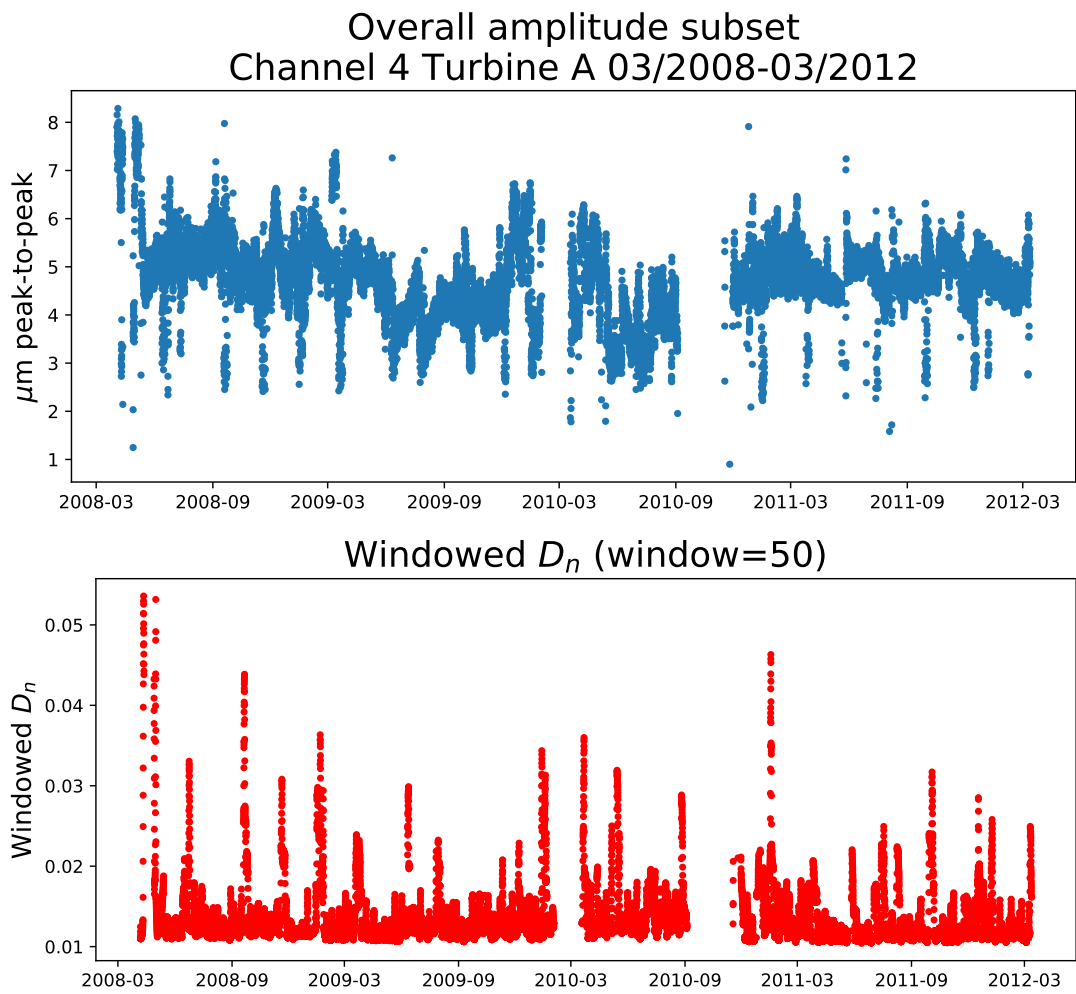


Figure 4.11: Example channel time series, with rolling KS-statistic series

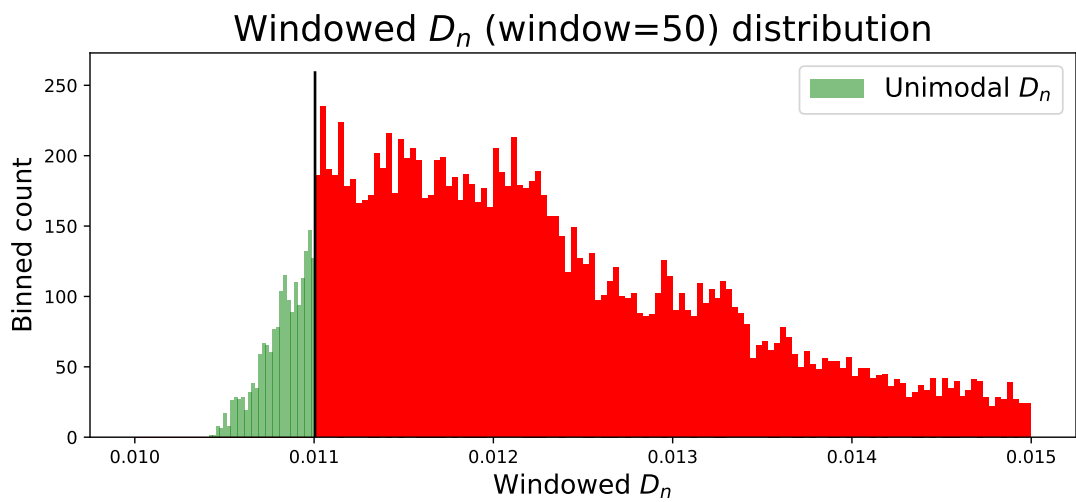


Figure 4.12: Distribution of KS-statistic values




Figure 4.13: Tuned distribution example

4.3.2 c) Machine state phase space

To ensure that empirically unimodal data corresponds to steady-state behaviours from the TG during the intended operating conditions, and not to stationary data recorded during periods of interim operation (e.g. at stepped duty cycles, or during low power operation regimes), the selection process also considers the operational conditions corresponding to the selected tuning distribution. Specifically, the generator load values exhibited during these periods is examined to identify non-transient conditions due to its contributory effects to turbine behaviour [MGiAR08].

Considering the example empirical tuning data distribution introduced in Fig. 4.13, the associated load values for this are illustrated in the vibration-load phase space in Fig. 4.14. This shows a small number of non-full load (< 600MW) states introducing vibration outliers which slightly skew the empirical tuning data. A more dramatic view of this is shown in Fig. 4.15, where the rolling dip unimodality filter is not applied. Without the pre-processing step, the distribution parameters extracted from the data shown in Fig. 4.15 would not correspond to a meaningful estimate of normal, steady-state behaviour. Ensuring that the tuning data for defining normality is only taken




Figure 4.14: Vibration-load phase space with D_n filter, with crisp load boundary




Figure 4.15: Vibration-load phase space without D_n filter, with crisp load boundary

from the relevant operational conditions helps keep the learned parameters as robust as possible.

4.3.2 d) Extracting the dominant distribution

Over and above the efforts outlined so far with empirical unimodality and the operational condition filter, periods of vibration behaviour can remain in the tuning data distribution which do not correspond to an ideal, unimodal distribution. This is especially true when extracting suitable data from an extended period of historical operation (corresponding to 3 years+ on particular channels in the dataset used for this study). The most common cause of this will be due to 'genuine' changes in the TG vibration response which cannot be at-

tributed to some sub-behaviour in the generator load values. An example of this is provided in Fig. 4.15.

To mitigate against this, the selected data is clustered using k -means ($k = 2$), with the most populous labelling being selected as the data to be used as part of the tuning distribution. This is based on the assumption that, post-filtering, there is likely to only be a single mode of outlier behaviour alongside the target mode for tuning. Multi-modal outlier behaviours would not be safe-guarded against using this method, and visual inspection of time series periods likely to be subject to these conditions is recommended.

4.3.3 Data selection - event-based

Selecting a representative collection of step changes that typify the characteristics of a machine is achieved through use of changepoint analysis techniques. These examples can be used to determine the validity of the selected channel profile parameters for step changes, and make any necessary adjustments using empirical evidence from the TG's historical behaviour.

Changepoint identification as an area of study has seen many applications in time series inference [CA17], and specifically in vibration monitoring [LZLL17] of machinery.

Examining step changes and their expected properties marks an alternative approach to reasoning about system normality beyond simplistic steady-state analysis. Two changepoint extraction techniques are introduced in this section (a simple standard deviation-based method, and a density-based method similar to the KS-test approaches applied in Section 4.3.2 b)) for comparative study.

The target parameters each approach extracts from the bulk time series are: minimum step change magnitude ($|S|$), the pre- (τ_{Lead}) and post-change (τ_{Tail}) stabilities in terms of standard deviation, and the period (T_{Step}) over

which changepoints emerge. Note that periods including zero'd data values are treated as sensor errors and are not considered in the empirically-defined populations of step changes. Each of these parameters are evaluated post-extraction (with both methods described in the next two sub-sections) as populations in order to ascertain the most representative values for the data. The standardised nature of the step change definition (a changepoint, book-ended by a before and after period) allows for the mean $|S|$ and pre- and post-changepoint standard deviations (τ_{Lead} and τ_{Tail}) to be calculated.

4.3.3 a) Standard deviation-based

A straightforward approach to step change extraction is the comparison of point-by-point delta values to the properties of the batch distribution of the entire time series. Defined as:

$$f_{Changepoint}(x(t)) = \begin{cases} 1, & : \|x(t) - x(t-1)\| > \mu(x_\Delta) + N_\sigma\sigma(x_\Delta) \\ 0, & : else \end{cases} \quad (4.4)$$

where x is the full time series, $x(t)$ is the value of x at a given time t , x_Δ is the first-order discrete difference of x and N_σ is a number of standard deviations $\sigma(x_\Delta)$ from the mean of the difference distribution $\mu(x_\Delta)$ defined as allowable change non-indicative of a step change. In other terms, any single point-wise change greater than $\mu(x_\Delta) + N\sigma(x_\Delta)$ will be labelled as a potential changepoint. Figure 4.16 provides an illustrative example of these. Typically, N_σ is set = 3 in accordance with the 99.7% rule [WH13] of statistical populations, where is assumed that 99.7% of all samples from a population will fall within three standard deviations of the mean.

The definition of step events introduced by the parametric view in Section 4.2.2 c) requires a level of stability exists both before and after the discrete

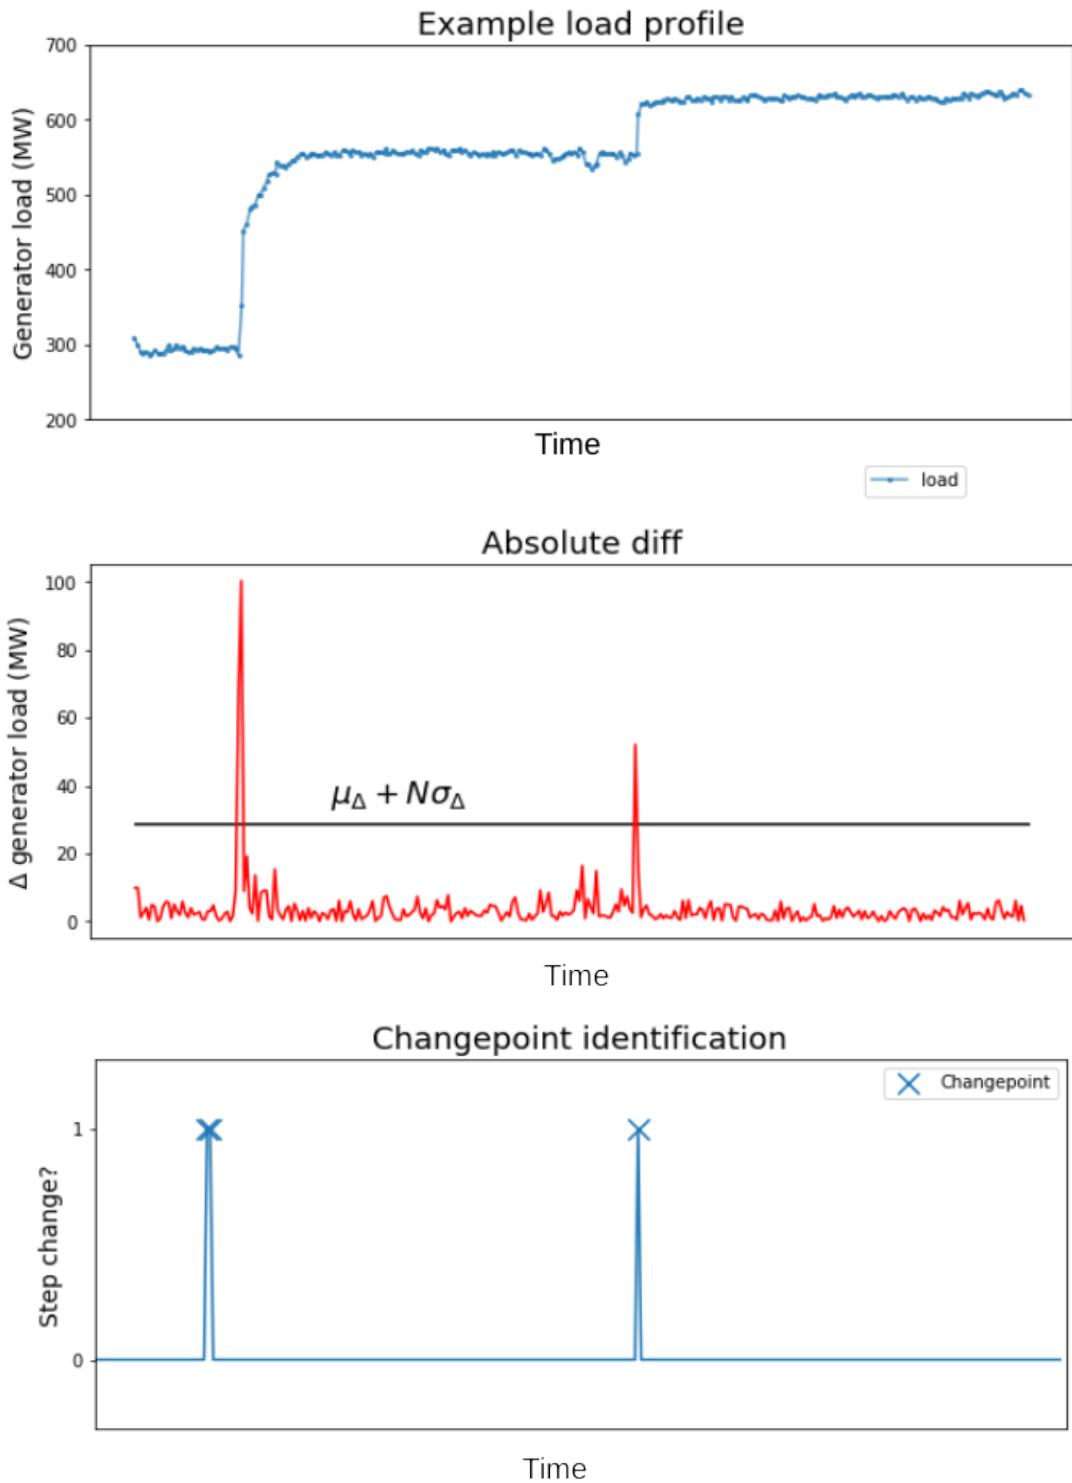


Figure 4.16: Overview of standard deviation-based identification of change-points in vibration time series data

change-point in order to be quantified as a repeatable step feature.

This process belongs to a larger family of similar techniques [CBK09] that

have seen wide adoption in a variety of application domains. The strength of such techniques lie in their ease of explanation and simplicity in implementation: meaning standard deviation-based anomaly or novelty detection approaches are often the first models tried in the identification of features of interest.

4.3.3 b) Density-based

Using the same statistical principles but with the opposite inferential target as the unimodality techniques outlined in Section 4.3.2 a), monitoring the rolling density and corresponding measures of the distribution allows for examples of bi- or multi-modality to be identified. Typical changes in the PDE under changepoint conditions are illustrated in Fig 4.17, showing the evolution from unimodal to bimodal behaviour.

From each window distribution, the rolling KS-statistic (as introduced in Section 4.3.2 b)) can then be calculated to reason about the modality of each time series snapshot. An example set of rolling dip figures for examining changes is provided in Fig. 4.18, showing higher D_n values around points of discrete change in the time series. The rolling D_n values above the 90th percentile are highlighted (in contrast with the unimodality considerations outlined in 4.3.2 b)).

Extracting changepoints is achieved by identifying groups of consecutive $D_n > 90^{th}$ percentile values and taking the first index of each of these groups. Example changepoints derived using this approach are annotated in the example provided in Fig. 4.18. It should be noted that this method successfully identifies the second changepoint in the time series despite their proximity in load values (approx. 550MW to 620MW). The sensitivity of this method in monitoring different time series observables for rotating machinery needs to be investigated in more depth to ensure all relevant changepoints are likely to

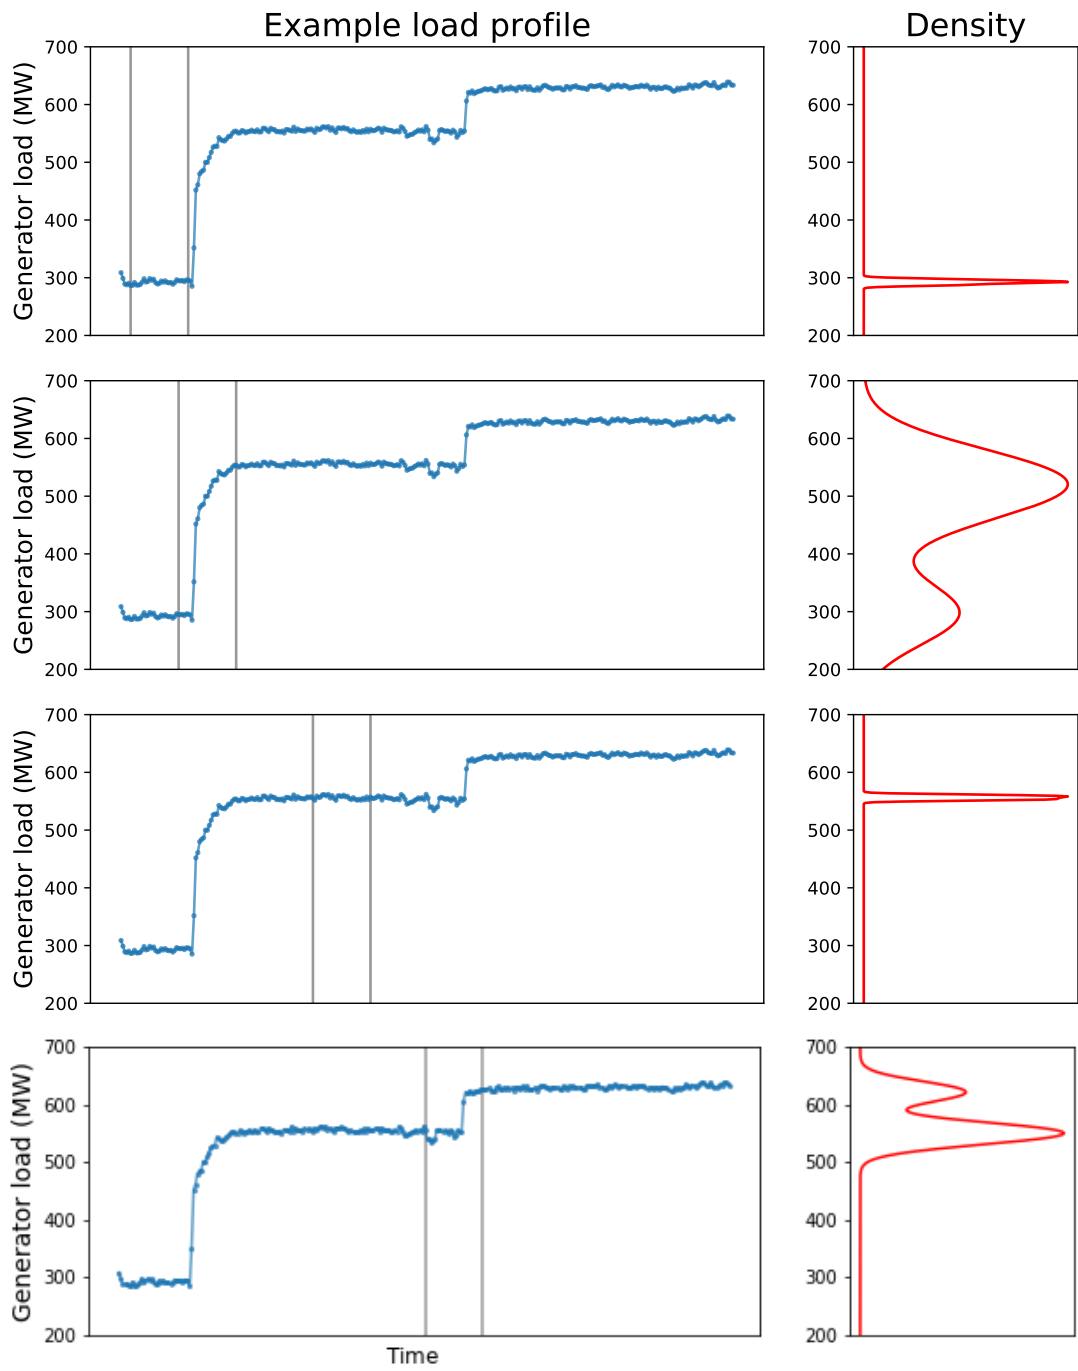


Figure 4.17: Evolution of PDE at various stages of an example time series with step conditions

be flagged.




Figure 4.18: Rolling KS-statistic applied to example load profile

4.4 Case study: Tuning channels

The techniques introduced in Section 4.3.2 are brought together into an automated tuning procedure, with the aim of identifying robust estimates of the required channel profile parameters from suitable historical data. This allows for previously 'un-tuned' turbine channels to benefit from the KBS approach to routine alarm analysis without the requirement for lengthy engineering consultation on a channel-by-channel basis.

For this study, the time series histories for 9 sensor channels (each corresponding to a particular TG component sensor) across two machines over the period 2008 - 2012 (at various points of the operational history) were extracted to provide a dataset for experimentation.

4.4.1 Channel profile vs. data

Comparing the preset channel profile values with typical operation data from Turbine A gives an insight into how the KBS uses its parameters to define normality when faced with typical conditions of a machine. An example period of operation with corresponding distribution of values is illustrated in Fig. 4.19, showing how overall amplitude behaviour deviates above and below the boundary Λ_{Lower} , resulting in the creation of `impulse` symbols in the STST

module along with the major mode of operation falling inside the expected bounds.

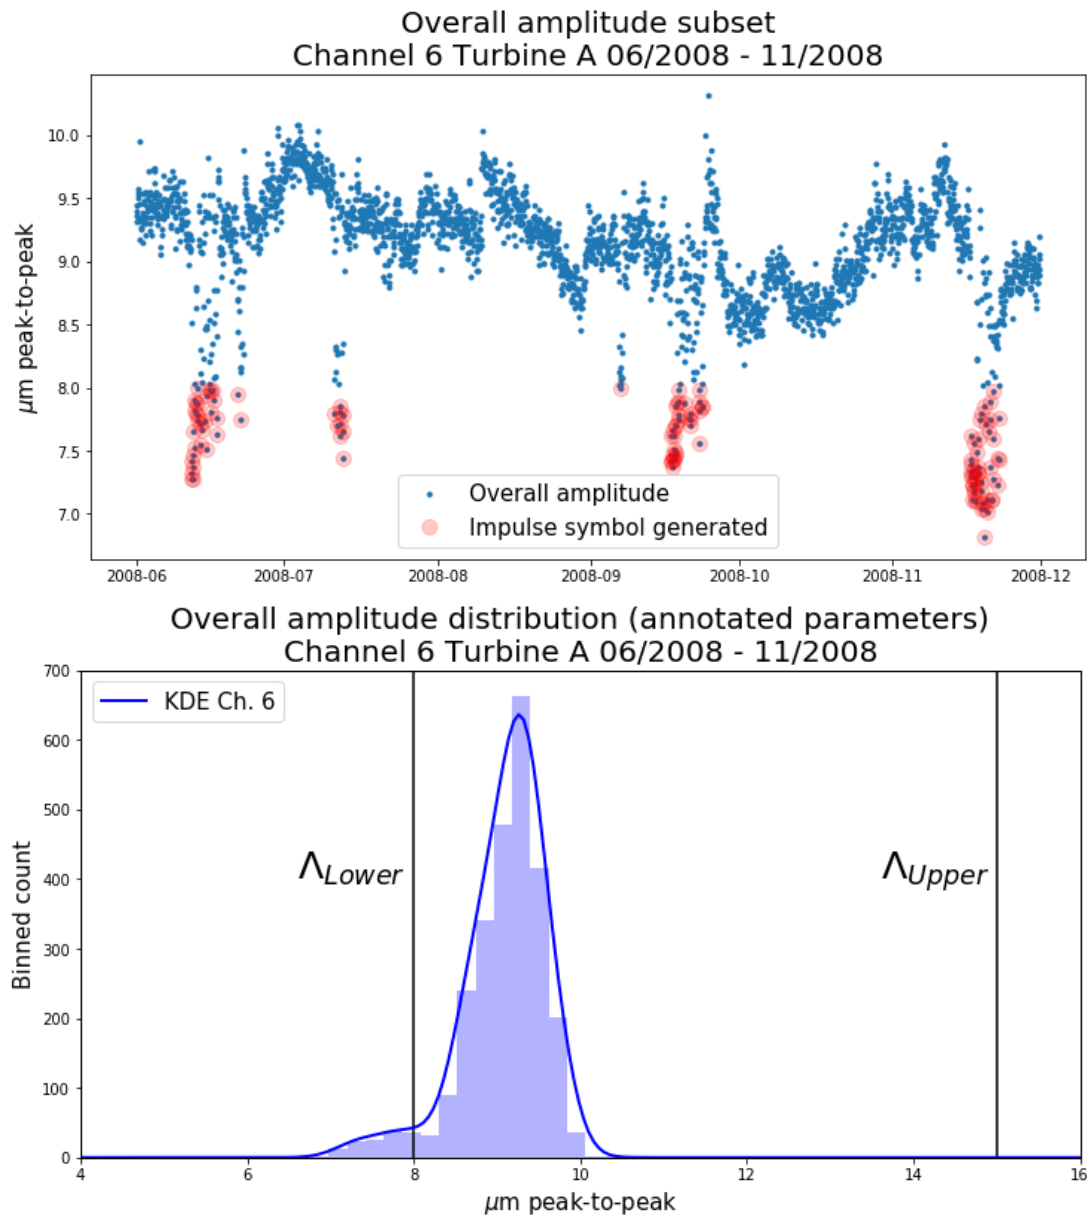


Figure 4.19: Time series and distribution of example overall amplitude subset, with annotated channel profile parameters

Such behaviour is generally consistent with the dynamic operation of a TG in the nuclear power context, and the highlighted `impulse` features given in Fig. 4.19 provide useful alerting in isolation (when cross-examined with operational deviations), and in the context of a KBS where rule-chaining can deliver

more complex conclusions.

In contrast, if this same channel profile is applied to the same observable (overall amplitude) but on a different channel (channel 7, the transducer mounted on the front bearing of the next turbine stage), the difference between the expected behaviour and observed data becomes apparent due to the difference in expected vibration level from these components. Fig. 4.20 shows the preset parameters are no longer applicable, with the expected vibration distribution differing from that on channel 6. This results in vast over-reporting of impulse symbols, due to the dominant behaviour now existing on the boundary of the rule threshold.

4.4.2 Tuning process and results

4.4.2 a) Empirical tuning distributions

The following empirical tuning distributions were derived from Turbine A data across the period of operation (where sufficient data was available for each channel). Fig. 4.21 provides each of these extracted vibration observable distributions. Note that the bandwidth selection method used for these density estimates was Silverman's [Sil87] rule. The corresponding empirical distributions for the 3 operational observables over the same data period are illustrated in Fig. 4.22. Both these collections of distributions illustrate the diversity of normal conditions exhibited across turbine channels that are otherwise not considered in the KBS machine profile without a statistical augmentation that the tuning methods provide.

These bulk calculations provide the basis for defining normality, bounds of expected behaviour from historical steady-state data and subsequently more applicable envelope values for use in the routine alarm KBS. Note that the x-axis for each of the distributions in these collections represents the parameter

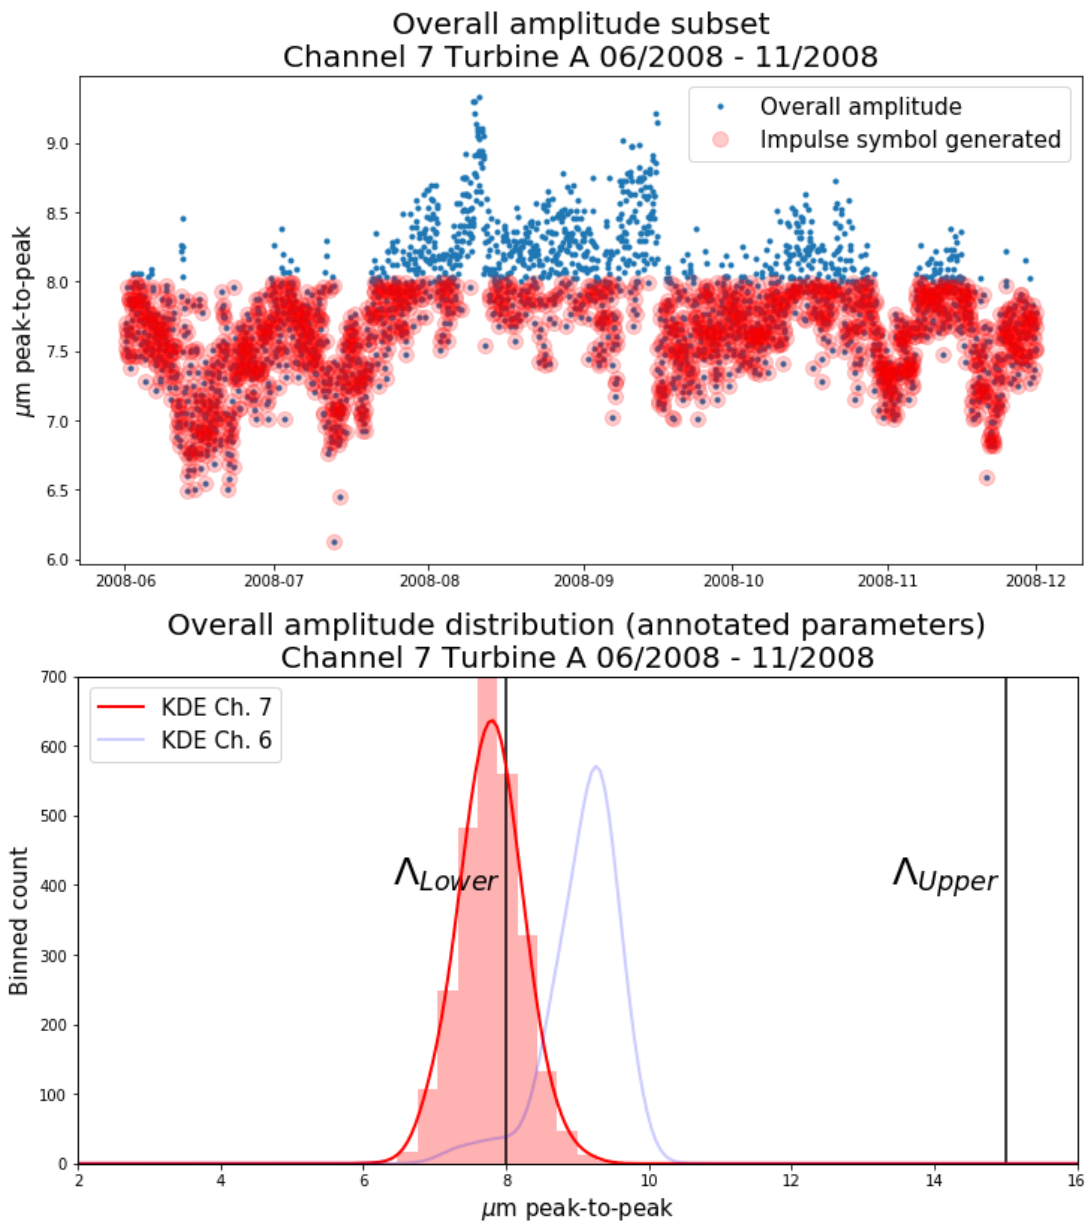


Figure 4.20: Time series and distribution of example overall amplitude subset, with annotated channel profile parameters (incorrect channel)

value itself, with the y-axis corresponding to probability density.

4.4.2 b) Selecting parameters

Where available, the preset channel profile envelope values can be compared directly with these empirically-derived distributions: both in the case of operational observables in Fig. 4.23, and for vibration observables in Fig. 4.24.

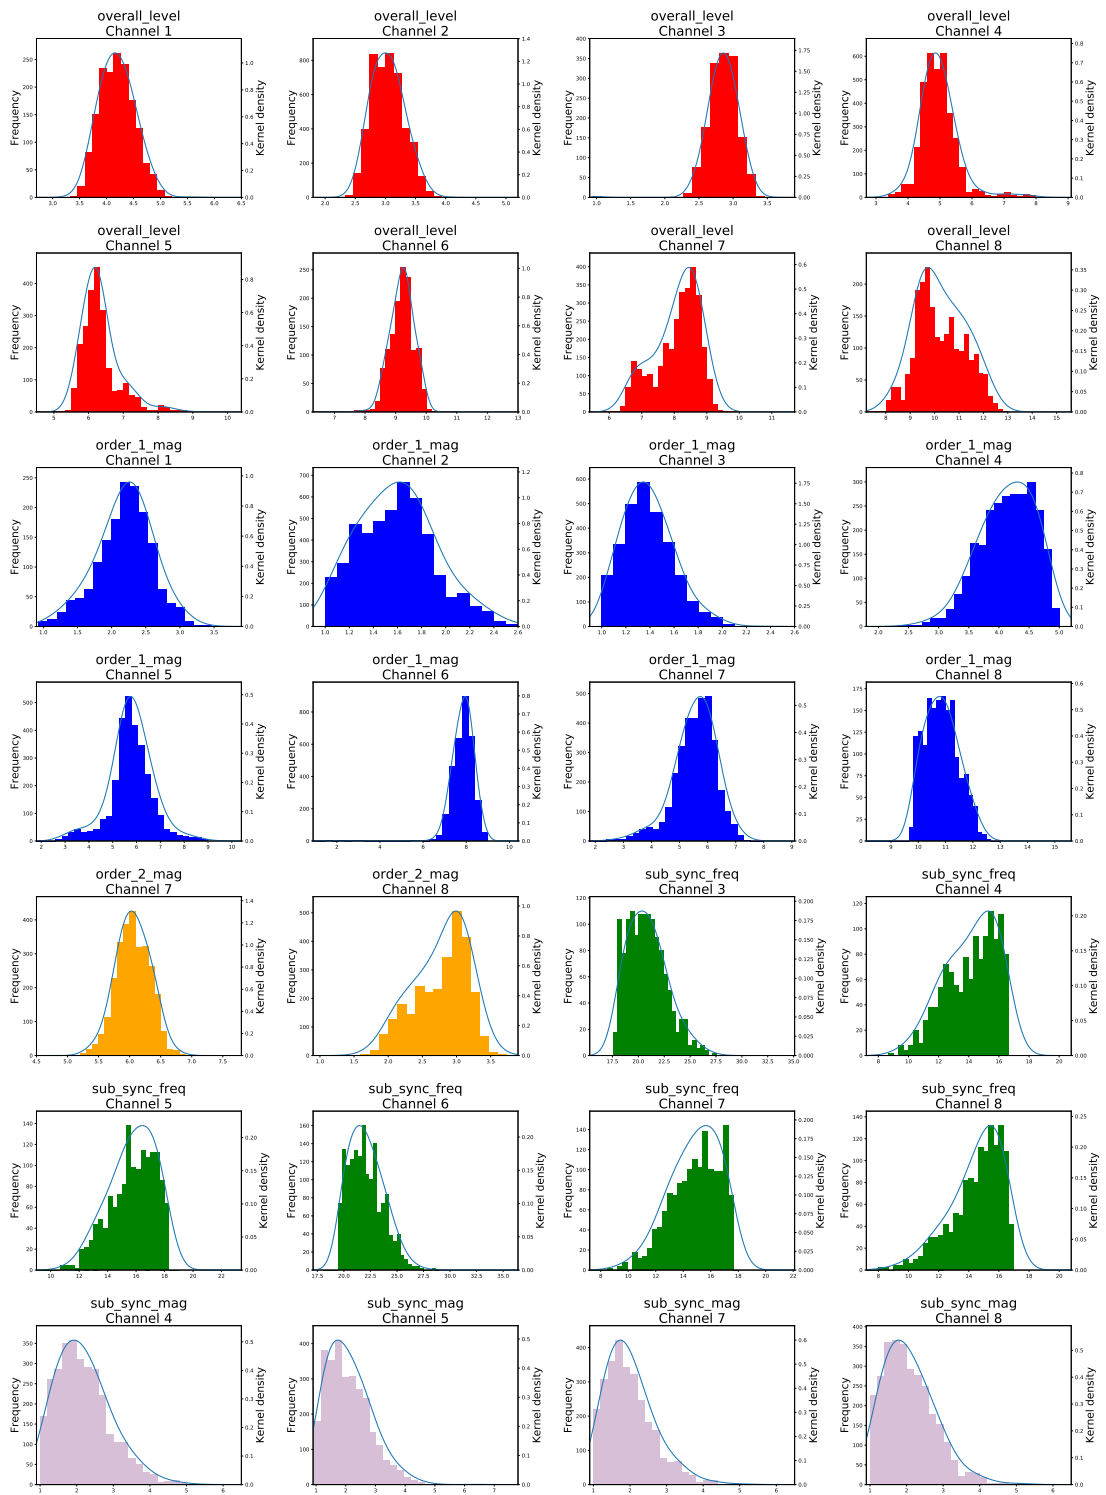


Figure 4.21: Empirical distributions for vibration observables

For the operational observables, both the rotor current and generator VARs exhibit differences between the initialised profile values for Λ_{Upper} and Λ_{Lower} and the empirical distribution. Since the nature of these observables is such




Figure 4.22: Empirical distributions for operational observables

that they are determined by the duty cycle and operational strategy of the reactor, there is less of a case for altering the bounds of operation without input from the reliability engineering staff. However, highlighting the differences in the expected duty cycle bounds and the empirically-derived distribution from the data has the potential to explain large volumes of operational change-based routine alarms, which can be clearly explained from the alterations made to the operational observables.

In the case of the vibration observables, the initialised values are largely in agreement with the distributions when simply considering encapsulation by the bounds (not considering the width of the envelope), except for examples in channel 9 overall level & first order magnitude, and sub-synchronous magnitude on both channels 5 & 9. These are clear candidates for amendments to be made to the selected values for Λ_{Upper} and Λ_{Lower} .

For such examples of disagreement or where the preset values do not already exist, the density function created by the KDE can be used to statistically define the new values for Λ_{Upper} and Λ_{Lower} . The FWHM is a useful descrip-




Figure 4.23: Empirical kernel density estimates for operational observables, with annotated existing channel profile limits

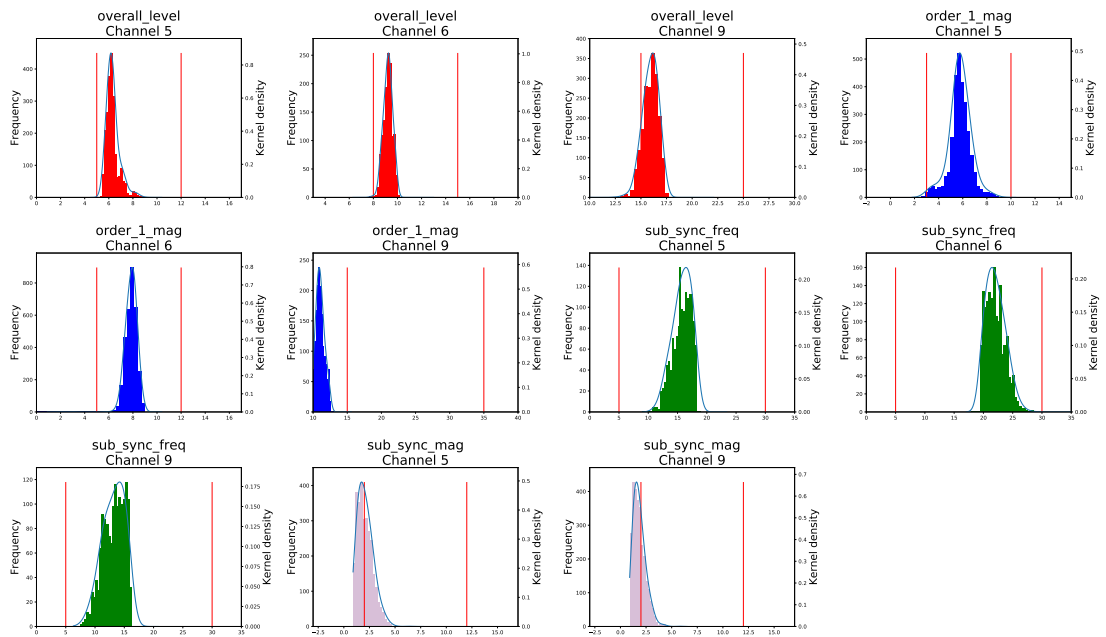


Figure 4.24: Empirical kernel density estimates for vibration observables, with annotated existing channel profile limits

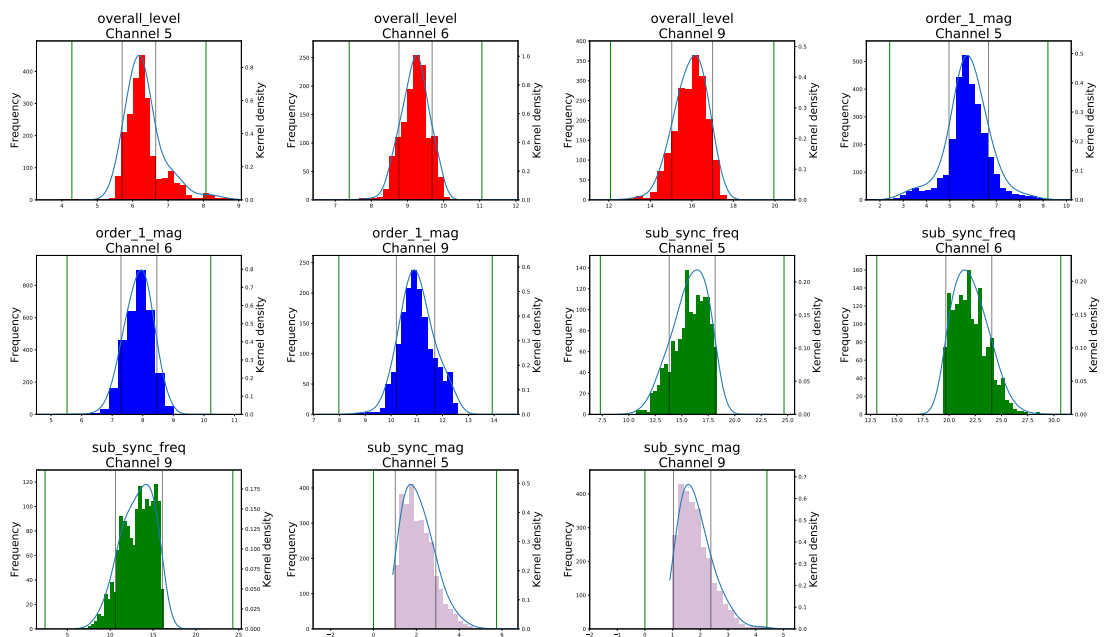


Figure 4.25: Density-based envelope setting of revised channel profile values

tive measure [Bar13] from the probability density around the dominant mode, from which a scaled factor can be taken to define bounds around which an envelope can be constructed ($FWHM \approx 2.355\sigma$). Fig. 4.26 provides an example of this metric.

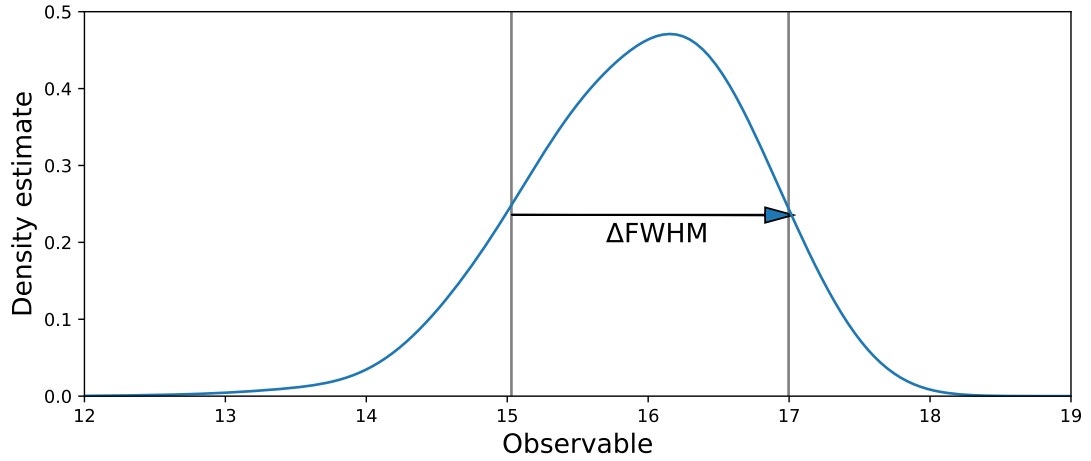


Figure 4.26: Illustration of FWHM taken from a density estimate

Based on the FWHM value of the dominant mode of each density estimate, Fig. 4.25 shows the previous examples from 4.24 with a self-tuning approach to selecting these channel profile parameters from the empirical distribution itself, using a factor of the FWHM value:

$$\Lambda_{Upper} = \bar{x} + N \frac{\Delta_{FWHM}}{2} \quad (4.5)$$

$$\Lambda_{Lower} = \max(0, \bar{x} - N \frac{\Delta_{FWHM}}{2}) \quad (4.6)$$

where \bar{x} is the dominant mode of the empirical distribution (argmax of the density function), Δ_{FWHM} is the FWHM value is the difference between the density values at the FWHM, and N is the envelope factor around which the thresholds are defined. Note the *max* expression keeps the lower bound from not falling below a value of 0. For the examples provided in 4.25, N is set to a value of 3 (as suggested in [WH13], consistent with the aforementioned 99.7% rule).




Figure 4.27: Standard deviation- & density-based changepoints on channel 9 time series

4.5 Case study: typical step changes

4.5.1 Technique selection

The techniques introduced in Section 4.3.3 are compared on the channel 9 overall level history of Turbine A, with the aim of ascertaining which is most suitable for bulk use for the remaining channels. Each of the identified changepoints are highlighted in Fig. 4.27 alongside the time series they were derived from, and summary statistics for both methods are provided in Table 4.6. As shown by the extracted step size and pre-/post- stabilities along with the number of instances identified, the standard deviation-based method tends to identify candidate step changes more often and less with greater pre-/post- instability than the density-based technique. The mean step profiles for both up- and down-steps for each of the changepoint detection methods are illustrated in Fig. 4.28, for both feature extraction and visual validation purposes.

Table 4.6: Identified steps per method

Method	No. steps	$\mu(S)$	$\mu(\tau_{Lead})$	$\mu(\tau_{Tail})$
Standard deviation-based	96	3.75	2.16	2.71
Density-based	39	2.71	1.55	1.56




Figure 4.28: Standard deviation- & density-based mean step profiles

Due to the windowed nature of the techniques, the potential for duplication of step instances in the returned results is a potential issue for both methods. Examining the uniqueness of the step example populations returned for each is accomplished by comparing their time series hierarchical clustering [Lia05] results, as illustrated in Fig. 4.29.

From this view, the standard-deviation method appears to exhibit a large redundancy effect in the labelled step instances for a significant proportion of the extracted examples, whereas the density-based method returns fewer steps but with improved cluster diversity. Considering each of the points highlighted in this section, the density-based changepoint extraction method is selected for use in bulk to continue the procedure of tuning the event-based channel profile parameters.





Figure 4.29: Standard deviation- & density-based hierarchical clustering

4.5.2 Tuning process and results

4.5.2 a) Empirically selected steps

For each of the Turbine A channel examples first examined in the results of Section 4.4, the step extraction method was applied to build a characteristic step profile for both up- and down-steps where sufficient data was available

(at least 5 examples of steps were required for a step profile to be constructed for a given channel-observable pair). The aggregate results for each of these are provided in Figs. 4.30 and 4.31, showing a wide variety of time series stabilities, changepoint magnitudes and general characteristics when considering typical step changes for each of the channels.

It should be noted that several of the channels included in these results for both up- and down-step examples do not appear to show sufficient stability or repeatability to begin examining their properties from a parametric perspective (e.g. sub synchronous up-step on channel 2). Also, several of the down-step examples include zero-level data which wouldn't otherwise be considered in routine alarm inference (e.g. order 1 & 2 magnitude downsteps in channel 6). These results were included to illustrate some of the edge cases associated with bulk step change inference in the vibration domain, and should be discarded when considering such approaches in a fully automated solution when deployed. This also highlights the importance of keeping engineering staff in-the-loop when making decisions about the profiling of channel characteristics: a key part of providing effective decision support.

4.5.2 b) Selecting parameters

The summary statistics for extracted steps are provided where enough data and examples were extracted from the channel history. The results for each of these are provided in Tables 4.7, 4.8 and 4.9.

The most striking element from these results is the inter-channel variability in the typical stage change example values extracted by the self-tuning methods. With the exception of the channel 6 first order magnitude extracted step, the mean step associated with each of the component channels is much less than the initialised value provided in the original KBS parameterisation (as shown in Section 4.2.4). This suggests that the flat, channel-agnostic setting

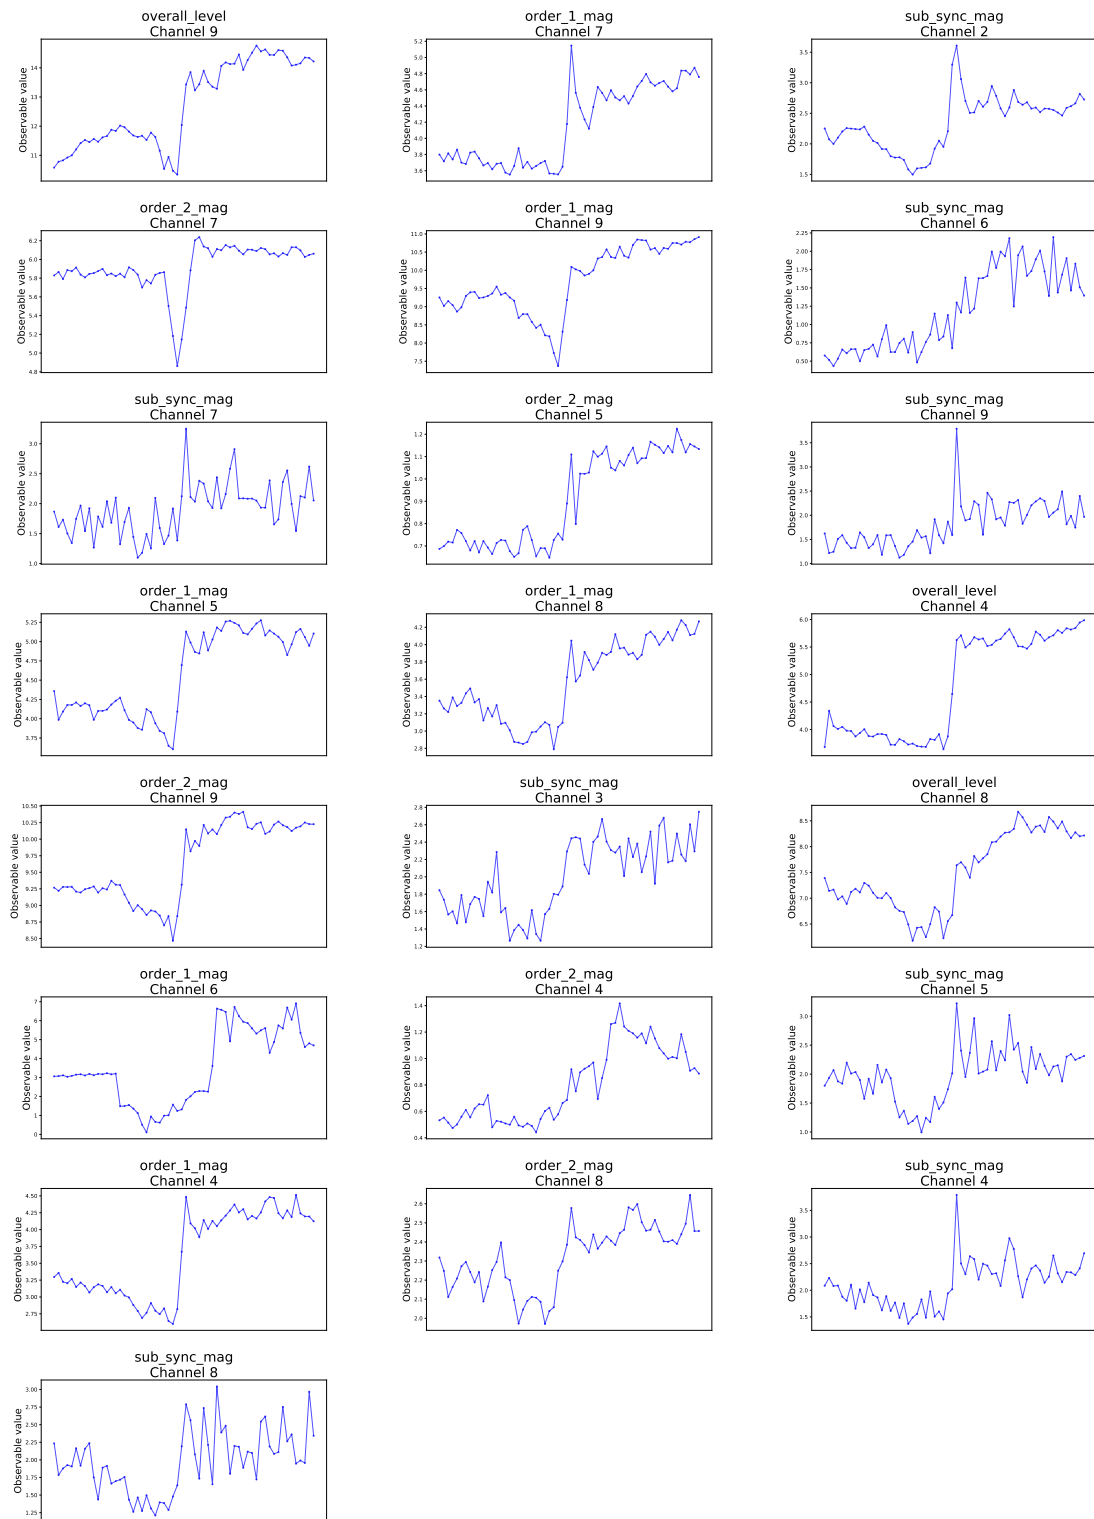


Figure 4.30: Upstep profiles for Turbine A vibration channels, extracted using the density-based methodf

for step change extraction is not sensitive enough to the differing conditions seen on each channel, and could be misdiagnosing large proportions of rou-

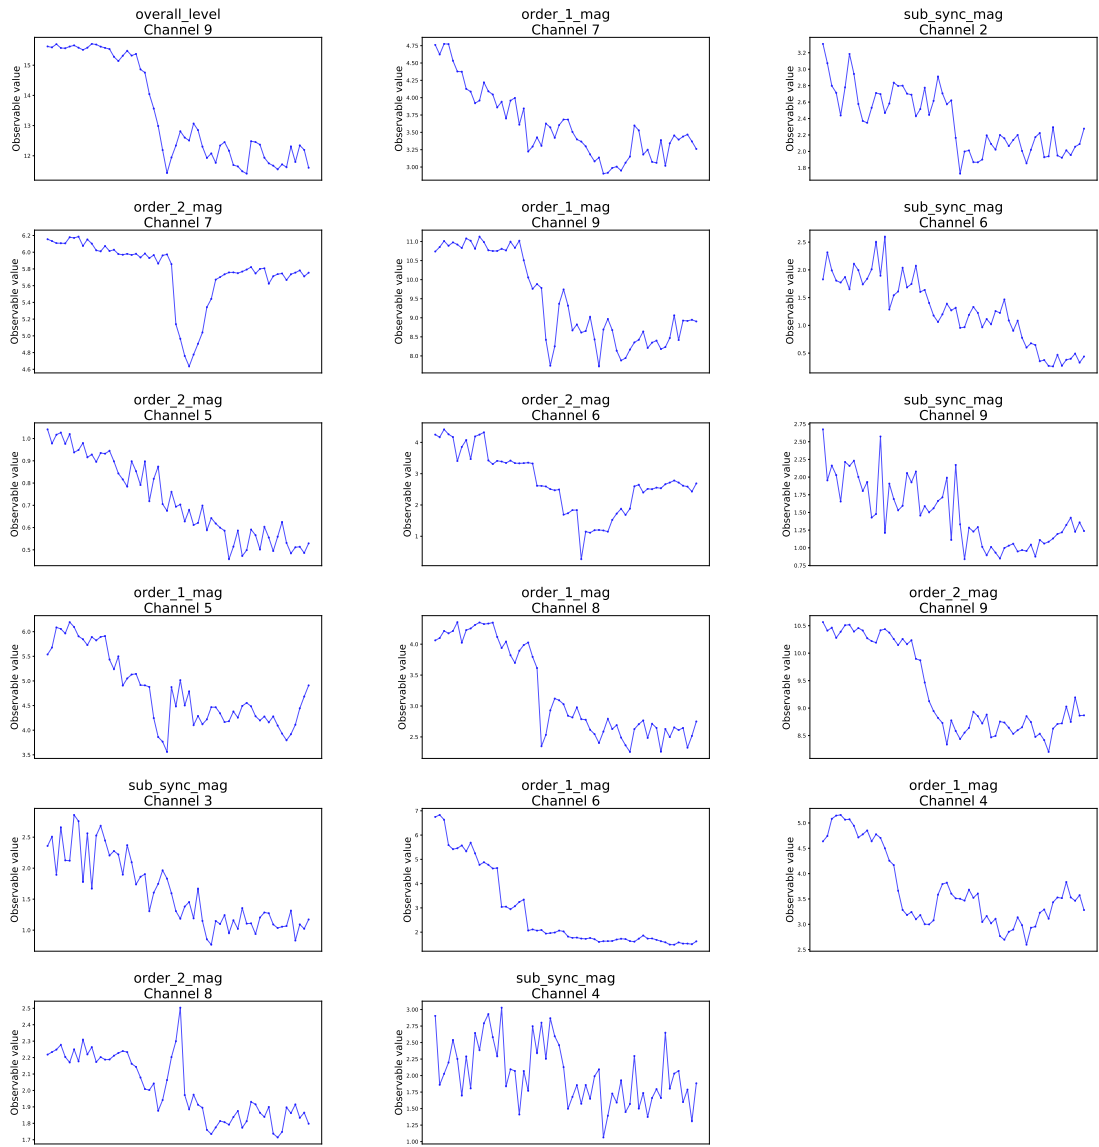


Figure 4.31: Downstep profiles for Turbine A vibration channels, extracted using the density-based method

Table 4.7: Turbine A overall level steps

Channel	No. steps	$\mu(S)$	$\mu(\tau_{Lead})$	$\mu(\tau_{Tail})$
4	9	1.51	0.84	0.56
5	5	1.66	0.51	0.55
6	3	1.42	1.23	0.76
7	3	1.21	0.99	0.53
8	9	1.41	0.92	0.57
9	39	2.72	1.55	1.56

Table 4.8: Turbine A first order magnitude steps

Channel	No. steps	$\mu(S)$	$\mu(\tau_{Lead})$	$\mu(\tau_{Tail})$
4	31	1.10	0.80	0.74
5	30	0.99	0.80	0.66
6	10	2.56	1.80	1.41
7	33	0.85	0.58	0.63
8	33	0.96	0.64	0.50
9	53	1.70	1.34	1.26

Table 4.9: Turbine A second order magnitude steps

Channel	No. steps	$\mu(S)$	$\mu(\tau_{Lead})$	$\mu(\tau_{Tail})$
4	14	0.43	0.20	0.28
5	20	0.35	0.19	0.17
6	8	1.24	1.07	0.75
7	24	0.38	0.37	0.40
8	32	0.28	0.26	0.26
9	89	1.15	0.88	0.88

tine alarm instances where an operational change has occurred.

4.6 Discussion

This chapter introduced a toolkit of statistical inference and data-driven techniques for the augmentation of a pre-existing TG KBS-based intelligent system, which allow for STST-base rulesets to be updated and self-tuned using automatic data selection and feature extraction methods. Further to providing a means to apply the knowledge base across more machine channels and examples of TGs, the distribution and parameter outputs from the approaches give insight into the empirical behaviours associated with steady-state and step change conditions for historical TG instances. Such methods can be used to aid in the decision making regarding the long-term health and state evolution of assets.

Chapter 5

A data-driven degradation model for circulators during refuelling

This chapter describes the second major contribution from this body of research: the investigation and development of a data-driven model for GC units during AGR refuelling events. The associated engineering problems, selected data-driven approach and potential applications are each discussed, with a focus on the value the modelling approach can provide to engineering staff to aid in GC health monitoring.

5.1 Problem definition

Gas circulators are monitored in a similar manner to turbine generators and other auxiliary rotating plant items on AGR sites: through the deployment of vibration monitoring transducers and surveillance systems. The mounting of these devices allows for streams of vibration measurements to be taken during the operation of the GCs, and give an insight into their behaviour and health. This is true for both steady state operation (as with the TGs discussed in Chapter 4), and for transient, dynamic conditions. This section will look into a key

transient event in the GC duty cycle: the low power refuelling event, for which online vibration data is available and potentially useful for reasoning about GC health.

5.1.1 Low power refuelling

One of the unique operational features associated with the AGR design, *low power refuelling* (LPR) is a regularly scheduled reactor event which allows for the replenishment of fuel channels [Non96] without a complete halting of generation. The channels not involved in the refuelling process are allowed to continue transferring energy to the TGs through the primary cycle which refuelling occurs. The main benefit of this is in increased uptime in generation hours when compared to offline refuelling reactor designs. Any MW output that can be safely maintained during fuel replenishment is desirable for the operator from an economic perspective.

From a monitoring perspective, regular LPR and other in-core interaction events also provide unique opportunities to make measurements of the refuelled channels to assess in-core condition and health. These structures are often unobservable to the operator without costly invasive inspections made while the NPP is offline, so any opportunity to collect data about channel state is worthwhile. The AGR design's flexibility in allowing for dynamic duty cycles and the process of partial refuelling is a feature that has been taken advantage of for numerous successful core health studies [WMT12, BWMR17]. Refuelling events typically occur for 5-8 fuel stringers per campaign over the period of days, with around 6-8 weeks separating campaigns [ZPC⁺04].

The LPR event is defined by an operational characteristic referred to as 'castling': the modification of the output reactor load to approximately 70% and 30% of full capacity at set intervals. This duty cycle provides alternate periods of medium and low power, and during the low phases it becomes

possible for the operator to interact with particular channels, restock fuel, and maintain or replace fuel stringer components. Castling gives the generator load time series during a single LPR campaign a distinctive profile; an example of which is provided in Fig. 5.1.




Figure 5.1: Typical generator load profile during AGR refuelling events with LPR highlighted

The example data in Fig. 5.1 for a single LPR is taken from a typical CM surveillance and interrogation hardware/software system used by reliability engineers responsible for the health of GCs. The variable fidelity of the time series in this instance can be explained by the common practice of taking increased granularity snapshots during periods of operational interest.

Correspondingly, the GCs are driven at variable circulatory output rates during each of the 'castle' steps in a refuelling campaign to regulate the cooling level required by that portion of the reactor. These changes of state are reflected in the vibration-based response of the GC data recorded from each unit, showing variation during each of the change periods. In common AGR practice, the castle features will total between 6-8 periods during a single LPR event. This is dependent on the replenishment requirements from an operational and fuel stock perspective.

An example of the vibration profile exhibited during a refuelling event is

provided in Fig. 5.2, with peak-to-peak overall amplitude measurements taken in both the horizontal and vertical axes at the drive and non-drive ends (DE, NDE) of each unit.




Figure 5.2: Typical vibration response to changing conditions during AGR re-fuelling events

5.1.2 Transient monitoring

‘Transient’ refers to periods of operation undergoing change or fluctuation, in contrast to ‘steady state’ behaviours where conditions are taken from a stationary distribution of values. This has historically been used as a catch-all term to describe unexpected changes, faults or temporary fluctuations, however it is argued that the definition can also be expanded to include non-steady state behaviours that are part of the normal operational duty cycle of a plant item.

Increased interest in transients as a potential source of diagnostic and prognostic information in recent years [BIAL⁺17] stems from the premise that engineering systems undergoing change can be subject to greater stresses, and their resulting data streams during transients can provide new information regarding the asset’s condition previously unattainable during steady state operation.

5.2 Strategy

5.2.1 Importance of the LPR

Monitoring GCs during refuelling represents a potentially useful source of information for a few reasons:

- **LPRs drive the circulators over a large dynamic range:** The effects of bringing the experienced load from full to 70% to 30% at regular intervals is argued to be of similar stressing factors, or at least comparable, to typical run up and run down conditions. Run up/run down events are already commonly examined [ZGMS04, RST17] in turbomachinery CM, providing useful data-based condition insights.
- **LPRs are regular:** During the lifetime of a circulator (the studied example has data over the course of 4 years), it can expect to experience numerous refuelling campaigns scheduled at roughly consistent intervals (approx. every 6-8 weeks). This repeated transient action could provide a wealth of previously unexamined information about the health and potential degradation of the machine class at key points in its lifecycle.
- **LPRs have a strong driver/response relationship:** Since the driving conditions are entirely observable (that is, we have measurement of the load alterations in the system), it becomes possible to construct an empirical relationship between the input and response observables.

The final point in particular is noteworthy; many engineering systems requiring reliability studies do not have a strong coupling between input and response, which may necessitate implicit measurements of latent, unobservable states for meaningful analysis. For example, the dynamic conditions of the AGR core itself are too complicated and numerous to map directly to a single operational driving parameter.

5.2.2 Data-driven model

The strategy for building a model of circulator operation during refuelling campaigns herein can be split into two main goals:

- 1. Construct a robust and accurate LPR time series classifier which automatically labels data consistent with refuelling events from historical records of circulator operation.

With eight instrumented GCs per reactor unit, the multiplicity of CM data streams introduced thus far is already a factor in examining the fleet-wide reliability of these primary cycle units. Extracting LPR and LPR-like transient data in an automated fashion from the large back catalogue of archived GC data opens the potential for analyses on a large subset of vibration signals previously unsegmented from regular operation.

- 2. Examine any temporal characteristics of the model's continued use to isolate potential prognostic metrics.

The discrete nature of the repeated refuelling campaigns applied to GCs provides a 'snapshot' of a system event at numerous moments of the operational history of the plant item. Any evolution in model parameters, orientation or empirical indicators in line with continued operation could provide an implicit mapping of any underlying degradation processes being experienced by the machine.

5.3 Model selection - theory

The best practice in training an ML model is to examine a number of suitable algorithms and select the approach which yields the best accuracy [Kot07, Mur12] with the training domain and test criteria at hand. Note that performance in this study refers to model accuracy, but this can also be quantified

other metrics such as precision or mutual information. A variety of approaches exist for this evaluation; *cross validation* [Koh95, CW12] (CV), model evidence comparison [Bis06] and likelihood ratios [Vuo89] encompass most of the model selection approach types for frequential- and Bayesian-style reasoning respectively.

The suitability of each of the earlier identified candidate algorithms for use in the LPR vibration data model is examined using a scoring method combining CV and out-of-sample test error (defined as the misclassification rate in the held-out test set) when compared to the full LPR 2006 - 2010 dataset.

5.3.1 *k*-fold cross validation

Splitting the dataset \mathcal{D} into training and test subsets while avoiding unintentionally over-fitting to a particular test collection can be a difficult process. A thorough shuffling and iterative train/test split procedure is often recommended to avoid over-training to a particular data orientation, which would reduce the applicability of the classifier in practice. Building a suitably general model which is robust with as wide a range of test inputs relies on minimising any inherent biases in the evaluation criteria; an issue that cross validation-type techniques seek to address.

k-fold CV splits the dataset into *k* evenly sized partitions, which are then each used in turn as the test set against a model trained by the remaining $k - 1$ folds. Each of these train/test splits are evaluated on their predictive accuracy (in the LPR model example, how well the model predicts the $\{online, upper, lower\}$ class membership for each item in each test set), as illustrated in Fig. 5.3. The average performance among the *k* splits is utilised as the overall CV performance of the model.

Selecting *k* is dependent on the size and nature of \mathcal{D} , and the requirements from the trained model. Since the LPR dataset is built from 21 individual re-




Figure 5.3: k -fold cross validation procedure

fuelling events, $k = 21$ is used in this example. This shuffles the full refuelling data into folds of approximately equivalent size to the LPR events, but removes any temporal- or event-based biases inherent from potential degradation from the circulator unit.

5.3.2 Hyperparameters

Learning algorithms often require a number of initial conditions to be set which govern the modelling technique's strategy when finding an optimal hypothesis. These *hyperparameters* differ for each ML approach, but their selection is often critical to the success of the final model: for both \mathcal{D} -classification accuracy and generalisability.

The systematic measurement of classification success with respect to a given dataset \mathcal{D} across a range of hyperparameter values for a learning technique is known as a *grid search* strategy [CWCCCJ03], and represents one of the most effective means to tune the hyperparameters of ML techniques.

5.3.3 Classifiers

A selection of linear and non-linear classifiers were investigated with the refuelling dataset in order to identify the most suitable.

5.3.3 a) Linear classifiers

A total of four linear classifiers were evaluated (with their corresponding reference section parenthesised):

- Perceptron-based linear model (3.3.2),
- Logistic linear model (3.3.3),
- Linear SVM with $L1$ -regularisation (3.3.4 a)),
- Linear SVM with $L2$ -regularisation.

The implementations of these techniques [PVG11, PA13] each have a single regularisation hyperparameter to be optimised; α for the perceptron and logistic linear models, and C for the linear SVM models. Note that this simplification when compared to historical versions of these methods (which can typically require a grid search of three or four hyperparameters) is due to improvements made in the `scikit-learn` model implementations, which use a combination of optimised search and hypothesis space navigation heuristics.

These four were selected over the other techniques outlined throughout Chapter 3 due to their suitability for low dimensional supervised learning problems, and their track record in application to this domain. Future research into other techniques (such as random forests, which often perform well in high-dimensional spaces) and with enriched feature vectors would be a worthwhile next step.

5.3.3 b) Higher-order classifiers

The non-linear RBF kernel SVM was also trained and tested with the refuelling dataset, in an effort to investigate potential non-linearities in the decision function hyperplanes separating the three load classes. An additional kernel hyperparameter γ is required to be optimised, making the grid search 2-dimensional.

5.3.4 Dataset overview

To gain a better perspective of the dataset at hand, this section provides a few explanatory visualisations and clarifies a number of characteristics inherent to the GC data. Forming a proper understanding of the available information by examining not only the relevant domain knowledge but also through techniques of data exploration is important in constructing robust data-driven models and useful systems built upon these.

The available observables for the GC asset class are broadly equivalent to the vibration-based machine views common across rotating plant items. In order to reason about the LPR and any potential effects it may impart on the GCs, any model requires a measure of the duty cycle and corresponding response of the machine. The LPR state is defined by variable load signatures (as introduced in Sec. 5.1.1), making the *generator load* operational observable the best choice to define state. For quantifying response to state, the *overall amplitude* vibration signals at both the DE and NDE on the horizontal and vertical axes are used.

A single circulator unit, which went on to experience an eventual inspection-based failure, was taken as the data source for the construction of the model. This data covers the period 2006 - 2010 and includes numerous LPRs, outages and ad hoc operational adjustments. This totals 18,147 4-dimensional time-stamped measurements. The experienced load duty cycle for this period is pre-




Figure 5.4: Full (2006-2010) duty cycle of the GC used to construct the model, with zoomed LPR instance

sented in Fig. 5.4.

From this period, a total of 21 labelled LPR events were selected to build the LPR model representation. The load periods for these are provided in the context of the full duty cycle in Fig. 5.5.

Fig. 5.6 provides the corresponding vibration response across each of these events in horizontal and vertical phase space for both the DE and NDE transducers mounted on the circulator, with different colours corresponding to separate LPRs. This visualisation provides the full range of vibration response behaviours to repeated LPR events, and shows initial visual evidence of some clustering behaviour - potentially around the three LPR states.

Fig. 5.7 illustrates a kernel density estimate of the generator load observ-




Figure 5.5: Duty cycle of the selected LPRs used to train the model

able across each of the selected LPR events, along with an average probability density function. Three dominant modes of load value are evident along with their probabilistic frequency - which maps intuitively to the LPR's casting characteristics.

The density estimates in Fig. 5.7 provides the kernel densities for each LPR, along with the average kernel density denoted by the thick blue line. This average is used by the GMM model to identify the suitable LPR boundary




Figure 5.6: Drive and non-drive end vibration phase space of the refuelling events

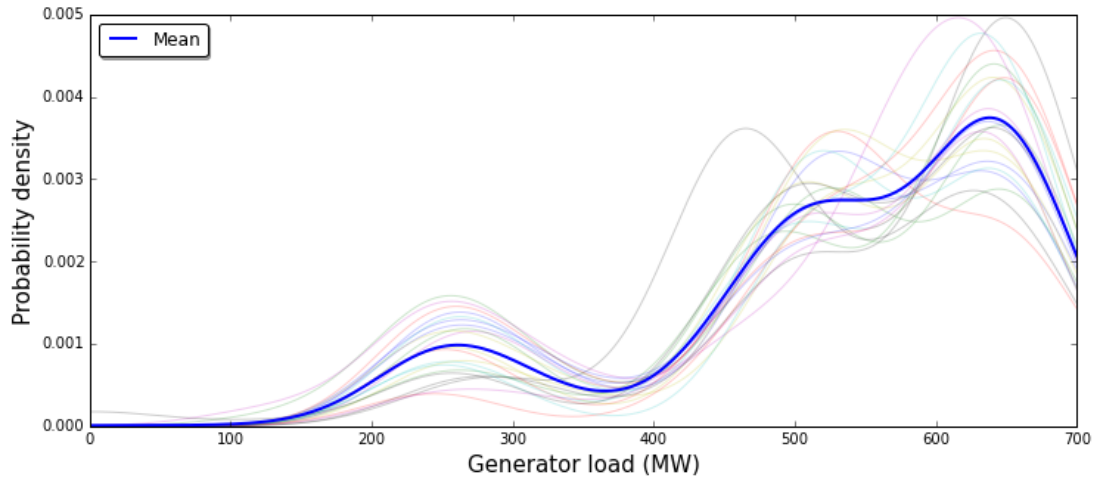


Figure 5.7: Kernel density estimate across the range of load values for the entire dataset

settings empirically.

5.3.5 Labelling LPR behaviour

Using an $N = 3$ Gaussian mixture model (as introduced in Section 3.3.6 b)), the classification boundaries which quantify membership to each of the LPR load states (formally $\{\text{Online}, \text{Upper}, \text{Lower}\}$, each of the load levels) are estimated empirically. This is achieved by taking the density minima between the identified distributions (image provided in Fig. 5.8) as the boundaries. The resulting piecewise labelling function is provided in (5.1).

$$f_{\text{State}}(x) = \begin{cases} \text{Online}, & : x > 603 \\ \text{Upper}, & : 386 \leq x \leq 603 \\ \text{Lower} & : 0 < x < 386 \end{cases} \quad (5.1)$$

The delimitation of the load values into three discrete behavioural categories allows for the labelling of the entire vibration dataset; providing the basis for a supervised classification problem as introduced in previously. With one of the three major driving behaviours of the LPR assigned to every time




Figure 5.8: Mixture of Gaussians ($N = 3$) model of load data, identifying the bounds of the three load-based behaviours

series observable, it becomes possible to generalise the vibration response to each of these behaviours in a sufficiently trained data-driven model and potentially reason about any variation in this response from an observed range of nominal values.

5.3.6 Training data feature vector and dimensionality

The data used in the supervised learning problem described in the following sections is explained in this short section to provide some clarity about the ML task at hand. The input feature vector is a 4-dimensional array of the horizontal and vertical drive-end and non-drive end peak-to-peak instantaneous vibration, while the output label is the load-based class taken from $[Online, Upper, Lower]$ (corresponding to the load values). Over the 21 LPR campaigns used in the training data, there were 2,366 examples, with an average of 113 measurements per refuelling event. For each class, 43% of the data points in LPRs were *Online*, 40% were *Lower* and 17% were *Upper*. The imbalance between each of these classes was deemed to be manageable due to the relative similarity in scale between steady-state (*Online*) and transient (*Upper* and *Lower*) as a whole.

5.4 Model selection - results

For a primary implementation of an LPR model with a focus on label accuracy and performance, each of the candidate models take all four vibration observables as range normalised real value inputs, as recommended by various practical ML engineering sources [CWCCCJ03]. This allows for any 4-tuple of vibration values of the form (DE horz, DE vert, NDE horz, NDE vert) to be used to predict the circulator state with respect to refuelling at that point. Note that the primary aim of this refuelling model is to predict the category of duty cycle behaviour (as defined by the load) from the vibration alone, which explains why load is omitted.

A final evaluation for each of the ‘tuned’ candidate models post-grid search is made on the full available operational records from which the refuelling events were selected. The classification accuracy across this extended period is used as the decision point in selecting the most suitable model for further implementation.

5.4.1 Grid search results

5.4.1 a) Linear classifiers

The CV grid search scores for the linear perceptron, linear logistic, linear $L1$ -SVM and linear $L2$ -SVM are provided in Fig. 5.9. These figures show the variation in CV accuracy (including the standard deviation of each data point highlighted by the blue envelope) when searching the hyperparameter space associated with each model type. The hyperparameter setting with the largest CV score for each model is used as the exemplar model for each. The test accuracy dropping when compared to the cross-validation results is to be expected as the model is tested across the fuller validation set. It is of interest that the $L2$ -SVM increases in accuracy when considering the full test set: there are a

Table 5.1: CV and test scores for linear classifiers

<i>Classifier</i>	<i>CV results</i>			<i>Test accuracy</i>
	<i>Hyperparameter</i>	<i>Value</i>	<i>CV score</i>	<i>Score</i>
Linear perceptron	α	0.1233	88.5%	76.05%
Linear logistic	α	$8E^{-5}$	87.8%	71.55%
Linear $L1$ -SVM	C	10.975	87.8%	75.69%
Linear $L2$-SVM	C	58.57	87.9%	88.23%

number of potential reasons for this, including that the training dataset being drawn entirely from dynamic LPR states was more challenging to model state more accurately than the state-steady examples in the full LPR history.

It should be noted that general shape of these results being approximately mirror images (perceptron/logistic vs. SVMs) of each other is to be expected - SVM convention dictates [KS03] that values of C are roughly proportional to $\frac{1}{\alpha}$, their gradient descent counterpart.

Each model (with the tuned hyperparameter settings from the grid search results) is then applied to the full 2006 - 2010 GC dataset. The percentage of correctly labelled states from the vibration input - the test accuracy - is used as the final selection criteria. The full results for the linear models are provided in Table 5.1, showing the linear SVM with $L2$ regularisation is the most accurate from the candidate algorithms.

5.4.1 b) Higher-order classifiers

Results of the two-dimensional grid search for the RBF SVM are provided in Fig. 5.10 and Table 5.2. One of the most prominent features of this search space is a high accuracy diagonal, where the proportional variation of C and γ maintain a high cross validation rate. This has been observed in a number of SVM models [Gri14] and can be considered normal - each of the candidate models in this diagonal should be a candidate for selection, with the highest test accuracy providing a useful measure for the most desirable selection. The range of results for both CV and test accuracy (in terms of correct labelling of




Figure 5.9: CV grid search scores for each candidate linear model

state) observed in this diagonal region is also provided in Fig. 5.10.

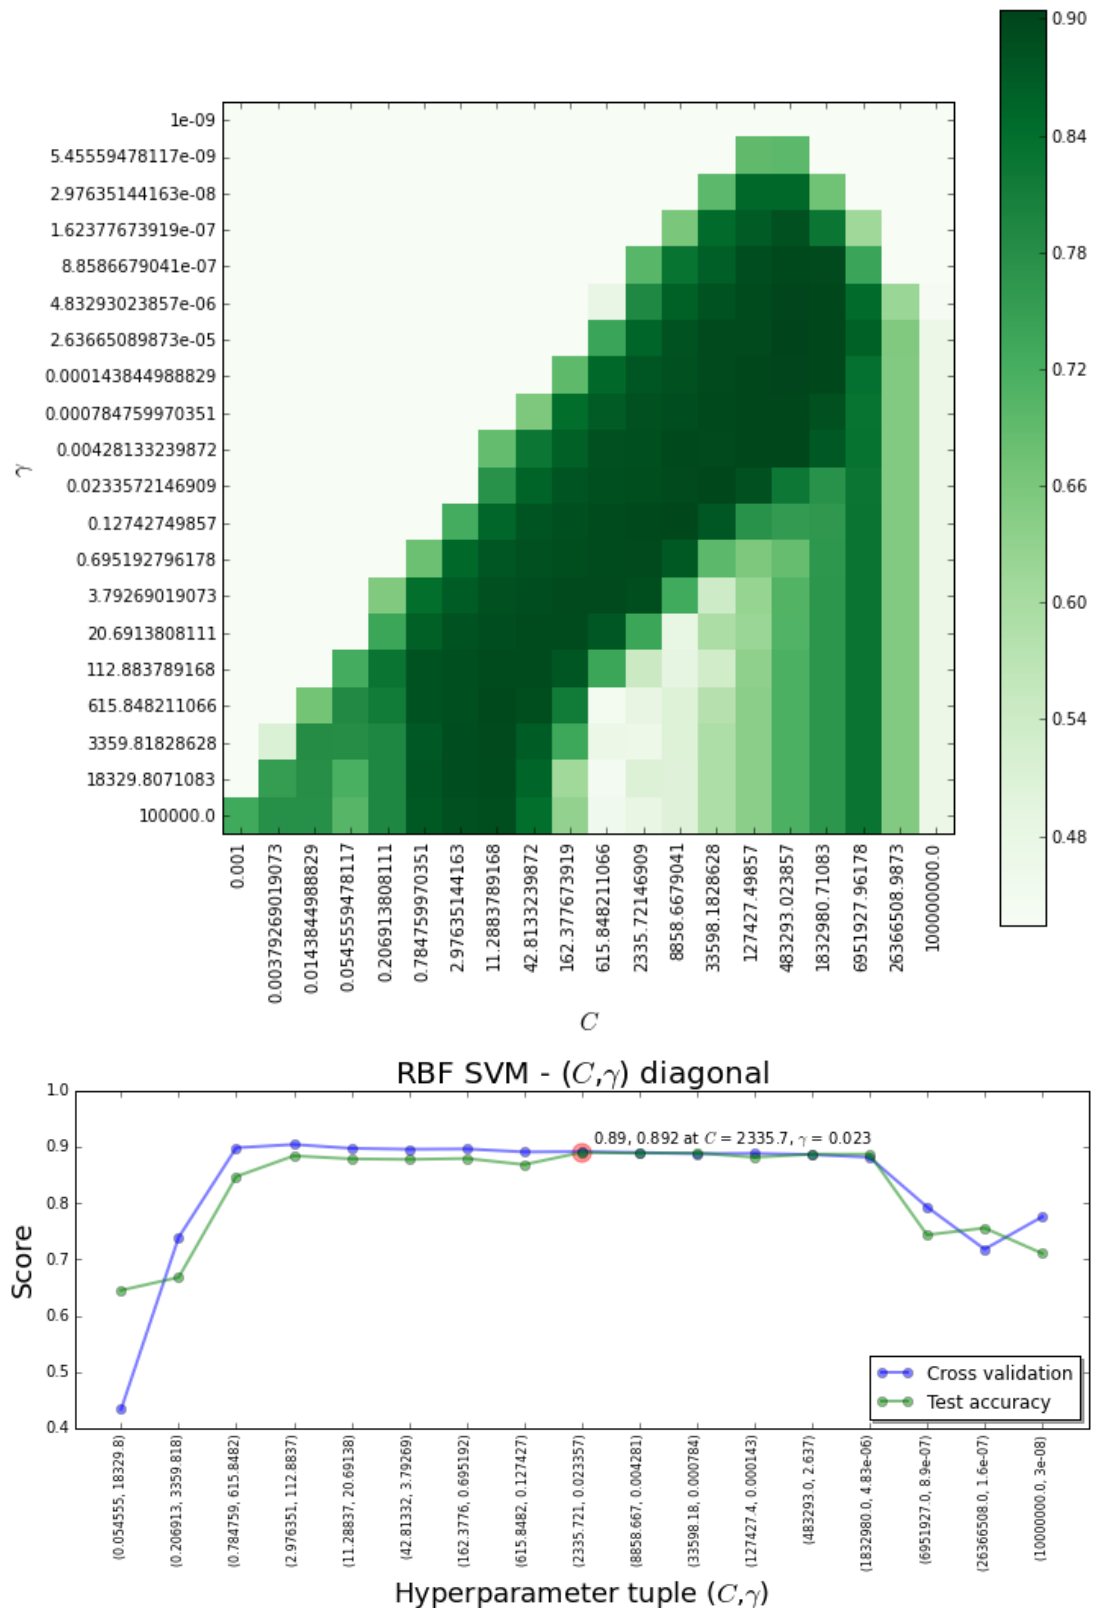


Figure 5.10: Classification accuracy of RBF SVM with varying C and γ

Table 5.2: CV and test scores for RBF SVM diagonal

<i>Hyperparameters</i>		<i>Validation</i>	
<i>C</i>	γ	<i>CV score</i>	<i>Test accuracy</i>
0.0545	18329.8071	43.4%	64.5%
0.2069	3359.8182	73.9%	66.8%
0.7847	615.8482	89.8%	84.7%
2.9763	112.8837	90.4%	88.4%
11.2883	20.6913	89.7%	87.8%
42.8133	3.7926	89.5%	87.7%
162.3776	0.6951	89.6%	87.9%
615.8482	0.1274	89.1%	86.8%
2335.7214	0.0233	89.1%	89.0%
8858.6679	0.0042	89.0%	88.8%
33598.1828	0.0007	88.7%	88.9%
127427.4985	$1.438E^{-4}$	88.8%	88.1%
483293.0238	$2.636E^{-5}$	88.6%	88.7%
1832980.7108	$4.832E^{-6}$	88.1%	88.7%
6951927.9617	$8.858E^{-7}$	79.3%	74.3%
26366508.9873	$1.623E^{-7}$	71.8%	75.5%
100000000	$2.976E^{-8}$	77.5%	71.1%

Note that the hyperparameters used in the results outlined in Table 5.2 were selected with a systematic grid-search exploration strategy, but without any parameter space exploration beyond these bounds. Using techniques such as Bayesian optimisation [CT07] to investigate this would be a worthwhile next step.

5.4.2 Discussion

The results provided in the previous section illustrate the effectiveness in using the SVM process when building a classification model for the LPR labelling problem, with the two best test scores coming from the linear $L2$ - (88.23%) and RBF (89.0%) SVMs. These procedures are in fact theoretically similar [KL03], with the linear SVM being a specific case of an RBF-type SVM - explaining their coincidence at the top of test accuracy rankings in this case.

Training times¹ for the RBF SVM were significantly longer than those ex-

¹Training on the LPR dataset was done on a Macbook Pro Retina 2.5GHz Intel Core i7 with

hibited by the linear $L2$ -SVM: taking an average of 57:05 per run (5 runs) compared with an average of 2:35 (5 runs) respectively. Despite the slight improvement in test accuracy when using the non-linear RBF approach, the associated training time is presented to be disadvantageous when considering the nature of the problem at hand. Successive training of models at various timesteps of the GC operational lifecycle combined with the number of units fleet-wide (8 per reactor) make the shorter training time of the linear $L2$ -regularised model preferable, despite the small loss in accuracy.

With these points considered, the **L2-regularised linear SVM** was selected as the most suitable model for the LPR classification task. This model provides the most accurate means of identifying LPR-type conditions from vibration data only, and enables the long-term modelling approaches introduced later in this chapter.

5.5 Decision support

The focus when developing intelligent and data-driven technologies for the CM engineering context should be on delivering actionable insights and information to end users. Putting any predictive capabilities to the side (discussed later in Section 5.6) for the moment, the model seeks to provide two key functionalities:

- **Automatic LPR data identification:** The model should allow for the easy categorisation of LPR-type data points from the large historical records available to the operator, without the requirement for logistic records or dates of campaigns to be input. This allows for the long-term behaviour of refuelling events to be documented and studied with minimal co-ordination and effort.

16GB DDR3 memory

- **Improved presentation of LPR as an event:** The model should provide a useful visualisation of the refuelling event considered as a whole, moving away from the time series-only approaches currently employed. This allows for the LPR to be reasoned about discretely - a step forward from the rudimentary analyses which don't currently consider nuances of any operational events which may be driving change.

This section examines the performance of the techniques employed in meeting these requirements and how they might be utilised in the reliability engineering environment.

5.5.1 Identifying LPR state data

The load-based class values and the predicted values (as annotated using 5.1) from the modelling approach are compared in Fig. 5.11 across the full operational history dataset, covering the GC's duty cycle between 2006 and 2010. Included in this diagram are the deltas between the predicted and actual LPR states: 0 refers to a correct result, 1 is a class difference of one (e.g. *Online* instead of *Upper*), 2 is a class difference of two (e.g. *Online* instead of *Lower*)

The prediction of the classifier is provided for each of the load values through the duty cycle in Fig. 5.12, which can be compared with the earlier Fig. 5.4. As demonstrated in Section 5.6.1, the model classifies the state correctly 88.23% of the time using the vibration data, with each of the discrepancies highlighted by the red deltas in Fig. 5.11. The vibration observables for each of the prediction-based segmentations at both the drive and non-drive ends of the GC are presented in Figs. 5.13 and 5.14. Note that the hue of points denotes density of instances in that part of the time series.

These groupings represent a data-driven estimation of the *Online*, *Upper* and *Lower* behaviours across the full operational dataset. This is an improvement from simply thresholding the load-based duty cycle for one key reason:




Figure 5.11: Class vs. predicted output, with deltas shown in red

assigning the state classes based on the experienced vibration profile takes an entirely response-based data view, which considers potential changes and evolution to the behaviour of the individual GC with continual use. This becomes particularly useful when considering the health monitoring of numerous circulators subject to differing duty cycles, conditions and maintenance regimes, where any circulator specific information regarding condition could prove use-




Figure 5.12: Load values of predicted class results




Figure 5.13: Load values of predicted class results for drive end vibration

ful.

Specifically for the LPR instances, the confusion matrix corresponding to the 2,366 measurements taking during refuelling campaigns is provided in Fig. 5.15.

5.5.2 Phase space view

Reliability engineers utilise phase plots in their examination of rotating machinery in order to ascertain properties regarding the magnitude and phase of vibration on their machines. This data view is distinct from the common time series analysis associated with much of the rest of CM (and predominantly the focus of this chapter thus far) - the temporal aspect of the data stream is largely abstracted out in favour of a holistic overview of the observed behaviours and




Figure 5.14: Load values of predicted class results for non-drive end vibration their context in the larger domain.

Phase plot-like approaches to visualisation are useful for considering the




Figure 5.15: Confusion matrix of model output from time series values taken from periods of refuelling

LPR behaviours in both the context of single refuelling events and the general behaviour of the machine over longer periods. Considering the drive- and non-drive-end orientations of the machine, the vibration states as presented in time series in Figs. 5.13 and 5.14 can be presented to illustrate their context in the larger data domain of the machine.




Figure 5.16: Phase space of drive end vibration, with predicted labels

The state segmentation at the drive end (Fig. 5.16) is the most prominent of the two phase space overviews, with a clear visual clustering of the three behaviours in the 20 - 50 μm pk-to-pk ranges on both observable axes. The non-drive end instance is much less discrete, with a much larger dynamic range

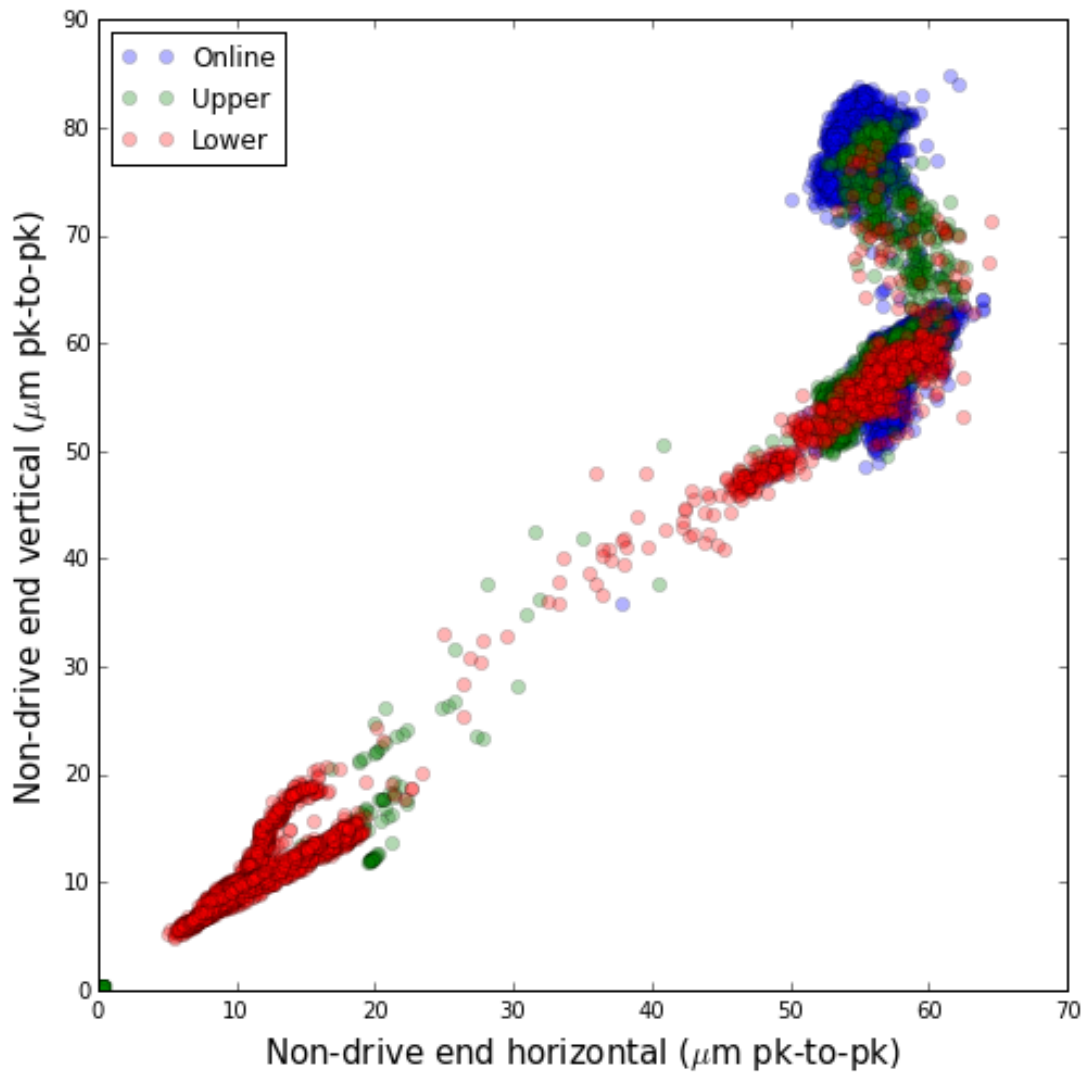


Figure 5.17: Phase space of drive end vibration, with predicted labels

without the near linear separability of the drive end data.

Both examples exhibit a further linear behaviour cluster corresponding to markedly lower vibration observable values, which are predictably dominated by the `Lower` behaviour. In the general scheme of circulator operation, these values are likely to correspond to run-up and run-down transient conditions experienced by the GC at the start and end of maintenance periods.

The potential usefulness of presenting the model classifications in phase space become more apparent when considering the characteristics of individual or groups of refuelling events for comparative purposes, as illustrated in

Fig. 5.18, which shows the output from 3 LPR events. Applying the same principles of hyperplane separation to the drive end subset of the GC vibration data allows presentation of the behaviour as sectors in phase space. These data-driven regions (and their associated accuracy properties) can provide an informative overview of the data at a particular end of the circulator to allow for comparison of behaviours by temporal groupings.




Figure 5.18: Multiple segmentations of the operational data in drive end space, providing an abstraction of the vibration data for examination

5.6 Predictive capabilities

A sufficiently trained LPR model represents a quantitative measure of normality for the given machine at that point in its lifecycle, which can then be compared with subsequent refuelling campaigns and re-trainings of the model with new GC data. This snapshot can provide the basis to reason about any evolution in the vibration response of the GC to continued operation, which opens the potential for prognostic measures of state change and remaining useful life (RUL) to be investigated. This section explores the opportunities the developed refuelled model provides in creating predictive signals to aid in determining future machine states.

5.6.1 Model output evolution

Considering the entire operational history of the GC used throughout this study, the misclassification rate of the `online` behaviour increases during the latter stages of the unit's operation. This is illustrated in Fig. 5.19, highlighting that as the GC progresses with use (with the assumption of some latent degradation process being experienced by the machine), the description of normality provided by the trained model becomes less applicable when considering the vibration response to `online` (non-LPR) behaviour. This feature can be used as an early indicator of potential state change in the monitoring of the GC: for example, if the windowed average of misclassified instances begins to increase then further investigative action can be taken. It is important to note that this rise in misclassification is present despite behaviour corresponding to the late-stage life of the GC being included in the overall training data.

5.6.2 State estimation

A key observation from the rising misclassification rate of the `online` behaviour shown in Section 5.6.1 is that the vibration response is quantitatively being altered enough for a discrepancy to be identified by the trained LPR model. Coupled with the context of the eventual failure of the GC used in




Figure 5.19: Misclassification of the `online` behaviour with operation

the study, this suggests the existence of a degradation process being captured explicitly or implicitly in the available circulator observables.

5.6.2 a) State classification

To investigate any change metric sequentially in the GC lifetime, each of the identified behaviours classified by the data-driven model (namely `online`, `upper` and `lower`) were delimited into four stages of operational progression: from early use through to operation just before the final inspection which brought the unit offline permanently. Each of these stages were labelled $\{\text{early}, \text{mid1}, \text{mid2}, \text{late}\}$ by order of progression.

Each of these data subsets were used as labelled inputs to a further classification model - allowing for the prediction of the point in the GC lifecycle a particular set of vibration observables corresponds to. The logistic linear model was used as it provides probabilistic outputs for each prediction, which allows for the state estimation process to be done continuously rather than discretely, allowing for the creation of a predictive function which can be updated with each new data instance. This approach follows a similar state estimation classification-based process to the prognostic model developed in systems such as [KTM⁺08].

The series in Fig. 5.20 presents the probabilistic output of the class membership for each of the behaviours labelled by the LPR model. Both the `online` and `upper` behaviours exhibit a clear ordering of each of the four latent states, while the `lower` behaviour is not as successful. This suggests that the most fruitful segmentations of the data for extracting a health indicator will be from either of the `online` or `upper` behaviours.

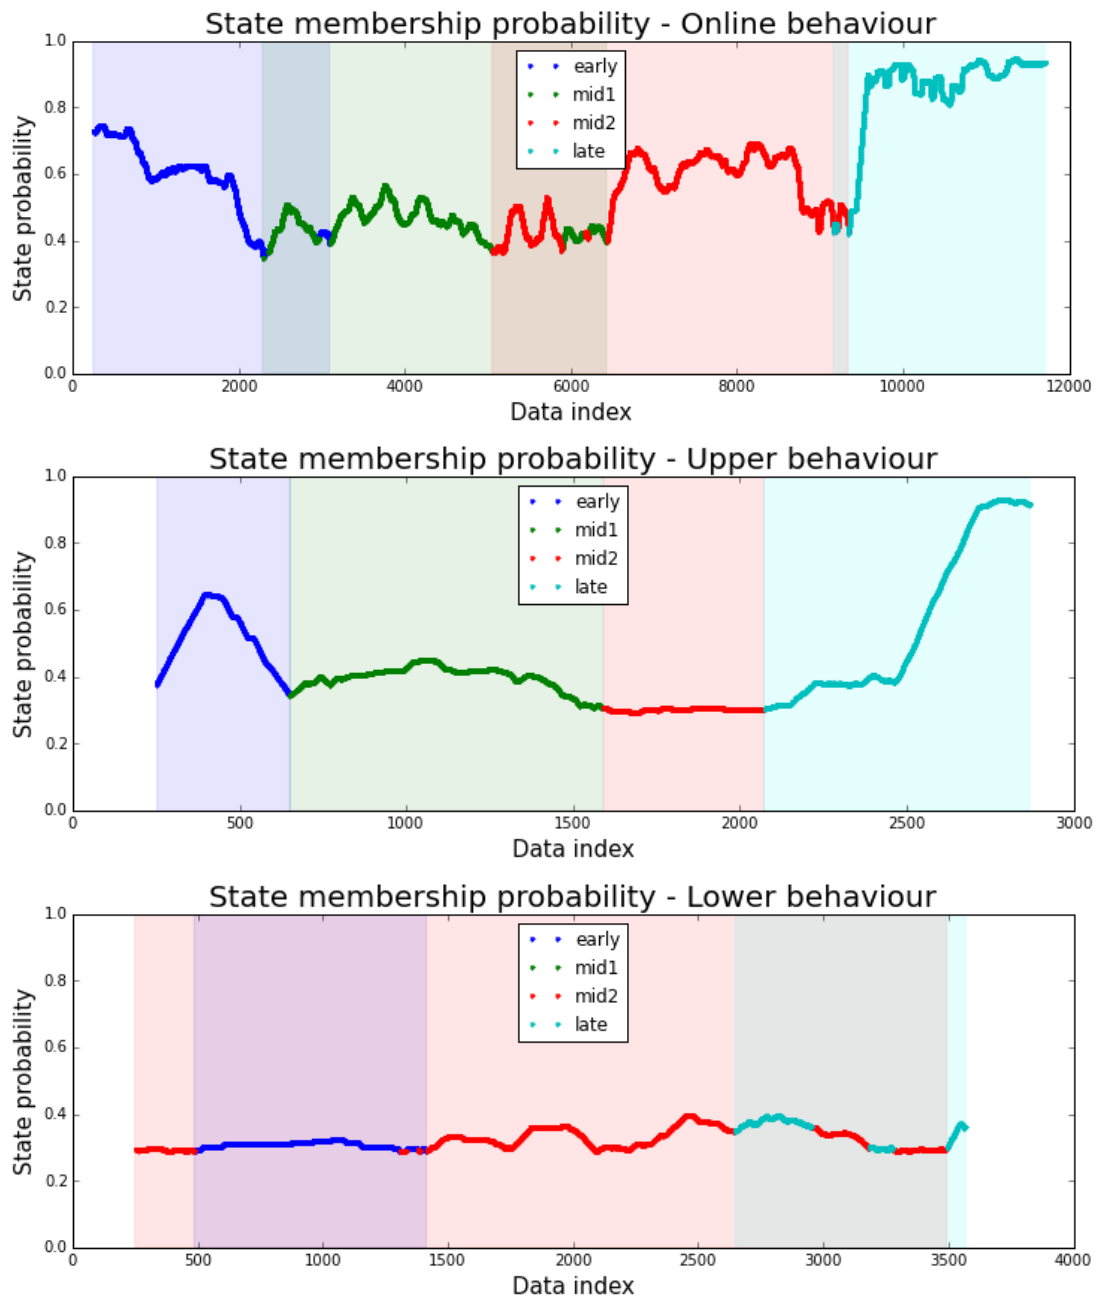


Figure 5.20: Most probable state estimates for each of the refuelling behaviours in chronological order

5.6.2 b) RUL estimation

Based on the assumption that the conclusion of the time series corresponds to the failure point in the life cycle of the GC, the RUL can be defined as:

$$RUL(t) = t_{Failure} - t \quad (5.2)$$

where $t_{Failure}$ is the end of the time series and t is any timestamp with corresponding data. Obviously the explicit RUL is unavailable at the point of operation, so the goal in building a prognostic system in this case is to use the available labelled RUL from this example to identify an implicit measure which can be used by engineers prior to failures. Note that the RUL is calculated at t for each individual data point, which explains why the RUL series differ in each figure of analysis.

Since the `late` behaviour corresponds to the response of the machine observed closest to the end of life, the probabilistic output of the LPR model for this behaviour represents a potentially useful measure for prediction. As the probability of `late` increases, the failure criteria being met is assumed to be sooner. The series' in Fig. 5.21 provide a comparison of the RUL (represented by the blue series) with $p(\text{late})$ for each of the LPR behaviours identified by the data-driven model. For comparison, the series' in Fig. 5.22 illustrate the results for the `mid2` classification; a label corresponding to earlier in the expected life cycle of the GC.

When evaluating these predictive measures, there are a number of considerations for determining suitability. Primarily, the accuracy when compared to the explicit RUL is most important. Also, as the operational history of the machine used in this study does not include any maintenance or repair records, any degradation measure is assumed to be largely monotonic [Cob10].

Considering these points, the `late` probabilities for the `Online` and `Upper` behaviours appear to map best to the RUL. While there are some limited breaches of monotonicity shown in the smoothed moving average for these, the prob-

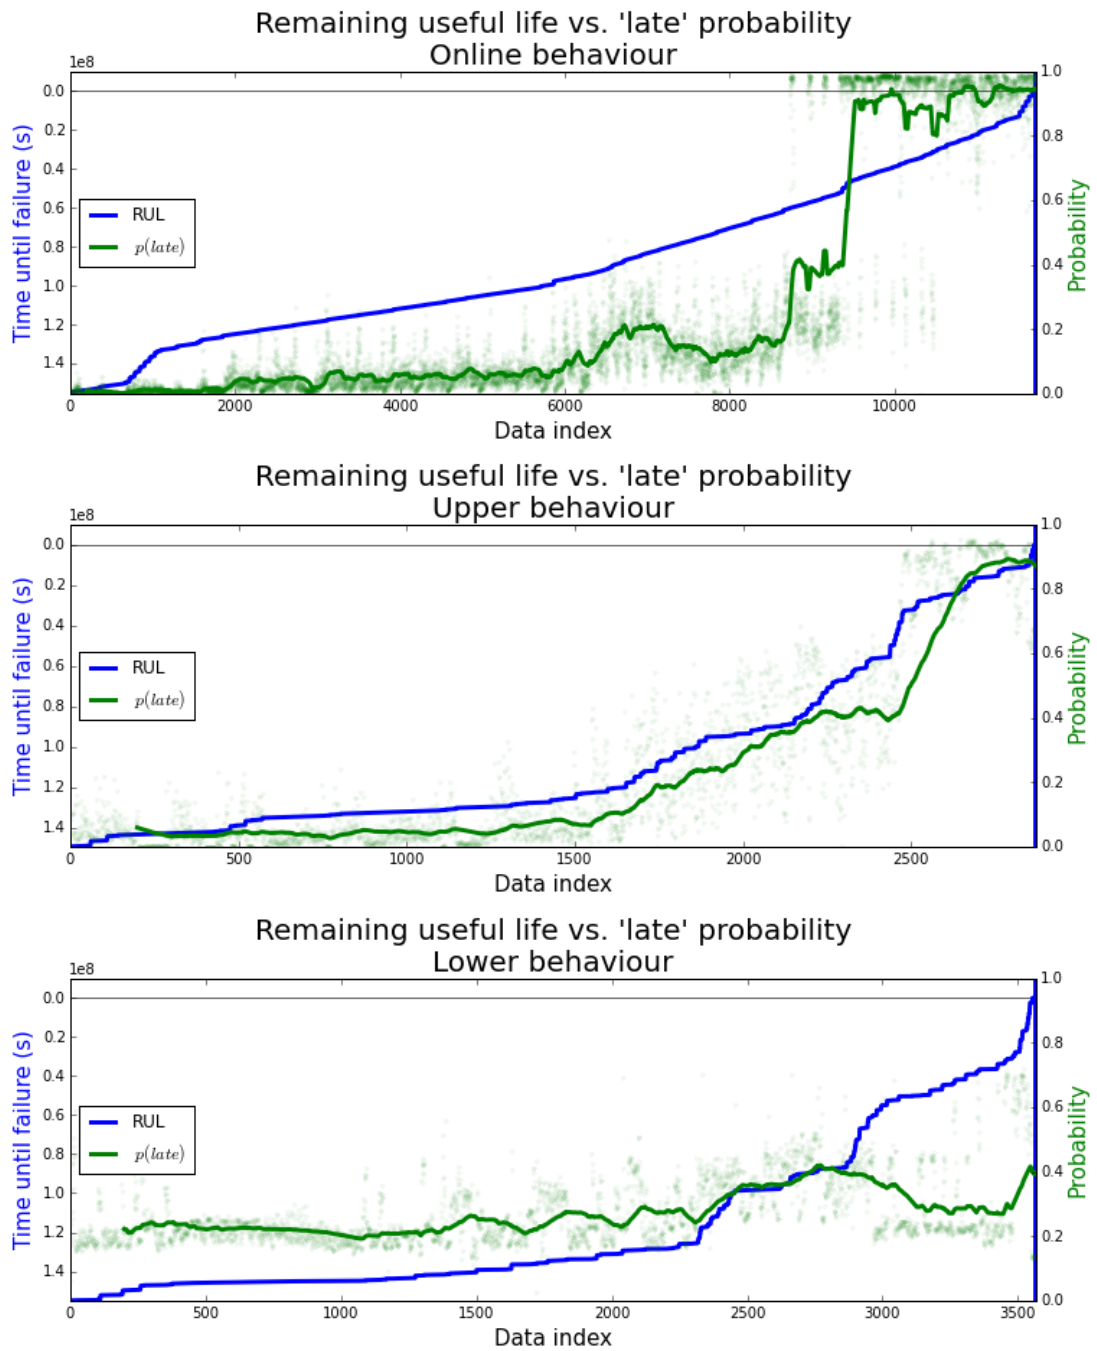


Figure 5.21: Remaining useful life compared with probabilistic output of LPR model for `late` class

abilities are largely trending upwards. The `Online` case has a much more discrete up-tick than the more gradual `Upper` measure, which suggests a combination of the two observables may be useful in identifying impending failures.

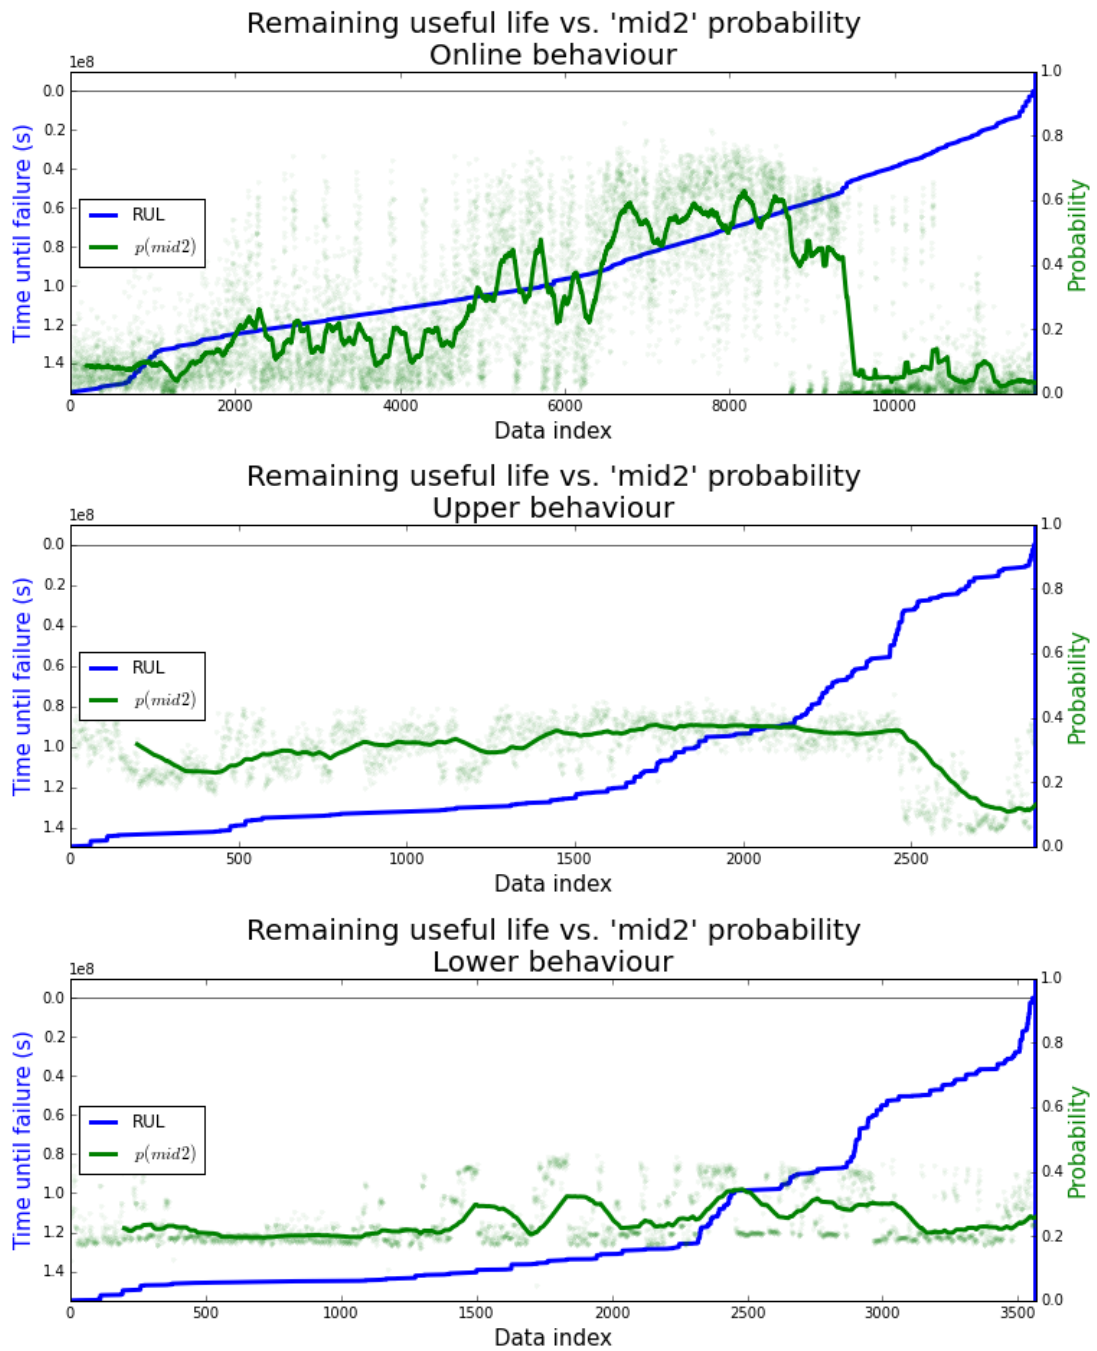


Figure 5.22: Remaining useful life compared with probabilistic output of LPR model for `mid2` class

The Pearson's correlation coefficients for both the class membership to `mid2` and late temporal states when compared to the true RUL are provided in Table 5.3. This shows that the vibration data labelled as `Upper` is strongly inversely correlated to the true RUL value. In more specific terms, vibration response

during 70% load transient periods appears to have an underlying temporal feature that is encapsulated by the ML model that can then be taken advantage of from a reliability engineering perspective by examining the likelihood of it belonging to a late-period class.

Table 5.3: Correlation between late temporal class & true RUL

<i>Temporal class</i>	<i>LPR behaviour</i>	<i>Pearson's corr.</i>
late	Online	-0.838
late	Upper	-0.938
late	Lower	-0.472
mid2	Online	-0.041
mid2	Upper	0.542
mid2	Lower	-0.146

The intended interpretation of the temporal class probability for reliability engineers is as a proxy for the remaining useful life. This should be presented in a more suitable context (e.g. 'predicted health metric' or similar). A rising trend in this metric between LPRs should be used as an early indicator that the GC is degrading, with higher probabilities corresponding to more likelihood of conditions associated with a machine failure.

5.7 Discussion

This chapter introduced a data-driven model for the health monitoring of AGR gas circulator units, detailing its construction, evaluation and potential for prognostics. For diagnostic classification problems, the model utilises a trained $L2$ -regularised linear SVM to identify vibration data corresponding to each of the states encountered during refuelling campaigns. This algorithm provides an 88.23% cross-validated accuracy, and was selected over the slightly more accurate RBF SVM due to its considerably shorter training time. As well as demonstrating high accuracy, this chapter also provided context on how the output of the model would be used by engineering staff to augment their exist-

ing CM processes surrounding GC units and their condition during refuelling events.

The selected techniques in this study are not an exhaustive set of the potentially applicable ML modelling approaches. It would be useful to continue this research by comparing the results from other supervised learning algorithms such as random forests [Bre01], gradient boosting machines [TGF18] and deep learning approaches.

Further to this, the model was shown to have some potential for predictive analytics when considering the overall health of the GCs. The model output was used as an input to a further classification metric (using the logistic linear algorithm, due to its probabilistic properties) which was trained based on temporal states in the GC's operational history. This showed promise in approximating the RUL of the circulator and could be used to schedule inspections and maintenance pre-emptively to avoid unexpected failures and outages.

An alternative approach to predicting the RUL could be in fitting a regression model to the RUL value at each timestep. The discrete classification technique outlined in Section 5.6.2 was selected due to the asynchronous nature of the time series data used in the study (i.e. the period between time series instances was not fixed). However, if a regular measurement could be guaranteed (from a single measurement stream) then taking a regression-based approach could potentially provide a much more fine-grained and accurate view of the health state.

It should be noted that this modelling approach does not currently consider the potential changes or uncertainties a scheduled or reactive maintenance action would introduce to the RUL calculation and prediction process. Integrating such events robustly into prognostic metrics is an open research challenge [DVLL15], and should be considered in future investigations.

Chapter 6

Conclusions & future work

As outlined in the opening of this thesis (Chapter 1 Section 1.2), there are a few central contributions to the domain from the work undertaken as a part of this research. This chapter will discuss and summarise these in more retrospective depth than previous sections, along with the implications each may have for future developments and areas of research. The contributions can be listed as:

- Augmentation of an existing knowledge-based intelligent system with ML and statistical inference techniques, providing an improved hybrid intelligent system tackling the engineering problem of routine alarm analysis in TGs,
- A self-tuning framework for vibration diagnostics, allowing for the application a routine alarm knowledge base across an entire asset family under a single maintenance regime,
- The use of techniques in statistical inference to automatically define periods of system normality and transient behaviour in rotomachinery vibration data,
- Construction of a data-driven classification model with the ability to accurately label historical periods of refuelling events from the vibration response data,

- Development of an empirical model mapping event state to remaining useful life, providing predictive metrics to anticipate future failures in GC units,
- Presentation of a new GC phase space view for refuelling campaigns, which visually provides feedback on typical vibration characteristics across the machine,

6.1 Conclusions

The research described falls into two related but distinct application areas of rotating machinery in NPPs: namely turbine generators and gas circulators. The methods employed for each of these asset classes are also different, as they attempt to solve differing challenges and problems associated with the health monitoring of each. Chapter 4 approaches the augmentation of existing knowledge-based approaches through statistical methods for TGs, while Chapter 5 introduces an empirical methodology to the exploration of predictive metrics for repeatable system events for overall GC health.

However, the unifying theme across both sub-sections is the use of intelligent techniques and data to provide value for the reliability engineer in their efforts to ensure key asset uptime is maximised. Both routine alarms and unplanned catastrophic failures (the two problem domains examined for TGs and GCs respectively) have a negative effect on the generation capacity of their associated NPPs, and any methods that provide insights or automatic approaches to tackling such problems will aid in mitigating the risks for both.

6.1.1 Turbine generator monitoring

The techniques presented throughout Chapter 4 provide data-driven approaches to help with a number of key problems associated with routine alarm analysis.

specifically on steam turbine generators in the nuclear-context.

- Ensuring the applicability and accuracy of the parametric rule-based approaches at the centre of the routine alarm expert system described in Chapter 4,
- Providing a re-usable empirical toolkit which can be applied to bulk historical turbine data in order to automatically extract the defining features of normality for a given machine instance, without the requirement for onerous manual labelling of data or behaviours,
- Extending the domain and applicability of the intelligent approach to previously uninvestigated machine channels and instances, without the same requirement for expert knowledge elicitation as with the original KBS development,
- Development of statistical techniques which automatically define periods and features of system normality and transient behaviour in rotomachinery vibration data.

Each of these points are defined by their use of data to augment an existing knowledge-based approach to increase the scope and flexibility of scenarios where crisp features such as rules, conditionals or inference trees can be used. The motivation behind a hybrid strategy like this is to take advantage of the ever increasing asset of historical machine data (as explored in Section 2.5.2 a)) while maintain the strong explicability and confidence associated with knowledge-based systems.

There were also a number of statistical methods unique in application to the domain:

- Use of the Kolmogorov-Smirnov (KS) statistic to identify periods of unimodality for steady state behaviour parameter selection on vibration data,

- Use of a windowed KS-statistic in order to extract periods of bi- and multimodality for changepoint and step change detection purposes on vibration data.

The strength of utilising density-based approaches (such as the KS-enabled methods enumerated above) lies in the explicability and communicability to engineering staff inherent in the extracted density functions. Diagrams such as those provided in Fig. 4.21 (created in the process of using the aforementioned statistical methods) can provide an intuitive view of how a distribution behaves and the characteristics of data used in initialising or updating decision support parameters (as with the expert system in Chapter 4).

The case studies outlined in Sections 4.4 and 4.5 cover important quantitative tasks associated with the data-driven augmentation of the KBS for routine alarm analysis, but also for general benchmarking and statistical reasoning about machine condition. The role that both steady state and transient conditions play in the day-to-day surveillance and asset management of rotating plant items is an important one: investigating both stable and changing machine conditions increases the range, variety and coverage of scenarios under scrutiny. The concept of event monitoring (repeatable changing conditions) is a theme explored in further depth in the subsequent chapter.

6.1.2 Gas circulator monitoring

Chapter 5 examines a specific AGR event - the low power refuelling (LPR) - and introduces an empirical modelling method utilising machine learning techniques to explore potential prognostic measures inherent in the observed gas circulator response to it. The philosophy behind this approach is to treat the LPR as a semi-regular stressor event on the circulator's long-term duty cycle, and use the vibration-based reaction to the change in conditions as a source of information on GC health itself.

As a result of the model development, there are a number of research outcomes associated with the construction and evaluation of this technique:

- Automatic classification of LPR state based on the vibration response of the circulator, allowing for bulk labelling of vibration data typical of re-fuelling periods,
- Presentation of a phase space view of machine state (exemplified by Figs. 5.16, 5.17 and 5.18) for GCs based on LPR state classification, visually providing feedback on typical vibration characteristics at each end of the circulator unit
- Development of a state membership probability-to-remaining useful life mapping, a metric that showed positive correlation with the true RUL of the example circulator and has the potential to be used as an early indicator of impending machine failure.

The development of the ML methods outlined throughout Chapter 5 follows a well-established pattern of train-test-cross validate seen in many ML-type problem solving scenarios. The source of novelty for this section lies in the combination of the application of the ML models to a specific sub-problem in the GC health monitoring sphere, and the delimitation of the LPR as a repeated event of interest for ongoing reliability engineering for nuclear-context rotating plant.

The combination of ML techniques used as part of the final model (*L2*-SVM with a logistic regression temporal step) is not unique in itself [CWY10] - the novelty in the outlined approach lies with specifics of the application area and the selected features and events the model does inference on. Specifically, considering non-stationery periods of dynamic reactor operation at key points of the plant item life cycle, and extracting quantitative metrics to map along potential degradation behaviours is novel in application to the AGR GC. The

LPR itself is a defining feature of the GC unit, and the ML model success is predicated on the LPR.

It should also be reiterated that the model was built on the operational history of a single circulator unit undergoing degradation, mainly for the purposes of metric exploration for prospective PHM and prognostic measures. While robust testing and cross-validation was undertaken to mitigate against overfitting, a wider set of circulator examples should be used to ensure the generalisability of the model features to GCs in refuelling campaigns as a whole.

6.2 Future work

The fields of intelligent systems and ML have continued to proceed rapidly in terms of technology and notable achievements, largely due to the ever-growing rates of data collection and reduction in costs associated with data warehousing and archiving. Accordingly, there is a variety of potential next steps the research outlined herein could be taken for improvement, or in exploration of other related areas of rotating plant health monitoring. This section provides some thoughts on where future efforts might best be spent in order to continue to take advantage of new approaches and improve upon the methods employed so far.

6.2.1 Routine alarm analysis

One of the strengths associated with the self-tuning methodology employed is the perceived explicability improvement when compared with more black-box techniques, which might require more ML or statistical understanding to fully grasp the outcomes from. The use of intuitive parameters such as step change magnitudes and visual distributions corresponding to steady state bounds of behaviour is argued to be more suitable for use by reliability professionals.

However, this needs to be tested and explored more fully in the development of an industrial deployment of such a system: as there has not been any feedback on the output of these methods in a formal or informal way from engineering staff.

The parameters associated with both the steady state/impulse and step change instances were explored, but the trend primitive of the KBS remains to be studied in greater detail. This is due to fewer examples of trend-like periods existing within the TG data used for exploration. Trending behaviours (as defined by the KBS rulebase) fall between steady state and step change features, and the definition of their feature extraction would be useful in extracting long-term changes that remain undetected due to their gradual emergence. A study of multiple TGs, or other rotating asset classes, with a focus on trend extraction would be a good next step to develop these methods.

Defining normality for the purposes of machine profile initialisation to utilise the routine alarm KBS, or simply benchmark typical behaviours of each component channel, was the primary aim of the techniques outlined in Chapter 4. One potential re-application of these approaches is to consider the self-tuning results at numerous points along a given TGs life cycle in order to track any changes consistent with machine degradation. Such a system could automatically re-tune after a given operation or time window, and provide a rolling update of the channel profile that could in turn be surveyed and alarmed on in the event of large enough a change in the extracted parameterisation and distributions. In a similar vein to the ML model explored in the following chapter, metrics from such a system could prospectively map onto the RUL of a TG and provide predictive decision support concerning the long-term strategy for operating a turbine unit. While this has not yet been observed, it would be a worthwhile research area for investigation get the similarities in the way in which TGs are monitored to GCs, and their deployment in the dynamic sce-

nario of the AGR.

6.2.2 Gas circulator monitoring

The inclusion of multiple GC instances in the training, testing and validation data would be the next step in investigating and developing the potential for LPR-based features from the model providing prognostic metrics. A gold standard process would be in keeping entire histories of GC operational data for examples of GC with and without catastrophic failures during their lifetimes held-out from any trained candidate model, and validating the efficacy of the late state membership feature-RUL mapping on previously unseen examples. This could be challenging, however, due to the relative rarity of full historical data to failure for circulator units.

Specific model improvements could also include deeper investigation of the temporal state fidelity, which was heuristically set to 5 time slices along the degradation history of the GC unit. The number of time slices itself could be increased to provide a more strongly-defined lifetime state to predict on, or alternative ML methods could be employed to deal with the sequential degradation aspect. Tools such as relevance vector machines (RVMs) [ZD12] have seen some success with these types of engineering problem.

6.2.3 Machine learning and reliability

Much of the techniques applied to the TGs and GCs have strong crossover potential for both asset classes (i.e. the techniques demonstrated on TGs could be further developed for use on GC monitoring, and vice versa). The next steps for building out the KBS-centred approach to cover circulator routine alarm analysis would be to apply the self-tuning methods demonstrated on the TG examples of Chapter 4 to instances of GC. The aim behind this suite

of methods was precisely for this purpose: allowing for the rule-base to be generalised beyond the original asset class it was developed on.

Introduced as an alternative technique in the context of unsupervised learning in Section 3.2.2 b), deep learning technologies have exploded in applicability to a variety of domain areas and kickstarted their own sub-field [GBC16] of ML research. 'Deep' techniques have already begun to see investigation in the vibration monitoring field [CLS15, OJJY18], and the growing data volumes for nuclear-context assets as discussed in this work lends itself to application of these techniques. Caution should be exercised, however, with strongly empirically weighted or 'black-box' style methods when working in safety critical domains such as nuclear, and densely built, deep neural networks are often difficult to provide validation or communicable decision support [Cas16] back to end-users with. The role of these methods in combination with knowledge-based approaches (as explored throughout Chapter 4) represents a rich area for further exploration and study: where the depth of the most advanced empirical methods such as NNs are reinforced with the clear explicability and human-readable reasoning associated with KBS-like solutions. There is an active research interest in explaining neural net outputs [FH17], and investigating such advanced in the context of rotating machinery could be highly valuable.

Bibliography

- [AAK17] H. S. Arel, E. Aydin, and S. D. Kore. Ageing management and life extension of concrete in nuclear power plants. *Powder Technology*, 321:390–408, 2017.
- [ABO10] E. Aronova, K. S. Baker, and N. Oreskes. Big science and big data in biology. *Historical Studies in Natural Sciences*, 40(2):183–224, 2010.
- [ACN15] B. W. Ang, W. L Choong, and T. S. Ng. Energy security: Definitions, dimensions and indexes. *Renewable and Sustainable Energy Reviews*, 42:1077–1093, 2015.
- [ACS⁺08] A. Arranz, A. Cruz, M. Sanzbobi, P. Ruiz, and J. Coutino. DADICC: Intelligent system for anomaly detection in a combined-cycle gas turbine plant. *Expert Systems with Applications*, 34(4):2267–2277, 2008.
- [Aea16] M. Abadi et al. Tensorflow: A system for large-scale machine learning. *Operating Systems Design and Implementation*, 16:265–283, 2016.
- [AHMR18] H. Akhavan-Hejazi and H. Mohsenian-Rad. Power systems big data analytics: An assessment of paradigm shift barriers and prospects. *Energy Reports*, 4:91–100, 2018.

- [Alt09] F. Altiparmak. A general neural network model for estimating telecommunications network reliability. *IEEE Transactions on Reliability*, 58(1):2–9, 2009.
- [Ama14] Y. Amano. Nuclear Technology Review. In *The 58th Regular Session of the IAEA General Conference*, 2014.
- [BAG⁺14] A. Ballesteros, E. Altstadt, F. Gillemot, H. Hein, J. Wagemans, J. Rouden, J. Barthelmes, K. Wilford, M. Serrano, M. Bru-movsky, and R. Chaouadi. Monitoring radiation embrittlement during life extension periods. *Nuclear Engineering and Design*, 267:197–206, 2014.
- [Bar13] R. J Barlow. A Guide to the use of Statistical Methods in the Physical Sciences. *John Wiley and Sons*, 2013.
- [BBC54] BBC. New authority for atomic energy. *On This Day: BBC News in Retrospect (news article)*, 1954.
- [BCFR12] S. Butler, F. O. Connor, D. Farren, and J. V. Ringwood. A feasibility study into prognostics for the main bearing of a wind turbine. In *IEEE International Conference on Control Applications*, pages 1092–1097, Dubrovnik, Croatia, 2012.
- [BD10] E. Byon and Y. Ding. Season-dependent condition-based maintenance for a wind turbine using a partially observed Markov decision process. *IEEE Transactions on Power Systems*, 25(4):1824–1834, 2010.
- [Ben96] H. Benson. *University Physics*. 1996.
- [BH02] D. E. Bently and C. T. Hatch. *Fundamentals of Rotating Machinery Diagnostics*. 2002.

- [BIAL⁺17] I. Bravo-Imaz, H.D. Ardakani, Z. Liu, A. García-Arribas, A. Arnaiz, and J. Lee. Motor current signature analysis for gearbox condition monitoring under transient speeds using wavelet analysis and dual-level time synchronous averaging. *Mechanical Systems and Signal Processing*, 94:73–84, 2017.
- [Bis06] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
- [BMZ12] P. Baraldi, F. Mangili, and E. Zio. Ensemble of bootstrapped models for the prediction of the remaining useful life of a creeping turbine blade. *2012 IEEE Conference on Prognostics and Health Management*, pages 1–8, 2012.
- [BMZ15] P. Baraldi, F. Mangili, and E. Zio. A prognostics approach to nuclear component degradation modeling based on Gaussian Process regression. *Progress in Nuclear Energy*, 78:141–154, 2015.
- [Bon06] P. Bonissone. Prognostics and health management at GE. *General Electric Global Research - Presentation*, 2006.
- [Bra10] M. J. Bradshaw. Global energy dilemmas: A geographical perspective. *The Geographical Journal*, 176(4), 2010.
- [Bre68] L. Breiman. *Probability*. Society for Industrial and Applied Mathematics, 1968.
- [Bre01] L. Breiman. Random Forests: A Review. *Machine Learning*, pages 5–32, 2001.
- [BRG02] C. S. Byington, M. J. Roemer, and T. Galie. Prognostic enhancements to diagnostic systems for improved condition-based

- maintenance. In *IEEE Aerospace Conference*, pages 2815–2824, Big Sky, MT, USA, 2002.
- [Bri05] British Energy. 10-year life extension at Dungeness B nuclear power station. *British Energy News*, 2005.
- [BRTL11] L. J. Bond, P. Ramuhalli, M. S. Tawfik, and N. J. Lybeck. Prognostics and life beyond 60 years for nuclear power plants. In *IEEE Conference on Prognostics and Health Management*, pages 1–7, 2011.
- [Bru11] G. Brumfiel. High-energy physics: Down the petabyte highway. *Nature*, 469:282–283, 2011.
- [BWMR17] C. Berry, G. West, S. D. J. McArthur, and A. Rudge. Semi-supervised learning approach for crack detection and identification in advanced gas-cooled reactor graphite bricks. In *10th International Topical Meeting on Nuclear Plant Instrumentation, Control and Human Machine Interface Technologies*, 2017.
- [BZS76] E. N. Bamberger, E. V. Zaretsky, and H. Signer. Endurance and failure characteristic of main-shaft jet engine bearing at 3×10^6 DN. *Journal of Tribology*, 98(4), 1976.
- [CA17] D. J. Cook and S Aminikhangahi. A survey of methods for time series changepoint detection. *Knowledge and Information Systems*, 51(2):339–367, 2017.
- [Cam02] Colin Campbell. Kernel methods: a survey of current techniques. *Neurocomputing*, 48(1-4):63–84, oct 2002.
- [Cas16] D. Castelvechi. Can we open the black box of AI? *Nature News*, 538(7623):20, 2016.

- [CBK09] V. Chandola, A. Banerjee, and V. Kumar. Anomaly detection: A survey. *ACM Computing Surveys*, 41(3):15, 2009.
- [CBT07] D. A. Clifton, P. Bannister, and L. Tarassenko. A framework for novelty detection in jet engine vibration data. *Key Engineering Materials*, 347:305–310, 2007.
- [CC04] G. M. Clarke and D. Cooke. *A Basic Course in Statistics*. Oxford University Press, fifth edit edition, 2004.
- [CF04] E Peter Carden and Paul Fanning. Vibration-based condition monitoring: A review. *Structural Health Monitoring*, 3(4):355–377, 2004.
- [CI17] K. J. Chalvatzis and A. Ioannidis. Energy supply security in the EU: Benchmarking diversity and dependence of primary energy. *Applied Energy*, 207:465–476, 2017.
- [CKY08] R. Caruana, N. Karampatziakis, and A. Yessenalina. An empirical evaluation of supervised learning in high dimensions. *Proceedings of the 25th International Conference on Machine Learning*, pages 96–103, 2008.
- [CLEP12] B. Chun, P. Lau, W. M M. Eden, and M. Pecht. Review of offshore wind turbine failures and fault prognostic methods. In *Prognostics & System Health Management Conference*, Beijing, China, 2012.
- [CLS15] Z. Chen, C. Li, and R. V. Sanchez. Gearbox fault identification and classification with convolutional neural networks. *Shock and Vibration*, 2015.

- [CNGT97] S. Caldara, S. Nuccio, G. R. Galluzzo, and M. Trapanese. A fuzzy diagnostic system: Application to linear induction motor drives. In *IEEE Instrumentation and Measurement Technology Conference*, pages 257–262, 1997.
- [Cob10] J. Coble. *Merging data sources to predict remaining useful life - An automated method to identify prognostic parameters*. PhD thesis, University of Tennessee, Knoxville, 2010.
- [COG09] COGENT. Power people: The civil nuclear workforce 2009-2025. *Renaissance Nuclear Skills Series*, 2009.
- [Cop84] A. Coppola. Reliability engineering of electronic equipment: A historical perspective. *IEEE Transactions on Reliability*, 33:29–35, 1984.
- [CPL96] J. Chen, R. J. Patton, and G. P. Liu. Optimal residual design for fault diagnosis using multi-objective optimization and genetic algorithms. *International Journal of Systems Science*, 27(6):567–576, 1996.
- [CSS⁺15] D Cabrera, F. Sancho, R. Sanchez, G. Zurita, M. Cerrada, L. Chuan, and R. Vasquez. Fault diagnosis of spur gearbox based on random forest and wavelet packet decomposition. *Frontiers of Mechanical Engineering*, pages 277–286, 2015.
- [CT07] G. C. Cawley and N. L. C. Talbot. Preventing over-fitting during model selection via Bayesian regularisation of the hyperparameters. *Journal of Machine Learning Research*, 8:841–861, 2007.
- [CTM⁺08] D. A. Clifton, L. Tarassenko, N. McGrogan, D. King, S. King, and P. Anuzis. Bayesian extreme value statistics for novelty

- detection in gas-turbine engines. In *IEEE Aerospace Conference*, 2008.
- [CV95] C. Cortes and V. Vapnik. Support Vector Networks. *Machine Learning*, 297:273–297, 1995.
- [CVS⁺11] A. Corner, D. Venables, A. Spence, W. Poortinga, C. Demski, and N. Pidgeon. Nuclear power, climate change and energy security: Exploring British public attitudes. *Energy Policy*, 39(9):4823–4833, sep 2011.
- [CW12] D. Conway and J. M. White. *Machine Learning for Hackers*. O’Reilly, Sebastopol, CA, USA, 2012.
- [CW14] E. Cambira and B. White. Jumping NLP curves: A review of natural language processing research. *IEEE Computational Intelligence Magazine*, pages 48–57, 2014.
- [CWCCCJ03] H. Chih-Wei, C. Chih-Chung, and L. Chih-Jen. *A Practical Guide to Support Vector Classification*, 2003.
- [CWT⁺11] W. Caesarendra, A. Widodo, P. H. Thom, B.-S. Yang, and J. D. Setiawan. Combined probability approach and indirect data-driven method for bearing degradation prognostics. *IEEE Transactions on Reliability*, 60(1):14–20, 2011.
- [CWY10] W. Caesarendra, A. Widodo, and B.-S. Yang. Application of relevance vector machine and logistic regression for machine degradation assessment. *Mechanical Systems and Signal Processing*, 24(4):1161–1171, 2010.

- [CX10] Z. Chen and L Xiangjiao. Fault diagnosis for valves of compressors based on support vector machine. In *Control and Decision Conference (CCDC)*, pages 1235–1238, 2010.
- [CZC⁺16] J. Chen, L. Ziping, G. Chen, Zi. Y., J. Yuan, B. Chen, and Z. He. Wavelet transform based on inner product in fault diagnosis of rotating machinery: A review. *Mechanical Systems and Signal Processing*, 70–71:1–35, 2016.
- [dB08] Y. de Boer. On Accounting of Emissions and Assigned Amount. *Kyoto Protocol: Reference Manual*, 2008.
- [DMM03] E. M. Davidson, S. D. J. McArthur, and J. R. McDonald. A toolset for applying model-based reasoning techniques to diagnostics for systems protection. *IEEE Power Engineering Society General Meeting*, page 642, 2003.
- [DMM⁺06] E. M. Davidson, S. D. J. McArthur, J. R. McDonald, R. James, T. Cumming, and I. Watt. Applying multi-agent system technology in practice: automated management and analysis of SCADA and digital fault recorder data. *IEEE Transactions on Power Systems*, 21(2):559–567, 2006.
- [dS12] C. W. de Silva. *Vibration Monitoring, Testing and Instrumentation*. CRC Press, 2012.
- [DVLL15] P. Do, A. Voisin, E. Levrat, and B. Lung. A proactive condition-based maintenance strategy with both perfect and imperfect maintenance actions. *Reliability Engineering and System Safety*, 133:22–32, 2015.
- [EB04] M. J. Embrechts and S. Benedek. Hybrid Identification of Nuclear Power Plant Transients With Artificial Neural Networks.

IEEE Transactions on Industrial Electronics, 51(3):686–693, jun 2004.

- [EMAQAG16] R. El-Mahayni, K. Al-Qahtani, and A. H. Al-Gheeth. Large synchronous motor failure investigation: Measurements, analysis and lessons learned. *IEEE Transactions on Industry Applications*, 52(6):5318–5326, 2016.
- [EP08] S. Ebersbach and Z. Peng. Expert system development for vibration analysis in machine condition monitoring. *Expert Systems with Applications*, 34(1):291–299, 2008.
- [Fer46] E. Fermi. The development of the first chain reaction pile. *Proceedings of the American Philosophical Society*, 90:20–24, 1946.
- [FH17] N. Frosst and G. Hinton. Distilling a neural network into a decision tree. *arXiv pre-prints*, 1711.09784, 2017.
- [FLL⁺16] J. Feng, Y. Lei, J. Lin, Z. Zin, and N. Lu. Deep neural networks: A promising tool for fault characteristic mining and intelligent diagnosis of rotating machine with massive data. *Mechanical Systems and Signal Processing*, 72–73:303–315, 2016.
- [FMJ⁺11] W. C. Flores, E. E. Mombello, J. A. Jardini, G. Rattá, and A. M. Corvo. Expert system for the assessment of power transformer insulation condition based on type-2 fuzzy logic systems. *Expert Systems with Applications*, 38(7):8119–8127, 2011.
- [Fos67] G. B. Foster. Recent developments in machine vibration monitoring. *IEEE Transactions on Industry and General Applications*, 3(2):149–159, 1967.

- [Fri01] J. H. Friedman. Greedy function approximation: a gradient boosting machine. *Annals of Statistics*, pages 1189–1232, 2001.
- [Gao14] J. Gao. Machine learning applications for data center optimization. <http://research.google.com>, 2014.
- [GBC16] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016.
- [GCCGP16] M. J. Gomez, C. Castejon, and J. C. Garcia-Prada. Automatic condition monitoring system for crack detection in rotating machinery. *Reliability Engineering and System Safety*, 152:239–247, 2016.
- [Gib16] E. Gibney. Google AI algorithm masters ancient game of Go. *Nature News*, 445, 2016.
- [GMW13] J. L. Godwin, P. Matthews, and C. Watson. Classification and detection of electrical control system faults through SCADA data analysis. *Chemical Engineering Transactions*, 33:985–990, 2013.
- [GP16] D. Goyal and B. S. Pabla. The vibration monitoring methods and signal processing techniques for structural health monitoring: A review. *Archives of Computational Methods in Engineering*, 23(4):585–594, 2016.
- [GPJW12] L. Gelman, I. Petrunin, I. K. Jennions, and M. Walters. Diagnostics of local tooth damage in gears by the wavelet technology. *International Journal of Prognostics and Health Management*, 5:2153–2648, 2012.

- [Gri14] O. Grisel. RBF SVM parameters - Tutorial. *scikit-learn Technical Notes*, 2014.
- [HACF05] A. Hess, V. A Arlington, G. Calvello, and P. Frith. Challenges, issues and lessons learned chasing the 'big P'. In *IEEE Aerospace Conference*, pages 3610–3619, 2005.
- [Hat12] Y. Hatamura. Final report of the investigating committee on the accident at the Fukushima nuclear power stations of Tokyo Electric Power Company (TEPCO). *Technical Report*, 2012.
- [Haz01] M. Hazewinkel. *Encyclopedia of Mathematics*. Springer Science & Business Media, 2001.
- [Heb49] D. Hebb. *The Organization of Behavior*. John Wiley & Sons, New York, USA, 1949.
- [HG08] HM-Government. Meeting the energy challenge: A white paper on nuclear power. *Department for Business Enterprise & Regulatory Reform (January)*, 2008.
- [HG13] HM-Government. The UK's Nuclear Future - Nuclear Industrial Strategy. *Department of Energy & Climate Change and Department for Business, Innovation & Skills*, 2013.
- [HH85] J. A. Hartigan and P. M. Hartigan. The Dip Test of Unimodality. *The Annals of Statistics*, 13(1):70–84, 1985.
- [HH13] M. Hayashi and L Hughes. The Fukushima nuclear accident and its effect on global energy security. *Energy Policy*, 59:102–111, 2013.
- [HKO04] H. Hyvarinen, J. Karhunen, and E. Oja. Independent component analysis. *John Wiley and Sons*, 2004.

- [Ho95] T. K. Ho. Random decision forests. *Proceedings of the 3rd International Conference on Document Analysis and Recognition*, 1:278–282, 1995.
- [Hop11] A. A. Hopgood. *Intelligent systems for engineers and scientists*. CRC Press, 2011.
- [HSTA00] P. Hayton, B. Scholkopf, L. Tarassenko, and P. Anuzis. Support vector novelty detection applied to jet engine vibration spectra. In *Proceedings of Neural Information Processing Systems*, 2000.
- [HW58] J. M. Harrer and E. A. Wimunc. EBWR turbine blade failure: Summary report. Technical report, Argonne National Laboratory, Lemont, Ill., 1958.
- [HWB06] D. He, S. Wu, and E. Bechhofer. Probabilistic model-based algorithms for prognostics. In *IEEE Aerospace Conference*, Big Sky, MT, USA, 2006.
- [HYHY10] Y. Huan, L. Yibing, A. Hongwen, and Z. Yanbing. Turbine vibration source separation based on independent component analysis. In *International Conference on Mechatronics and Automation*, pages 1728–1731, Xi’an, China, 2010.
- [HZTM09] A. Heng, S. Zhang, A. C. C. Tan, and J. Mathew. Rotating machinery prognostics: State of the art, challenges and opportunities. *Mechanical Systems and Signal Processing*, 23(3):724–739, 2009.
- [IAT13] A. Ikonomopoulos, M. Alamaniotis, and H. Tsouakals. Gaussian processes for state identification in pressurized water reactors. *Nuclear Technology*, 182:1–12, 2013.

- [IEA17] IEA. *World Energy Outlook 2017 - Executive Summary*. 2017.
- [IOU+12] I. Aho-Mantila, O.Cronvall, U.Ehrnstén, H.Keinänen, R.Rintamaa, A.Saarenheimo, K.Simola, and E.Vesikari. Lifetime prediction techniques for nuclear power plant systems. *Nuclear Corrosion Science and Engineering*, pages 449–470, 2012.
- [iso09] ISO 22266-1:20009: Mechanical vibration - Torsional vibration of rotating machinery - Part 1: Land-based steam and gas turbine generator sets in excess of 50MW. *International Organization for Standardization*, 2009.
- [JE56] K. Jay and B. Edmund. *Calder Hall: The story of Britian's first atomic power station*. 1956.
- [JLB06] A. K. S. Jardine, D. Lin, and D. Banjevic. A review on machinery diagnostics and prognostics implementing condition-based maintenance, oct 2006.
- [JLD+91] H. E Johnson, D. J. Littler, E. J. Davies, F. Kirkby, P. B Myerscough, and W. Wright. *Modern Power Station Practice Volume C: Turbines, Generators and Associated Plant*. Pergamon Press, 1991.
- [JM15] M. I. Jordan and T. M. Mitchell. Machine learning: Trends, perspectives and prospects. *Science*, 349(6245):255–260, 2015.
- [JMS96] C. M. Jones, J. S. Marron, and S. J. Sheather. A brief survey of bandwidth selection for density estimation. *Journal of the American Statistical Association*, 91:401–407, 1996.

- [JN00] L B Jack and A K Nandi. Genetic algorithms for feature selection in machine condition monitoring with vibration signals. *IEEE Proceedings - Vision, Image, and Signal Processing*, 147(3):205–212, 2000.
- [Jol11] Ian Jolliffe. *Principal Component Analysis*. Springer Berlin Heidelberg, 2011.
- [JSB04] D. B. Jarrell, D. R. Sisk, and L. J. Bond. Prognostics and condition-based maintenance: A new approach to precursive metrics. *Nuclear Technology - Nuclear Plant Oerations and Control*, (145):275–286, 2004.
- [Kal72] D. Kalderon. Steam turbine failure at Hinkley Point A. In *Proceedings of the Institution of Mechanical Engineers*, pages 341–377, 1972.
- [KCMT13] A. D. Kenyon, V. M. Catterson, S. D. J. McArthur, and J. Twiddle. An agent-based implementation of hidden Markov models for gas turbine condition monitoring. *IEEE Transactions on Systems, Man and Cybernetics*, 99, 2013.
- [KHV06] R. Kothamasu, S. H. Huang, and W. H. VerDuin. System health monitoring and prognostics a review of current paradigms and practices. *The International Journal of Advanced Manufacturing Technology*, 28(9-10):1012–1024, jul 2006.
- [KL03] S. S. Keerthi and C.-J. Lin. Asymptotic behaviors of support vector machines with Gaussian kernel. *Neural computation*, 15(7):1667–1689, 2003.
- [KL11] A. Kusiak and W. Li. The prediction and diagnosis of wind turbine faults. *Renewable Energy*, 36(1):16–23, 2011.

- [KLO⁺12] E. M. Kan, M. H. Lim, Y. S. Ong, A. H. Tan, and S. P. Yeo. Extreme learning machine terrain-based navigation for unmanned aerial vehicles. *Neural Computing and Applications*, 22(3-4):469–477, feb 2012.
- [KMF⁺16] D. Kwon, Hodkiewicz M., J. Fan, T. Shibutani, and M. G. Pecht. IoT-based prognostics and systems health management for industrial applications. *IEEE Access*, 4:3659–3670, 2016.
- [Koh95] R. Kohavi. A Study of Cross-Validation and Bootstrap for Accuracy Estimation and Model Selection. *International Joint Conference on Artificial Intelligence*, 14(12):1137–1143, 1995.
- [Kot07] S. B. Kotsiantis. Supervised machine learning: A review of classification techniques. *Informatica*, 31(249-268), 2007.
- [Kri03] K. Krishnakumar. Intelligent systems for aerospace engineering - An overview. *National Aeronautics and Space Administration Ames Research Centre - Report*, pages 1–15, 2003.
- [KS03] Y. Koshiba and A. Shigeo. Comparison of L1 and L2 support vector machines. *Proceedings of the International Joint Conference on Neural Networks*, 2003.
- [KSB⁺18] A. Kusupati, M. Sing, K. Bhatia, A. Kumar, P. Jain, and M. Varma. FastGRNN: A fast, accurate, stable and tiny kilobyte-sized gated recurrent neural network. *32nd Conference on Neural Information Processing Systems (NeurIPS), Montreal*, 2018.
- [KTM⁺08] H.-E. Kim, A. C. C. Tan, J Mathew, E Y H Kim, and B.-K. Choi. Machine prognostics based on health state estimation using

- SVM. In *Proceedings of Third World Congress on Engineering Asset Management and Intelligent Maintenance Systems Conference*, pages 834–845, Beijing, China, 2008.
- [KWM92] N. Karunanithi, D. Whitley, and Y. K. Malaiya. Using neural networks in reliability prediction. *IEEE Software*, pages 53–59, 1992.
- [Lal14] C. Lalenne. *Random Vibration: Mechanical Vibration and Shock Analysis*. John Wiley & Sons, 2014.
- [Leg05] A.-M. Legendre. *Nouvelles Méthodes pour la Détermination des Orbites des Comètes*. Didot Publications, 1805.
- [Lia05] Shu-hsien Liao. Expert system methodologies and applications - A decade review from 1995 to 2004. *Expert Systems with Applications*, 28(1):93–103, jan 2005.
- [LK12] J. Lin and A. Kolcz. Large scale machine learning at Twitter. In *ACM SIGMOD International Conference on Management of Data*, page 793804, 2012.
- [LLBK13] J. Lee, E. Lapira, B. Bagheri, and H. Kao. Recent advances and trends in predictive manufacturing systems in big data environment. *Manufacturing Letters*, 1(1):38–41, oct 2013.
- [LMR10] A. Löschel, U. Moslener, and D. T. G. Rübhelke. Indicators of energy security in industrialised countries. *Energy Policy*, 38(4):1665–1671, apr 2010.
- [Loh12] S. Lohr. *The Age of Big Data*, 2012.

- [LSG⁺16] C. Li, R-V. Snchez, Zurita G., Cerrada M., and Cabrera D. Fault diagnosis for rotating machinery using vibration measurement deep statistical feature learning. *Sensors*, 16:1–35, 2016.
- [LSoO11] N. Laouti, N. Sheibat-othman, and S. Othman. Support vector machines for fault detection in wind turbines. In *18th IFAC World Congress*, pages 7067–7072, Milano, Italy, 2011.
- [LWDG17] T. Lardner, G. M. West, G. Dobie, and A. Gachagan. Automated sizing and classification of defects in CANDU pressure tubes. *Nuclear Engineering and Design*, 325:25–32, 2017.
- [LWZ⁺14] J. Lee, F. Wu, W. Zhao, M. Ghaffari, and L. Liao. Prognostics and health management design for rotary machinery systems Reviews , methodology and applications. 42:314–334, 2014.
- [Lyo13] A. Lyon. Why are normal distributions normal? *The British Journal for the Philosophy of Science*, 65(3):621–649, 2013.
- [LYZC18] R. Liu, B. Yang, E. Zio, and X. Chen. Artificial intelligence for fault diagnosis of rotating machinery: A review. *Mechanical Systems and Signal Processing*, 108, 2018.
- [LZLL17] G. Lu, Y Zhou, C. Lu, and X. Li. A novel framework of change-point detection for machine monitoring. *Mechanical Systems and Signal Processing*, 83:533–548, 2017.
- [MAK14] N. Mahantesh, P. Aditya, and U. Kumar. Integrated machine health monitoring: a knowledge-based approach. *International Journal of System Assurance Engineering and Management*, 5(3):371–382, 2014.

- [Mas51] F. J. Massey. The Kolmogorov-Smirnov test for goodness of fit. *American Statistical Association*, 46(253):68–78, 1951.
- [MGiAR08] Z. Mazur, R. Garcia-illescas, and J. Aguirre-Romano. Steam turbine blade failure analysis. *Engineering Failure Analysis*, (15):129–141, 2008.
- [MH08] L. v. d. Maaten and G. Hinton. Visualising data using t-SNE. *Journal of Machine Learning Research*, 2008.
- [MH14] W. Q. Meeker and Y. Hong. Reliability meets Big Data: Opportunities and challenges. *Quality Engineering*, 26(1):102–116, 2014.
- [MHA18] I. MacLeay, K. Harris, and A. Annut. Digest of United Kingdom energy statistics (DUKES) 2018. *Department of Energy & Climate Change - Digest*, 2018.
- [Mic01] R. Michal. Fifty years ago in December: Atomic reactor EBR-1 produced first electricity. *ANS Nuclear News - Operations*, 2001.
- [Min88] M. Minsky. *The Society of Mind*. Simon & Schuster, New York, USA, 1988.
- [Mit81] J. S. Mitchell. *Machinery analysis and monitoring*. Penn. Well, 1981.
- [Mit97] T. M. Mitchell. *Machine Learning*. McGraw-Hill International Editions, 1997.
- [MLY17] S. Min, B. Lee, and S. Yoon. Deep learning in bio-informatics. *Briefings in bio-informatics*, 18(5):851–869, 2017.
- [MM07] Q. Miao and V. Makis. Condition monitoring and classification of rotating machinery using wavelets and hidden Markov

- models. *Mechanical Systems and Signal Processing*, 21(2):840–855, 2007.
- [Moh11] L. Mohrbach. The defence-in-depth safety concept: Comparison between the Fukushima Daiichi units and German nuclear power units. *International Journal for Nuclear Power*, 56(4-5), 2011.
- [Moh18] A. R. Mohanty. *Machinery condition monitoring: Principles and practices*. CRC Press, 2018.
- [Moo96] T. K Moon. The expectation-maximisation algorithm. *IEEE Signal Processing Magazine*, 13(6):47–60, 1996.
- [MRS10] P. Maul, P. Robinson, and A. Steer. Understanding AGR graphite brick cracking using physical understanding and statistical modelling. In *Securing the Safe Performance of Graphite Reactor Cores - Royal Society of Chemistry*, pages 103–110, 2010.
- [MSR⁺97] K. R. Muller, A. Smola, G. Ratsch, B. Scholkopf, J. Kohlmorgen, and V. Vapnik. Predicting time series with support vector machines. In *International Conference on Artificial Neural Networks*, page 999, 1997.
- [Mur12] K. P. Murphy. *Machine Learning: A Probabilistic Perspective*. MIT Press, Cambridge, Mass., 2012.
- [MWMM16] P. Murray, G. West, S. Marshall, and S McArthur. Automated in-core image generation from video to aid visual inspection of nuclear power plant cores. *Nuclear Engineering and Design*, 300:57–66, 2016.

- [NAS] NASA. Prognostics Centre of Excellence (PCoE). *Ames Research Centre, CA*.
- [Nat13] S. Nathan. March 1963: The Windscale AGR. *The Engineer*, 2013.
- [NBTT12] M. N. Nguyen, C. Bao, K. L. Tew, and S. D. Teddy. Ensemble Based Real-Time Adaptive Classification System for Intelligent Sensing Machine Diagnostics. *IEEE Transactions on Reliability*, 61(2):303–313, 2012.
- [Nea06] M. Neale. Learning from rotating machinery failures around the world. *Institution of Mechanical Engineers - The President's Choice (August)*, 2006.
- [NEI12] NEI. EDF plans longer life extensions for UK AGRs. Technical report, 2012.
- [NEI13] NEI. Nuclear accounts for 19{ } of UK electricity generation in 2012. Technical report, 2013.
- [NK03] M. P. Norton and D. G. Kargzub. *Fundamentals of Noise and Vibration Analysis for Engineers*. Cambridge University Press, 2003.
- [NM11] S. Natti and K. Mladen. Assessing circuit breaker performance using condition-based data and a Bayesian approach. *Electric Power Systems Research*, 81(9):1796–1804, 2011.
- [Non96] E. Nonbol. Description of the Advanced Gas Cooled Type of Reactor (AGR). *Riso National Laboratory Report*, (November 1996), 1996.

- [NRC] NRC. <http://www.nrc.gov/images/reading-rm/photo-gallery/20071115-058.jpg>.
- [OFG97] E. Osuna, R. Freund, and F. Girosi. Training support vector machines: an application to face detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 130–136, 1997.
- [OJBH06] O. A. Omitaomu, M. K. Jeong, A. B. Badiru, and J. W. Hines. Online prediction of motor shaft misalignment using Fast Fourier Transform generated spectra data and support vector regression. *Journal of Manufacturing Science and Engineering*, 128, 2006.
- [OJJY18] H. Oh, J. H. Jung, B. C. Jeon, and B. D. Youn. Scalable and unsupervised feature engineering using vibration-imaging and deep learning for rotor system diagnosis. *IEEE Transactions on Industrial Electronics*, 65(4):3539–3549, 2018.
- [OMS⁺15] L. Olatomiwa, S. Mekhilef, S. Shamshirband, K. Mohammadi, D. Petkovic, and C. Sudheer. A support vector machine-firefly algorithm-based model for global solar radiation prediction. *Solar Energy*, 115:632–544, 2015.
- [OP11] J. O. Ogutu and T. Piepho, H.-P. abnd Schulz-Streeck. A comparison of random forests, boosting and support vector machines for genomic selection. *BMC Proceedings, BioMed Central*, 5(3):1–11, 2011.
- [PA13] F. Pedregosa and Et Al. *Scikit-learn : Machine Learning in Python*. 2013.
- [PAH⁺17] M. Pehl, A. Arvesen, F. Humpenoder, A. Popp, E. G. Hertwich, and G. Luderer. Understanding future emissions from low-

carbon power systems by integration of life-cycle assessment and integrated energy modelling. *Nature Energy*, 2:939–945, 2017.

- [Par11] C. Parsons. The Steam Turbine - Rede Lecture. *Cambridge University Press*, 1911.
- [PCLH15] J Phillips, E Cripps, John W Lau, and M R Hodkiewicz. Classifying machinery condition using oil samples and binary logistic regression. *Mechanical Systems and Signal Processing*, 60-61:316–325, 2015.
- [PFKC08] S. N. Pakzad, G. L. Fenves, S. Kim, and D. E. Culler. Design and implementation of scalable wireless sensor networks for structural monitoring. *Infrastructure Systems*, pages 89–101, 2008.
- [PK15] S. Pfenniinger and J. Keristead. Renewables, nuclear, or fossil fuels? Scenarios for Great Britains power system considering costs, emissions and energy security. *Applied Energy*, 152:83–93, 2015.
- [PNA⁺02] S Poyhonen, M Negrea, a. Arkkio, H Hyotyniemi, and H Koivo. Support vector classification for fault diagnostics of an electrical machine. *6th International Conference on Signal Processing, 2002.*, 2:373–378, 2002.
- [Pom89] D. A. Pomerleau. ALVINN: An autonomous land vehicle in a neural network. *Carnegie Mellon Research Showcase*, 1989.
- [POS11] POST. Carbon footprint of electricity generation. *Parliamentary Office of Science and Technology, UK Houses of Parliament*, (383), 2011.

- [PP08] T. Paek and R. Pieraccini. Automating spoken dialogue management design using machine learning: An industry perspective. *Speech Recognition*, 59(8-9):716–729, 2008.
- [PRB03] O. Pietsch, W. Reimche, and F.-W. Bach. Transient vibration signatures at steam and gas turbines. In *American Conference for Nondestructive Testing*, 2003.
- [PSJ⁺13] R. Pfister, K. A. Schwarz, M. Janczyk, R. Dale, and J. Freeman. Good things peak in pairs: a note on the bimodality coefficient. *Frontiers in Psychology*, 4(2):700, 2013.
- [PVG11] F. Pedregosa, G. Varoquaux, A. Gramfort, and V. Michel. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [Rag18] E. Raguseo. Big data technologies: An empirical investigation of their adoption, benefits and risks for companies. *International Journal of Information Management*, 38:187–195, 2018.
- [Ran04] R. B. Randall. State of the art in monitoring rotating machinery Part 1. *Sound and Vibration*, (March):14–20, 2004.
- [RCMJ11] S. E. Rudd, V. M. Catterson, S. D. J. McArthur, and C. Johnstone. Circuit breaker prognostics using SF6 data. *Power and Energy General Meeting*, 2011.
- [RMS15] A. Radford, L. Metz, and Chintala S. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprints*, 2015.

- [RN88] J. L. Rodgers and W. A. Nicewander. Thirteen ways to look at the correlation coefficient. *The American Statistician*, 42(1):59–66, 1988.
- [Ros58] F. Rosenblatt. The Perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6), 1958.
- [RST17] I. Rossi, A. Sorce, and A. Traverso. Gas turbine combined cycle start-up and stress evaluation: A simplified dynamic approach. *Applied Energy*, 190:880–890, 2017.
- [RSWZ18] L. Ren, Y. Sun, H. Wang, and L. Zhang. Prediction of bearing remaining useful life with deep convolutional neural network. *IEEE Access*, 6:13041–13049, 2018.
- [Rus14] R. Rusaw. Big data meets nuclear power. *Nuclear Engineering International*, 2014.
- [SAA⁺00] G. Schreiber, H. Akkermans, A. Anjewierden, R. de Hoog, N. Shadbolt, W. Van De Velde, and Bob Wielinga. *Knowledge Engineering and Management*. 2000.
- [SAT03] S. Seker, E. Ayaz, and E. Turkan. Elman’s recurrent neural network applications to condition monitoring in nuclear power plant and rotating machinery. *Engineering Applications of Artificial Intelligence*, 16:647–656, 2003.
- [SBH⁺05] A. Shamshad, M. A. Bawadi, W. M. A. Hussin, T. A. Majid, and S. A. M. Sanusi. First and second order Markov chain models for synthetic generation of wind speed time series. *Energy*, 30(5):693–708, 2005.

- [Sch04] C. Scheffer. *Practical machinery vibration analysis and predictive maintenance*. Newnes Publications, 2004.
- [SF07] J. K. Shultis and R. E. Faw. *Fundamentals of Nuclear Science and Engineering*. CRC Press, second edition, 2007.
- [SF11] M. Schlechtingen and I. Ferreira Santos. Comparative analysis of neural network and regression based condition monitoring approaches for wind turbine fault detection. *Mechanical Systems and Signal Processing*, 25(5):1849–1875, jul 2011.
- [SGPC09] B. Saha, K. Goebel, S. Poll, and J. Christopherson. Prognostics methods for battery health monitoring using a Bayesian framework. *Instrumentation and Measurement, IEEE Transactions on*, 58(2):291–296, 2009.
- [SGS⁺18] S. Schelter, S. Grafberger, P. Schmidt, T. Rukat, M. Kiessling, A. Taptunov, F. Biessmann, and D. Lange. Deequ - Data Quality Validation for Machine Learning Pipelines. *32nd Conference on Neural Information Processing Systems (NeurIPS), Montreal*, 2018.
- [SHM11] J. Z. Sikorska, M. Hodkiewicz, and L. Ma. Prognostic modelling options for remaining useful life estimation by industry. *Mechanical Systems and Signal Processing*, pages 1803–1836, 2011.
- [Sil81] B. W. Silverman. Using kernel density estimates to investigate multimodality. *J. R. Statist. Soc*, 43(1):97–99, 1981.
- [Sil87] B. W. Silverman. Density estimation for statistical and data analysis. *Monographs on Statistics and Applied Probability*, 1951:1–22, 1987.

- [SJZW17] H. Shao, H. Jiang, H. Zhao, and F. Wang. A novel deep autoencoder feature learning method for rotating machinery fault diagnosis. *Mechanical Systems and Signal Processing*, 95:187–204, 2017.
- [SKW09] J. W. Sheppard, M. A. Kaufman, and T. J. Wilmering. IEEE standards for prognostics and health management. *IEEE A{&}E Systems Magazine*, (September):34–41, 2009.
- [SP07] V. A. Sotiris and M. Pecht. Support Vector Prognostics Analysis of Electronic Products and Systems. In *Artificial Intelligence for Prognostics*, pages 120–127, 2007.
- [SRM⁺08] S. Strachan, S. Rudd, S. D. J. McArthur, M. Judd, S. Meijer, and E. Gulski. Knowledge-based diagnosis of partial discharges in power transformers. *IEEE Transactions on Dielectrics and Electrical Insulation*, 15(1):259–268, 2008.
- [SSR08] V. Sugumaran, G. R. Sabareesh, and K. I. Ramachandran. Fault diagnostics of roller bearing using kernel based neighborhood score multi-class support vector machine. *Expert Systems with Applications*, 34:3090–3098, 2008.
- [ST17] Siri-Team. Hey Siri: An on-device DNN-powered voice trigger for Apple’s personal assistant. *Apple Machine Learning Journal*, 1(6), 2017.
- [SWdH⁺83] G. Schreiber, B. Wielinga, R. de Hoog, H. Akkermans, and W. Van De Velde. CommonKADS: A Comprehensive Methodology for KBS Development. *IEEE Expert Intelligent Systems And Their Applications*, 1983.

- [SZ15] F. K. Shaikh and S. Zeadally. Extended Kalman Filtering for Remaining Useful Life Estimation of Bearings. *IEEE Transactions on Industrial Electronics*, 62:1781–1790, 2015.
- [SZ16] F. K. Shaikh and S. Zeadally. Energy harvesting in wireless sensor networks: A comprehensive review. *Renewable and Sustainable Energy Reviews*, 55:1041–1054, 2016.
- [Tap08] R. L. Tapping. Materials performance in CANDU reactors: The first 30 years and the prognosis for life extension and new designs. *Journal of Nuclear Materials*, 383(1-2):1–8, 2008.
- [Tav08] P. J. Tavner. Review of condition monitoring of rotating electrical machines. 2(4):215–247, 2008.
- [TBT15] W. W. Tiddens, A. Braaksma, and T. Tinga. The adoption of prognostic technologies in maintenance decision making: a multiple case study. *Procedia CIRP*, 38:171–176, 2015.
- [TGF18] S. Touzani, J. Granderson, and S. Fernandes. Gradient boosting machine for modelling the energy consumption of commercial buildings. *Energy and Buildings*, pages 1533–1543, 2018.
- [Tho96] W. T. Thompson. *Theory of Vibrations*. Nelson Thornes, 1996.
- [Tip00] M. E. Tipping. The relevance vector machine. *Advances in neural information processing systems*, pages 652–658, 2000.
- [TM92] L. Travé-Massuyès. Gas-Turbine Condition Monitoring Using Qualitative Model-Based Diagnosis. *Knowledge Creation Diffusion Utilization*, 1992.
- [TMMS07] M. Todd, S. D. J. McArthur, J. R McDonald, and S. Shaw. A semiautomatic approach to deriving turbine generator diag-

- nostic knowledge. *IEEE Transactions of Systems, Man and Cybernetics, Part C (Applications and Reviews)*, 37(5):979–992, 2007.
- [TMMZT12] D. A. Tobon-Mejia, K. Medjaher, N. Zerhouni, and G. Tripot. A data-driven failure prognostics method based on mixture of Gaussian hidden Markov models. *IEEE Transactions on Reliability*, 61(2):491–503, jun 2012.
- [Tod09] M. Todd. *Combining knowledge-based systems and machine learning for turbine condition monitoring*. PhD thesis, University of Strathclyde, 2009.
- [Ton02] B. H. Tongue. *Principles of Vibration*. Oxford University Press, second edition, 2002.
- [Tru79] G. V. Trunk. A problem of dimensionality: A simple example. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 3:306–307, 1979.
- [UKC08] UKCT. Carbon footprinting: The next step in reducing your emissions. *UK Carbon Trust - Press Release*, 2008.
- [VBKR03] S. G. Vinod, A. K. Babar, H. S. Kushwaha, and V. Raj. Symptom based diagnostic system for nuclear power plant operations using artificial neural networks. *Reliability Engineering & System Safety*, 82:33–40, 2003.
- [Vee80] H. J. M Veendrick. The behaviour of flip-flops as synchronizers and prediction of their failure rate. *IEEE Journal of Solid State Circuits*, 15(2):169–176, 1980.
- [VR11] P. Vaidya and M. Rausand. Remaining useful life, technical health and life extension. *Proceedings of the Institution of Mechan-*

ical Engineers, Part O: Journal of Risk and Reliability, 225(2):219–231, 2011.

- [Vuo89] Q. H. Vuong. Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica: Journal of the Econometrics Society*, 1:307–333, 1989.
- [Wes11] Westinghouse. Oxide Fuels Complex Overview. *Springfield Nuclear Sites*, 2011.
- [WGC⁺18] A. Wan, F. Gu, J. Chen, L. Zheng, P. Hall, Y. Ji, and X. Gu. Prognostics of gas turbine: A condition-based maintenance approach based on multi-environmental time similarity. *Mechanical Systems and Signal Processing*, 109:150–165, 2018.
- [WH13] P. Westfall and K. S. S. Henning. *Understanding Advanced Statistical Methods*. CRC Press, 2013.
- [Win12] C. Winzer. Conceptualizing energy security. *Energy Policy*, 46:36–48, 2012.
- [Wis00] W. H. Wiser. *Energy Resources: Occurrence, Production, Conversion, Use*. 2000.
- [WJM⁺06] G. M. West, G. J. Jahn, S. D. J. McArthur, J. R. McDonald, and J. Reed. Data mining reactor fuel grab load trace data to support nuclear core condition monitoring. *IEEE Transactions on Nuclear Science*, 53(3):1494–1503, jun 2006.
- [WMT09] G. M. West, S. D. J. McArthur, and D. Towle. *BETA: A System for Automated Intelligent Analysis of Fuel Grab Load Trace Data for Graphite Core Condition Monitoring*. 2009.

- [WMT12] G. M. West, S. D J McArthur, and D. Towle. Industrial implementation of intelligent system techniques for nuclear power plant condition monitoring. *Expert Systems with Applications*, 39(8):7432–7440, 2012.
- [WNA13] WNA. Nuclear Power in Germany - World Nuclear Institute. *WNA Country Profiles*, 2013.
- [WNA18a] WNA. Nuclear Power in the United Kingdom. *World Nuclear Association - Country Profiles, United Kingdom*, 2018.
- [WNA18b] WNA. Nuclear Power in the World Today. *World Nuclear Association - Information Library, Current and Future Generation*, 2018.
- [WNA19] WNA. World Nuclear Power Reactors & Uranium Requirements. Technical report, 2019.
- [WRM12] J. Weston, F. Ratle, and H. Mobahi. Deep learning via semi-supervised embedding. *Neural Networks: Tricks of the Trade*, pages 639–655, 2012.
- [WWJ⁺10] G M West, C J Wallace, G J Jahn, S D J McArthur, George Street, and Barnett Way. Predicting the Ageing of Advanced Gas-Cooled Reactor (Agr) Graphite Bricks. *Interface*, pages 845–855, 2010.
- [WWJM10] C. J. Wallace, G. M. West, G. J. Jahn, and S. D. J. McArthur. Control rod monitoring of Advanced Gas-cooled Reactor. In *Nuclear Plant Instrumentation, Control and Human-Machine Interface Technologies*, pages 254–263, Las Vegas, Nevada, 2010.
- [WWS13] A. J. Wood, B. F. Wollenberg, and G. B. Sheble. *Power generation, operation and control*. John Wiley & Sons, 2013.

- [WY07] A. Widodo and B.-S. Yang. Support vector machine in machine condition monitoring and fault diagnosis. *Mechanical Systems and Signal Processing*, 21:2560–2574, 2007.
- [Xin07] C. Xindi. Time series prediction with recurrent neural networks trained by a hybrid PSOEA algorithm. *Neurocomputing*, 70(13):2342–2353, 2007.
- [XSW⁺17] Y. Xu, Y. Sun, J. Wan, X. Lui, and Z. Song. Industrial Big Data for Fault Diagnosis: Taxonomy, Review and Applications. *IEEE Access*, 5:17368–17380, 2017.
- [Yan06] W. Yan. Application of random forest to aircraft engine fault diagnosis. *Computational Engineering in Systems Applications*, pages 468–475, 2006.
- [YJH16] W. Yuxiang, S. Jiqiang, and Z. He. Study on new considerations of defence-in-depth strategy for nuclear power plants. In *International Pacific Basin Nuclear Conference*, pages 647–661, 2016.
- [YLT05] B. Yang, D. Lim, and A. Tan. VIBEX: an expert system for vibration fault diagnosis of rotating machinery using decision tree and decision table. *Expert Systems with Applications*, 28(4):735–742, 2005.
- [YM13] S. K. Yadav and G. C. Mishra. Global energy demand consequences vs. greenhouse gases emission. *International Journal of Engineering Research and Technology*, 6(6):781–788, 2013.
- [YS02] W. Y. W. Yan and H. S. Shao. Application of support vector machine nonlinear classifier to fault diagnoses. *Proceedings of the 4th World Congress on Intelligent Control and Automation (Cat. No.02EX527)*, 4:2697–2700, 2002.

- [YZX⁺11] N. Yang, S. Zhang, Y. Xu, W. Zhao, and Y. Zhu. An Expert System for Vibration Fault Diagnosis of Large Steam Turbine Generator Set. In *International Conference on Computer Research and Development*, 2011.
- [Zad93] L. A. Zadeh. Fuzzy logic, neural networks and soft computing. In *Seventh Annual Workshop on Space Operation Applications and Research (SOAR)*, 1993.
- [ZCB10] Q. Zhang, L. Cheng, and R. Boutaba. Cloud computing: state-of-the-art and research challenges. *Journal of Internet Services and Applications*, 1(1):7–18, 2010.
- [ZD10] E. Zio and F. Di Maio. A data-driven fuzzy approach for prediction the remaining useful life in dynamic failure scenarios of a nuclear system. *Reliability Engineering & System Safety*, 95(1):49–57, 2010.
- [ZD12] E. Zio and F. Di Maio. Fatigue crack growth estimation by relevance vector machine. *Expert Systems with Applications*, 39(12):10681–10692, 2012.
- [ZGMS04] C. Zuluaga-Giraldo, D. Mba, and M. Smart. Acoustic emission during run-up and run-down of a power generation turbine. *Tribology International*, 37(5):415–422, 2004.
- [Zio09] E. Zio. Reliability engineering: Old problems and new challenges. *Reliability Engineering & System Safety*, 94(2):125–141, 2009.
- [ZLCM14] N Zhao, S. Li, Y. Cao, and H. Meng. Remote intelligent expert system for operation state of marine gas turbine engine.

Proceedings of the 11th World Congress on Intelligent Control and Automation, pages 3210–3215, 2014.

- [ZMHF16] M. A. Zaidan, A. R. Mills, R. F. Harrison, and P. J. Fleming. Gas turbine engine prognostics using Bayesian hierarchical models: A variational approach. *Mechanical Systems and Signal Processing*, 70:120–140, 2016.
- [ZPC⁺04] A. K. Ziver, C. C. Pain, J. N. Carter, C. R. E. de Oliviera, A. J. H. Goddard, and R. S. Overton. Genetic algorithms and artificial neural networks for loading pattern optimisation of advanced gas-cooled reactors. *Annals of Nuclear Energy*, 31:431–457, 2004.
- [ZZXC10] C. Zou, E. Zheng, H.-W. Xu, and L. Chen. SVM-based multi-class cost-sensitive classification with reject option for fault diagnosis of steam turbine generator. *2010 Second International Conference on Machine Learning and Computing*, pages 66–70, 2010.