


# Automated Structural Damage Detection Using One Class Machine Learning

by

James Long

Bachelor of Engineering in Civil Engineering  
University of Limerick (2012)

**ARCHIVES**



Submitted to the Department of Civil and Environmental Engineering  
in partial fulfillment of the requirements for the degree of  
Master of Science in Civil and Environmental Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2014

© 2014 Massachusetts Institute of Technology. All rights reserved.

  
Signature redacted

Author .....  
Department of Civil and Environmental Engineering  
May 18, 2014

  
Signature redacted

Certified by .....  
Oral Buyukozturk  
Professor of Civil and Environmental Engineering  
Thesis Supervisor

  
Signature redacted

Accepted by .....  
Heidi M. Nepf  
Chair, Departmental Committee for Graduate Students



# **Automated Structural Damage Detection Using One Class Machine Learning**

by

James Long

Submitted to the Department of Civil and Environmental Engineering  
on May 18, 2014, in partial fulfillment of the  
requirements for the degree of  
Master of Science in Civil and Environmental Engineering

## **Abstract**

Measuring and analysing the vibration of structures using sensors can help identify and detect damage, potentially prolonging the life of structures and preventing disasters. Wireless sensor systems promise to make this technology more affordable and more widely applicable. Data driven structural health monitoring methodologies take raw signals obtained from sensor networks, and process them to obtain damage sensitive features. New measurements are then compared with baselines to detect damage. Because damage-sensitive features also exhibit variation due to environmental and operational changes, these comparisons are not always straightforward and sophisticated statistical analysis is necessary in order to detect abnormal changes in the damage sensitive features. In this thesis, an automated methodology which uses the one-class support vector machine (OCSVM) for damage detection and localisation is proposed. The OCSVM is a nonparametric machine learning method which can accurately classify new data points based only on data from the baseline condition of the structure. This methodology combines feature extraction, by means of autoregressive modeling, and wavelet analysis, with statistical pattern recognition using the OCSVM. The potential for embedding this damage detection methodology at the sensor level is also discussed. Efficacy is demonstrated using real experimental data from a steel frame laboratory structure, for various damage locations and scenarios.

Thesis Supervisor: Oral Buyukozturk

Title: Professor of Civil and Environmental Engineering





## Acknowledgments

My sincere thanks go to my advisor, Professor Oral Buyukozturk for his guidance and support. I am grateful also for the assistance provided by my colleagues Justin Chen, Reza Mohammadi and Dr. Young-Jin Cha, as well as Dr. Michael Feng and the team at Draper Laboratory. To Kris Kipp, Kiley Clapper, Rebecca Fowler and everyone at the MIT Department of Civil and Environmental Engineering, thank you for all your help.

I would also like to express my gratitude to Edward H. Linde (1962) and the Linde Family Foundation for their generosity. Thanks also for the support given by Royal Dutch Shell through the MIT Energy Initiative, which makes this research possible.

Finally, to Peter and Siobhán Long, thank you for teaching, guiding, listening, encouraging, even proofreading. I couldn't ask for more patient, supportive and loving parents.

# Contents

<b>1</b>	<b>Introduction</b>	<b>13</b>
1.1	Motivation for Structural Health Monitoring . . . . .	13
1.2	Research Objectives and Review of Related Work . . . . .	14
1.3	Chapter Organisation . . . . .	16
1.4	Fundamentals of SHM . . . . .	17
1.4.1	Sensor Technology . . . . .	19
1.4.2	Model based vs Data based . . . . .	21
<b>2</b>	<b>Problem Statement</b>	<b>25</b>
<b>3</b>	<b>Feature Extraction</b>	<b>29</b>
3.1	Review of Feature Extraction Methods . . . . .	31
3.1.1	Modal Properties . . . . .	31
3.1.2	Non-Stationary Frequency Domain Methods . . . . .	33
3.1.3	Nonlinear Methods . . . . .	33
3.1.4	Time series analysis . . . . .	34
3.2	Auto Regressive Time Series Analysis . . . . .	35
3.3	Continuous Wavelet Transform . . . . .	36
3.4	Data Fusion . . . . .	38
3.5	Chapter Summary . . . . .	39
<b>4</b>	<b>Statistical Pattern Recognition</b>	<b>41</b>
4.1	One Class Support Vector Machine . . . . .	43

4.1.1	The Kernel Trick . . . . .	43
4.1.2	Derivation of the One-Class Support Vector Machine . . . . .	46
4.1.3	Hyperparameter selection for the One-Class SVM using the Gaussian Kernel . . . . .	49
4.2	Damage Detection Methodology . . . . .	53
4.3	Damage Localisation . . . . .	56
4.3.1	OCSVM based decision function . . . . .	57
4.3.2	One Class Regularised Least Squares based localisation index	57
4.3.3	Chapter Summary . . . . .	60
<b>5</b>	<b>Experimental Results</b>	<b>63</b>
5.1	Evaluating Damage Detection Performance . . . . .	65
5.2	Hyperparameter Selection Method . . . . .	66
5.3	Damage Detection Results . . . . .	68
5.4	Damage Localisation Results . . . . .	69
5.4.1	Damage at Sensor 1 . . . . .	69
5.4.2	Damage at Sensor 17 . . . . .	71
5.4.3	Combined Damage at Sensor 1 and Sensor 17 . . . . .	73
5.5	Data Visualisation . . . . .	75
5.5.1	Visualisation of Damage Scenarios at Sensor 1 . . . . .	75
5.5.2	Visualisation of Damage Scenarios at Sensor 17 . . . . .	77
5.6	Sensitivity to Sample Size . . . . .	79
5.7	Effect of Data Acquisition Parameters . . . . .	81
5.7.1	Sampling Frequency . . . . .	81
5.7.2	Sampling Resolution . . . . .	81
5.7.3	Chapter Summary . . . . .	86
<b>6</b>	<b>Conclusions</b>	<b>87</b>
6.1	Future Work . . . . .	88
<b>A</b>	<b>Damage Detection Results</b>	<b>91</b>

# List of Figures

3-1	Toy example illustrating benefits of wavelet analysis . . . . .	38
4-1	Toy example illustrating the kernel trick . . . . .	45
4-2	Example of an OCSVM decision boundary . . . . .	48
4-3	Effect of hyperparameter selection on decision boundary . . . . .	50
4-4	Example of hyperparameter selection based on the training error rate	51
4-5	Illustration of the iterative hyperparameter selection method . . . . .	54
4-6	Comparison of OCSVM and RLS Localisation Index contours . . . . .	61
5-1	Schematic diagram and photograph of instrumented lab structure . .	64
5-2	Comparison of damage detection performance using different hyperparameter selection methods . . . . .	67
5-3	Detection of Damage at Sensor 17 . . . . .	68
5-4	Localisation of Damage at Sensor 1 using OCSVM LI . . . . .	70
5-5	Localisation of Damage at Sensor 1 using RLS LI . . . . .	71
5-6	Localisation of Damage at Sensor 17 using OCSVM LI . . . . .	72
5-7	Localisation of Damage at Sensor 1 using RLS LI . . . . .	73
5-8	Localisation of Combined Damage using OCSVM LI . . . . .	74
5-9	Localisation of Damage at Sensor 1 using RLS LI . . . . .	74
5-10	2D visualisation of data and decision boundary from Sensor 1, for the damage scenario at Sensor 1 . . . . .	76
5-11	2D visualisation of data and decision boundary from Sensor 4, for the damage scenario at Sensor 1 . . . . .	77

5-12	2D visualisation of data and decision boundary from Sensor 17, for the damage scenario at Sensor 17 . . . . .	78
5-13	2D visualisation of data and decision boundary from Sensor 9, for the damage scenario at Sensor 17 . . . . .	79
5-14	Effect of sample size on damage detection performance . . . . .	80
5-15	Illustration of data quantisation . . . . .	82
5-16	Effect of bit rate and sampling frequency on damage detection performance . . . . .	83
5-17	Comparison between 24 bit data and 8 bit data . . . . .	84
5-18	Comparison between 24 bit data and 8 bit data in the frequency domain	85
A-1	Damage detection for sensor 1 damage scenarios using geometric hyperparameter selection. . . . .	92
A-2	Damage detection for sensor 17 damage scenarios using geometric hyperparameter selection. . . . .	93
A-3	Damage detection for multiple location damage scenarios using geometric hyperparameter selection. . . . .	94
A-4	Damage detection for sensor 1 damage scenarios using iterative hyperparameter selection. . . . .	95
A-5	Damage detection for sensor 17 damage scenarios using iterative hyperparameter selection. . . . .	96
A-6	Damage detection for multiple location damage scenarios using iterative hyperparameter selection. . . . .	97
A-7	Damage detection for sensor 1 damage scenarios using training error based hyperparameter selection. . . . .	98
A-8	Damage detection for sensor 17 damage scenarios using training error based hyperparameter selection. . . . .	99

A-9	Damage detection for multiple location damage scenarios using training error based hyperparamter selection. . . . .	100
-----	---------------------------------------------------------------------------------------------------------------------	-----

## List of Tables

5.1	Experimental Damage Scenarios and Locations . . . . .	64
-----	-------------------------------------------------------	----

# Nomenclature

## Symbols

Symbol	Description
$a$	Wavelet scale number
$b$	Wavelet time shift
$C$	Regularised Least Squares penalisation parameter
$d_{max}$	Maximum distance between training examples
$\mathbf{e}$	Vector of ones
$e_x(t)$	Autoregressive residual
$E_a$	Wavelet energy at scale a
Hz	Hertz
$\mathbf{I}$	Identity Matrix
$j$	The imaginary unit
$N$	Number of training examples given to classifier
$\mathbb{R}^n$	Real coordinate space of n dimensions
sgn	Signum function
$K(x, y)$	Kernel function
$t$	time
$x_1, x_2$	Dimensions of 2D real coordinate space
$X(a, b)$	Wavelet transform at scale a and time b
$\mathbf{w}$	Coefficients describing hyperplane in OCSVM derivation
$x(t)$	Vibration time series
$\alpha_j$	Lagrangian multiplier in OCSVM derivation

$\mu(t)$	Mean of vibration time series
$\nu$	OCSVM parameter controlling outlier fraction
$\xi_i$	Slack variable in OCSVM derivation
$\rho$	Constant bias term of hyperplane in OCSVM derivation
$\sigma$	Gaussian kernel parameter
$\sigma(t)$	Standard deviation of vibration time series
$\varphi_k$	Autoregressive model coefficients
$\phi(x)$	Feature Space induced by kernel
$\psi(t)$	Wavelet function
$\omega_0$	Morlet mother wavelet central frequency

## Acronyms

Acronym	Description
AR	Autoregressive
ARMA	Autoregressive Moving Average
ASCE	American Society of Civil Engineers
CWT	Continuous Wavelet Transform
DSF	Damage Sensitive Feature
FFT	Fast Fourier Transform
LI	Localisation Index
MEMS	Microelectromechanical Systems
OCSVM	One Class Support Vector Machine
PCA	Principle Component Analysis
RLS	Regularised Least Squares
SHM	Structural Health Monitoring
SPR	Statistical Pattern Recognition
SVM	Support Vector Machine
SVs	Support Vectors



# Chapter 1

## Introduction

### 1.1 Motivation for Structural Health Monitoring

The aim of Structural Health Monitoring (SHM) is to provide warning when the condition of a structure deteriorates, by collecting and analysing data from sensors measuring the structure. A system which reliably achieves this goal has the obvious benefits of preventing disastrous structural collapses, extending infrastructure life time, and reducing maintenance costs. The idea of monitoring the condition of structures is not new. Almost all structures experience deterioration over their lifetimes, and therefore periodic visual inspection and non-destructive testing of structures has long been practised. Advances in sensor technology and computing capability promise to vastly improve on these techniques. SHM has many potential benefits: Fundamentally, the use of sensors allows for earlier and more reliable detection of structural damage in comparison to visual inspection or non-destructive testing. It also provides a method for the assessment of components which can not be visibly inspected, for example aerospace components concealed by heat shields, or which are costly or difficult to inspect manually such as offshore structures. Automated SHM approaches would also enable us to monitor a much greater number of structures, on a much more frequent basis, than existing manual inspection methods. Given that there are more than 600,000 bridges in the United States, with an average age of 42 years, one in nine of which have been identified as structurally deficient, the benefit is clear.

## 1.2 Research Objectives and Review of Related Work

The primary goal of this thesis is to develop and verify an automated vibration based damage identification methodology for wireless SHM which can be embedded on board a sensor equipped with computing power. This would enable a network of SHM sensors to act autonomously, acquiring and analysing data without the need for intervention from a human expert, and only transmitting data when an alarm is raised. The advantages of this approach are numerous, potentially lowering the cost of deploying a SHM system, reducing the amount of power required and increasing the reliability of data transmission. Further discussion of the motivation for this decentralised, automated approach will be discussed in detail in Chapter 2.

Previous research which has focused on the goal of developing automated, 'smart sensing' SHM technology, as described above, has adopted a data-based approach to damage identification, where time series of data acquired from sensors are processed into damage-sensitive features, and then statistical pattern recognition is carried out, comparing damage-sensitive features with baseline values to identify damage. For many applications, particularly in infrastructure monitoring, this comparison is a novelty detection problem, meaning new tests are compared only with nominally 'normal' baseline values. This approach is sensible and attractive for smart sensing, and will be adopted in this thesis.

However, existing approaches to the novelty detection problem in the SHM literature have a number of drawbacks, and in this thesis non-parametric, kernel methods from the field of machine learning will be proposed as an improvement to the existing state of the art in SHM. We will now describe a selection of related state of the art research efforts, and outline the drawbacks of these approaches which will be addressed by the research in this thesis.

Lynch [14] developed functioning smart sensing hardware which has the capability to extract autoregressive based features from time series, on board a sensor equipped with a microcontroller. The damage sensitive feature used is the standard deviation

of the residual error of an ARX model. For statistical pattern recognition, first a prescreening step is carried out to identify which entry in the baseline database has the most similar AR model, to reduce the effect of operational variation. Once the most similar baseline measurement is identified, the ratio between the baseline and new damage sensitive features is computed. A threshold value for this ratio is chosen by 'engineering judgement'. This approach, while suitable for on board computation and quite sophisticated in its prescreening phase, requires the ad-hoc selection of thresholds, and is not suitable for the use of multivariate vectors of damage sensitive features which may capture a greater range of damage types and scenarios.

Worden et al [28] published one of the first SHM studies to explicitly address the issue of statistical pattern recognition as part of a comprehensive damage detection methodology. In this work, the Mahalanobis squared distance is used as a discordancy measure between new features and baseline values. A threshold value of this measure, above which new tests are classified as damaged, is calculated by conducting a Monte Carlo simulation and retaining the value above which only some small percentage of the simulated values fall. This outlier analysis procedure has been adopted by many other SHM researchers, for example Gul and Catbas [9]. The drawback of this approach is that the Monte Carlo procedure makes assumptions on the probability distribution of the data. As will be shown later in this thesis using experimental data, these assumptions may not always hold true, and in some cases may lead to a failure to detect damage.

Nair et al [18], combine Autoregressive coefficients into a single damage sensitive feature. Statistical pattern recognition is then conducted by using a t hypothesis test to establish the statistical significance of the difference in mean between the undamaged baseline DSFs, and the new damaged DSFs. This approach also has drawbacks. The mean value of a DSF representing a damaged condition may be very close to the mean value of the baseline data. Consider the case of a bimodal distribution of baseline damage sensitive features: The mean of such a distribution lies far from the actual baseline data and would likely suggest the occurrence of damage, but using this methodology new data which lies near this mean value would

be considered to be from the undamaged structure.

Nair and Kiremidjian [16] suggest the use of Gaussian Mixtures Modelling for statistical pattern recognition. Gaussian Mixtures Modelling models data as the superposition of Gaussian distributions. In this work, it is reasoned that once damage occurs the number of mixtures in the model will be altered. This approach overcomes many of the issues with parametric methods. However, for novelty detection problems in other fields, Gaussian Mixtures Modelling has been outperformed by one class kernel methods. In this thesis we will investigate the use of two such kernel methods, the One Class Support Vector Machine, and the One Class Least Squares SVM, for damage detection and damage localisation. A comprehensive damage detection and localisation methodology will be developed using these approaches for the statistical pattern recognition phase, and the methodologies will be investigated using real experimental data from a steel frame laboratory structure.

## 1.3 Chapter Organisation

In the previous section we have outlined the primary research objectives of this thesis, reviewed the related literature and identified what improvements we hope to make on the existing state of the art research in data based, wireless structural health monitoring. For context, the remainder of this chapter will provide a broader overview of the structural health monitoring problem and available technology, not limited to data based, smart sensing methodologies.

Chapter 2 will provide a more detailed discussion on the motivation for adopting the smart sensing, data based approach, and an overview of the methodology developed in this thesis.

Chapter 3 will begin with a discussion of the specific requirements for feature extraction in a wireless system, before an overview of common feature extraction methods in the SHM literature and a discussion of their merits. A number of candidate damage-sensitive features for the methodology presented in this thesis will then be identified and discussed in detail.

Chapter 4 will review previous approaches to the statistical pattern recognition element of SHM, before discussing the advantages of the OCSVM and providing its theoretical formulation. Three methods of automated hyperparameter selection for the OCSVM will be suggested. A detailed discussion of the timeline and implementation of the full damage detection methodology will then follow. Finally, the problem of damage localisation within the decentralised smart sensing framework will be addressed, and two methods for defining a localisation index will be presented.

Experimental verification of all algorithms developed in this thesis will be presented in Chapter 5. A description of the steel frame laboratory structure, its instrumentation and the experimental procedures used to acquire data will be provided, as well as results demonstrating the efficacy of the damage detection and localisation algorithms. Studies on the effect of the volume of training data from the baseline structure used, as well as the effect of sampling frequency and resolution, key sensor parameters, are also conducted and discussed. Finally, Chapter 6 will provide a summary of the work carried out for this thesis, the main conclusions reached, and a suggestion for what future work will entail.

## 1.4 Fundamentals of SHM

Over the past few decades SHM has attracted a great deal of attention, with the establishment of dedicated academic conferences, journals and an active research community. This section will provide an introduction to the basic problem of SHM, and the various technologies and strategies that have been suggested to meet this challenge. Vibration based damage detection is the field of SHM which seeks to identify damage in structures by measuring their vibration response; the thinking being, that any change in the structural system will manifest itself as a change in the vibration characteristics. While there are other ways of approaching the SHM problem, this thesis is focused on vibration based damage detection.

Worden et al [27] state that the goal of SHM is to establish the existence of damage, the location of damage(s), the type of damage and the severity of damage.

In order to achieve any of these goals, the following steps, identified by Farrar et al [5] (and slightly modified and generalised here) are required as part of any SHM methodology:

- **Operational evaluation:** The first step in developing a SHM system is establishing what type of damage is likely to occur, and of interest for a given structure, what conditions the structure operates in, and what constraints exist on how data can be collected. This entails making high level decisions on what the SHM methodology should consist of, and what will be monitored. It is possible that these higher level decisions can be made not for an individual structure, but for a class of structures, for example concrete highway bridges in a given geographical region. Clearly, the value, importance and uniqueness of the structure will all impact the degree to which operational evaluation is customised.
- **Data acquisition, cleansing and normalisation:** Before useful data can be acquired from a structure, a number of decisions are required: What type of sensors are appropriate, where should they be located, how often will they collect data, how will the data be transmitted or stored? Again, economic considerations are important. In particular, the type of sensors chosen, and method of data transmission can greatly influence the cost of the system, and the structural health monitoring methodology. There are many knock on effects on the next two steps in the general SHM methodology presented here, feature extraction, and baseline comparison. A discussion of these effects and the advantages and disadvantages of different sensor technology will follow this overview of SHM methodologies.
- **Feature extraction and information condensation:** As Farrar [5], observes, the definition of damage sensitive features, which condense time series measurements and retain sensitivity to damage, has been the focus of much of the research published on SHM. Which features are most useful will depend on the sensor technology, the application, and also on the method of baseline

comparison used. An overview of popular feature extraction algorithms will be provided in Chapter 3.

- **Comparison with baseline:** Once the measured data has been processed into damage-sensitive features, these features need to be compared to a baseline in order to establish whether damage has occurred. This comparison can be conducted in two ways: Model-based or Data-based. In the model-based approach new data is compared with a model of the structure (e.g. finite element model), while in the data-based approach new data is simply compared with existing data from the same structure. The choice of one of these approaches has many implications in the broader picture of SHM, and will guide the entire SHM methodology. Sharp distinctions between these two approaches have been drawn in the literature, and their relative strengths and weaknesses will be further discussed later in this chapter.

Any method for vibration based monitoring of structures must address these topics. While none of these steps exist in a vacuum, there are clear divides in the SHM literature in how each step is approached. Some of the most important distinctions between SHM methodologies will now be discussed within the context of the above framework.

### 1.4.1 Sensor Technology

A wide variety of sensors are available for use in SHM systems. To capture environmental and operational variations it may be desirable to measure temperature, humidity, wind direction or other salient features. The incorporation of these sensors into a holistic SHM sensor system is important and likely to improve the diagnostic capability of damage detection algorithms. Beyond these operational factors, one of the key considerations in designing a sensor network for vibration based damage detection is how to measure the movement of the structure. Fundamentally, vibration sensors convert changes in displacement, velocity or acceleration to changes in an electric signal. This analog electric signal is then converted to a digital represen-

tation. How closely spaced the discretised digital representation is, and how much resolution is provided are important parameters with implications for damage detection performance. This issue will be examined in Chapter 5 of this thesis. A number of technologies, and therefore sensor types, exist for the conversion of mechanical vibration to an electric signal. A brief review of some of the most popular sensor types will be presented here. Perhaps the most popular sensor used in the SHM literature is the piezoelectric accelerometer. This sensor consists of a known mass spring system to which a piezoelectric material is attached. As the piezoelectric material deforms, it experiences stress, which causes it to output a voltage due to the piezoelectric effect. Sensors of this type have proven to be reliable and accurate [6], but are relatively expensive. Force balance accelerometers consist of a known mass at the end of a pendulum positioned between magnetic coils. As a force is applied to the sensor, the movement of the mass is resisted by means of a current passing through the coils, and the mass is kept in its original position. The greater the force acting on the mass, the more current is required to keep it in position, and thus the acceleration is measured by the current passing through the coils. Sensors of this type are extremely sensitive, but are large and most suitable for low frequency applications, hence their popularity for measuring seismic movement. Micro electrical-mechanical systems (MEMS) accelerometers have become increasingly popular, due to their small size and affordability. Typically manufactured from silicon, these sensors consist of a movable plate of known mass with extending fingers, as well as complementary stationary fingers. As force acts on the sensor, the plate moves, and the attached fingers move relative to the stationary fingers. Because these 'fingers' make up a differential capacitor, this movement results in the flow of current. MEMS sensors can be as much as an order of magnitude less expensive than traditional piezoelectric accelerometers, and as such represent an opportunity to make widespread application of SHM systems a realistic prospect. Once mechanical vibration has been converted to an electric signal, and then to a digital representation, the question of what to do with this data arises. Wired sensor networks transmit the data via wired connections to a computer for storage and further processing. For a dense network of sensors



the installation of such a system quickly becomes a non-trivial task and can consume significant time and money. Furthermore, wired networks rely on the existence of electricity, telecommunications and an appropriate location to house a computer. In many scenarios this is not realistic. Wireless networks promise to ameliorate many of these issues, but come with their own problems. A single typical test for damage detection purposes will contain thousands of entries. When considering a network of hundreds or thousands of sensors, the transmission of this test data is problematic, potentially saturating wireless bandwidth. Even in normal operating conditions wirelessly transmitted data can be subject to lost packets, requiring the data series to be resent. In addition, wireless transmission of data uses substantial amounts of power. Given these issues, there has been a growing trend towards using micro-controllers embedded at the sensor location to process data before transmission. Lynch et al [14] reported that by processing acceleration data into damage-sensitive features prior to transmission a reduction in power usage of approximately 50 % was achieved.

#### **1.4.2 Model based vs Data based**

The necessity to perform comparisons with a baseline was identified by Worden and Farrar [27] as a fundamental axiom of structural health monitoring. How this comparison is conducted effectively divides all SHM technologies. The distinction between different approaches most often drawn in the SHM literature is that between model-based and data-based approaches. These, fundamentally, are different ways to approach comparing new data with a baseline.

##### **Model Based comparisons**

In the model-based approach the comparison can be mathematically formalised as an optimisation problem, where the model itself is changed, or updated, to minimise the differences between the predicted response and the measured response. This assumes that a model of the structure, typically a finite element model, is available. The model-based approach has the advantage of potentially identifying not only the oc-

currence and location of damage, but also the severity. This is usually achieved by the solution of an inverse problem, where the error between measured modal parameters and those predicted by the model, through the solution of an eigenproblem, is minimised by altering the mass, stiffness or damping matrices of the model. Friswell [7], provides an overview of some common methods for solving this inverse problem. Despite the advantage of localising and quantifying damage, the model-based approach is fraught with difficulty, both practical and technical. Because the number of available measurements, for example the number of natural frequencies, is almost always less than the number of parameters for potential updating, the inverse problem suffers from non-unique solutions. As outlined by Friswell [7], this problem can be overcome by using physical intuition to parameterise the finite element model. For example, we may stipulate that only the locations most likely to incur damage be allowed to vary in the updating procedure. Alternatively, a regularisation approach can be taken, for instance introducing penalising the norm of the solution to ensure uniqueness. Still, difficulties persist. If the type of damage anticipated can not be accurately described by a linear reduction in stiffness or mass, a nonlinear finite element model must be used. In addition, before the damage detection process can begin, the model must be calibrated to ensure that it accurately represents the baseline condition of the structure. This in itself is a difficult task, complicated by potential imperfections in boundary conditions and uncertainties and variations in the physical properties of the structure. There is also an economic consideration, as if a model of the structure does not already exist, development may be an expensive prospect. Lastly, as discussed in Section 1.4.1, there is a growing trend towards wireless technology in SHM, and towards the embedding of the damage detection process on board microcontrollers. The complexity and global nature of finite element models makes it difficult to imagine a scenario where it would be feasible to use a model-based approach as part of a wireless smart sensing system.

## Data Based comparisons

Farrar [5] introduced a paradigm for data based damage detection which has been widely adopted in SHM studies. A number of studies which use this approach were discussed earlier in this chapter, in Section 1.2. Further discussion of the literature will be conducted in Chapter 4. In the data-based approach, comparisons of features from new tests with baseline values are statistical in nature. This comparison can be conducted as a novelty detection problem, where the baseline is taken only from the existing condition of the structure, or it can be a fully supervised problem, where the database contains data from the existing condition of the structure, as well as various damaged configurations. Generally speaking, novelty detection approaches can potentially detect and localise damage, but fully supervised algorithms are required to quantify the type and severity of damage. These two approaches will be further discussed in Chapter 4. Regardless of which is chosen, the data-based approach relies on the collection of a sufficiently representative database of features to describe the condition of the structure. To provide reliable performance, this database should incorporate data representing the variations in environmental and operating conditions the structure is subjected to. The acquisition of this dataset may seem prohibitive, but can be done in an automated manner in many applications, and compared to the development and calibration of a finite element model may prove relatively inexpensive. The data-based approach allows the use of any damage-sensitive feature, and is not limited to modal parameters easily tractable for model updating procedures. This flexibility potentially allows for more sophisticated damage detection, incorporating nonlinear and nonstationary effects. If we carefully consider the choice of damage-sensitive features the data-based approach has one major advantage over the model-based approach: It can be adapted for use in a smart sensing system, where computation is incorporated into the sensor technology. For these reasons, this thesis will adopt a data-based damage detection philosophy. The motivations for this choice, and an outline of the specific problems that will be addressed in this thesis are provided in the following chapter.



# Chapter 2

## Problem Statement

Recent developments in MEMS accelerometer technology and wireless sensing protocols promise to dramatically reduce the cost of SHM systems, potentially making the widespread deployment of dense networks of sensors feasible. As mentioned in Section 1.4.1, MEMS accelerometers can be as much as an order of magnitude less expensive than conventional piezoelectric sensors, while wireless transmission of data allows the monitoring of structures in locations without existing telecommunications and electrical infrastructure. In addition, even in locations with these amenities, the installation of a wired sensor network is costly, time-consuming and difficult. New wireless sensor technologies present a massive opportunity to transform SHM from a largely academic endeavour, to a technology which is widely implemented and seen as a vital component of infrastructure systems. However, the goal of achieving low-cost, easily adaptable sensor based damage detection systems requires methodologies which account for the constraints of wireless sensing. In particular, algorithms which mitigate the need to transmit full acceleration series are desirable for a number of reasons. The transmission of full vibration time series is likely to saturate wireless bandwidth, in addition to quickly depleting battery life. While energy harvesting techniques may mitigate the need to replace batteries manually, the reduction of power use is still a desirable goal. By carrying out some, or all, of the damage detection methodology prior to transmission, the full time series can be condensed to a set of damage sensitive features, or even a single scalar number indicating a decision.

Lynch et al [14] report that by processing data into damage sensitive features on a micro-controller at the sensor level prior to data transmission, dramatic reductions in power consumption can be achieved. This decentralised, parallel computing approach not only reduces power consumption, but also bandwidth usage, thereby improving the reliability of transmission. Additionally, in the event that wireless transmission is interrupted, the reduction in the volume of relevant data facilitates temporary storage until transmission can be successfully completed. If a fully automated methodology can be implemented on board, the sensor may not need to transmit any data other than a check in to confirm it is working, unless damage is detected.

The desire for embedded computation naturally leads to the choice of a data-based damage detection methodology (described in Section 1.4.2), as model-based approaches are likely to require more computational effort, and being global in nature are less suitable for decentralised computation. To recap, in the data-based approach, raw acceleration histories are acquired and then processed into damage sensitive features (DSFs) prior to statistical comparison with a baseline. This process, referred to as feature extraction, often compresses the data, while remaining sensitive to changes in the underlying structural system. While the literature provides many suggestions for appropriate DSFs, embedding of the feature extraction process on a micro-controller places many additional constraints on the choice of DSFs. These constraints are discussed in Chapter 3, and a survey and evaluation of the literature is conducted. On the basis of this survey, features derived from an autoregressive model and from the continuous wavelet transform are chosen as suitable candidates.

Once the acceleration signals have been processed into damage-sensitive features, data-based damage detection methodologies typically then use statistical methods to compare new features with a database of baseline features. In Chapter 4, we will argue that, for many SHM applications, this statistical comparison requires a 'one-class', or novelty detection approach. The one-class approach consists of developing a model of data taken only from the baseline condition of the structure, assumed to be intact, and then using this model to evaluate data taken from new tests. The

obvious drawback of this approach is the assumption that the structure starts the monitoring process in an intact state. However, in many cases, we are not interested in whether or not the structure is in exactly the same as originally constructed, but rather whether it is degrading from its current state. Indeed, even new structures often contain defects or behave slightly differently than intended by designers. Rather than attempt to quantify and model every deviation from the designer’s intention, the data-based approach tries to detect changes from the current state, by constructing a statistical model of data from the structure’s real behaviour.

In 2000, Farrar [5], noted that almost none of hundreds of SHM studies conducted to that point had made use of any statistical methods, but since then growing attention has focused on this step of the damage detection process. However, approaches to one-class classification in the structural health monitoring literature have focused almost exclusively on parametric approaches, where damage-sensitive features from the baseline structure are assumed to adhere to a known probability distribution. In this thesis a non-parametric, state of the art machine learning algorithm called the one-class support vector machine (OCSVM) [20] will be discussed and applied to the problem of structural damage detection in a manner which is suitable for embedded processing on board a micro-controller. The OCSVM, and related kernel methods, have consistently outperformed traditional parametric methods on analogous novelty detection problems in fields outside of structural health monitoring. The one-class support vector machine is trained on data from the baseline condition of the structure, and can then evaluate new points by comparing them with the baseline in a sparse manner; meaning that only a fraction of the training data needs to be stored. A non-parametric decision boundary, separating normal values of feature vectors from anomalous values, is generated, offering a much more flexible method for novelty detection than the parametric methods popular in the SHM literature. Additionally, we conduct the entire methodology in a fully decentralised manner, where each sensor acquires data locally, and processes this data without the need for external communication, before returning a binary decision indicating whether damage has been detected at the sensor location or not.

Due to their ability to generate accurate, non-linear decision boundaries at relatively low computational cost, support vector machine methods have attracted attention from researchers working on embedded sensing technology in other fields. Lee and Verma [12] demonstrated a low-power implementation of the support vector machine for detecting the onset of seizures and arrhythmia in patients. Feature-extraction for data-based damage detection has been successfully embedded on a sensor mote by Lynch et al [14], using autoregressive time-series analysis, but to date, little attention has been given to embedding statistical pattern recognition algorithms for performing damage detection at the sensor level prior to data transmission. The OCSVM approach developed in this thesis is an ideal candidate for powerful, low-cost statistical pattern recognition for embedded structural damage detection. A sensor mote which combines existing feature-extraction algorithms with novelty detection by means of the OCSVM would provide the capability for automated, sophisticated damage detection at the sensor level, reducing the required wireless transmission of data to a single scalar for each sensor.

Identifying the location of damage is usually viewed as a more challenging problem than simply detecting the occurrence of damage. Many researchers have relied either on a model-based approach, or a data-based approach which incorporates data from several damaged configurations of the structure, to successfully localise damage. Discussion in the literature of data-based, smart-sensing applicable algorithms for localisation, which do not require data from damaged scenarios a priori, is extremely limited. Nair et al [18], present such an approach, which is decentralised and suitable for embedded computation. In this approach a localisation index is calculated which describes how far test data deviates from the baseline training data. The hypothesis is that this localisation index will be highest closest to the damage location. A similar methodology for damage localisation is proposed in this thesis, but here, two localisation indices based on the OCSVM are proposed. Because these localisation indices are non-parametric, they offer the potential for more accurate description of distance from the training data, and therefore better localisation, while still operating within the decentralised methodology.



# Chapter 3

## Feature Extraction

The requirements of a low-cost, low-power damage detection system, place a number of constraints on the feature-extraction process. While a significant research effort has been dedicated to the suggestion of DSFs for structural damage detection, many of these features are not appropriate for embedding at the sensor level. Specifically, in addition to being sensitive to damage, and robust with respect to benign ambient variations, features should also satisfy some, or all, of the following criteria:

- Output only: This requirement stipulates that to extract the damage-sensitive feature, we do not need to know explicitly the input excitation to the structure. Some SHM strategies depend on the ability to provide or measure the input excitation, for example with an impact hammer, or by relying on measurable seismic excitations. This information is undoubtedly useful, but hard to come by. The methodology and algorithms described in this thesis are developed with the goal of widely applicable damage detection, and therefore it is desirable to be able to extract damage-sensitive features without explicit measurement of the input. Of course, to extract useful damage-sensitive features, there must be energy to cause the structure to vibrate; we anticipate that for a real life application, this would be from ambient sources such as wind, moving vehicles, waves or machinery. In this thesis we will assume that there is an excitation source available, and that the frequency content of this source excitation is

relatively consistent. This assumption is certainly not universally true, but it is reasonable in many cases. For example, an offshore structure is excited by ocean waves, which are almost always present, and while stochastic in nature, are generally consistent in frequency content. Similarly, the excitation provided to a bridge by a moving vehicle is dependable and relatively consistent. Further work is required to ensure that test data is acquired in a manner which satisfies this assumption, for example triggering the collection of data on a bridge when a single vehicle passes a certain point in the span, on an offshore structure when waves are coming from a certain direction.

- Single sensor only: To allow for decentralised, embedded execution of the damage detection algorithm, the extraction of damage-sensitive features should be conducted only using data from the location of the sensor in question. As previously discussed, wireless transmission of time-series is a costly prospect. This is not limited to communication from sensors to a base station; it also applies to communication between sensors. Therefore, in the decentralised methodology, it is highly desirable that damage-sensitive features can be extracted locally, rather than globally.
- Low computational effort: As referenced in Chapter 2, Lynch [14] has reported 50% reductions in power consumption by processing data into features prior to wireless transmission. Of course, as the feature extraction process becomes more and more computationally expensive, the advantage of on-board processing diminishes. If the algorithm for feature extraction is extremely complex, the benefit of embedded computation may cease to exist.
- Condenses volume of raw data: Generally, the power saving obtained by on-board processing is due to a reduction in the volume of data to be transmitted. The methodology presented in this thesis proposes that in addition to feature extraction, statistical classification should be conducted on board a micro-processor. Although this will result ultimately in the condensation of data to a single value, there are still benefits to condensing the data at the fea-

ture extraction stage. The OCSVM classification algorithm used in this thesis requires the storage of baseline features at the sensor location. If the dimensions of the features are large, this will not only require more memory to store a single example, but for algorithmic reasons, the OCSVM may require more examples to be stored. Therefore there is a significant benefit in condensing the time series to a smaller number of features.

- Sensitive with proximity to damage: In chapter 4, an algorithm for identifying the location of damage will be introduced. The hypothesis this algorithm is based on, is that damage-sensitive features become more sensitive the closer they are to the damage location. In the case of global features, for example the frequency of the fundamental mode of the structure, this will almost certainly not be true. For locally defined features, this hypothesis is perhaps more reasonable.

## 3.1 Review of Feature Extraction Methods

Based on the criteria outlined above, a review of the existing literature will now be conducted. The suitability of popular feature extraction methods for embedded computation will be discussed. As noted by Ying et al [30], no one feature is likely to be universally sensitive to all potential damage scenarios, robust to all operational variations, for all types of structure; and therefore instead of trying to find one optimal feature, the strategy should be to identify a number of potentially useful features.

### 3.1.1 Modal Properties

Many of the early studies in structural health monitoring focused on the fact that changes in structure can manifest themselves as changes in natural frequency, mode shape, or modal damping. In 1975, Vandiver [25], used changes in natural frequency to detect structural failure of an offshore fixed platform. Yuen [31], conducted a numerical study of the effect of damage in a cantilever beam on its mode shapes.

Pandey et al [19], used mode shape curvature, the second derivative of mode shape, in a simulation showing the potential for damage detection and localisation in a beam. Early studies used either numerical models, for which an eigenproblem can directly be solved, or alternatively a simple manual peak-picking and curve fitting scheme. Advances in the highly related field of modal analysis have provided algorithms which can identify in an automated fashion, the natural frequency, mode shapes, and damping using only output from sensors, to a higher degree of accuracy than peak-picking. Stochastic subspace identification, developed by Van Overschee and De Moor [24] and Frequency Domain Decomposition, developed by Brincker et al [2] are two such algorithms which can identify modal parameters in an accurate and automated fashion. However, modal parameters are generally considered to be relatively insensitive to damage, and their identification requires output from multiple sensor locations. Therefore the use of modal parameters as damage-sensitive features will not be considered further in this thesis, as it does not meet the criteria outlined at the beginning of this chapter. Instead of attempting to extract global features, such as mode shapes, we can instead analyse the frequency content of individual sensor signals. The most common method of conducting frequency analysis is to use the ubiquitous Fast Fourier Transform (FFT), to transform time series to the frequency domain. This may provide insight into the nature of the signal, but it only condenses the data to half the original time series length. Therefore additional processing is usually used to further condense the frequency domain representation of the data to a more reasonable dimension. An example of this approach is shown by Ying et al [30], who define various different features, such as the amplitude of the largest peak in the frequency domain divided by its location, or the number of peaks greater than 60% of the largest amplitude. Due to the efficiency of FFT algorithms this approach is one of the least computationally expensive available for the extraction of damage-sensitive features. However, as with modal parameters, FFT based damage-sensitive features may not show sufficient sensitivity to damage due to assumptions of linearity and stationarity.

### 3.1.2 Non-Stationary Frequency Domain Methods

One of the limitations of FFT based methods of feature extraction is the assumption of stationarity. Time-frequency and time-scale methods, which allow the frequency representation of the signal to vary over the duration of the signal, promise to overcome this drawback. A comprehensive overview of these methods and their application to structural health monitoring is provided by Staszewski and Robertson [22]. Among the most popular time-frequency methods is the short-time Fourier transform, an extension of the Fourier transform which uses a moving window to allow for non-stationarity. The need to use a window function arises due to the fact that fourier components are of infinite duration. Time-scale analysis, on the other hand, decomposes a signal into wavelet components, which are localised in both time and frequency, and therefore does not require windowing. The use of features derived from the continuous wavelet transform (CWT), for low power, embedded feature extraction was suggested by Nair and Kiremidjian [17], and the efficacy of this method was verified on data from the ASCE benchmark structure. Noh et al demonstrated the use of CWT features for damage detection using strong motion data. Based on this, and the advantages of non-stationary time-scale analysis over traditional Fourier analysis, the CWT approach will be adopted as one of the feature extraction methods in this thesis.

### 3.1.3 Nonlinear Methods

All feature extraction methods discussed thus far are inherently linear. Almost every structure displays some degree of nonlinearity, but this effect is usually small in undamaged structures, and linearisation is acceptable. Damage in a structure due to cracking, yielding or connection-loosening may cause an increase in nonlinearity. If this effect can be captured from the measured sensor data it may enable earlier and more accurate detection of damage. A number of feature extraction methods which try to use this concept have been developed. For example, Todd [23] uses nonlinear dynamics concepts to define the local attractor variance as a damage-sensitive feature.

The structure is provided with a chaotic excitation, and the response is analysed in state-space, where it is hypothesised that damage will cause changes to the steady state trajectory, or 'attractor'. This is potentially a very powerful methodology, but unfortunately requires precise measurement of the input excitation, which may not be possible for many real-life structures.

Empirical Mode Decomposition, and the related Hilbert-Huang Transform are nonlinear signal processing techniques which have been used for SHM feature extraction. Unlike Fourier or wavelet transformation, empirical mode decomposition does not specify basis functions for signal expansion, but rather uses an empirical, iterative algorithm to expand the signal into well-behaved 'intrinsic modes'. This adaptive method preserves nonlinearity in the signal and allows the 'modes' to have time-varying frequency and amplitude. While some very promising applications of this algorithm to damage detection have been demonstrated in the literature [13], [29], the computational effort required is high, likely making this approach unsuitable for wireless application.

### **3.1.4 Time series analysis**

Autoregressive (AR) time series models have been proposed as a means for extracting DSFs in a number of SHM studies [21], [16]. The autoregressive model assumes that the output at the current time in the model is a weighted linear combination of the output of a number of the previous time steps. How many of the previous time steps are included is referred to as the order of the model. An Autoregressive Moving Average (ARMA) model, extends this by relaxing the assumption of stationarity. The AR model has been implemented by Lynch [14] on board a microprocessor, and was demonstrated to deliver a dramatic reduction in power consumption compared to transmitting the original time series in full. The AR model is output only, requires only information from a single sensor, is computationally efficient, and condenses the time series data significantly. While the AR model is linear and stationary, and as such may suffer from the same drawbacks in sensitivity as FFT based methods, it does provide a much greater condensation of data than the FFT, and the definition

of damage-sensitive features is thus clearer. Instead of arbitrarily defining specific features of the frequency spectrum, we can simply use the AR coefficients as damage-sensitive features.

A comprehensive review of feature extraction methods has been presented in this section and the merits of these methods for wireless sensing applications have been discussed. Based on this review, we have selected the continuous wavelet transform method, and the AR coefficient method, as suitable feature extraction techniques. These methods will now be presented and discussed further.

## 3.2 Auto Regressive Time Series Analysis

As already stated, the simple Autoregressive model is a linear, stationary model which models the evolution of the time series as being dependent in a linear fashion on the previous  $p$  time steps. All time series are normalised to have zero mean and unit standard deviation prior to fitting of the model

$$x(t) = \frac{x(t) - \mu(t)}{\sigma(t)} \quad (3.1)$$

The AR model is described by the following equation:

$$x(t) = \sum_{k=1}^p \varphi_k x(t-k) + e_x(t) \quad (3.2)$$

where  $x(t)$  is a single acceleration time series,  $\varphi_k$  are the autoregression coefficients,  $p$  is the order of the model, and  $e_x(t)$  is the residual error. The Burg method for estimating the AR model is chosen in this thesis. The first five coefficients of the AR model are chosen as the damage-sensitive features. Due to the linear nature of the model these features are most likely to capture changes such as reductions in stiffness, cross-section or mass.

### 3.3 Continuous Wavelet Transform

Another promising method for extraction of DSFs is the use of wavelet analysis. Unlike traditional Fourier analysis, or the time series approach described in the previous section, wavelet analysis is nonstationary, providing information in both the time and scale domains, which may yield more insightful information on transient phenomena. Nair and Kiremidjian [17] also demonstrated that the energy of the wavelet coefficients obtained from a continuous wavelet transform are directly related to the modal properties of the structure, and can therefore be used as a damage-sensitive feature. The continuous wavelet transform (CWT) at scale  $a$ , and time  $b$  is given by:

$$X(a, b) = \frac{1}{\sqrt{a}} \int_{-\infty}^{+\infty} x(t) \psi\left(\frac{t-b}{a}\right) dt \quad (3.3)$$

where  $\psi\left(\frac{t-b}{a}\right)$  is the mother wavelet function which has been scaled by the factor  $a$ , and translated by  $b$ . The mother wavelet function must have the following properties:

$$\int_{\mathbb{R}} \psi(t) dt = 0 \quad (3.4)$$

$$\| \psi(t) \| = 1 \quad (3.5)$$

One such mother wavelet function, which will be used in this thesis, is the Morlet wavelet:

$$\psi(t) = e^{j\omega_0 t} e^{-\frac{t^2}{2}} \quad (3.6)$$

where  $j$  is the imaginary unit,  $\omega_0$  is the central frequency of the mother wavelet and  $t$  is time. Nair and Kiremidjian [17] defined the damage-sensitive features as the total energy contained at each wavelet scale. In this thesis we will use 5 wavelet based damage-sensitive features, corresponding with the total energy contained at each of the first five  $a^{th}$  scales.

$$E_a = \sum_{b=1}^T |X(a, b)|^2 \quad (3.7)$$

The advantages of the non-stationary nature of the CWT are illustrated in Fig-



ure 3-1. Here, a toy example showing a signal consisting of a 60Hz wave, with a delta function superimposed at the half way point is shown. Underneath the time series, the continuous wavelet transform of the signal is shown. The CWT, due to its non-stationary nature, is able to decompose the signal into the constant sinusoid component, and also a number of high frequency components localised in time around the location of the delta function. In contrast, the FFT, because it is stationary, does not detect the delta function. This is obviously an unrealistic toy example, but we can easily imagine damage such as a crack or loosened connection manifesting itself as quickly decaying high frequency content.

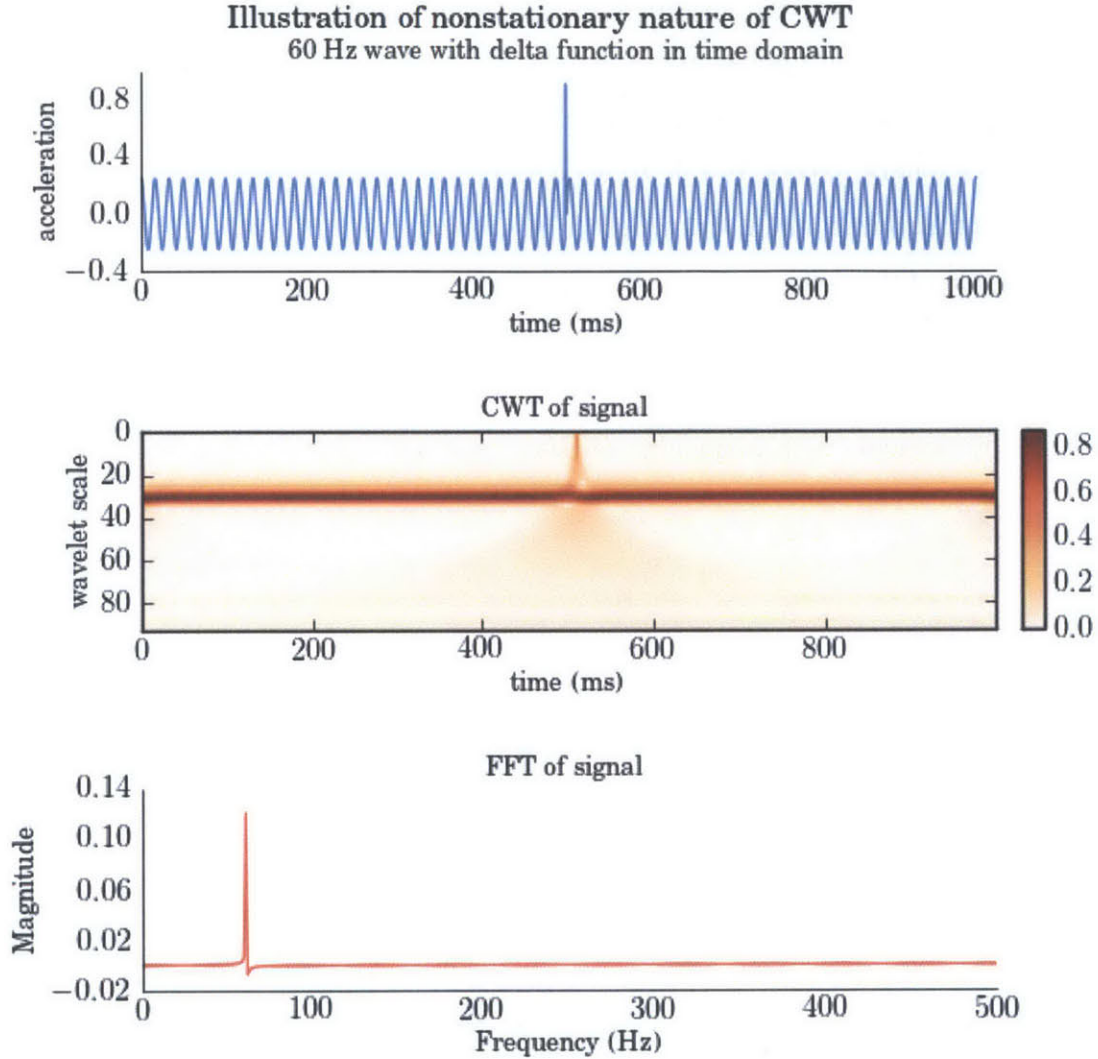


Figure 3-1: Toy example illustrating benefits of wavelet analysis

### 3.4 Data Fusion

The feature extraction stage is completed by fusing the features into a vector which provides more comprehensive information than any one individual feature. Triaxial sensors are used in this study, providing acceleration time series in the three axes for each test. Combining the 5 AR coefficients with the wavelet coefficient energies for the first five scales gives 10 features for each direction, for a total of 30 features for each

test. In the future we hope to combine DSFs with salient, measurable environmental factors such as temperature and humidity in the data fusion process.

## **3.5 Chapter Summary**

This chapter discussed the specific requirements wireless sensing places on the feature extraction process. A survey of the existing literature dealing with feature extraction for SHM was conducted, and a critical discussion of the existing methods' suitability for wireless applications was presented. On the basis of this discussion the Continuous Wavelet Transform, and the Autoregressive time series model were chosen as promising feature extraction algorithms. In data-based damage detection methodology feature extraction is followed by statistical pattern recognition. In the following chapter novel methods for statistical pattern recognition using support vector machines will be presented.



## Chapter 4

# Statistical Pattern Recognition

In a purely deterministic world the selection of an appropriate set of features would be all that is necessary for damage detection. However, even the most sophisticated features exhibit random variations, and thus, in the data-based damage detection paradigm, feature extraction is followed by statistical pattern recognition to detect and localise damage. For the vast majority of structural health monitoring applications, the statistical pattern recognition problem is fundamentally one of novelty detection. That is, distinguishing whether features extracted from a new signal are normal or abnormal based only on information from the normal condition of the structure. This is in contrast with many binary or multi-class classification problems presented in the literature [30], which train a classifier with data from the undamaged state of the structure, in addition to data from several damaged configurations, in order to detect and localise damage given new test data. This is a viable strategy in the case of a small, mass produced component, where it may be possible to acquire examples of data from damaged conditions by inducing cracks, notches and holes in a test component and undertaking extensive testing. This multi-class approach has several benefits over novelty detection, allowing for quantification of damage severity and type, as well as improving damage localisation. Unfortunately, in the vast majority of cases, particularly for the monitoring of existing large-scale infrastructure, it is difficult to imagine a scenario where data from the damaged structure would be available before it occurs. Therefore, despite the allure of the multi-class

approach, it is more practical and useful to concentrate on the one-class, or novelty detection approach to statistical pattern recognition. Much of the research in the field of data-based damage detection has employed Mahalanobis distance based approaches for one-class statistical pattern recognition. Worden et al [28] developed an outlier analysis procedure based on the Mahalanobis distance which was also used by Gul and Catbas [9]. Nair et al [18] hypothesis testing on the Mahalanobis distance as a means for damage detection. While these approaches can produce good results, the drawback is in the assumption that collections of features, or feature vectors, are normally distributed. This assumption may be reasonable in some cases, but factors such as the choice of features, or changing environmental conditions can cause feature vectors to have non-Gaussian or even multi-modal distributions, as noted by Farrar [6]. To combat this issue, more flexible methods such as Gaussian Mixture Models [16], or fully non-parametric methods such as k-means clustering [10], have been used for pattern recognition in SHM. Outside of the realm of SHM, nonparametric kernel methods for novelty detection have consistently shown excellent performance. In particular, the one-class support vector machine (OCSVM), developed by Scholkopf [20] has performed well on novelty detection for hyperspectral imagery problems [1], document classification [15], and seizure detection [8]. Das et al [4] used the OCSVM to classify different damage signatures in composite plates. The OCSVM is trained on a set of baseline feature vectors, from what is assumed to be a safe condition of the structure. The training process generates a nonlinear decision boundary, which can then be used to output a binary decision, when given new feature vectors, indicating whether they are from the normal condition of the structure, or from some anomalous condition. Lee and Verma [12], demonstrated that support vector machine methods can be implemented in a low power, low cost fashion on microcontroller devices. We believe that the one-class support vector machine can not only outperform existing statistical pattern recognition methods in SHM, but can do so in a low cost, embedded manner as part of a smart sensor network. This thesis presents a method for decentralized, automated structural damage detection using the one-class support vector machine for statistical pattern recognition.

## 4.1 One Class Support Vector Machine

Support vector machine methods have attracted some attention from researchers in SHM, but in the context of binary or multi-class classification, where data from the expected damaged state of the structure is available during the training stage. In contrast, the one-class support vector machine (OCSVM) requires only examples of data from the normal, baseline state of the structure to decide if a new test point is abnormal. As previously discussed, we cannot realistically expect to be able to generate examples of data from the expected damage scenarios for most civil structures, and thus the one-class approach is more practical. The OCSVM has a number of advantages over other one-class statistical pattern recognition methods presented in the SHM literature. It is fully nonparametric, enabling accurate description of datasets which do not adhere to a known probability distribution function. Additionally, once trained, the OCSVM decision function requires the retention of only a subset of the training data, and has been shown in other applications to be suitable for low-power embedded computation. The support vector machine's power lies in its ability to generate non-linear, flexible decision boundaries at low cost. This ability is due to the use of kernel functions, and is commonly referred to as 'the kernel trick'. Because the OCSVM is derived by a simple alteration to the traditional binary classification SVM (although with major implications for SHM), it is instructive to discuss briefly the kernel trick in the context of binary classification before proceeding to a discussion of the theoretical aspects of the OCSVM.

### 4.1.1 The Kernel Trick

Given examples of data from two different categories, the binary SVM transforms the data from its original dimensions,  $X$ , into a feature space,  $\phi(X)$  often nonlinear and higher dimensional, where the data can be better separated by a linear function. This concept is illustrated in Figure 4-1 below, where we wish to separate the purple data

from the green data. The original data has dimensions:

$$X_{\in \mathbb{R}^2} = (x_1, x_2) \quad (4.1)$$

where  $X$  is an Euclidean space with dimensions  $x_1$  and  $x_2$ . This data has an obvious pattern, but one which cannot be explained with a linear separator. The data is transformed to the feature space where the data can clearly now be separated much more easily.

$$\phi(X)_{\in \mathbb{R}^3} = (x_1^2, \sqrt{2} x_1 x_2, x_2^2) \quad (4.2)$$


where  $\phi(X)$  is now a three dimensional space. This data has an obvious pattern, but one which cannot be explained with a linear separator. The data is transformed to the feature space where the data can clearly now be separated much more easily.

Of course, this mapping to a higher dimensional feature space quickly becomes very expensive as the dimensions of the feature space increase. Significantly, the calculation of the decision boundary in the SVM algorithm is formulated in such a way that it requires the inner products in feature space, and not the explicit map,  $\phi(X)$ . This means that for any two training examples,  $x$  and  $y$ , we need to know the inner product in the feature space. For the example given in Figure 4-1, where  $x$  and  $y$  are both two dimensional vectors, their inner product in feature space is as follows: follows:


$$\begin{aligned} \langle \phi(x), \phi(y) \rangle &= x_1^2 y_1^2 + 2x_1 y_1 x_2 y_2 + x_2^2 y_2^2 \\ \langle \phi(x), \phi(y) \rangle &= (x_1 y_1 + x_2 y_2)^2 \\ K(x, y) &= \langle \phi(x), \phi(y) \rangle = (\langle x, y \rangle)^2 \end{aligned} \quad (4.3)$$

The inner product in the higher dimensional space is given simply by the square of the inner product in the original space. Clearly this is less expensive than performing calculations explicitly in the feature space. Functions such as  $K(x, y)$  in equation 4.1.1 are known as kernels. The ability to compute inner products in high dimensional space at low cost is extremely powerful and has led to the widespread use of kernel methods





(a) Data in original dimensions



(b) Data in higher dimensional feature space

Figure 4-1: Toy example illustrating the kernel trick

in machine learning. The usefulness of 'the kernel trick' becomes more apparent when we consider the basic task of many machine learning algorithms: That is, to calculate how similar a new data point is to an existing point whose label is known. Perhaps the most intuitive mathematical way to calculate similarity between vectors is to compute their inner product, and indeed many classifications algorithms, including the linear support vector machine, depend on the calculation of inner products. However, as in this example, similarity in Euclidean space does not always capture important patterns in the data. The kernel trick allows us to calculate inner products, at low cost in a nonlinear space, (a Reproducing Kernel Hilbert Space) where the patterns are more apparent. The above calculation is just one example of how inner products in higher dimensional feature spaces can be calculated with the kernel trick. When and how the kernel trick can be used, and what is a suitable kernel function, is a topic in its own right. One commonly used kernel, which we will use for the OCSVM, is the Gaussian kernel:

$$K(x, y) = \exp \left( - \frac{\|x - y\|^2}{\sigma^2} \right) \quad (4.4)$$

where  $x$  and  $y$  are vectors describing single training examples, and  $\sigma$  is a free parameter, which will be discussed later in Section 4.1.3. Now that we have described the kernel function and its use, we will provide the formulation of the one-class support vector machine.

### 4.1.2 Derivation of the One-Class Support Vector Machine

As mentioned previously, the binary version of SVM aims to separate two classes of training data by transforming them to a nonlinear feature space and finding a linear function which optimally separates them. The OCSVM, on the other hand, does not have any data for the second class, and so instead seeks to find a hyperplane which separates the training data from the origin, as a proxy for the second class. This is formulated by the following convex optimisation problem:

$$\min_{w \in \mathbb{R}^2, \xi \in \mathbb{R}^N, \rho \in \mathbb{R}} \frac{1}{2} \|w\|^2 + \frac{1}{\nu N} \sum_i \xi_i - \rho \quad (4.5)$$

subject to  $\mathbf{w} \cdot \phi(x_i) \geq \rho - \xi_i, \xi_i \geq 0$

where  $\nu$  is a bound on the fraction of training examples classified as outliers,  $N$  is the number of training examples,  $\xi_i$  are slack variables, and  $\mathbf{w} \cdot \phi(x_i) - \rho$  is the separating hyperplane in feature space. The decision function for a new test point  $x$  is then given by:

$$f(x) = \text{sgn}(\mathbf{w} \cdot \phi(x) - \rho) \quad (4.6)$$

Using Lagrangian multipliers, and the kernel trick, the dual problem can be formulated as:

$$\min_{\alpha} \frac{1}{2} \sum_{i,j} \alpha_j \alpha_i K(x_j, x_i) \quad (4.7)$$

subject to  $0 \leq \alpha_i \leq \frac{1}{N}, \sum \alpha_i = 1$

Most of the Lagrange multipliers,  $\alpha_i$ , in equation 4.7, will evaluate to zero. Any non-zero  $\alpha_i$  is called a support vector. Only the support vectors are required to evaluate the decision function for a new point  $x$ . Once the convex minimization problem in equation 4.7 has been solved, the OCSVM is ready to evaluate new test points. The binary decision function on new test points,  $x$  is given by:

$$f(x) = \text{sgn} \left( \sum_i^{\text{no. SVs}} \alpha_i K(x, x_i) - \rho \right) \quad (4.8)$$

where  $\rho$  can be recovered as:

$$\rho = \sum_j \alpha_j K(x_k, x_j), \alpha_j \in \left( 0, \frac{1}{\nu N} \right) \quad (4.9)$$

To illustrate the OCSVM, a typical dataset with decision boundary and support




Figure 4-2: Example of an OCSVM decision boundary

vectors is shown below in Figure 4-2. The evaluation of new test points depends only on the support vectors, which are a fraction of the original data used for training.

The main cost associated with the decision function is the evaluation of the kernel function between the new test point and each support vector. Clearly, if the number of support vectors is low, the decision function can be evaluated at low cost. The nature of this decision function suggests the potential for embedding the statistical pattern recognition phase of the damage detection methodology at the sensor level. Because of the sparsity, and low required computational effort, and based on successes in embedding the decision function for a two-class SVM on a low-power microcontroller [12], we believe this OCSVM decision function can be embedded at the sensor level for damage detection and localisation purposes.


### 4.1.3 Hyperparameter selection for the One-Class SVM using the Gaussian Kernel

The OCSVM contains two free parameters which must be selected by the user. The first parameter,  $\nu$ , controls the fraction of training examples which can be classified as outliers, and by extension, the false positive rate. Because  $\nu$  controls the false positive rate, typically a low value is desirable and this parameter is often set to 0.01. The second parameter, the Gaussian kernel parameter,  $\sigma$ , can dramatically affect the performance of the OCSVM. Small values of  $\sigma$  can lead to overfitting of the data, and thus, poor generalisation. Large values of  $\sigma$  underfit the data, resulting in an inability to detect non-trivial patterns. Typically free parameters are chosen by cross-validation: A fraction of the data is withheld from training, and the parameter which gives the best results in classifying this withheld data is selected. For the same reason that we are forced to use one-class methods, we cannot use cross-validation; namely, we do not have data to withhold which represents the damaged state of the structure.


In lieu of cross-validation an alternative approach to selecting the Gaussian kernel parameter,  $\sigma$  is required. Three potential methods for selecting this parameter are discussed in the following sections, and subsequently evaluated on acceleration data from an experimental steel frame laboratory structure later in this thesis.

#### Training Error Based Hyperparameter selection

Although true cross validation cannot be performed due to a lack of data from the damaged state, the following approach, similar to that used by Das [4] is conceptually related to cross-validation. At very small values of  $\sigma$  the OCSVM tends to overfit the data, essentially defining a decision boundary tightly surrounding each individual training point. This results in unseen points from the same distribution as the training data being classified as anomalous. As the value of  $\sigma$  increases the decision boundary begins to fit the data better, and unseen points from the same distribution will be classified as such. This is shown in Figure 4-3. Because  $\nu$  is a theoretical upper bound



(a) Example of overfitted data



(b) Example of well fitted data

Figure 4-3: Effect of hyperparameter selection on decision boundary




Figure 4-4: Example of hyperparameter selection based on the training error rate

on the fraction of normal data which will be classified as anomalous, as  $\sigma$  increases the decision boundary will eventually tend to classify more and more of the withheld data as part of the training distribution, eventually classifying only the exact fraction  $\nu$  of the withheld data as anomalous. Reasoning that we want to fit the data as closely as possible without overfitting, we can select the first value of  $\sigma$  at which the classification accuracy on the withheld data reaches  $1 - \nu$ . An example of this method is shown below in Figure 4-4.

### Geometry Based Hyperparameter selection

Motivated by the use of OCSVM for anomaly detection in hyperspectral imagery, Khazai et al [11] propose a method for automatic calculation of the optimal Gaussian kernel parameter,  $\sigma$ . The argument made in this paper is briefly summarised in this section. The Gaussian kernel, as shown in equation 4.4, can be viewed as a measure of how similar two data points are. Identical training points yield the maximum value,



1, of the Gaussian kernel. The smallest value of the kernel for the training examples will occur for the two least similar training points, i.e. those which are farthest apart:

$$\delta = \exp\left(\frac{-d_{max}^2}{\sigma}\right) \quad (4.10)$$

$$\text{where } d_{max} = \max_{i,j} \|x_i - x_j\|^2 \quad (4.11)$$

Presumably a new test point which is an outlier would be farther away than  $d_{max}$  from some of the training points. Therefore the kernel value for this new test point will fall between 0 and  $\delta$ . Using Equation 4.10 we can estimate  $\sigma$  by the following equation:

$$\sigma^2 = \frac{d_{max}}{\sqrt{-\ln(\delta)}} \quad (4.12)$$

The problem is to find an appropriate, small, value for  $\delta$ . Using a geometric argument Khazai et al [11] propose that an appropriate value for  $\delta$  can be estimated by:

$$\delta = \frac{1}{N(1 - \nu) + 1} \quad (4.13)$$

where  $N$  is the total number of training examples, and  $\nu$  is the fraction of training examples we permit to be outside the decision boundary. The optimal value of  $\sigma$  can then be calculated by combining equation 4.10, equation 4.12 and equation 4.13.

### Iterative Method for Hyperparameter selection

Wang et al [26] propose an iterative method which converges on a solution for the decision boundary which is judged neither too loose (underfitting) or too tight (overfitting). The basic algorithm is as follows:

1. Initialise by choosing an upper bound,  $\sigma_u$  and a lower bound  $\sigma_l$
2. Set  $\sigma = (\sigma_u + \sigma_l)/2$  i.e. binary search
3. Generate a decision boundary by training the OC-SVM using a pre-determined  $\nu$ , and the  $\sigma$  from step 2.
4. Evaluate whether the decision boundary is loose or tight.



5. If the boundary is tight set  $\sigma_l = \sigma$ , if loose set  $\sigma_u = \sigma$
6. Stop if  $\sigma_u - \sigma_l$  is less than a threshold value, or if the boundary is neither loose nor tight. Else return to 2.


Clearly a key part of this algorithm is the evaluation of whether the decision boundary is loose or tight in Step 3. As presented in [26] this determination is made like so: The data is split into two sets of pairs of training points. The first set is comprised of pairs of training points which are closer together than a certain distance  $D$ . The second contains pairs which are farther apart than  $D$ . The boundary is judged loose if for any midpoint of the second set of pairs, which is inside the decision boundary, there is no training point within a hypersphere of radius  $D$ : This physically means that there is a substantial hole inside the decision boundary devoid of any data. The boundary is judged to be tight, if any midpoint of the first set of pairs falls outside the decision boundary, indicating that the boundary has become concave between the two closely spaced points. This is illustrated below in Figure 4-5 where examples of both loose and tight boundaries are shown.

One drawback of this method, is that the presence of closely spaced outliers in the training examples can cause the algorithm to generate a boundary which is too loose. This is due to the fact that the midpoint of two data points, which are almost always outside the decision boundary generated by the OCSVM, will also almost always be outside the decision boundary. This will result in the algorithm presented in [26] judging the decision boundary to be tight at every iteration. As a remedy for this issue we suggest that the decision boundary be judged tight only if the midpoint of two training examples belonging to the first group of pairs, who themselves are within the decision boundary, falls outside the decision boundary.

## 4.2 Damage Detection Methodology


An overview of support vector machine methods, as well as the formulation of the one-class support vector machine, and three potential methods for OCSVM hyperparam-

Iterative method at a step where boundary is judged as loose



(a) Step in iterative algorithm where boundary is judged loose

Iterative method at a step where boundary is judged as tight



(b) Step in iterative algorithm where boundary is judged tight

Figure 4-5: Illustration of the iterative hyperparameter selection method

eter selection has been presented in the previous section. We will now outline an automated, decentralised damage detection methodology based on the OCSVM, which will later be experimentally verified. In the data-based damage detection methodology raw signals are acquired from the sensors and then processed into damage-sensitive features. Once a sufficient baseline dataset, describing the normal operating conditions of the structure has been accumulated, statistical pattern recognition can be used to compare new test data to the baseline. As we have discussed in this chapter, the OCSVM is a promising new method for SHM novelty detection which has a number of advantages over the popular approaches in the literature. As previously discussed, advances in MEMS sensor technology and wireless sensing present an opportunity to realise the ambition of 'smart sensing'. The goal of the methodology presented in this thesis is to develop the capability for automated sensors which collect and analyse data before transmission, extending battery life, eliminating bottlenecks in data transmission, and providing redundancy in the sensor network. We propose a decentralised approach to damage detection, where each sensor operates independently and locally, first acquiring data, then processing the data into a vector of damage sensitive features, and then finally returning a decision using the one-class support vector machine. A two phase approach is suggested to achieve this goal. Phase 1 consists of acquiring data and extracting features at the sensor level until a sufficient data base describing the normal operating conditions of the structure has been created. Once a sufficient database of features is available and has been transmitted to a base station, an OCSVM classifier is trained offline for each sensor location. The output of this training process, namely the coefficients,  $\alpha_i$ , the support vectors,  $x_i$ , the scalar offset term  $\rho$  and the kernel parameter  $\sigma$ , are then transmitted to the individual sensors. Evaluation of new test points can then be conducted in a fully automated, embedded manner, at the sensor level. A new test point at any sensor location is deemed to be damaged if the decision function at the sensor outputs a value of 1. The process of training the OCSVM requires the solution of a quadratic programming problem. While this is not especially computationally intensive for small to moderate sized data sets, and could be conducted on board the sensor, it is likely

that an evaluation period would be desired, to tweak the data acquisition and feature extraction processes. As the field of SHM matures and monitoring procedures become more standardised, it may be desirable to embed the SVM training stage on board sensor nodes. For now, the two stage approach provides an initial trial period to establish the baseline and train the OCSVM, followed by a fully automated, live test period where data is acquired, processed and evaluated using the OCSVM, yielding a single scalar decision on each new test, at each sensor location. This methodology is a sophisticated, automated approach to structural damage detection, incorporating state of the art nonparametric novelty detection by means of the OCSVM.

### 4.3 Damage Localisation

Once damage has been detected by the system, the next challenge is to identify where the damage has occurred. The OCSVM outputs a binary decision for each new test, at each sensor location, indicating whether the test is anomalous or not. As we will see later, these binary decisions do not typically provide any information on the location of the damage. In this section two methods for defining a localisation index (LI) will be presented. Operating within the decentralised, automated framework discussed above, the goal is to allow the output of a scalar localisation index, as well as a binary decision, which is likely to be highest at the sensor closest to the damage. This approach is similar in spirit to that proposed by Nair [18], who defined a localisation index based on the Euclidean distance between damaged and undamaged clouds of auto-regressive feature vectors at each sensor. Here, two localisation indices which are closely related to the OCSVM, and can therefore better account for the shape of the data than a Euclidean distance, are proposed. It should be noted that the efficacy of this approach is highly dependent on the sensitivity of the damage-sensitive features with proximity to damage.

### 4.3.1 OCSVM based decision function

A minor modification to the OCSVM decision function allows us to output a distance metric, rather than a binary decision. By simply dropping the sign function in front of equation 4.8, we obtain this distance:

$$f(x) = \sum_i^{\text{no. SVs}} \alpha_i K(x, x_i) - \rho \quad (4.14)$$

Inside the decision boundary, this function outputs positive numbers, with higher numbers being closest to the highest density area of the training data. As we move farther away from the training data, and outside the decision boundary, the distance metric becomes more strongly negative. This OCSVM based distance is a more accurate descriptor of proximity to the training data than an Euclidean distance from the mean of the training data, or a Mahalanobis distance which imposes an elliptical shape on the data. To define an intuitive localisation index we simply invert the sign in equation 4.14, so that large positive numbers indicate larger deviations from the training data. We therefore expect that the sensor location closest to the damage will output the largest value of the localisation index. If effective, this localisation index has an obvious benefit: It requires no additional analysis or memory on top of that already needed for the damage detection process.

### 4.3.2 One Class Regularised Least Squares based localisation index

The goal of the OCSVM is to define a decision boundary which best encloses a high percentage of the training data. Although the distance metric described above does provide a measure of proximity to the training data, this measure may be distorted by the requirement to enclose an exact fraction of the data within the decision boundary. A closely related algorithm, the Least-Squares SVM, developed by Choi [3], has been developed specifically to describe proximity to the training data, and has outperformed the OCSVM for relevance ranking tasks. This algorithm seeks to find

a function, in a feature space described by a kernel function, which minimises the quadratic error between the training data, and a hyperplane  $\rho - \mathbf{w} \cdot \phi(x)$ . This can be described by the following equation:

$$\min \sum_i (\rho - \mathbf{w} \cdot \phi(x_i))^2 \quad (4.15)$$

The solution of this minimisation problem will tend to overfit the data, and thus we introduce a penalisation term to regularise the minimisation problem. The weight we place on the penalisation is controlled by the parameter  $C$ .

$$\min \frac{1}{2} \|\mathbf{w}\|^2 - \rho + \frac{C}{2} \sum_i (\rho - \mathbf{w} \cdot \phi(x_i))^2 \quad (4.16)$$

Introducing Lagrange multipliers  $\alpha_i$ , and introducing  $\xi_i = \rho - \mathbf{w} \cdot \phi(x_i)$  we can rewrite equation 4.16 as follows:

$$L = \frac{1}{2} \|\mathbf{w}\|^2 - \rho + \frac{C}{2} \sum_i \xi_i^2 - \sum_i \alpha_i (\phi(\mathbf{x}) \cdot \mathbf{w} + \xi_i - \rho) \quad (4.17)$$

Taking the first derivative of equation 4.17 with respect to  $\mathbf{w}$ ,  $\xi_j$ ,  $\rho$  and  $\alpha_i$ , and setting them equal to zero, yields the following:

$$\begin{aligned} \mathbf{w} &= \sum_i \alpha_i \phi(\mathbf{x}_i) \\ \xi_i &= \frac{\alpha_i}{C} \\ \sum_i \alpha_i &= 1 \\ \phi(\mathbf{x}) \cdot \mathbf{w} + \xi_i - \rho &= 0 \end{aligned} \quad (4.18)$$

Combining these equations gives:

$$\sum_i \alpha_i \phi(\mathbf{x}_i) \cdot \phi(\mathbf{x}_i) - \rho + \alpha_i/C = 0 \quad (4.19)$$

Recalling that  $\phi(\mathbf{x}_i) \cdot \phi(\mathbf{x}_i) = \mathbf{K}$ , and using the fact that  $\sum_i \alpha_i = 1$  from equation 4.18 the problem reduces to the solution of the following linear system:

$$\begin{bmatrix} 0 & \mathbf{e}' \\ \mathbf{e} & \mathbf{K} + \mathbf{I}/C \end{bmatrix} = \begin{bmatrix} 1 \\ \mathbf{0} \end{bmatrix} \quad (4.20)$$

where  $\mathbf{e}$  is the vector of all ones,  $\mathbf{e}'$  denotes the transpose,  $\mathbf{I}$  is the identity matrix and  $\mathbf{K}$  is the kernel matrix. The solution of this system can be found using block matrix inversion and yields:

$$\rho = \frac{1}{\mathbf{e}' (\mathbf{K} + \mathbf{I}/C)^{-1} \mathbf{e}} \quad (4.21)$$

$$\boldsymbol{\alpha} = \frac{(\mathbf{K} + \mathbf{I}/C)^{-1} \mathbf{e}}{\mathbf{e}' (\mathbf{K} + \mathbf{I}/C)^{-1} \mathbf{e}} \quad (4.22)$$

This solution finds the optimal hyperplane  $\rho - \mathbf{w} \cdot \phi(x)$ , by minimising the objection function with respect to both  $\mathbf{w}$  and  $\rho$ , the constant bias term. In many cases, especially when the dimension of the data is large, the bias term is unlikely to be important. Therefore we can fix the bias term to a constant value, for example 1, and optimise with respect to  $\mathbf{w}$  only. This is the approach that we will take in this thesis to define a regularised least squares based localisation index. This approach simplifies the solution of the original minimisation problem to the single linear system:

$$\boldsymbol{\alpha} = (\mathbf{K} + \mathbf{I}/C)^{-1} \mathbf{e} \quad (4.23)$$

where the bias term has been replaced with  $\mathbf{e}$ , the vector of ones. The decision function for new points is then given by:

$$f(x) = \sum_i \alpha_i K(x, x_i) - 1 \quad (4.24)$$

For new points very similar to the training data, this function will approach a value of 1. As we move away from the training data, as is expected when damage is incurred, the value of the function drops towards 0. To make the localisation index more intuitive, we perform the following operation on the decision function shown in

equation 4.24:


$$LI(x) = -1(f(x) - 1) \quad (4.25)$$

The regularised least squares (RLS) decision function is of a very similar form to the OCSVM decision function, but with one crucial difference: The function is no longer sparse and now depends on all of the points in the training set. This, clearly has implications for memory use and computational efficiency. Nonetheless, the one-class RLS formulation has shown excellent performance for other applications, and is worth investigating. A comparison of the RLS based distance, and the OCSVM based distance is shown in Figure 4-6. Here we can see that the OCSVM, by virtue of the hard requirement to encapsulate a prespecified fraction of the training data with the 0 level contour, is more sensitive to outliers. Qualitatively, at least, the RLS based localisation index fits the shape of the data better. Whether any performance advantage the RLS approach may have is worth the increased memory requirement is unclear. A comparison between the two suggested localisation index algorithm's ability to localise damage is provided as part of the experimental results.

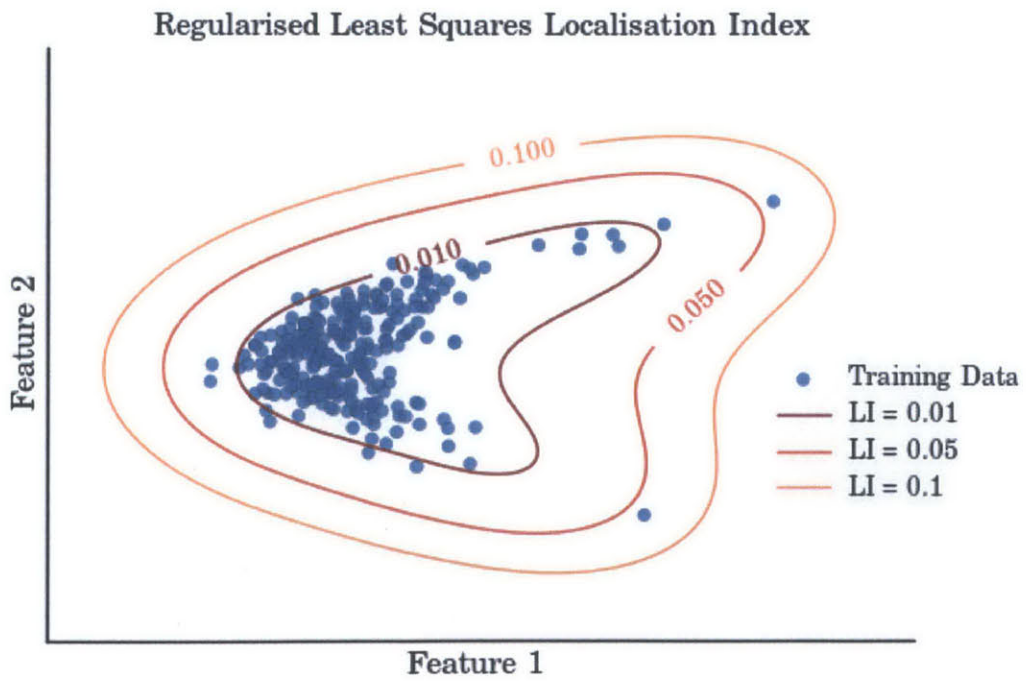
### 4.3.3 Chapter Summary

In this chapter we have discussed the need for statistical pattern recognition in data-based SHM. The advantage of the one class approach to statistical pattern recognition over multi-class methods in the context of SHM was outlined. A review of existing literature on this topic was presented, with a discussion of the limitations of these approaches. Based on success in other fields, the one-class support vector machine was proposed as a potential way to overcome these deficiencies. Motivated by the demands of wireless sensing, and a desire to perform analysis on board sensors prior to data transmission, an automated, decentralised methodology, incorporating the one-class support vector machine was described. Finally, the extension of this method to damage localisation was discussed, and two potential methods to achieve this goal were proposed. We will now proceed to a description of the extensive experimentation undertaken to evaluate and verify the methodology developed in this chapter.





(a) Contours of OCSVM based localisation index



(b) Contours of RLS based localisation index

Figure 4-6: Comparison of OCSVM and RLS Localisation Index contours




# Chapter 5

## Experimental Results

The efficacy of all methods discussed thus far in this thesis are evaluated based on data acquired from an experimental laboratory structure which has been developed specifically for the purposes of testing and validating SHM algorithms. The structure is a modular, 3 story, 2 bay steel frame structure with bolted connections. 18 tri-axial accelerometers measure the vibration response of the structure. The locations of these sensors are illustrated below in Figure 5-1. The structure is excited by a shaker, located in close proximity to sensor no. 18. This shaker inputs random, white Gaussian noise from a spectrum between 5 Hz and 350 Hz, to the structure.

As previously discussed, any data based damage detection methodology requires comparison with data from the baseline condition of the structure in order to detect damage. With this in mind, an initial database of 258 ten second samples was taken at a sampling rate of 3000 Hz at each sensor location. A test dataset was then acquired which includes data from the same, intact structure, as well as from 5 different damaged configurations of the structure. These scenarios are shown in Table 5.1. The structure is a steel frame, consisting of beam and column elements. These elements are connected with a bolted connection containing four bolts. Damage is induced at a sensor location by loosening the bolts at the connection nearest the sensor. Minor damage corresponds with loosening two of the four bolts, while for major damage all four bolts are loosened.



(a) 1a Diagram of sensor locations on lab structure



(b) 1b Photograph of lab structure

Figure 5-1: Schematic diagram and photograph of instrumented lab structure

Table 5.1: Experimental Damage Scenarios and Locations

Test Number	Damage Type	Damage Location
01-45	No Damage	-
46-75	2 bolts loosened at connection	Sensor 1
76-105	4 bolts loosened at connection	Sensor 1
106-135	2 bolts loosened at connection	Sensor 17
136-165	4 bolts loosened at connection	Sensor 17
165-195	4 bolts loosened at both connections	Sensor 17

## 5.1 Evaluating Damage Detection Performance

The detection of damage is typically regarded as the lowest level in the damage identification process, but it is perhaps the fundamental task of SHM. The ability to remotely and autonomously detect damage in a system using only a sensor network has tremendous value. In this section we will present results demonstrating the damage detection performance of the methodology presented in this thesis on the data shown in Table 5.1.

We will first quickly recap the damage detection methodology. In the first phase of the damage detection methodology, baseline data is acquired, and raw acceleration time series are processed into a 30 dimensional feature vector for every sample, at each sensor location. The feature vector is comprised of the first five wavelet coefficients for all three axes of the sensor, in addition to the five AR coefficients for each sensor axis, to give 30 total features for each individual test. Once a sufficient amount of training data has been acquired the OCSVM is then trained on a  $n \times 30$  matrix of training data where  $n$  is the number of training tests at an individual sensor location. We then enter the second, live testing phase of the methodology. New test data is acquired and processed into the same 30 dimensional feature vector before being inputted to the trained OCSVM which returns a binary decision indicating whether the test is anomalous or not.

To allow easy quantitative comparisons for parametric studies, we will define a metric for damage detection performance. Classification algorithms can make two types of errors. Type 1 errors, or false positives, occur when a test from the healthy condition is classified as damaged. Type 2 errors occur when tests from the damaged condition are classified as healthy. Both of these errors are important and need to be accounted for. Type 1 errors have economic implications, and also broader effects on how the damage detection system is viewed. If the Type 1 error rate is too high, we risk the SHM system being viewed as 'the boy who cried wolf'. The effect of type 2 errors is obvious. Because our system has 18 individual components acting independently we need to clarify exactly what the error rate is. We will define the

Type 1 error over all 18 individual sensor components, but the Type 2 error rate will be calculated only at the sensor location closest to the damage. The false positive rate in this experimental study is defined as the total fraction of tests from Scenario 1 (the undamaged test scenario), which were classified as damaged, at all sensor locations. The negative rate, or type 2 error rate, is defined as the fraction of tests at the sensor location closest to the damage, which were classified as healthy. So for Scenario 1, this quantity is only calculated for Sensor 1. We can use the F1 score as a metric which accounts for both the false positive and false negative rates:

$$F_1 = 2 \cdot \frac{\text{sensitivity} \cdot \text{specificity}}{\text{sensitivity} + \text{specificity}} \quad (5.1)$$

$$\text{where sensitivity} = \sum_{\text{damage scenarios}} \text{True positive rate at damage location}$$

$$\text{specificity} = \text{True negative rate over all sensors}$$

A high sensitivity score indicates the ability to correctly identify damage at the sensor location, while a high specificity score indicates the ability to correctly classify healthy tests. The F1 score is the harmonic mean of these two scores, and a higher F1 score indicates better performance of the methodology.

## 5.2 Hyperparameter Selection Method

Recall from Section 4.1.3, that the OCSVM has a number of free parameters which need to be specified. To encourage low false positive rates, the OCSVM parameter,  $\nu$  is set to 0.05 for all experiments. A detailed discussion on three potential methods for selecting the Gaussian kernel parameter  $\sigma$  was presented in Section 4.1.3. These three methods will now be compared using the F1 score as a performance metric. The full damage detection methodology was executed three times, each using a different hyperparameter selection method, and the results are shown in Figure 5-2. As we can see, there is almost no difference in performance between the three methods.




Figure 5-2: Comparison of damage detection performance using different hyperparameter selection methods

The geometric method is the most simple and inexpensive to calculate as it does not require any iteration, and therefore we will use this method for all other results in this thesis.

### 5.3 Damage Detection Results

Representative results of the damage detection performance are shown in this section for discussion. Figure 5-3 below shows the fraction of tests from both the major and minor damage scenario at sensor 17, which were classified as damaged for each sensor location. For comparison, results from the undamaged test scenario (Test 01-45) are plotted also to illustrate the true negative rate.




Figure 5-3: Detection of Damage at Sensor 17

The damage detection performance is excellent, for both major and minor damage scenarios, with 100% detection rates at the damage location, and false positive rates below 5% at every sensor location. We can also observe that the major damage scenario, which is at sensor 17, is detected by all sensor locations, while the minor



damage detection rate drops at sensor locations farther from the damage location. However, we can see that the simple damage detection process does not provide any real information on the location of the damage. Damage detection results all three hyperparameter selection methods, and for all damage scenarios are provided in the Appendix and show similar performance.

## 5.4 Damage Localisation Results

As we have seen in Figure 5-3, the binary damage detection methodology does not typically provide us with any information on the location of the damage. In Section 4.3 we introduced and discussed two potential methods for localising damage using a localisation index. To briefly recap, the first localisation index was defined by a simple alteration to the OCSVM, and has the advantage of not requiring any computation that has not already been performed for detection, as well as preserving a sparse decision function, meaning that only a fraction of the training data needs to be saved. We will refer to this method as the OCSVM LI method in this section.

The second method was derived from a regularised least squares formulation, and may provide a more accurate distance metric than the OCSVM LI method, with the downside that the decision function is not sparse, and the training and evaluation is all additional to that required for detection. This method will be referred to as the RLS LI method.

Localisation results from all damage scenarios are shown in the following sections for the purpose of discussion. These results are, as with the detection results, compared with data from the blind undamaged case. Both localisation index methods show very promising results. As we will show, the RLS LI outperforms the OCSVM LI, but the difference in performance is not dramatic.

### 5.4.1 Damage at Sensor 1

Below in Figure 5-4, the average value of the OCSVM LI for the minor and major damage scenario at Sensor 1 are shown. There are 30 individual tests in each damaged

scenario, so we average the values of the localisation index returned for each test at each sensor location and the plot them as a column chart. In the case of the minor damage the value of the LI is highest at Sensor 2, followed by sensor 4, 7, and 1. Referring to Figure 5-1, we can see that locations 2 and 4 are adjacent to the damage location at sensor 1 and thus the LI provides valuable information on the location of the damage. In the case of the major damage, the LI is highest at sensor 1, effectively localising the damage.




Figure 5-4: Localisation of Damage at Sensor 1 using OCSVM LI

For comparison, the values of the RLS LI for the same damage scenarios are shown in Figure 5-5. Here we see a similar pattern, with the largest value of the LI for the major damage being at the damage location. Again for the minor damage, the highest value of the localisation index is at Sensor 2. Qualitatively, it seems that the RLS LI is providing clearer localisation, as the LI values are comparatively higher near the damage locations; that is to say, the difference between the locations near the damage and farther from the damage is more distinct.




Figure 5-5: Localisation of Damage at Sensor 1 using RLS LI

#### 5.4.2 Damage at Sensor 17

The results of the damage localisation methodology, using both the OCSVM LI and the RLS LI, for the damage scenarios at Sensor 17 are shown in this section. In Figure 5-6, the average value of the OCSVM LI for both the major and minor damage scenarios at sensor 17, as well as the undamaged scenario are shown. In comparison to the OCSVM LI results for the damage scenario at Sensor 1, shown in Figure 5-4, the localisation is much less clear. While the value of the LI is high at Sensor 17, it is not distinctly higher than many other locations. For the major damage, Sensor 1, and Sensor 10, neither of which is adjacent to Sensor 17, show values just as high as Sensor 17. Similarly, for the minor damage scenario Sensor 11, and Sensor 2 show values of the LI comparable with that at Sensor 17. Although some information on the damage location is provided, the performance is less than ideal.




Figure 5-6: Localisation of Damage at Sensor 17 using OCSVM LI

In contrast, the localisation results for the Sensor 17 damage scenarios using the RLS LI, demonstrate much more effective performance. These results are shown in Figure 5-7. Here we see that, for the major damage scenario, the value of the LI is highest at Sensor 17. The next highest values are at Sensor 16 and Sensor 18, and all three are much higher than the other locations. The damage is unambiguously indicated to be in the vicinity of Sensor 17. The minor damage scenarios, too show much better performance than the OCSVM LI. Here, the values are highest at Sensor 17, Sensor 18 and Sensor 2, and all are much higher than the other locations. While Sensor 2 is not adjacent to the damage location, this is still a marked improvement over the results of the OCSVM LI method, shown in Figure 5-6.






Figure 5-7: Localisation of Damage at Sensor 1 using RLS LI

### 5.4.3 Combined Damage at Sensor 1 and Sensor 17

To create a damage scenario with multiple damage locations, four bolts were loosened at the connection nearest to both Sensor 1 and Sensor 17. Referring to Figure 5-1, we can see that these damage locations are on diagonally opposite corners of the 3 story 2 bay steel frame structure, making the localisation of these damages a difficult proposition. Plotted below in Figure 5-8, are the results of the OCSVM LI for this combined damage scenario, as well as the undamaged scenario for comparison. Here we see that the values of the localisation index are high at almost every sensor location, with Sensor 1 showing a slightly higher value than elsewhere. The localisation of damage is not clear.




Figure 5-8: Localisation of Combined Damage using OCSVM LI




Figure 5-9: Localisation of Damage at Sensor 1 using RLS LI

Figure 5-9 shows the localisation results using the RLS LI. Here, although the values are still relatively high at many locations, we can see a clear 'V' pattern showing that the LI values are highest as we approach Sensor 1 and Sensor 17. This, again, is a marked improvement over the OCSVM method.

In summary, localisation results for all damage scenarios, using both the RLS LI and the OCSVM LI have been shown and discussed. For both the damage scenarios at Sensor 17, and the multiple damage scenario, the RLS LI outperforms the OCSVM LI, providing clearer distinctions between the damage vicinity and other sensor locations than the OCSVM LI. To gain more insight into the damage detection and localisation process, and the nature of the data we are analysing, we will now proceed to show select visualisations of the data.

## 5.5 Data Visualisation

To visualise how the data is distributed, and the nature of the OCSVM decision boundary, we perform principle component analysis (PCA) on the 30 dimensional feature vectors, retaining the two largest principle components. We then fit a OCSVM to this 2 dimensional data to allow visualisation of the decision boundary.

### 5.5.1 Visualisation of Damage Scenarios at Sensor 1

Baseline training data for sensor 1, and the decision boundary generated by the OCSVM, are plotted below in Figure 5-10. Data from the undamaged test case, the minor and major damage scenarios at sensor 1 are also shown. This visualisation provides insight into how and why the damage detection and localisation process works. Referring to Figure 5-10, we can see that the baseline data and undamaged test data are in good agreement, clearly coming from the same underlying system, leading to low false positive rates in damage detection. In addition we see that both major and minor damage are well detected by the decision boundary, and that the major damage is clearly much further from the baseline data.

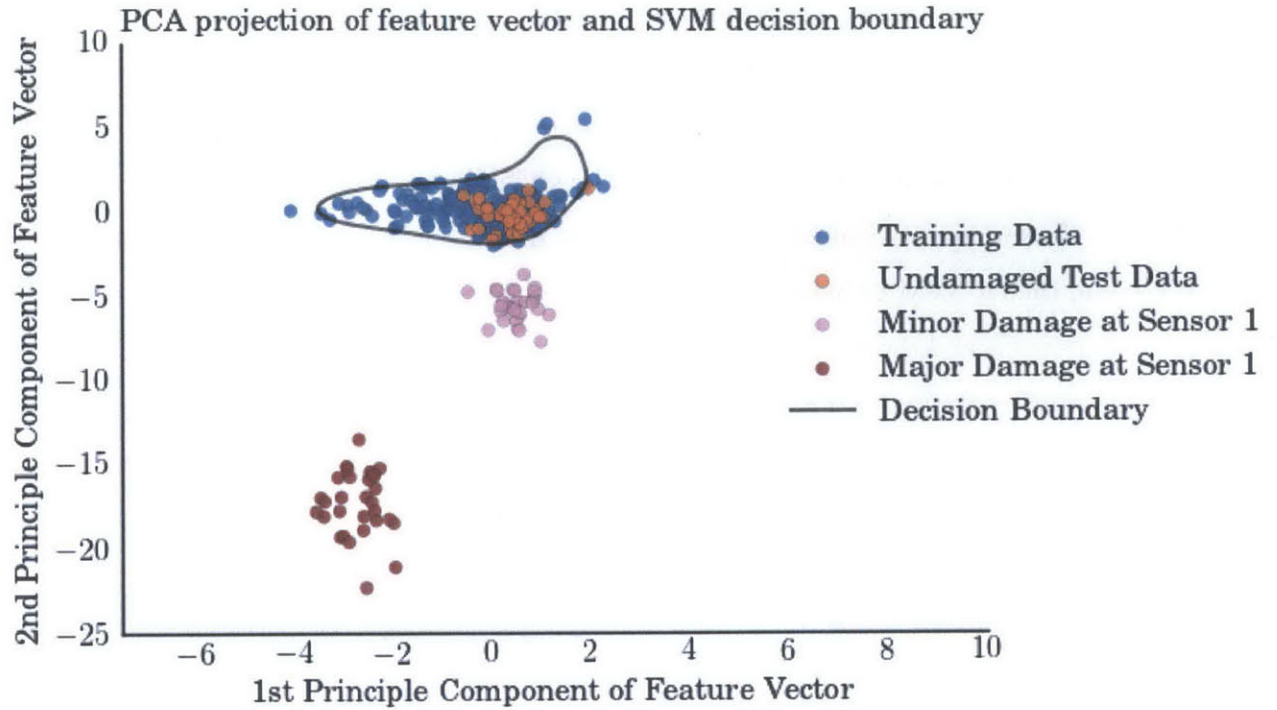


Figure 5-10: 2D visualisation of data and decision boundary from Sensor 1, for the damage scenario at Sensor 1

Data from sensor 4 is plotted in Figure 5-11 below. This data, from a sensor farther from the damage location, allows us to understand how damage can be localised. In contrast to the data from sensor 1, we can see that the major damage scenario data is much closer to the baseline data. This leads to a higher value of the LI at sensor 1, and clear localisation of damage for the major damage scenario. Figure 5-11 also illustrates some of the advantages the OCSVM method demonstrated in this paper has over previously reported SHM pattern recognition algorithms. The data from sensor 4 is clearly asymmetric and non-Gaussian, exhibiting a long tail to the negative x direction, and a steep drop off in density to the positive y direction from the region of highest probability density. The data from the damaged scenarios lies quite close to the baseline data in a Euclidean space, but clearly in a region of low probability density. The OCSVM, because it is a nonparametric kernel method, allows us to identify these patterns, and correctly classify the data from the damaged scenarios.



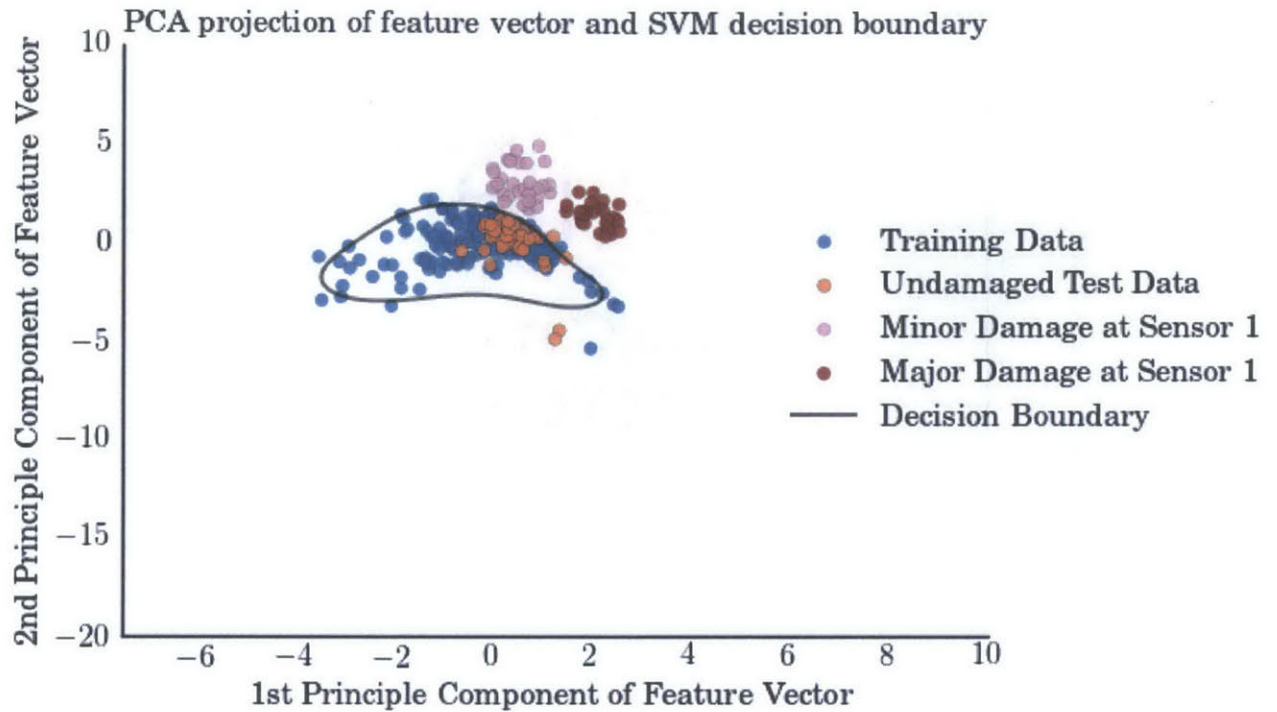


Figure 5-11: 2D visualisation of data and decision boundary from Sensor 4, for the damage scenario at Sensor 1

### 5.5.2 Visualisation of Damage Scenarios at Sensor 17

Data from sensor 17 is plotted below, for the baseline, undamaged test, and major and minor damage scenarios at sensor 17. As we would expect at the location of the damage, the data corresponding to the major damage has migrated significantly from the baseline data. Compared to the data from Sensor 1 in the previous section, the minor damage data is much closer to the baseline, perhaps explaining why the performance of the localisation algorithm was better on the damage scenarios at Sensor 1.




Figure 5-12: 2D visualisation of data and decision boundary from Sensor 17, for the damage scenario at Sensor 17

In order to illustrate the concept on which the localisation index method is built, plotted below is data from Sensor 9, containing the baseline data as well as data corresponding to the damage scenarios at Sensor 17. If the localisation algorithm is to work, the data from the damaged scenarios must be relatively closer to the baseline data than it was in Figure 5-12 above. Here we see that this in fact the case, and both the minor and major damage scenario data is much closer to the baseline data for the Sensor 9 data plotted below. Again, as with Figure 5-11 we can see the benefits of the OCSVM methodology here, as the major damage scenarios are only detected by virtue of the decision boundary accounting for the asymmetric nature of the baseline data which has a long tail in the positive y direction, but a very sharp drop in density in the negative y direction. The data from the major damage scenario is located in close proximity to the baseline data, but because the OCSVM accounts for the steep decline in baseline density in this region, a reasonable fraction of the damage points are classified correctly.




Figure 5-13: 2D visualisation of data and decision boundary from Sensor 9, for the damage scenario at Sensor 17

## 5.6 Sensitivity to Sample Size

Thus far, all results presented have been generated using a baseline training dataset of 258 samples. In this section we will investigate what effect, if any, the size of the training dataset has on the damage detection performance. The importance of this is clear if we recall the damage detection strategy outlined in Section 4.2. Briefly recapping this method, the first phase of the monitoring program consists of acquiring a sufficient database of feature vectors from the baseline condition of the structure. Once this database is acquired, the OCSVM is trained offline, and the results of the training communicated to the sensors via radio. The monitoring program then enters the second phase, where all processing is conducted on board the sensor, and data is only transmitted in the event of the detection of damage. Clearly, the sooner the monitoring program can enter the live phase, the sooner we can start providing real time information on the health of the structure, and the more cost



effective it will be. To investigate what sample size we require before reliable damage detection results are obtained, the damage detection algorithm was executed using training sample sizes varying from 1 to 258. Additionally, to investigate if the effect of sample sizes depended on the dimension of the feature vector, this analysis was carried out for three different feature vector types: a 15 dimensional feature vector of AR features, a 30 dimensional feature vector of combined AR and Wavelet features, and a 45 dimensional feature vector, again of combined AR and Wavelet features. The damage detection performance was evaluated using the F1 metric, as described in Section 5.1. The results of this analysis are shown in Figure 5-14



Figure 5-14: Effect of sample size on damage detection performance

We can see that at sample size above approximately 50, the performance plateaus and is very stable. Once the sample size goes below 50 the performance falls off a cliff. The effect of feature vector dimension is not pronounced, but somewhat expectedly, the 45 dimensional features require a slightly bigger sample size before reaching a stable, reliable damage detection performance level.

## 5.7 Effect of Data Acquisition Parameters

In Section 1.4.1 we reviewed some of the most popular types of accelerometers. Regardless of exactly what accelerometer technology is chosen, a number of decisions on how they will operate are required. Two important decisions, with respect to the efficiency of the sensor, are the sampling frequency, and the sampling resolution. When a signal moves from the analog domain to the digital domain it must be discretised, and quantised. How closely spaced the discretisation is controls the sampling frequency, while how many numbers are available to represent the signal in the quantisation process controls the sampling resolution. Both lower sampling frequency and lower sampling resolution will lead to lower power devices, but potentially with a loss in performance. Thus far, all results presented in this Chapter were obtained from a signal sampled at 3000 Hz with a sampling resolution of 24 bits. A parametric study on the effect of varying both the sampling frequency and resolution on damage detection performance is presented in this section.

### 5.7.1 Sampling Frequency

The rate at which the acceleration signal is sampled dictates the range of frequencies which we can accurately analyse. The well known Shannon-Nyquist theorem states that to avoid loss of information in the sampling process, the sampling rate must be twice that of the highest frequency contained in the signal. Therefore, if there is frequency content at 1000 Hz which is of importance to the damage detection process, we must sample at at least 2000 Hz in order to avoid losing some information about this content. The implications of the sampling rate on damage detection performance are clear: If our damage sensitive features depend on frequency content higher than half the sampling rate, there will be a degradation in performance.

### 5.7.2 Sampling Resolution

There are a finite number of values which can be used to represent signals in the digital domain. The bit rate dictates exactly how many integer numbers are available

to represent the signal: If the sampling resolution is 24 bits, there are  $2^{24}$  integers available. The combination of sampling resolution and dynamic range controls the sensitivity of the sensor. For example if the range of values the sensor can measure is between 0 and 1, and the resolution is 4 bit, then the sensitivity is  $1/2^4 = 1/16$ . In effect this means that all values measured by the sensor are rounded to the nearest  $1/16^{th}$ . To illustrate this concept a sine wave which varies between 0 and 1 is shown in Figure 5-15. Overlaid on the 'continuous' sine wave, is the 4 bit representation.




Figure 5-15: Illustration of data quantisation

To investigate the effect of the resolution rate and sampling frequency on damage detection performance, the original 3000 Hz, 24 bit data was resampled and requantised. First, the 3000 Hz data was downsampled to 2000, 1500, 1000, 750, 600, 500, 400, 300, 200, 150, 120 and 100 Hz. All of these downsampled datasets were then requantised from 24 bit resolution to 20, 16, 12 and 8 bits. The damage detection algorithm was then executed on all combinations of downsampled and requantised data. The F1 score was used to evaluate damage detection performance. Results are

shown in Figure 5-16.

**F1 classification score under varying sampling rates and resolutions**




Figure 5-16: Effect of bit rate and sampling frequency on damage detection performance

Figure 5-16 shows a heat map where the colour represents the F1 damage detection score, as shown in the colour bar to the right of the figure. The effect of sampling frequency on the damage detection performance is clear. Damage detection performance appears to degrade significantly as the sampling rate drops. Performance degrades slowly from 3000 Hz to 1500 Hz, before dropping precipitously. This is not very surprising, and demonstrates that content above 750 Hz must be providing information on the occurrence of damage. Much more surprisingly, the bit rate does not appear to have any significant effect on the ability to detect damage. We see no real difference in the damage detection performance as the resolution drops from 24 bit to 8 bit. This is an interesting result, and suggests that the 24 bit resolution provided by the accelerometer is not necessary, and may be consuming power and memory without providing any advantage. To investigate further, an example of real data from our experimental structure is plotted below in both 24 bit and 8 bit.






Figure 5-17: Comparison between 24 bit data and 8 bit data

Visually, the 8 bit and 24 bit data are similar and the calculated difference between the two signals is small, although the effect of decreased sensitivity can be seen in the 8 bit data sets where some high frequency content appears more 'square'. This effect is highlighted in the zoomed inset to the right hand side of the figure. To investigate deeper, below the Fourier Transform of both signals is plotted. Here we



can see that the difference in the frequency domain is negligible, and has the sole effect of a marginal increase in noise.

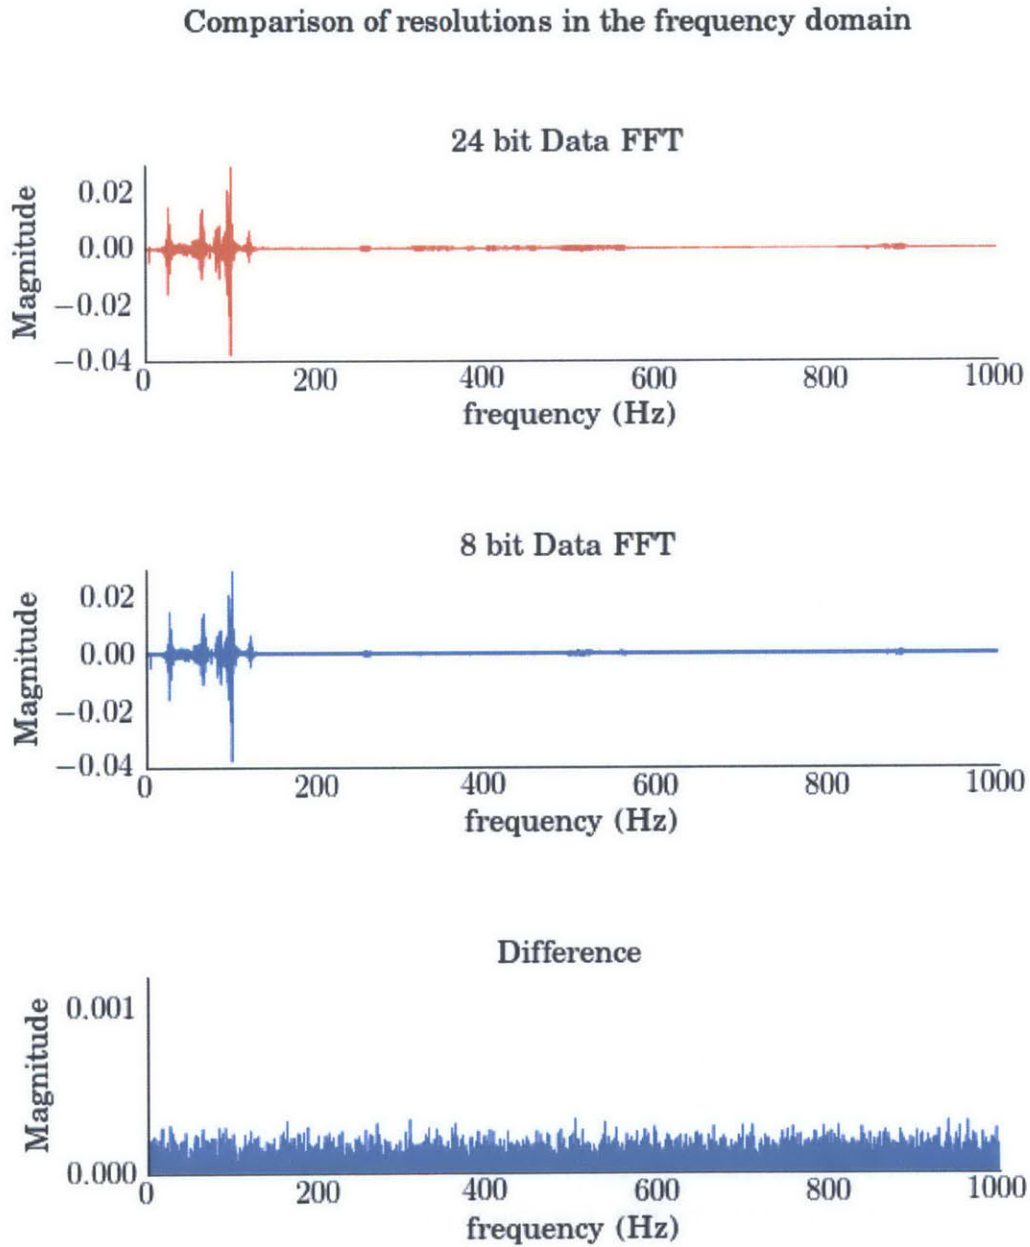


Figure 5-18: Comparison between 24 bit data and 8 bit data in the frequency domain

The dynamic range of the accelerometers used in this study is  $\pm 10g$ . The maximum value encountered in the analysed data sets is  $3.79g$ , while the minimum is

$-3.24g$  (unsurprisingly both at sensor 18, closest to the shaker and farthest from the ground). Given this poor match between the dynamic range of the sensors, and the measured behaviour of the laboratory structure, it is even more surprising that reducing the resolution from 24 bit to 8 bit does not degrade damage detection. Geometry and operating conditions play an important role in the range of accelerations an accelerometer will experience. In the lab structure described in this chapter, the accelerations experienced at the first floor level, for example at Sensor 1, are much lower than those experience at the highest level. Potentially a lower dynamic range could be used for the sensors which experience less intense accelerations, while preserving the same sensitivity. Whether this level of customisation is economically advantageous is less clear, but perhaps warrants further investigation.

### 5.7.3 Chapter Summary

A detailed description of the experimental procedures and testing regime used for verification and analysis of the SHM algorithms described in this thesis has been provided in this chapter. The performance of three potential methods for automatic selection of one-class support vector machine hyperparameters were compared using data acquired from the modular steel frame laboratory structure described. Representative damage detection results were presented from the most promising of these methods. Observing that the binary damage detection output does not typically provide any useful information for the localisation of damage, the two localisation indices proposed in Chapter 4 were implemented and tested on the same data set and their performance compared. Noting that in a real-life SHM application, the size of the training data set is an important consideration, an investigation into the effect this has on damage detection performance was conducted. Additionally, two sensor system parameters, sampling frequency and sampling resolution, which have important effects on the efficiency of the sensors, were investigated. The original data sets were downsampled and requantised, and analysis of the effect of this degradation in data quality was conducted. We will now proceed to present the conclusions arrived at through the experimental work conducted in this thesis.

# Chapter 6

## Conclusions

A comprehensive, fully automated, decentralised damage detection and localisation strategy for structural health monitoring has been presented in this thesis. Raw acceleration data obtained from sensors placed on a structure is first processed into damage-sensitive features and then state of the art machine learning algorithms are applied to compare these features with baseline values from the intact structure. The efficacy of the algorithms has been demonstrated using data taken from a steel frame laboratory structure, under various different damage scenarios. In light of recent trends towards smart sensing networks which are capable of executing damage detection algorithms on board sensors, careful consideration of computational requirements has been given. A review of the SHM literature was conducted in order to select suitable feature extraction methods for smart sensing. Autoregressive and wavelet based features were chosen for feature extraction in this thesis.

Once the acceleration data has been processed into these damage-sensitive features damage detection is accomplished using the one-class support vector machine, which returns a binary value indicating whether the structure is damaged or not. The one-class support vector machine has numerous advantages over previous approaches in the SHM literature, and has been shown in other fields to be suitable for embedding on a microcontroller. The OCSVM requires tuning by means of a free hyperparameter. Three potential methods for automated tuning of this parameter have been discussed and evaluated in this thesis. All three perform adequately, but the geometry-based

method developed by Khazai et al [11] requires less computation, and is therefore preferable in this case. For all damage scenarios contained in the test data from the steel laboratory structure, damage is detected extremely reliably, with a low false positive rate.

The effect of the sample size of baseline data was also investigated. 258 samples from the intact laboratory structure were initially collected and used to train the OCSVM. To investigate whether damage detection performance degrades when this sample size is smaller, the damage detection algorithm was carried out on various smaller training data sets. Performance was stable above a sample size of 50, but quickly deteriorated below this number. Additional parametric studies were carried out to investigate the effect of sampling frequency and resolution. The initial data was acquired at a sampling frequency of 3000 Hz, and a 24 bit resolution. This data was downsampled and requantised to lower frequencies and resolutions. As the sampling frequency of the data decreased, the damage detection performance deteriorated significantly, but a reduction from 24 bit to 8 bit resolution had no effect.

A binary decision is returned at each location by the OCSVM. This typically does not provide any information on the location of damage, and therefore two different algorithms for damage localisation have been developed and presented in this thesis. These methods return a localisation index, instead of a binary decision, at each sensor location. The first method is based on a slight modification to the one-class support vector machine, the second on a one-class regularised least squares formulation. These methods were also evaluated using the experimental data from the 3 story, 2 bay steel frame structure. Both methods showed promising results, but the regularised least squares method outperformed the OCSVM based method.

## 6.1 Future Work

Future work should address the issue of how to rigorously account for environmental and operational variations. An additional data normalisation procedure is required

to alleviate the effect of varying excitation in a real world application of the algorithms presented in this thesis. To account for variations in environment, data from additional temperature and humidity sensors should be incorporated as part of an integrated SHM sensor network. The OCSVM algorithm presented in this thesis provides sufficient flexibility to adopt a data fusion approach, and include salient environmental measurements as part of the feature vector.

To realise the goal of embedded smart sensing, the algorithms presented in this paper should be implemented on board a microcontroller, and empirical studies of the energy saved by performing the full damage detection process prior to data transmission conducted to verify the benefits of this strategy. Based on the results of the parametric studies carried out on sampling resolution and frequency, further investigation into lowering the resolution of the sensors should be conducted. The methodology presented in this thesis uses a batch machine learning strategy, where the full database of features from the baseline structure is acquired, and then the OCSVM is trained on a once off basis. Future work on the development of this methodology should consider the use of an online strategy, where the classifier can learn incrementally as each new baseline test is added to the database.

As the field of SHM moves more and more towards wireless application, the identification and development of suitable DSFs for embedded computation is a key concern. The test facility described in this paper can be used, in conjunction with a comprehensive software library of feature extraction methods to investigate the performance of different feature extraction methods in identifying different types of structural damage. The computational effort required can then be weighed to choose features for embedding in a smart-sensing system. If SHM technology is to progress and mature into a useful tool for many different types of structures, a greater understanding of how data based approaches should be tailored to the structure type and its operating environment, is required.

In conclusion, an automated, decentralised data-based damage detection and localisation methodology has been developed in this thesis. This methodology uses machine learning algorithms for statistical pattern recognition which have a num-

ber of advantages over previous approaches. The efficacy of these methods has been evaluated on real experimental data from a steel frame structure instrumented with accelerometers and induced with several different damage scenarios. The algorithms detect damage reliably, provide useful information on the location of damage and have a very low false alarm rate. Suggestions on the future development of this work have been provided, in order to realise the goal of a low cost, flexible, automated, intelligent SHM system.

# Appendix A

## Damage Detection Results

Full damage detection results using all three hyperparameter selection methods are presented in this appendix. For each hyperparameter selection methods, results of the binary damage detection strategy are shown for all test scenarios (undamaged, minor damage at sensor 1, major damage at sensor 1, minor damage at sensor 17, major damage at sensor 17, combined major damage at sensor 1 and at sensor 17. These results are presented as column charts, where the x axis shows the sensor locations, and the y axis represents the fraction of tests for the given damage scenario which are classified as damaged.




Figure A-1: Damage detection for sensor 1 damage scenarios using geometric hyperparameter selection.






Figure A-2: Damage detection for sensor 17 damage scenarios using geometric hyperparameter selection.




Figure A-3: Damage detection for multiple location damage scenarios using geometric hyperparameter selection.




Figure A-4: Damage detection for sensor 1 damage scenarios using iterative hyperparameter selection.




Figure A-5: Damage detection for sensor 17 damage scenarios using iterative hyperparameter selection.




Figure A-6: Damage detection for multiple location damage scenarios using iterative hyperparameter selection.




Figure A-7: Damage detection for sensor 1 damage scenarios using training error based hyperparameter selection.






Figure A-8: Damage detection for sensor 17 damage scenarios using training error based hyperparameter selection.




Figure A-9: Damage detection for multiple location damage scenarios using training error based hyperparameter selection.



# Bibliography

- [1] Amit Banerjee, Philippe Burlina, and Chris Diehl. A support vector method for anomaly detection in hyperspectral imagery. *Geoscience and Remote Sensing, IEEE Transactions on*, 44(8):2282–2291, 2006.
- [2] Rune Brincker, Lingmi Zhang, and Palle Andersen. Modal identification of output-only systems using frequency domain decomposition. *Smart materials and structures*, 10(3):441, 2001.
- [3] Young-Sik Choi. Least squares one-class support vector machine. *Pattern Recognition Letters*, 30(13):1236–1240, 2009.
- [4] Santanu Das, Ashok N Srivastava, and Aditi Chattopadhyay. Classification of damage signatures in composite plates using one-class svms. In *Aerospace Conference, 2007 IEEE*, pages 1–19. IEEE, 2007.
- [5] Charles R Farrar, Thomas A Duffey, Scott W Doebling, and David A Nix. A statistical pattern recognition paradigm for vibration-based structural health monitoring. *Structural Health Monitoring*, 2000:764–773, 1999.
- [6] Charles R Farrar and Keith Worden. *Structural health monitoring: A machine learning perspective*. John Wiley & Sons, 2012.
- [7] Michael I Friswell. Damage identification using inverse methods. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 365(1851):393–410, 2007.
- [8] Andrew B Gardner, Abba M Krieger, George Vachtsevanos, and Brian Litt. One-class novelty detection for seizure analysis from intracranial eeg. *The Journal of Machine Learning Research*, 7:1025–1044, 2006.
- [9] Mustafa Gul and F Necati Catbas. Statistical pattern recognition for structural health monitoring using time series modeling: Theory and experimental verifications. *Mechanical Systems and Signal Processing*, 23(7):2192–2204, 2009.
- [10] Krishnan Nair Kesavan and Anne S Kiremidjian. A wavelet-based damage diagnosis algorithm using principal component analysis. *Structural Control and Health Monitoring*, 19(8):672–685, 2012.

- [11] Safa Khazai, Saeid Homayouni, Abdolreza Safari, and Barat Mojaradi. Anomaly detection in hyperspectral images based on an adaptive support vector method. *Geoscience and Remote Sensing Letters, IEEE*, 8(4):646–650, 2011.
- [12] Kyong Ho Lee and Naveen Verma. A low-power processor with configurable embedded machine-learning accelerators for high-order and adaptive analysis of medical-sensor signals. 2013.
- [13] SJ Loutridis. Damage detection in gear systems using empirical mode decomposition. *Engineering Structures*, 26(12):1833–1841, 2004.
- [14] Jerome Peter Lynch, Arvind Sundararajan, Kincho H Law, Anne S Kiremidjian, and Ed Carryer. Embedding damage detection algorithms in a wireless sensing unit for operational power efficiency. *Smart Materials and Structures*, 13(4):800, 2004.
- [15] Larry M Manevitz and Malik Yousef. One-class svms for document classification. *the Journal of machine Learning research*, 2:139–154, 2002.
- [16] K Krishnan Nair and Anne S Kiremidjian. Time series based structural damage detection algorithm using gaussian mixtures modeling. *Journal of Dynamic Systems, Measurement, and Control*, 129(3):285–293, 2007.
- [17] K Krishnan Nair and Anne S Kiremidjian. Derivation of a damage sensitive feature using the haar wavelet transform. *Journal of Applied Mechanics*, 76(6):061015, 2009.
- [18] K Krishnan Nair, Anne S Kiremidjian, and Kincho H Law. Time series-based damage detection and localization algorithm with application to the asce benchmark structure. *Journal of Sound and Vibration*, 291(1):349–368, 2006.
- [19] AK Pandey, M Biswas, and MM Samman. Damage detection from changes in curvature mode shapes. *Journal of sound and vibration*, 145(2):321–332, 1991.
- [20] Bernhard Schölkopf, John C Platt, John Shawe-Taylor, Alex J Smola, and Robert C Williamson. Estimating the support of a high-dimensional distribution. *Neural computation*, 13(7):1443–1471, 2001.
- [21] Hoon Sohn and Charles R Farrar. Damage diagnosis using time series analysis of vibration signals. *Smart materials and structures*, 10(3):446, 2001.
- [22] Wiesław J Staszewski and Amy N Robertson. Time–frequency and time–scale analyses for structural health monitoring. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 365(1851):449–477, 2007.
- [23] MD Todd, JM Nichols, LM Pecora, and LN Virgin. Vibration-based damage assessment utilizing state space geometry changes: local attractor variance ratio. *Smart Materials and Structures*, 10(5):1000, 2001.

- [24] Peter Van Overschee and Bart De Moor. Subspace identification for linear systems: theory, implementation, applications. *status: published*, 1996.
- [25] J Kim Vandiver et al. Detection of structural failure on fixed platforms by measurement of dynamic response. In *Offshore Technology Conference*. Offshore Technology Conference, 1975.
- [26] Huangang Wang, Lin Zhang, Yingchao Xiao, and Wenli Xu. An approach to choosing gaussian kernel parameter for one-class svms via tightness detecting. In *Intelligent Human-Machine Systems and Cybernetics (IHMSC), 2012 4th International Conference on*, volume 2, pages 318–323. IEEE, 2012.
- [27] Keith Worden, Charles R Farrar, Graeme Manson, and Gyuhae Park. The fundamental axioms of structural health monitoring. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science*, 463(2082):1639–1664, 2007.
- [28] Keith Worden, Graeme Manson, and NRJ Fieller. Damage detection using outlier analysis. *Journal of Sound and Vibration*, 229(3):647–667, 2000.
- [29] YL Xu and J Chen. Structural damage detection using empirical mode decomposition: experimental investigation. *Journal of engineering mechanics*, 130(11):1279–1288, 2004.
- [30] Yujie Ying, James H Garrett Jr, Irving J Oppenheim, Lucio Soibelman, Joel B Harley, Jun Shi, and Yuanwei Jin. Toward data-driven structural health monitoring: Application of machine learning and signal processing to damage detection. *Journal of Computing in Civil Engineering*, 27(6):667–680, 2012.
- [31] MMF Yuen. A numerical study of the eigenparameters of a damaged cantilever. *Journal of sound and vibration*, 103(3):301–310, 1985.