On Model- and Data-based Approaches to Structural Health Monitoring

Robert James Barthorpe

A thesis submitted to the University of Sheffield for the degree of Doctor of Philosophy in the Faculty of Engineering

Department of Mechanical Engineering

November 2010

Abstract

Structural Heath Monitoring (SHM) is the term applied to the process of periodically monitoring the state of a structural system with the aim of diagnosing damage in the structure. Over the course of the past several decades there has been ongoing interest in approaches to the problem of SHM. This attention has been sustained by the belief that SHM will allow substantial economic and life-safety benefits to be realised across a wide range of applications. Several numerical and laboratory implementations have been successfully demonstrated. However, despite this research effort, real-world applications of SHM as originally envisaged are somewhat rare. Numerous technical barriers to the broader application of SHM methods have been identified, namely: severe restrictions on the availability of damaged-state data in real-world scenarios; difficulties associated with the numerical modelling of physical systems; and limited understanding of the physical effect of system inputs (including environmental and operational loads). This thesis focuses on the roles of law-based and data-based modelling in current applications of. First, established approaches to model-based SHM are introduced, with the aid of an exemplar 'wingbox' structure. The study highlights the degree of difficulty associated with applying model-updating-based methods and with producing numerical models capable of accurately predicting changes in structural response due to damage. These difficulties motivate the investigation of non-deterministic, predictive modelling of structural responses taking into account both experimental and modelling uncertainties. Secondly, a data-based approach to multiple-site damage location is introduced, which may allow the quantity of experimental data required for classifier training to be drastically reduced. A conclusion of the above research is the identification of hybrid approaches, in which a forward-mode law-based model informs a data-based damage identification scheme, as an area for future work.

Acknowledgments

First of all I would like express my gratitude to my supervisors, Keith Worden & Graeme Manson, for invaluable technical and moral support throughout the course of the work culminating in this thesis. They may not at all times have been quite aware of the directions the work was taking, but I couldn't imagine a more friendly and supportive environment in which to conduct a PhD.

Secondly, I would like to thank all the friends, housemates, and members of the research group, past & present, who have contributed in ways both great and small to maintaining the sanity of the author over the past several years. Research, like life, offers up a mish-mash of colourful surprises, and perhaps without always realising it you've all played a part in guiding me through the highs and the lows and giving me the inspiration to keep on going. Thank you!

For the work presented here, particular thanks goes to Vaggelis Papatheou for support, entertainment and technical discussion throughout, and to Daryl Hickey for his help and advice in the experimental aspects of the work.

I'd also like to extend my gratitude to Chuck Farrar, Francois Hemez and Gyuhai Park at Los Alamos National Laboratory, New Mexico, and Michael Link and Matthias Weiland at the University of Kassel, Germany, for their help and hospitality during visits to their respective institutions. Your help has been invaluable in conducting the work described.

Finally, I'd like to express my deep gratitude to my parents and brother for their dedication and support over so many years. I know I don't say this enough – thank you!

Table of Contents

1 Introduction

1.1 Structural Health Monitoring

Structural Health Monitoring (SHM) refers to a process of measuring and interpreting data from a system of sensors distributed about a structural system in order to objectively quantify the condition of the structure. The broad aim of an SHM system is to be able to identify, at an early stage, occurrences of damage that may ultimately lead to failure of the individual component or system. The damage identification outcomes made through monitoring may subsequently be used to inform decisions on remedial work.

SHM is distinct from the field of *Non-Destructive Evaluation* (NDE) in that it uses permanently installed sensors, is a 'global' technique that makes use of indirect measurements to assess structural state, and may be employed in an on-line mode. NDE methods, for example ultrasound or eddy current inspection, are inherently 'local' in nature and typically require the structure to taken off-line in order for inspection to be conducted. Additionally, inspection can be limited by issues of inaccessibility, and is typically reliant upon human expert interpretation of measured data.

SHM has the potential to offer enormous benefits for aerospace and civil structures in applications as diverse as aircraft, bridges and offshore structures: at the time of writing the advent of large-scale offshore wind farms and the commissioning of large passenger aircraft featuring composite wings appear to open up new fields for application of SHM technology. The three motivating aspects for SHM that are recurrently mentioned in the literature are [1]:

1) The life-safety benefits achievable through being able to continuously monitor critical components.

2) The economic benefits achievable by (i) avoiding unplanned down-time due system of structural failure and (ii) increasing the efficiency of inspection.
3) The optimisation of light-weight structures for which the current condition is known through monitoring.

1.1.1 Damage identification hierarchy

The damage identification problem can usefully be considered as a hierarchical process, developed from that discussed by Rytter [2]. The five levels of damage identification are:

1) **Detection** - is damage present in the system?
2) **Localisation** - where has the damage occurred?
3) **Classification** - what kind of damage is present?
4) **Assessment** - what is the extent of the damage?
5) **Prediction** - what is the residual life of the system, in light of the damage?

Success at a given level of the hierarchy above is largely contingent upon success having been achieved at the lower levels: one cannot assess the extent of damage without locating it first, nor will it be possible to make a prediction of the future extent of damage without having first identified the type of damage that has occurred. Identifying the level of identification sought is a key decision to be made when seeking to apply SHM methods and will have a major bearing on the approach adopted.

1.1.2 Features for SHM

Damage identification is reliant upon the existence of low-dimensional *features* that are indicative of damage. The identification of features is a key step in the application of SHM. The premise of vibration-based SHM is that damage to the structure will lead to changes in the dynamic response of the structure. Certain features of the recorded response will demonstrate a greater sensitivity to these changes than others. Feature extraction is the process of identifying these damage-sensitive features of the dynamic response and using them to create feature vectors. The focus in this thesis is damage identification based on low-frequency vibration. A summary of the most commonly applied features is given below.

- Natural frequencies
- Modeshapes
- Modeshape curvature/Modal strain energy
- Operational deflection shapes (ODS)
- Frequency response functions (FRF) and Transmissibilities
- Dynamic flexibility
- Damping coefficients

A challenge faced by all SHM approaches is that in addition to displaying sensitivity to damage, the features listed above will typically also display sensitivity to other factors, for example operational and environmental conditions [3].

1.2 SHM reviews

Structural Health Monitoring is a complex problem that has attracted considerable research interest. The literature is extensive; however a number of excellent reviews are available, including those produced by the Los Alamos National Laboratory Structural Health Monitoring research group.

An initial review [4] provides a history of developments in the field, before focusing on a broad range of relevant vibration-based damage detection methods. A section on practical applications of the technology, both for individual components and for entire civil and mechanical structures, is also included. The review concludes with a summary of the critical issues within the field, as a basis for future work. The issues highlighted include the dependence of many techniques on the availability of *a priori* analytical models, disagreement over the sensitivity of vibration-based methods, a lack of knowledge regarding optimal sensor placement and sensor quantity, and a lack of surveys that directly compare potential methodologies. It is concluded that sufficient evidence exists for the viability of vibration-based techniques in structural damage identification, but that future work should be directed at the application of the technology to specific applications and to real-world structures in their operating environment.

An updated review [5] covering advances from 1996-2001 was greatly expanded and restructured to encompass an even broader range of literature. Damage detection methodologies were separated into the steps defined by the Statistical Pattern Recognition Paradigm discussed in [6], and extensive sections on operational

evaluation, data acquisition and signal processing were included for the first time. The review of damage identification methods that constituted much of the previous review was separated into two sections. The first of these covered feature extraction and information condensation - essentially the generation of reduced-dimension damage sensitive features from the gathered vibration data. A significant number of new techniques that had been investigated in the intervening period between reviews were included, and an update of progress in the fields previously covered was provided. The second section, covering the statistical treatment of the calculated features, saw a similar expansion to include new techniques and updates on established methods. The review ended once again with a summary of outstanding issues.

In addition to the above, a special issue of the Philosophical Transactions of the Royal Society (prefaced by [7]) provides an updated overview of the field. Recent years have seen the publication of at least two monographs dedicated entirely to SHM [8, 9], and [10] must be considered the most comprehensive guide to the subject that has been published to date.

1.3 Approaches to SHM

It is commonly stated in the structural health monitoring literature that there are two approaches to the vibration-based SHM task: the *data-based* approach and the *model-based* approach. The data-based approaches are built entirely upon experimental data, without recourse to law-driven models and typically employing principles of statistical pattern recognition. The term 'model-based approach' is often taken as referring exclusively to a finite element model updating (FMU) task, with an initial model continually updated on the basis of newly-presented structural data in order to identify damage.

A final category of approach - and one that has received relatively little attention in the SHM literature - is referred to in this thesis as the *hybrid approach*. Hybrid approaches are those that lie somewhere between the 'extremes' of purely data-driven and purely model-driven approaches. In this thesis, the term 'hybrid approach' is taken to refer specifically to a method that makes use of predictive modelling to supplement experimental data, coupled with statistical modelling for damage discrimination. The three approaches to SHM are thus defined in this thesis as:

1. **Model-Based** Approaches
2. **Data-Based** Approaches
3. **Hybrid** Approaches

All three approaches are addressed to some extent in the work presented. The model-based and data-based approaches are briefly introduced below, including key references and an overview of the advantages and drawbacks of each. Additional references are introduced as appropriate throughout the thesis.

1.4 Model-based SHM

Model-based approaches are commonly-based upon the development of a Finite Element (FE) model representing the structure in question. After an initial FE model has been developed, experimental data is used to update its dynamic property matrices (mass, stiffness and damping) to allow the analytical model to more accurately represent the experimental structure. Damage detection, location and quantification are then possible through the solution of an inverse problem using data collected from the structure under test. Additionally, predictive law-based modelling is a requisite element in damage prognosis [11].

The application of law-based models in an inverse mode has received a great deal of attention in the SHM literature with several notable successes. In [12] an updating approach was applied to identify damage introduced in to the Z24 highway bridge. In [13] damage identification is applied to a 7-storey reinforced concrete structure, with the uncertainty of damage identification results assessed.

A challenge to be overcome in applying model updating-based methods is that the inverse problem is often ill-posed and requires careful regularisation [14, 15]. Conditioning may be improved by reducing the dimension of the inverse problem through subset selection [16]. Minimum rank perturbation theory has been applied for damage detection, location and extent assessment [17].

Model updating methods are susceptible to the effects of uncertainty attributable to measurement variability and model-form discrepancies. The modelling of complex substructures such as joints represents a non-trivial task, as modelling simplifications must inevitably be made. Methods for parameterising joints in a manner amenable to updating are discussed in [18]. Of equally great concern is the modelling of the damage introduced into the structure. For the example of crack

modelling, damage can be parameterised using a broad variety of approaches from simply reducing the stiffness of an element at the appropriate location to applying complex and highly parameterised fatigue crack models [19]. The use of simplified generic elements for damage identification is introduced in [20]. In [21] distinction is drawn between two classes of error in the model parameters: *damage parameters* and *other parameters*. The 'other' parameters are those attributable to modelling and experimental error. Calibrating the parameters of an initial, physically-representative model to normal-condition data from the experimental structure allows the portion of the test-FE metric attributable to the main effects of the 'other' parameters to be minimised. However, higher-order interaction effects may remain.

1.5 Data-based SHM

The alternative approach to diagnostics in SHM also involves the construction of a model, but this model is statistical rather than law-based in nature. The model is established by applying a machine learning or pattern recognition approach to experimentally acquired data. As the approach is based entirely upon the data, the requirement to develop and validate a numerical model is avoided. Further, uncertainty arising from measurement variability is automatically accommodated by the approach.

Data-based methods have been shown to be capable of successfully tackling damage identification problems at Levels 1-4 of Rytter's hierarchy as presented i.e. up to the level of assessment [22]. Progressing to damage prognosis requires that the physical mechanisms explaining the progress of damage be modelled; this is beyond the scope of a purely data-based model.

Statistical pattern recognition may be conducted in either a supervised or unsupervised mode. The key distinctions for damage identification are the level of identification that may be achieved, and the data that is required. It is broadly accepted that unsupervised learning methods are restricted to providing a level 1 diagnostic; however, they only require data from the undamaged structure for training. Supervised learning methods are capable of achieving higher levels of identification, but require training data from the structure in both its undamaged and damaged states. Such data will rarely be available for the large-scale, high-value structures for which monitoring is sought. The lack-of-data problem is perhaps the greatest challenge in applying pattern recognition methods.

Damage detection via unsupervised learning has been demonstrated using techniques that include outlier analysis [23] and control chart methods [24]. In [25] factor analysis is employed prior to the application of control charts in order to mitigate environmental effects. Supervised learning has been demonstrated using Artificial Neural Networks (ANNs), Support Vector Machines (SVMs) [26] and Kernel Discriminant Analysis (KDA) [27] among many others.

1.6 Objectives

The contributions of this thesis are to consider aspects of both the model-based and data-based approaches to SHM, as presented above; to identify the advantages and disadvantages of each; and to investigate whether a synthesis of the approaches can be found that overcomes some of the disadvantages of the individual approaches.

Approaches to SHM are exemplified using three studies.

- Application of a FE model-updating-based approach to an experimental structure.
- Investigation of methods for applying predictive law-based modelling to the SHM task.
- Application of a data-based approach to a challenging problem on a complex experimental structure.

By considering all three approaches to some extent, it is hoped that a broad picture of the challenges facing the SHM field may be presented.

1.7 Summary of Chapters

The thesis is structured as follows.

Chapter 2

The first damage identification approach to be considered here uses the updating of law-based models. In Chapter 2 the theoretical basis for the approach is introduced.

Chapter 3

Chapter 3 describes the application of the model-updating approach to an experimental structure for the purposes of damage location. The structure used is an Aluminium 'wingbox' which includes complicating factors such as riveted and

bolted joints. This simple structure allows an application of model-based SHM to be demonstrated, and the performance of the model-updating approach in handling random measurement noise, systematic measurement variability and model-form discrepancy to be investigated. Several of the pros and cons that have been identified in the literature manifested themselves, and these are discussed.

Chapter 4

Chapter 4 describes applications of law-based models in the forward mode. This approach forms the first step in a possible hybrid approach to SHM, whereby a predictive, law-based modelling task is used to inform the development of a discriminative statistical model. Practical tools from the field of probabilistic verification and validation (V&V) are introduced for the predictive task. Possible applications of forward modelling are illustrated via two case studies, employing the wingbox structure introduced in Chapter 3.

Chapter 5

In Chapters 5 to 7 the focus moves to data-driven approaches. Data-driven approaches are exemplified here by tackling a challenging SHM task: the identification of damage occurring at multiple locations. The possibility of constructing a data-driven classifier using observations of single-site damage, and consequently using this classifier to identify multiple-site damage is investigated. In Chapter 5 the problem of multiple damage location is introduced and the approach developed in Chapters 5 to 7 is set out. The first stage in this approach, experimental data acquisition, is also described, with particular attention being paid to test sequencing in order that the effects of systematic and random test variability are well-represented in the gathered data. The subject of the multiple damage location study is a section of aircraft wing sourced from a Piper Tomahawk trainer aircraft.

Chapter 6

Chapter 6 describes the feature selection process employed to reduce the high-dimensional data recorded in Chapter 5 into a set of low-dimensional features suitable for the training of statistical classifiers. The hypothesis in Chapter 6 is that it should be possible to find features that offer a good level of discrimination between multiple-site damage states, despite only single-site damage states being available for feature selection. The ability of individual features from the chosen feature set to

discriminate between damaged and undamaged data is evaluated as an intermediate step prior to the training of statistical classifiers.

Chapter 7

Chapter 7 describes the application of statistical discrimination for the multiple-site damage location problem. The method adopted is Support Vector Machine (SVM) classification, and outline theory relating to statistical learning theory and SVM classification is provided. A classifier is developed and trained using the feature set identified in Chapter 6 and a dataset that includes observations of normal condition and single-site damage data only. The performance of the classifier in classifying multiple-site damage is evaluated. The performance of the developed classifier is compared to a 'best case' scenario whereby the same feature set is employed, but the training dataset includes observations of multiple-site damage.

Chapter 8

The approaches presented and discussed throughout the thesis are drawn together in Chapter 8. The relative strengths and weaknesses of law-driven and data-driven approaches to the diagnostic element of Structural Health Monitoring (SHM) are discussed, with the intention of indicating where they are best used in practice. Discussion progresses from the individual approaches towards how they may be combined.

Chapter 9

In Chapter 9 conclusions are summarised, and areas for future work identified.

2 Parameter Estimation Theory

2.1 Introduction

The first damage identification approach to be considered here is the updating of law-based models. Model updating approaches were originally developed for the improvement of structural models. In recent years, model updating methods have been extensively applied to the problem of damage identification - several examples are presented in [14, 16, 28]. Numerous successes have been reported as a result of this research effort, but it is acknowledged that the application of updating approaches remains a challenging task. In this chapter the theoretical basis of the approach is introduced, drawing on the work presented in [29-31]. In the following chapter, model updating is applied to an experimental case study, allowing several of the pros and cons of the approach to be exemplified.

Damage identification through numerical model updating represents an inverse problem, with the added challenge that this problem will usually be ill-conditioned. Core to the solution of this problem is a parameter estimation task with least-squares estimation (LSE) being the principal tool applied. First, the classical least-squares problem is introduced in abstract form, with reference made to weighting of responses and the use of regularisation. Applications of least-squares estimation to model correlation and damage identification are introduced in turn. Two alternative approaches are identified for model updating-based damage identification, with use of updating to improve the correlation of response deviations advocated over a two-stage correlation approach. The advocated approach is applied to a case study in Chapter 3.

2.2 Parameter estimation

2.2.1 Least-squares parameter estimation

Consider the system of linear equations,

$$\boldsymbol{Ax} = \boldsymbol{b} \tag{2.1}$$

where $\boldsymbol{x}$ is a vector of parameters of length n, $\boldsymbol{b}$ is a vector of observations of length m and $\boldsymbol{A}$ is a matrix of size (m x n). The objective of the parameter estimation task is, for given $\boldsymbol{A}$ and $\boldsymbol{b}$, to determine the value of the parameter vector $\boldsymbol{x}$. The estimate of the 'true' parameter vector $\boldsymbol{x}$ is denoted herein as $\widehat{\boldsymbol{x}}$.

The parameter estimation problem may be formulated in a least-squares sense as,

$$\widehat{\boldsymbol{x}} = \min_{\boldsymbol{x}} \|\boldsymbol{Ax} - \boldsymbol{b}\|_2^2 \tag{2.2}$$

where $\|\ldots\|_2$ indicates the vector 2-norm. The minimisation in (2.2) may be expressed in full as,

$$\widehat{\boldsymbol{x}} = \min_{\boldsymbol{x}} \Big((\boldsymbol{Ax} - \boldsymbol{b})^{\mathrm{T}} (\boldsymbol{Ax} - \boldsymbol{b}) \Big) \tag{2.3}$$

2.2.2 Weighted least-squares estimation

A development of (2.3) is that each residual term may be individually weighted, for example to reflect the confidence ascribed to the accuracy of the observed measurements. The resulting *weighted least-squares* formulation may be expressed as,

$$\widehat{\boldsymbol{x}} = \min_{\boldsymbol{x}} \Big((\boldsymbol{Ax} - \boldsymbol{b})^{\mathrm{T}} \boldsymbol{W} (\boldsymbol{Ax} - \boldsymbol{b}) \Big) \tag{2.4}$$

where $\boldsymbol{W}$ is a weighting matrix of size (m x m) with $\boldsymbol{W} = diag(w_1, w_2, \ldots, w_m)$. If the statistical properties of the measurement uncertainties are known, then these properties may be used in the specification of $\boldsymbol{W}$. For example, in the case that the observations are normally-distributed and are uncorrelated with variances σ_i^2 (for $i = 1 \ldots m$), it can be shown [29] that adopting $\boldsymbol{W} = diag(1/\sigma_1^2, 1/\sigma_2^2, \ldots, 1/\sigma_m^2)$ leads to unbiased estimates with covariances given by,

$$var(\widehat{\boldsymbol{x}}) = \left(\boldsymbol{A}^{\mathrm{T}} \boldsymbol{W} \boldsymbol{A} \right)^{-1} \tag{2.5}$$

2.2.3 Solution of the least-squares problem

The method pursued for solving the least-squares problem is dependent upon the characteristics of the matrix $\boldsymbol{A}$. Three distinct cases present themselves:

i) $\boldsymbol{A}$ is full-rank (i.e. $rank(\boldsymbol{A}) = n$) and well-conditioned
ii) $\boldsymbol{A}$ is full-rank (i.e. $rank(\boldsymbol{A}) = n$) but ill-conditioned
iii) $\boldsymbol{A}$ is rank-deficient (i.e. $rank(\boldsymbol{A}) < n$)

If $\boldsymbol{A}$ is full-rank and well-conditioned, (2.4) may be solved directly as a set of simultaneous equations, or by the minimisation of a cost function, and will have a unique solution. Direct solution may be applied using Gaussian elimination. However, direct solution is not recommended for model-updating applications as in the case that $\boldsymbol{A}$ is ill-conditioned the outcomes may not be physically meaningful. Instead, a cost function approach may be adopted.

The cost function approach seeks the minimum of,

$$\boldsymbol{J}(\boldsymbol{x}) = \left\| \boldsymbol{W}^{\frac{1}{2}}(\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b}) \right\|_2^2 \tag{2.6}$$

or in full,

$$\boldsymbol{J}(\boldsymbol{x}) = (\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b})^{\mathrm{T}} \boldsymbol{W} (\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b}) \tag{2.7}$$

Iterative minimisation of the cost function in the context of model updating is described in Section 2.3.

If $\boldsymbol{A}$ is full-rank (i.e. $rank(\boldsymbol{A}) = n$) but ill-conditioned, small deviations in either the observation vector $\boldsymbol{b}$ or the matrix $\boldsymbol{A}$ can lead to large errors in the parameter estimate $\hat{\boldsymbol{x}}$. The condition number κ of the matrix is calculated as,

$$\kappa(\boldsymbol{A}) = \left| \frac{\lambda_{max}(\boldsymbol{A})}{\lambda_{min}(\boldsymbol{A})} \right| \tag{2.8}$$

It can be shown [29] that, for the full rank problem, the error in the parameter estimate $\hat{\boldsymbol{x}}$ has an upper bound $\kappa(\rho_{\boldsymbol{A}} + \rho_{\boldsymbol{b}})$ where $\rho_{\boldsymbol{A}}$ and $\rho_{\boldsymbol{b}}$ are the relative errors in $\boldsymbol{A}$ and $\boldsymbol{b}$ respectively.

If $\boldsymbol{A}$ is not full-rank then the problem in (2.2) is under-determined and an infinite number of solutions exist. A unique minimum-norm solution may be arrived at by

applying the Moore-Penrose pseudo-inverse, or through application of singular-value decomposition. Unfortunately, it is unlikely that the minimum-norm solution arrived at will represent a physically meaningful outcome in model-updating applications.

For both the rank-deficient and the ill-conditioned cases, *regularisation* may be applied in order to improve the conditioning of the solution. Side constraint regularisation is of particular relevance for damage identification applications and is introduced below. An alternative (and complementary) method for tackling problems of ill-conditioning and rank deficiency is to reduce the dimension of the parameter vector $\boldsymbol{x}$. *Parameter subset selection* with regard to damage identification is discussed further in [16].

2.2.4 Side constraint regularisation

The purpose of regularisation is to include further information about the desired solution of the least-squares problem, such that a stable solution may be achieved. One such additional piece of information may be to require that the deviation of the parameters changes from initial estimates be small. Alternatively it may be the case that certain of the parameters should take similar values, and thus the difference between them be minimised. This information can be incorporated as a side constraint of the form,

$$\boldsymbol{\Omega}(\boldsymbol{x}) = \left\| \boldsymbol{W}_{\boldsymbol{\theta}}^{\frac{1}{2}} \delta \boldsymbol{x} \right\|_2^2 = \delta \boldsymbol{x}^{\mathrm{T}} \boldsymbol{W}_{\boldsymbol{\theta}} \delta \boldsymbol{x} \tag{2.9}$$

where $\boldsymbol{W}_{\boldsymbol{\theta}}$ is a parameter weighting matrix and $\delta \boldsymbol{x}$ is a vector of parameter differences. Depending upon the solution scheme, $\delta \boldsymbol{x}$ may chosen to be,

$$\delta \boldsymbol{x} = \boldsymbol{x} - \boldsymbol{x}_{init} \tag{2.10}$$

where $\boldsymbol{x}_{init}$ is an initial parameter vector estimate, or in an iterative scheme,

$$\delta \boldsymbol{x} = \boldsymbol{x} - \boldsymbol{x}_j \tag{2.11}$$

where $\boldsymbol{x}_j$ is the parameter vector estimate at iteration step j. The matrix $\boldsymbol{W}_{\boldsymbol{\theta}}$ controls the type of constraint applied. Setting $\boldsymbol{W}_{\boldsymbol{\theta}}$ as the identity matrix $\boldsymbol{I}_{\boldsymbol{n}}$ weights the parameter changes equally, and has the effect of constraining the parameter changes to be small. Alternatively, $\boldsymbol{W}_{\boldsymbol{\theta}}$ may be chosen to reflect that particular groups of parameters should be nominally equal. $(\boldsymbol{x} - \boldsymbol{x}_{init})$

The regularised equivalent of the cost function (2.5) takes the form,

$$\boldsymbol{J}(\boldsymbol{x}) = \left\| \boldsymbol{W}^{\frac{1}{2}}(\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b}) \right\|_2^2 + \boldsymbol{\lambda}^2 \left\| \boldsymbol{W}_{\boldsymbol{\theta}}^{\frac{1}{2}} \delta \boldsymbol{x} \right\|_2^2 \tag{2.12}$$

or in full,

$$\boldsymbol{J}(\boldsymbol{x}) = (\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b})^{\mathrm{T}} \boldsymbol{W} (\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b}) + \boldsymbol{\lambda}^2 \delta \boldsymbol{x}^{\mathrm{T}} \boldsymbol{W}_{\theta} \delta \boldsymbol{x} \tag{2.13}$$

The two terms on the right hand side of (2.8) are referred to as the *residual norm* and the *side constraint* respectively. The coefficient λ in (2.8) is a scalar value that allows the gross influence of the two terms to be balanced. Appropriate values for λ may be arrived at through cross-validation, or by using an L-curve approach that seeks to estimate the optimal value λ_{opt}. The L-curve approach will be described further, with an example, in Chapter 3. The application of the L-curve approach to the general least-squares estimation problem and consideration of other forms of regularisation are discussed in [32]. A discussion and application of cross validation for model updating is given in [33].

Regularisation is a major complicating issue in model updating, and careful application of regularisation is required in order to achieve meaningful updating outcomes.

2.3 Model updating

2.3.1 Parameters and responses

The objective of model updating is to minimise the difference between the predictions of the analytical model $\boldsymbol{z}(\theta)$ and the corresponding measurements from the physical structure $\boldsymbol{z}_m$. The predictions of the analytical model are dependent on the vector of unknown parameters θ. The vectors $\boldsymbol{z}(\theta)$ and $\boldsymbol{z}_m$ are of length m and the vector θ is of length l. The problem may be formulated in a least-squares sense as the minimisation of the objective function,

$$J(\theta) = \|\boldsymbol{z}_m - \boldsymbol{z}(\theta)\|_2^2 = \epsilon^{\mathrm{T}} \epsilon \tag{2.14}$$

where ϵ is the residual vector given by,

$$\epsilon = \boldsymbol{z}_m - \boldsymbol{z}(\theta) \tag{2.15}$$

The measurements compared in vectors $\boldsymbol{z}_m$ and $\boldsymbol{z}(\theta)$ are usually modal responses. Natural frequencies are predominantly employed for comparison due to the relatively low level of associated measurement variability that may be achieved. Modeshapes have also been extensively investigated, although the higher levels of variability associated with their measurement can be detrimental to the outcome of the updating process. It is more common for the modeshapes to be employed solely for matching modes from the physical and analytical structures.

In cases where damping is included in the analytical model, FRFs may in principle be considered as measurements for comparison to allow the parameters of the damping model to be updated. However, comparison of frequency spectra is somewhat more involved than comparison of natural frequencies or modeshapes. Given that the spectra may contain several thousand spectral lines, cross-correlation-based criteria developed along the same lines as the MAC may be applied in order to reduce the dimension of the data, as in [34]. Such criteria may have many local minima, severely complicating the optimisation task. Further, the inclusion of a damping model increases the number of parameters to be considered in θ. Finally, given that it is possible to reconstruct the FRFs from the modal properties (assuming the effect of out-of-range modes is not significant) it is not apparent that FRFs contain any additional information to aid updating beyond estimation of the damping values.

Having formed the objective function $J(\theta)$, some method of optimisation must be applied in order to find its minima and thus the optimal estimate of the unknown parameters. Iterative sensitivity-based methods are commonly applied and the theory underlying this approach is given below.

2.3.2 Sensitivity-based iterative solutions

The cost function $J(\theta)$ given in equation (2.14) is a nonlinear function of the unknown parameters θ and there will potentially be many local minima. A solution may be sought by linearising the solution using as a truncated Taylor series and solving iteratively. For each iteration step j the full Taylor series approximation at point $\theta = \theta_j$ is given by,

$$\boldsymbol{z}(\theta) = \boldsymbol{z}_j + \boldsymbol{S}_j \delta\theta_j + \text{ higher order terms} \tag{2.16}$$

where,

$$\boldsymbol{z}_j = \boldsymbol{z}(\theta_j) \tag{2.17}$$

is the vector of m model predictions resulting from the parameter vector θ_j,

$$\delta\theta_j = \theta - \theta_j \tag{2.18}$$

is the vector of l parameter differences between the current parameter values and θ_j, and,

$$\boldsymbol{S}_j = \left.\frac{\partial \boldsymbol{z}}{\partial \theta}\right|_{\theta=\theta_j} \tag{2.19}$$

is an (m x l) sensitivity matrix formed from the first derivatives of the predicted responses with respect to the parameters.

Truncating the Taylor series expansion in (2.16) after the first-order terms allows a linear approximation to be made,

$$\boldsymbol{z}(\theta) = \boldsymbol{z}_j + \boldsymbol{S}_j\delta\theta_j \tag{2.20}$$

Introducing (2.20) into (2.15) gives,

$$\epsilon = \boldsymbol{S}_j\delta\theta_j - \delta\boldsymbol{z}_j \tag{2.21}$$

where,

$$\delta\boldsymbol{z}_j = \boldsymbol{z}_m - \boldsymbol{z}(\theta_j) \tag{2.22}$$

The cost function to be minimised at the j^{th} iteration is formed as,

$$J(\delta\boldsymbol{\theta}) = \left\|\boldsymbol{S}_j\delta\theta_j - \delta\boldsymbol{z}_j\right\|_2^2 = \epsilon^{\mathrm{T}}\epsilon \tag{2.23}$$

or, in full and augmented with weighting matrices and side constraint regularisation,

$$J(\delta\theta) = \left(\boldsymbol{S}_j\delta\theta_j - \delta\boldsymbol{z}_j\right)^{\mathrm{T}}\boldsymbol{W}\left(\boldsymbol{S}_j\delta\theta_j - \delta\boldsymbol{z}_j\right) + \boldsymbol{\lambda^2}\delta\theta^{\mathrm{T}}\boldsymbol{W}_\theta\delta\theta \tag{2.24}$$

Minimisation of the cost function (2.24) at each iteration involves calculating the derivative of the $J(\delta\theta)$ with respect to each parameter in turn and setting the result to zero, i.e.

$$\frac{\partial J(\delta\theta)}{\partial \delta\theta} = 0 \tag{2.25}$$

As shown in [29] this leads to a set of linear equations with a solution,

$$\delta\theta_j = \left(\boldsymbol{S}_j^{\mathrm{T}}\boldsymbol{W}\boldsymbol{S}_j + \boldsymbol{\lambda}^{\boldsymbol{2}}\boldsymbol{W}_{\theta}\right)^{-\mathbf{1}}\boldsymbol{S}_j^{\mathrm{T}}\boldsymbol{W}\left(\boldsymbol{z}_m - \boldsymbol{z}(\theta)\right) \tag{2.26}$$

The initial parameter estimate θ_{init} is iteratively updated as,

$$\theta = \theta_{init} + \delta\theta_1 + \delta\theta_2 \ldots \tag{2.27}$$

If regularisation is carefully applied and the parameter changes at each step are small, then stable convergence of the parameter estimates can be achieved.

2.4 Model updating for damage identification

Attention is now turned to how model-updating methods may be applied to the damage identification problem. Two principal approaches are available: *two-stage correlation* and *correlation of response deviations*.

2.4.1 Two-stage correlation

The first approach, considered in this section, involves two stages of updating: the first to produce a reliable model of the undamaged structure; the second to update this reliable model using damaged state data.

The structure of the approach is as follows:

1) To update the parameters of an initial model using data from the undamaged experimental structure $\boldsymbol{z}_m^{Undam}$, to produce a correlated model with parameters θ^{Undam} and responses $\boldsymbol{z}\left(\theta^{Undam}\right)$.
2) To update the parameters of the correlated model using data from the possibly damaged experimental structure $\boldsymbol{z}_m^{Dam}$, to produce a damaged-state model with estimated damage parameters θ^{Dam} and responses $\boldsymbol{z}(\theta^{Dam})$.

The particular application of the two-stage approach can be tailored through partition of the parameter vector θ into those parameters selected for improving the model correlation in the undamaged state θ_{Corr}, and those selected for identifying damage θ_{ID}. For the first stage of updating, the parameter vector is defined as either,

$$\theta = \begin{bmatrix} \theta_{Corr} \\ \theta_{ID} \end{bmatrix} \quad \text{or} \quad \theta = \theta_{Corr} \tag{2.28}$$

For the second stage, the parameter vector is defined as either,

$$\theta = \begin{bmatrix} \theta_{Corr} \\ \theta_{ID} \end{bmatrix} \quad \text{or} \quad \theta = \theta_{ID} \tag{2.29}$$

Solution of the parameter estimation problem for both stages proceeds as for equation (2.24).

An advantage of the method is that an initial round of correlation is undertaken which should, if conducted successfully, serve to increase confidence in the model predictions. Validation of this 'reliable' model is possible by comparing predictions to a validation dataset not used for updating.

There are, however, numerous factors that make application of the two-stage approach challenging in practice. The outputs of the model will typically display far greater sensitivity to the correlation parameters θ_{Corr} than to the parameters used for damage identification θ_{ID}. Further, the sensitivity of the outputs to model-form errors will typically be many times greater than to θ_{ID}. A final consideration is that as the dimension of the parameter vector is increased in order to consider both θ_{Corr} and θ_{ID}, there will be an adverse effect on the conditioning of the problem.

Despite the issues identified, several successful demonstrations of the two-stage approach have been reported - the approach was successfully applied to a steel frame structure in [28].

2.4.2 Correlation of response deviations

The second approach is to match the response deviations from the undamaged state to those for the structure. To the first-order, any systematic error that is present in both the damaged and the undamaged models will be removed. The cost function takes the form,

$$J(\theta) = \|\delta \mathbf{z}_m - \delta \mathbf{z}(\theta)\|^2 = \epsilon^{\mathrm{T}} \epsilon \tag{2.30}$$

where,

$$\delta \mathbf{z}_m = \mathbf{z}_m^{Dam} - \mathbf{z}_m^{Undam} \tag{2.31}$$

and,

$$\delta \boldsymbol{z}(\theta) = \boldsymbol{z}(\theta) - \boldsymbol{z}\left(\theta^{Undam}\right) \tag{2.32}$$

or in full,

$$\epsilon = \left(\boldsymbol{z}_m^{Dam} - \boldsymbol{z}_m^{Undam}\right) - \left(\boldsymbol{z}(\theta) - \boldsymbol{z}\left(\theta^{Undam}\right)\right) \tag{2.33}$$

Where an iterative scheme is employed, the linearisation of (2.33) that is the equivalent of (2.21) is,

$$\epsilon = \boldsymbol{S}_j \delta\theta_j - \left(\delta \boldsymbol{z}_m - \delta \boldsymbol{z}\left(\theta_j\right)\right) \tag{2.34}$$

where,

$$\delta \boldsymbol{z}\left(\theta_j\right) = \boldsymbol{z}\left(\theta_j\right) - \boldsymbol{z}\left(\theta^{Undam}\right) \tag{2.35}$$

The corresponding cost function augmented with weighting matrices and side constraint regularisation is,

$$J\left(\delta\theta_j\right) = \epsilon^{\mathrm{T}} \boldsymbol{W} \epsilon + \boldsymbol{\lambda_j}^{\boldsymbol{2}} {\delta\theta_j}^{\mathrm{T}} \boldsymbol{W}_\theta \delta\theta_j \tag{2.36}$$

A further option is to normalise the response deviations through division by the undamaged values. An example of when normalisation may be considered is the case that the response vector contains natural frequencies from a broad range of frequencies. A 5% deviation in a mode at 10Hz may be of comparable significance to a 5% deviation in a mode at 1000Hz; however this is not reflected in the respective absolute deviations of 0.5Hz and 50Hz. Equations (2.31) and (2.32) become,

$$\delta \boldsymbol{z}_m = diag\left(\boldsymbol{z}_m^{Undam}\right)^{-1}\left(\boldsymbol{z}_m^{Dam} - \boldsymbol{z}_m^{Undam}\right) \tag{2.37}$$

and,

$$\delta \boldsymbol{z}(\theta) = diag\left(\boldsymbol{z}\left(\theta^{Undam}\right)\right)^{-1}\left(\boldsymbol{z}(\theta) - \boldsymbol{z}\left(\theta^{Undam}\right)\right) \tag{2.38}$$

where $diag(\ldots)$ indicates a square matrix with the elements of the bracketed vector on its diagonal. The major advantage of the approach is the removal of first-order model-form errors present in the both the damaged and the undamaged models.

This goes some way towards overcoming a major hurdle presented by the two-stage correlation approach.

2.5 Summary

In this chapter, the theoretical basis of the model-updating approach has been introduced. The theory presented is not extensive - for a more thorough treatment, the reader is directed to the references provided. The theory presented here is, however, sufficient for the study presented in the following chapter. In Chapter 3, the updating of law-based models for the purposes of damage identification is demonstrated through an experimental case study.

3 Updating of Law-Based Models for SHM

3.1 Introduction

Methods based upon the principle of finite element model updating are the *de facto* standard for law-based structural health monitoring. However, despite an extended period of research interest, application of the approach is often found to be challenging, and successful applications of the approach to real-world problems remain rare.

In this chapter the application of, and difficulties related to, model-updating approaches for SHM are demonstrated using an example structure. The structure adopted is a laboratory-based 'wingbox'. The particular case considered is *damage location* on the basis of natural frequency deviations. The performance of the model-updating approach in handling random measurement noise, systematic measurement variability and model-form discrepancy are investigated. The features employed for updating are the natural frequencies of ten modes. The updating method employed for damage detection is the *correlation of response deviations* presented in Chapter 2. A characteristic of this approach is that first-order model-form effects are removed, mitigating the effect of model-form error to some extent. As such it was not deemed essential to pursue a first round of 'model-improvement' updating in this instance, although it is acknowledged such a step would in general be recommended in order to reduce model-form discrepancy.

The layout of the chapter is as follows: the experimental structure, data acquisition process and FE model are introduced; the particular model updating method employed is described; and damage location outcomes are presented for the cases of random measurement noise, systematic measurement variability and model-form

discrepancy. The chapter concludes with a discussion of the pros and cons of the model-updating approach. The difficulties encountered in applying the approach motivate the investigation of alternative applications of law-based models in Chapter 4, and serve as a counterpoint to the data-driven approach investigated in Chapters 5-7.

3.2 The wingbox structure

The structure under consideration is a stiffened Aluminium panel intended to replicate an aircraft wingbox. The top sheet of the wingbox is a 750 x 500 x 3mm Aluminium sheet. This is stiffened by the addition of two ribs composed of lengths of C-channel riveted to the shorter edges and two stiffening stringers composed of angle section that are bolted along the length of the sheet as shown in Figure 3-1. Free-free boundary conditions are approximated experimentally by suspending the wingbox from a substantial frame using springs and nylon line attached at the corners of the top sheet. The structure is very lightly damped and no problems were presented in realising the experimental modeshapes.

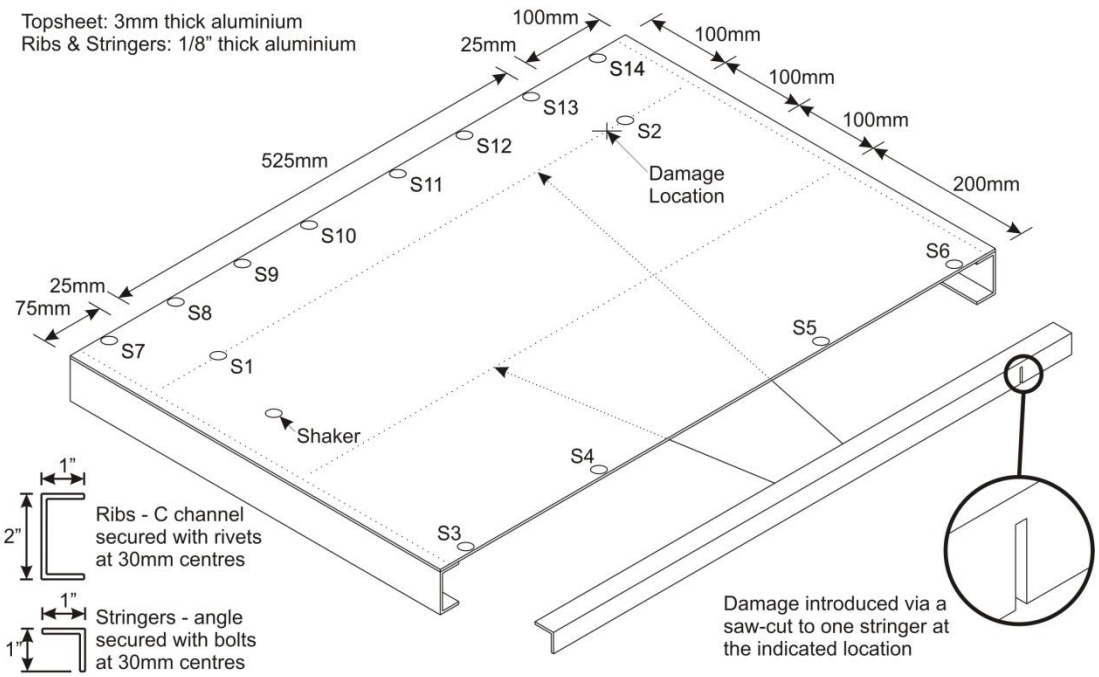

Figure 3-1 Schematic of the wingbox structure

3.3 Experimental data acquisition

3.3.1 Sensor placement

Responses were taken at 14 locations on the surface of the wingbox. Previous work, on a nominally identical structure [35], had found that transmissibilities between responses at locations on the stringer to be damaged were sufficiently sensitive for training a data-driven damage detection algorithm. Consistent with this, sensors S1 and S2 were placed on the stringer to be damaged, with the remainder distributed along the front (sensors S3-S6) and rear (sensors S7-S14) unconstrained edges of the wing box as illustrated in Figure 3-1. The choice of sensor position was constrained to locations corresponding to nodes of the FE model, with the locations of sensors S3-S14 informed by an Effective Independence – Average Driving Point Residue (EI-DPR) sensor placement algorithm [36] in an effort to best capture the modal response of the structure. The EI-DPR algorithm and further sensor placement optimisation options for SHM are provided in Appendix A. A dense sensor network was employed to provide a wealth of features for the following chapter and for future work. For the study in this chapter, it was found that a single sensor was sufficient for identifying the natural frequencies of interest.

Figure 3-2 Experimental structure

3.3.2 Damage introduction

In this study, damage cases of a single type and location but of differing severity are considered. The damage took the form of a saw-cut in one of the stiffening stringers, 125mm from the edge of the panel as shown in Figure 3-1. It is acknowledged that this form of open crack is not representative of the damage found in real structures. However, the relative ease with which the saw-cut may be modelled and the ability to introduce controlled amounts of 'damage' in situ were seen to be advantageous for the current study. The stringer was damaged to 25%, 50%, 75% and 100% of its flange depth, corresponding to saw-cuts 5.5mm, 11 mm, 16.5 mm and 22.5 mm deep and 2mm wide.

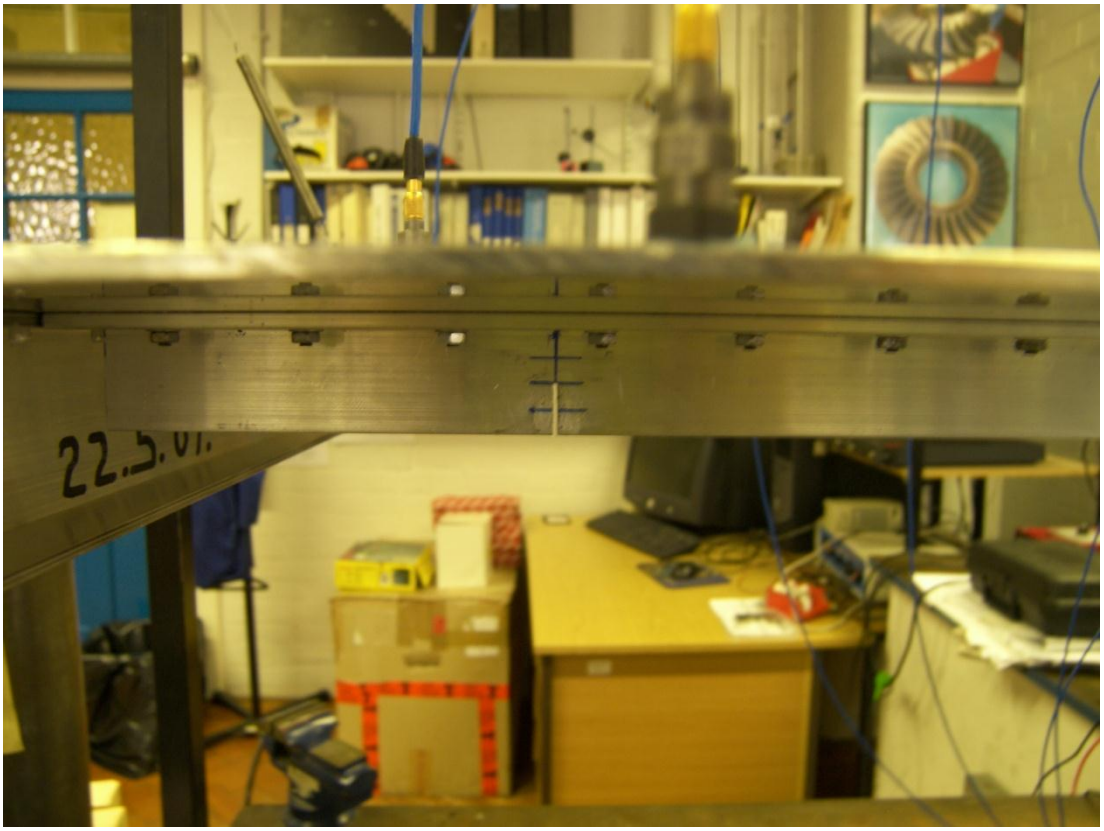

Figure 3-3 Saw-cut damage (50%)

3.3.3 Equipment and test sequence

The data acquisition system used during the test was a DIFA SCADAS III unit controlled by LMS software running on a desktop PC. All measurements were recorded within a frequency range of 0-2048 Hz with a resolution of 0.5 Hz. The structure was randomly-excited with a band-limited white Gaussian signal in the range 0-4096 Hz at the location specified in Figure 3-1. Both real and imaginary parts

of the FRFs were recorded at all response locations. For the normal condition, 200 samples of two-average data were recorded. For each of the four damage states a further 100 x two-average samples were taken. Two-average samples were found to offer a suitable balance between noise reduction and acquisition time.

A curve-fitting approach was employed for the estimation of natural frequency values from the large experimental dataset. A viscous modal damping model was used to fit a curve to the discretely sampled FRFs to allow estimates of natural frequency and modal damping to be made. In this case, the modal damping values were discarded. The availability of multiple response locations meant that those DOFs that best distinguished the modes of interest could be preferentially selected for curve-fitting. It was found that location S9 gave well-defined resonant peaks for all the modes of interest.

3.4 Finite element modelling

A numerical model of the wingbox was developed using the non-commercial MATFEM software [37]. The top plate was modelled using shell elements, and the stiffening components (ribs and stringers) are modelled using beam elements. The fastening components (rivets and bolts) are modelled using two-node spring elements to represent each rivet or bolt. The advantage of employing two-node spring elements is that the stiffness coefficients about each axis may be individually adjusted. In total, the model comprised 80 spring elements, 84 beam elements and 560 shell elements. As damage is to be simulated as a loss of stiffness, no modifications of the mesh for the purposes of modelling damage are necessary. The effects of sensor and shaker attachment are disregarded. Free-free boundary conditions are specified. The undeformed structure is illustrated in Figure 3-4.

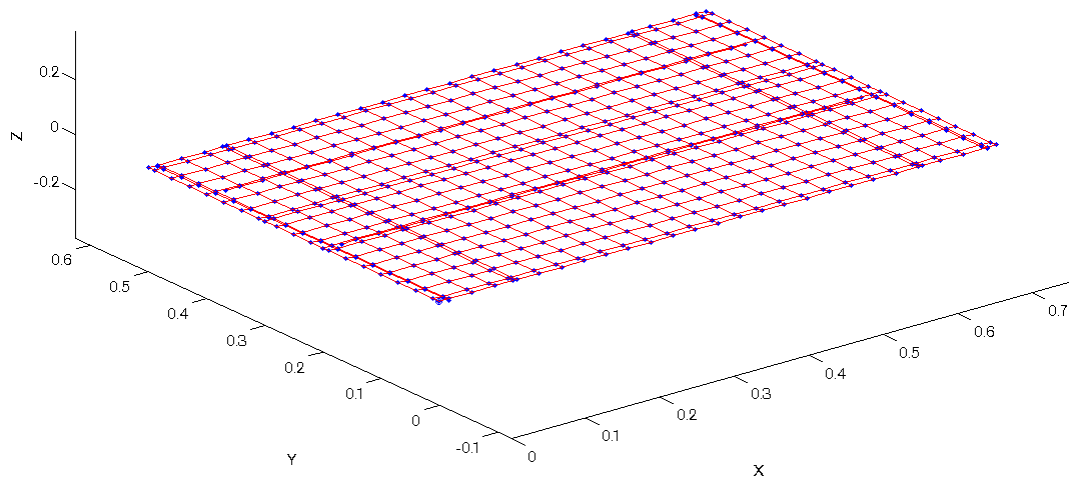

Figure 3-4 Undeformed finite element model

3.4.1 Evaluation of FE model

Agreement between the developed model and the 14DOF test data was assessed based on MAC and frequency correlations. The modal assurance criterion (MAC) for comparison of modeshape vectors $\boldsymbol{X}$ and $\boldsymbol{Y}$ of equal length is given by,

$$\mathrm{MAC} = \frac{\left(\boldsymbol{X}^{\mathrm{T}}\boldsymbol{Y}\right)}{\left(\boldsymbol{X}^{\mathrm{T}}\boldsymbol{X}\boldsymbol{Y}^{\mathrm{T}}\boldsymbol{Y}\right)}\ 100\% \tag{3.1}$$

Ten matched mode pairs were identified in the range of interest, as detailed in Table 3-1. While the overall correlation between the identified mode pairs was deemed to be 'good', low MAC values and large frequency deviations were observed for several of the modes. The low correlation values may be attributable in part to the relative complexity of the structure, and to the number and location of the measurement DOFs.

Fourteen accelerometers were distributed upon the structure. While this number of measured DOFs may be considered somewhat onerous from the perspective of a practicable damage identification methodology on a structure of this size, 14 DOFs is a very low value for modeshape analysis of a structure of the relative complexity of the wingbox. The availability of a limited set of responses places the onus on optimal placement of the available sensors. An objective of the EI-DPR algorithm employed is that the independence of the measured modeshapes be maximised. However, placing the sensors by applying the EI-DPR algorithm to the model-predicted modeshapes led to the two unconstrained edges of the structure being prioritised as measurement locations. The MAC value was thus highly dependent upon the level

of correlation recorded in just two regions of the structure. For mode pairs [E7, A16] and [E9, A17], 'good' correlation for the region monitored by sensors S3-S6 was counterbalanced by small deviations in the region monitored by sensors S7-S14, leading to MAC values that may be unrepresentatively low. For each of the 10 identified modes, the natural frequency of the i^{th} mode as predicted by the FE model $\omega_i(\theta)$ was lower than the measured natural frequency $\omega_{m,i}$. This may be indicative of a systematic bias in the model predictions arising from model-form decisions.

In this case, the mode pairings detailed in Table 3-1 were in agreement with concurrent tests on the same structure [38], that employed an impact hammer excitation and enabled the inclusion of a far greater number of response locations. On this basis, the mode pairings were deemed accurate. However, the results in Table 3-1 give an indication of the difficulty that may be encountered in achieving confidence in mode pairing for non-trivial structures, particularly where only a limited set of responses is available. Issues of mode pairing are returned to in the second case study in Chapter 4.

Mode No		Measured frequency ω_m [Hz]	Predicted frequency $\omega(\theta)$ [Hz]	Frequency deviation $\Delta\omega$ [%]	MAC
Exp	Model				
E2	A8	70.05	64.44	-8.70	90.79
E3	A9	104.39	101.65	-2.69	95.21
E4	A10	116.16	105.65	-9.95	65.73
E5	A13	186.07	180.08	-3.33	94.23
E6	A14	197.32	190.73	-3.46	89.83
E7	A16	236.32	219.84	-7.50	32.72
E9	A17	294.02	251.81	-16.76	54.68
E11	A22	346.8	343.58	-0.94	85.47
E12	A24	391.83	366.79	-6.83	60.06
E15	A27	457.44	436.93	-4.69	74.52

Table 3-1 Test-FE natural frequency and MAC comparison

3.4.2 Experimental results

The distributions of the normal and damaged state data resulting from the curve-fitting process are summarised in Table 3-2. The assumption is made that the data are normally-distributed.

From Table 3-2 it may be observed that the variability of each measured response appears to be relatively consistent across damage levels, and that some of the responses display markedly greater sensitivity to the introduced damage than

others. The behaviour of the responses with damage is more clearly visualised in Figure 3-5 to Figure 3-14. All observations at each damage state are included, with summary statistics (mean ± 2 standard deviations) overlaid upon the data.

Mode No	Measured frequency ω_m [Hz] (mean ± 1 std) [Hz]				
	0%	25%	50%	75%	100%
E2	70.05 ±0.25	70.02 ±0.24	70.01 ±0.22	70.00 ±0.21	70.00 ±0.24
E3	104.39 ±0.29	104.51 ±0.33	104.16 ±0.31	102.86 ±0.31	99.13 ±0.45
E4	116.16 ±0.33	116.07 ±0.40	115.84 ±0.36	116.23 ±0.29	116.13 ±0.29
E5	186.07 ±0.28	185.77 ±0.25	184.94 ±0.20	182.51 ±0.23	176.51 ±0.26
E6	197.32 ±0.27	197.16 ±0.33	196.88 ±0.32	196.46 ±0.38	195.97 ±0.57
E7	236.32 ±0.52	235.80 ±0.41	234.23 ±0.33	230.82 ±0.38	223.38 ±0.28
E9	294.02 ±0.38	293.77 ±0.42	293.40 ±0.32	293.49 ±0.32	292.84 ±0.30
E11	346.8 ±0.32	346.73 ±0.32	346.68 ±0.31	346.36 ±0.29	346.06 ±0.36
E12	391.83 ±0.74	392.00 ±0.87	391.77 ±0.79	391.94 ±0.80	390.73 ±1.03
E15	457.44 ±0.41	456.51 ±0.38	454.32 ±0.36	451.05 ±0.47	448.62 ±0.62

Table 3-2 Mean and standard deviations of experimentally identified modes: saw-cut damage

Prior to embarking upon the update-based damage location exercise, it is worth commenting on the behaviour of the individual features with damage. If a change in the mean response value such that it lies outside of the range of the undamaged mean ± 2 standard deviations is taken as being indicative of damage, then only mode E15 is indicative of damage for every investigated state. Two modes (E5 and E7) indicate damage for the 50%-100% damage states; two modes (E3 and E6) indicate damage at the 75%-100% damage states; and two modes (E9 and E11) indicate damage for the 100% damage state only. Three modes (E2, E4 and E12) are not by this measure indicative of damage for the states investigated, though mode E12 does display a reduction of between one and two standard deviations for the 100% case.

The analysis of individual features gives little insight into how the damage identification approach will perform as it is the pattern of changes, and in particular differences in the pattern of changes between different damage locations, that is of primary concern. However, the similarity in the order of the response deviations due to damage and to random measurement error is worthy of note. It is important, too, to note that this data is not available *a priori*- in practical applications the behaviour of the features with damage would be unknown during the development of the modelling system. *Forward approaches,* which make use of the developed numerical model to predict feature changes with damage, are considered in the next chapter.

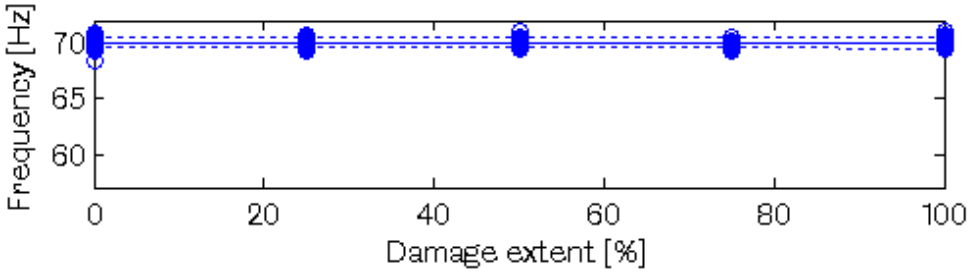

Figure 3-5 Natural frequencies, experimental mode E2

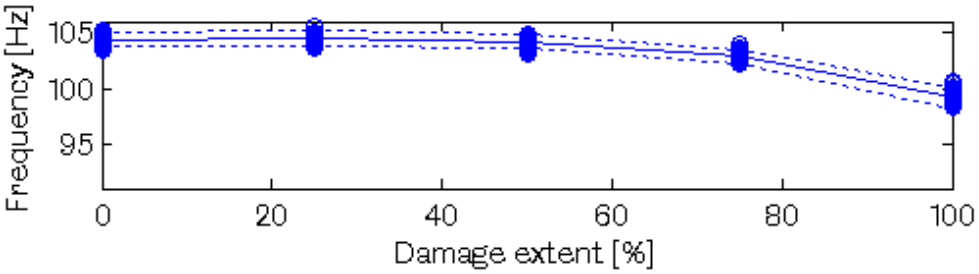

Figure 3-6 Natural frequencies, experimental mode E3

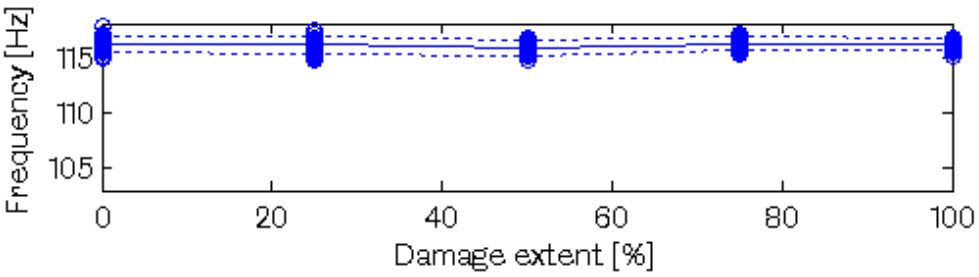

Figure 3-7 Natural frequencies, experimental mode E4

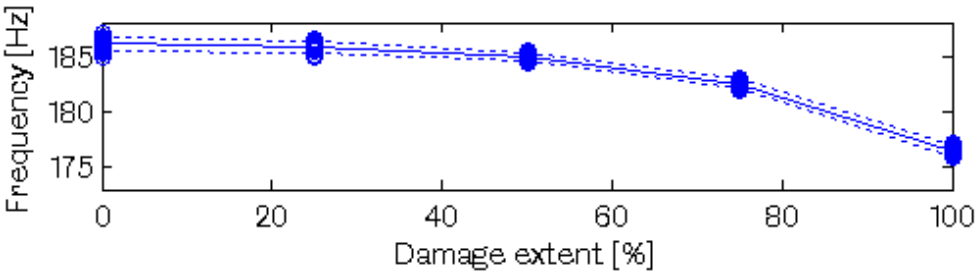

Figure 3-8 Natural frequencies, experimental mode E5

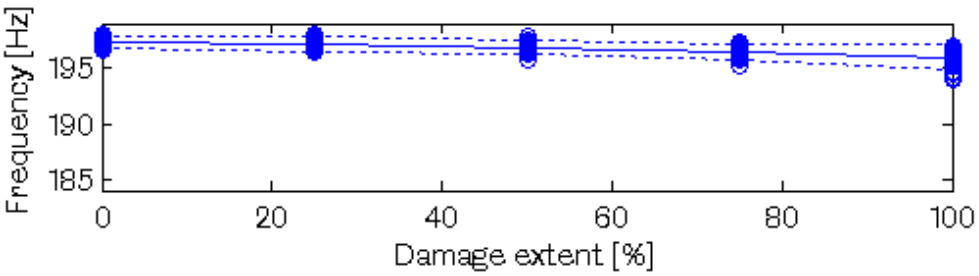

Figure 3-9 Natural frequencies, experimental mode E6

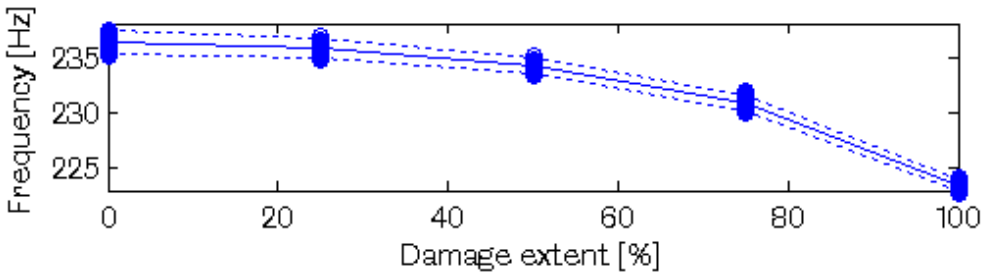

Figure 3-10 Natural frequencies, experimental mode E7

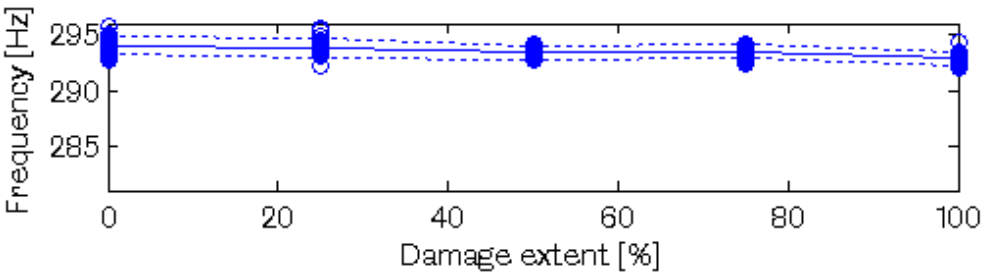

Figure 3-11 Natural frequencies, experimental mode E9

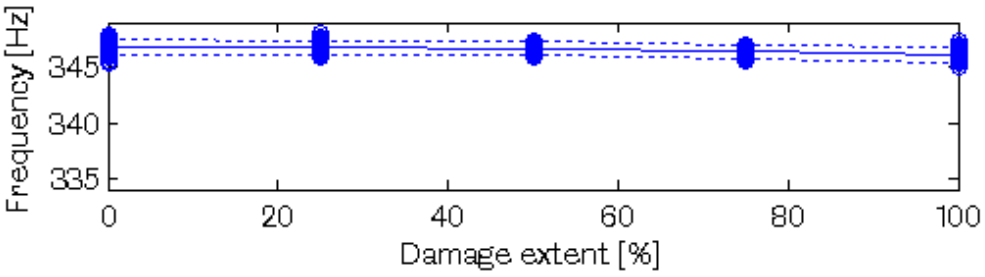

Figure 3-12 Natural frequencies, experimental mode E11

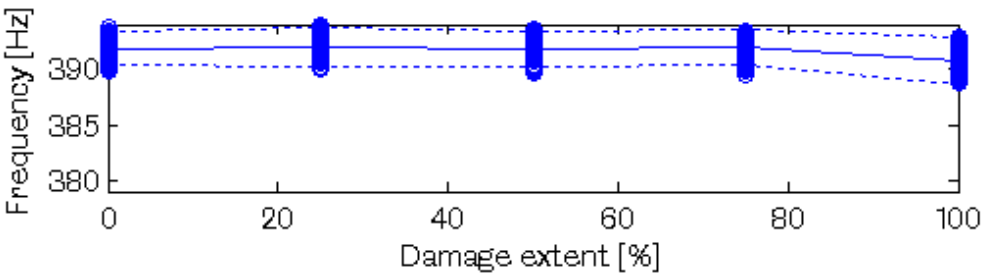

Figure 3-13 Natural frequencies, experimental mode E12

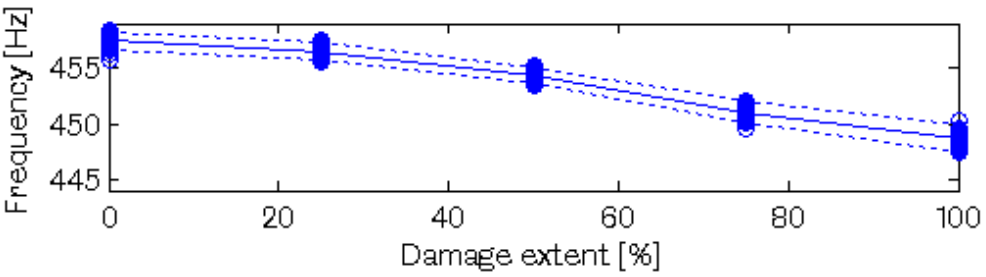

Figure 3-14 Natural frequencies, experimental mode E15

3.5 Model updating for damage location

3.5.1 Approach

The approach taken is to minimise the difference between the experimentally measured modal frequency deviations $\delta\boldsymbol{\omega}_m$ and the corresponding model-predicted modal frequency deviations $\delta\boldsymbol{\omega}(\theta)$. Damage is parameterised as a reduction in element stiffness. Stiffness reduction is judged to be suitable for damage location. For damage assessment, establishment of a relationship between stiffness and damage extent would be required – this is pursued as a case study in Chapter 4.

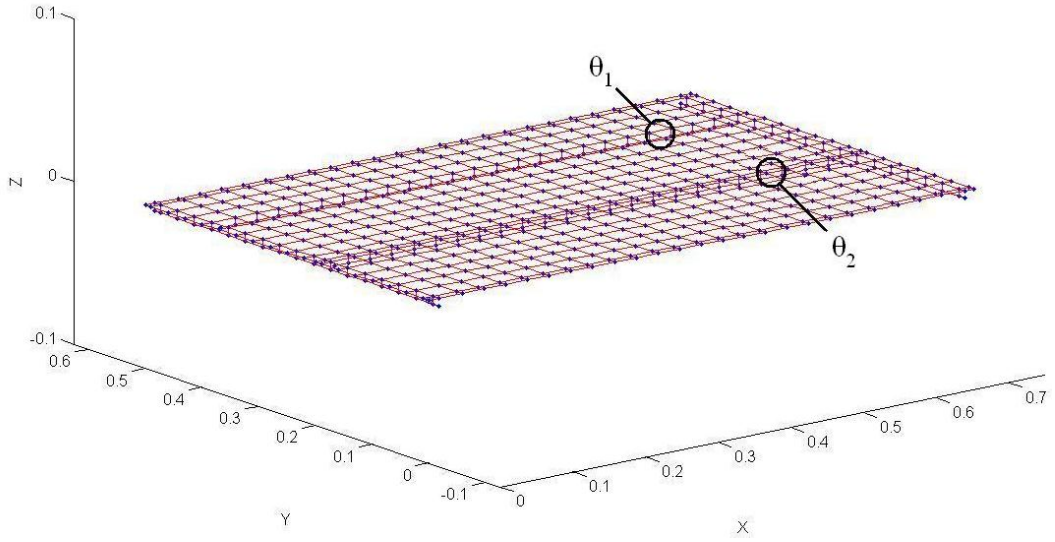

Figure 3-15 Locations of element groups governed by parameters θ_1 and θ_2

A deviation is made here from the classic updating approach. In the classic approach described in Chapter 2, the parameter change vector is arrived at iteratively. At each iteration step the sensitivity matrix $\boldsymbol{S}$ is calculated, providing an approximation of the model outputs that is a linear function of the parameter changes $\delta\theta$. From (2.34), at each iteration the solution of the cost function,

$$J(\delta\theta_j) = \epsilon^{\mathrm{T}}\boldsymbol{W}\epsilon + \boldsymbol{\lambda_j}^{\boldsymbol{2}}\delta\theta_j{}^{\mathrm{T}}\boldsymbol{W}_{\theta}\delta\theta_j \tag{3.1}$$

is sought, applying the linearisation,

$$\epsilon = \boldsymbol{S}_j\delta\theta_j - \left(\delta\boldsymbol{\omega}_m - \delta\boldsymbol{\omega}(\theta_j)\right) \tag{3.2}$$

Note that in the interests of clarity, the general response notation $\boldsymbol{z}$ used in Chapter 2 is replaced by the natural frequency notation $\boldsymbol{\omega}$. The regularisation parameter λ_j is

calculated at each iteration step using an appropriate method (for example using the L-curve approach or generalised cross-validation). The requirement that λ_j be assessed at each step becomes onerous in cases where a large number of updating cycles are to be executed.

In order to allow consideration of a large number of observations in a computationally efficient manner, a response surface approach was instead adopted. The model is executed for a structured sample of parameter values, allowing the response vector $\boldsymbol{\omega}(\theta)$ to be estimated at any point θ via interpolation. The parameter change estimate is arrived at in a single step by finding the minimum of a cost function,

$$J(\delta\theta) = \epsilon^{\mathrm{T}}\boldsymbol{W}\epsilon + \boldsymbol{\lambda^2}\delta\theta^{\mathrm{T}}\boldsymbol{W}_{\theta}\delta\theta \tag{3.3}$$

where,

$$\epsilon = \delta\boldsymbol{\omega}_m - \delta\boldsymbol{\omega}(\theta) \tag{3.4}$$

Incorporating the normalisation detailed in (2.37) and (2.38) the response deviations are,

$$\delta\boldsymbol{\omega}_m = diag\left(\boldsymbol{\omega}_m^{Undam}\right)^{-1}\left(\boldsymbol{\omega}_m^{Dam} - \boldsymbol{\omega}_m^{Undam}\right) \tag{3.5}$$

and,

$$\delta\boldsymbol{\omega}(\theta) = diag\left(\boldsymbol{\omega}\left(\theta^{Undam}\right)\right)^{-1}\left(\boldsymbol{\omega}(\theta) - \boldsymbol{\omega}\left(\theta^{Undam}\right)\right) \tag{3.6}$$

If non-normalised frequency shifts are used, the larger frequency shifts associated with higher frequency modes may dominate. The normalisation in (3.5) and (3.6) counters this potential bias.

Considering the parameter estimation in this way avoids the use of iterative linearisation as the full, nonlinear model output $\boldsymbol{\omega}(\theta)$ is estimated at all points. The compromise made is that control over regularisation is reduced - the regularisation parameter λ_j that would be set at each iterative step is held as a constant λ. The damage location outcomes are thus sub-optimal, but serve to illustrate several issues relating to model updating for damage identification.

In order to allow visualisation of the cost function over a range of parameters, only two parameters are investigated (i.e. number of parameters m=2). The parameters to be updated are indicated in Figure 3-15. Each parameter controls the Young's modulus of a group of three elements representing the stringer at the location shown in Figure 3-15. The element group governed by parameter θ_1 is located at the same position as the damage that was introduced into the experimental structure – it may be considered to be the 'true' location of damage. The element group governed by parameter θ_2 is located on the other stringer, at the same offset from the edge of the wingbox. The parameters are multiplying factors of the initial Young's modulus value of the elements $E_{init} = 69.9\text{GPa}$. The values of parameters θ_1 and θ_2 for the model in its undamaged state are 1. It is acknowledged that adjusting the stiffness of a group of parameters, rather than the stiffness of a single parameter, may introduce some degree of model-form discrepancy into the damaged-state predictions, as it may not be possible for the model to fully reconcile the changes introduced into the experimental structure. However, this set of element groups proved adequate to illustrate aspects of the model-updating approach.

All 10 paired natural frequencies are adopted as responses (i.e. number of responses n=10). In order to evaluate the cost function as a response surface, model predictions across the range of parameters are required. The FE model is executed across a grid of parameter values from 0.05 to 2 times the initial value of the parameters, requiring 1600 executions of the model. An estimate of the response vector $\boldsymbol{\omega}(\theta)$ at any parameter point θ within the given range is available via interpolation. It is acknowledged that the computational expense of sampling at the relatively high resolution used here is unnecessarily high. *Experimental design* may have been employed to reduce the number of model runs made.

Consideration of the response surface is useful in illustrating aspects of the damage location problem relating to the cost function. It is acknowledged that the solution of the cost function presents further issues, with the selection and size of the parameter set to be updated being central. While issues related to conditioning and subset selection are not investigated in this study, they are discussed at the end of the chapter.

3.5.2 Weighting matrices W_z and W_θ

The form and values of the response weighting matrix $\boldsymbol{W}$, the parameter weighting matrix $\boldsymbol{W}_\theta$, and the regularisation parameter λ must be set.

The response weighting matrix $\boldsymbol{W}_z$ may be set based on the 'confidence' held in the responses. This confidence that may be ascribed to the responses is taken here as the reciprocal of the variance of the measured responses. The response weighting matrix is thus specified as a square matrix of size $(n \times n)$ with diagonals,

$$[\boldsymbol{W}_z]_{ii} = \frac{1}{\sigma_i^2} \quad for\ i = 1 \dots n \tag{3.7}$$

and zero off-diagonal elements.

The choice of the parameter weighting matrix $\boldsymbol{W}_\theta$ is typically challenging. The sensitivity of the responses to the parameters can vary markedly, as can the initial values of the parameters. There may additionally be a need to introduce constraints to improve the conditioning of the problem. Engineering judgement must be exercised in setting the weighting matrix.

For the damage location case investigated no further constraint is required. The problem is set up such that the initial values of the two parameters considered are equal. As such an $(m \times m)$ identity matrix is deemed appropriate,

$$\boldsymbol{W}_\theta = \boldsymbol{I}_2 \tag{3.8}$$

3.5.3 Regularisation parameter λ

The regularisation parameter λ is used to achieve a balance between the gross weightings of the two terms in (3.3) in order to identify a unique and stable solution of the cost function. The parameter may be set using an L-curve approach, or on the basis of a cross-validation exercise. Here, the L-curve approach is considered.

The L-curve requires the cost function to be minimised for varying values of λ, and the norms of the resulting response and parameter terms in (3.1) to be plotted. The optimal value the regularisation parameter λ_{opt} is in the vicinity of the 'corner' of the L-curve, at the point of maximum gradient. Note that the value of λ_j would typically set at each iteration step. The setting of the regularisation parameter at each iteration step adds a computational cost to the updating process. In the subsequent

analysis, the updating procedure is effectively conducted in a single iterative step. This allows a large number of experimental observations to be assessed in a computationally efficient manner at the expense of a degree of accuracy.

Applying the L-curve methodology to find λ_{opt} for the 2-parameter case investigated was found to be challenging as the 'corner' of the L-curve plot is indistinct. An example for 25% damage is given in Figure 3-16.

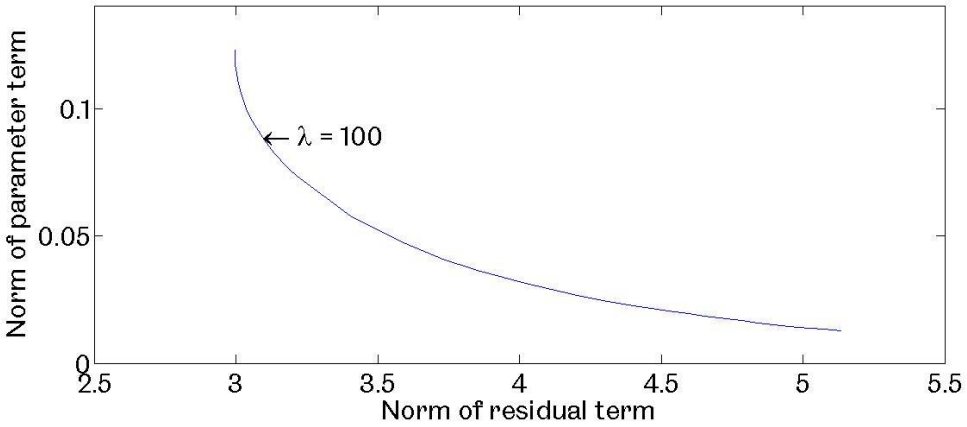

Figure 3-16 Diagram of L-curve for 25% damage

Following consideration of the L-curve for each damage states, a value of λ = 100 is adopted for all cases in the analysis below. This approach would not be recommended over the iterative approach in practice, but is deemed sufficient for illustration in this study.

3.6 Application of damage location approach

Parameter estimates were ultimately made for every experimental observation at each damage level - 600 observations in all. Initially, the mean results for each damage level are presented.

Equation (3.3) is visually interpreted for the 50% damage case (using the mean responses of the 100 observations) in Figure 3-17. The residual vector,

$$\epsilon^{\mathrm{T}}\boldsymbol{W}\epsilon \tag{3.9}$$

Is superimposed with the side constraint,

$$\boldsymbol{\lambda^2}\delta\theta^{\mathrm{T}}\boldsymbol{W}_{\theta}\delta\theta \tag{3.10}$$

to give the cost function $J(\delta\theta)$. The minimum of the response surface is adopted as the parameter estimate. Cubic spline interpolation across the cost function surface is applied to increase the resolution of the estimate.

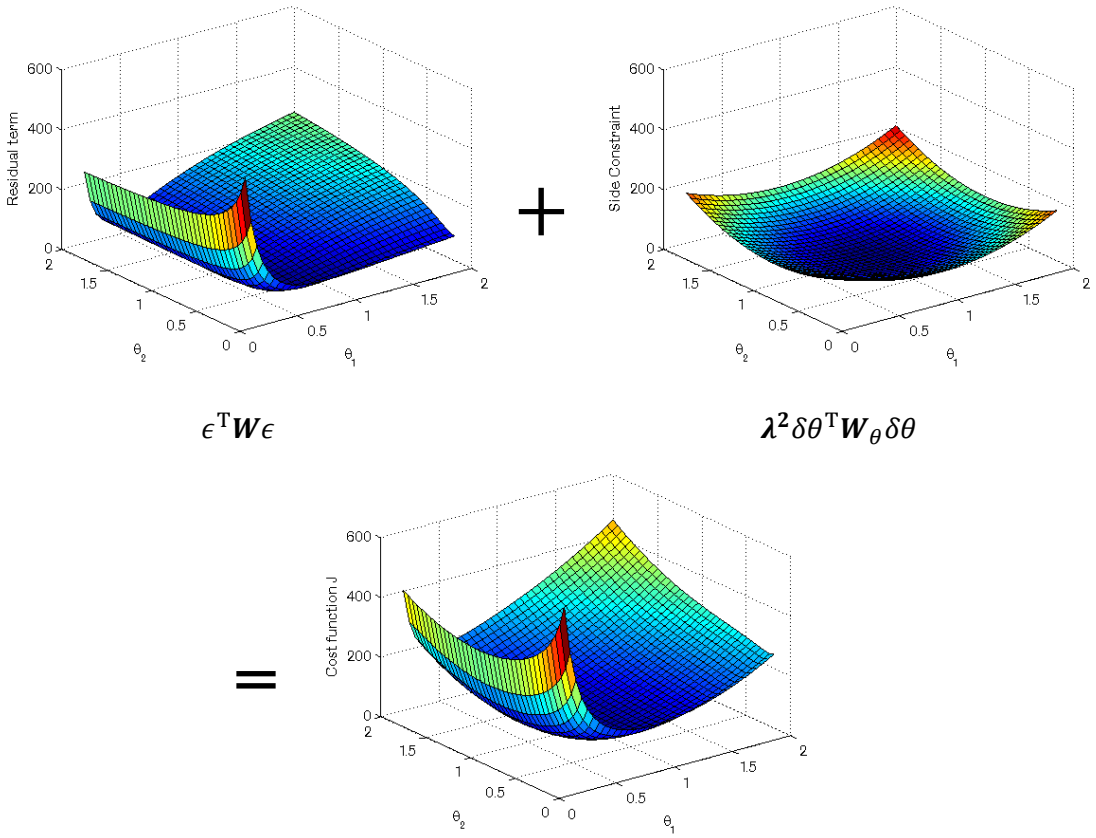

$\epsilon^{\mathrm{T}}\boldsymbol{W}\epsilon$ $\boldsymbol{\lambda^2}\delta\theta^{\mathrm{T}}\boldsymbol{W}_{\theta}\delta\theta$

$\epsilon^{\mathrm{T}}\boldsymbol{W}\epsilon + \boldsymbol{\lambda^2}\delta\theta^{\mathrm{T}}\boldsymbol{W}_{\theta}\delta\theta$

Figure 3-17 Superposition of residual and side constraint terms

The mean results for the 25%, 50%, 75% and 100% damage cases are presented in Figure 3-18 to Figure 3-25. For each state, the surface plot is given, alongside a contour plot on which the minima are marked in red. The initial parameter estimate in all cases is $\theta_{init} = [1\ 1]^{\mathrm{T}}$. The ranges over which the response figures are plotted are adjusted in some cases to aid visualisation.

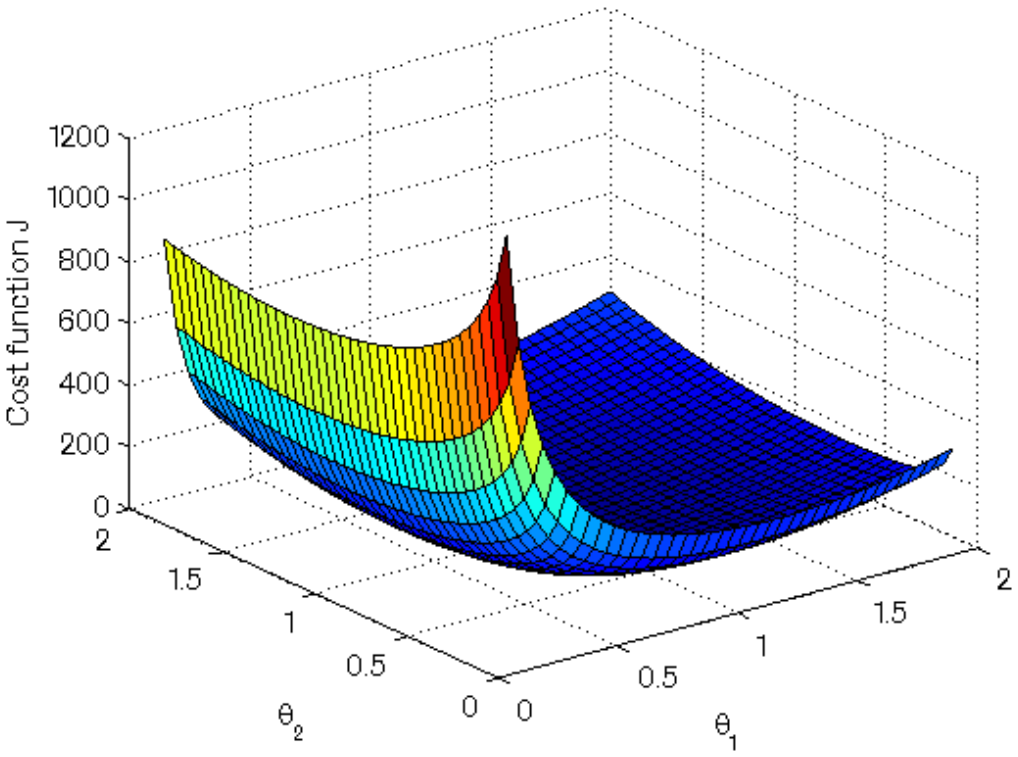

Figure 3-18 Cost function surface for 25% damage

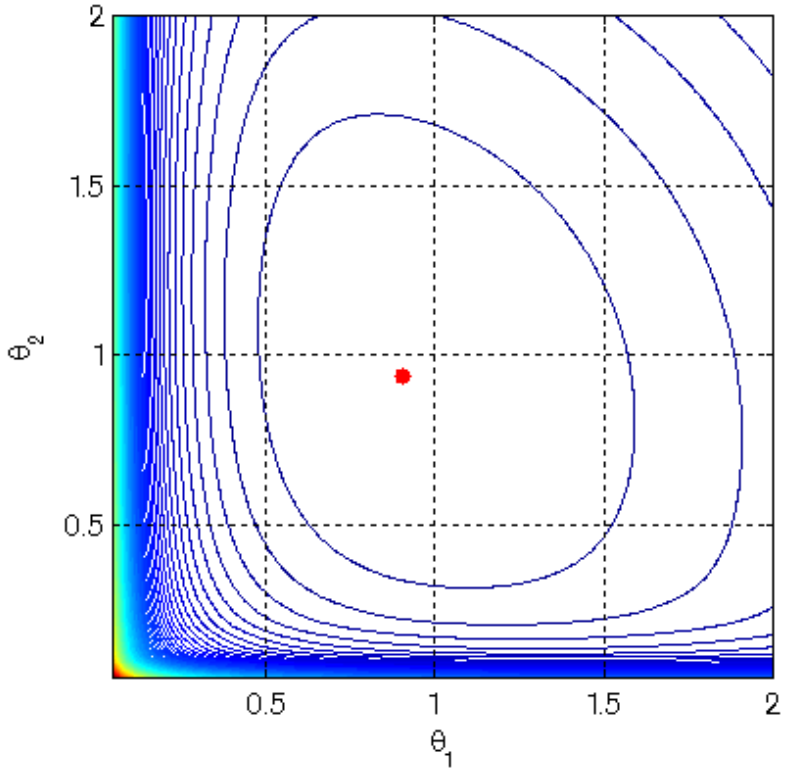

Figure 3-19 Cost function contour plot for 25% damage

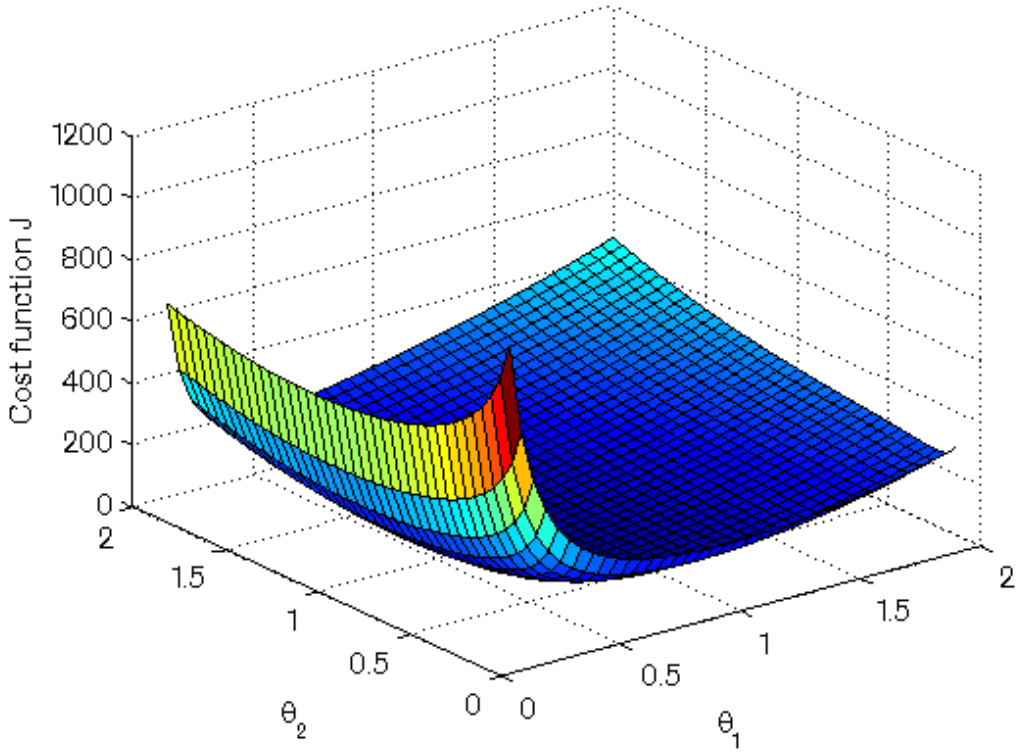

Figure 3-20 Cost function surface for 50% damage

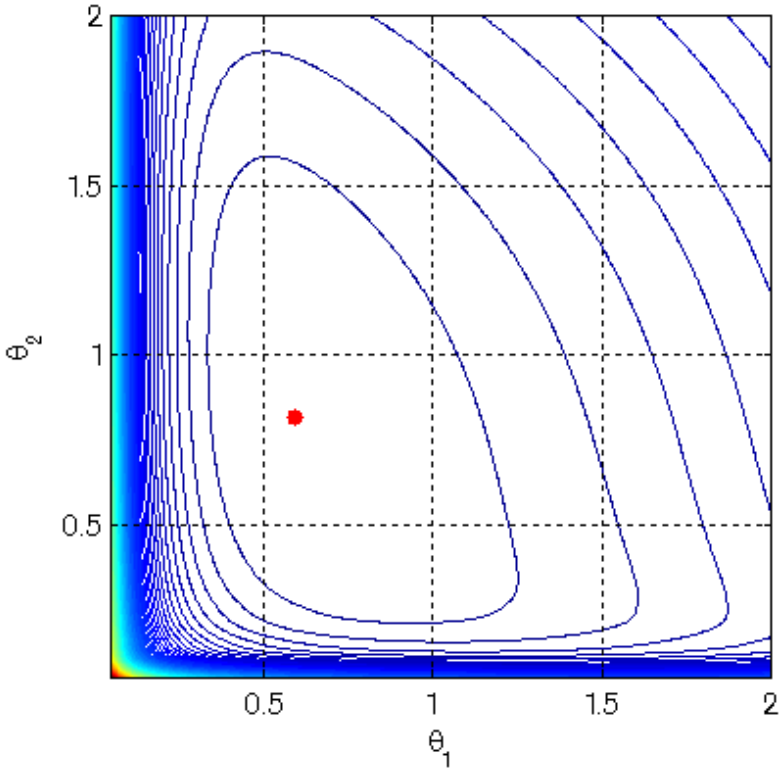

Figure 3-21 Cost function contour plot for 50% damage

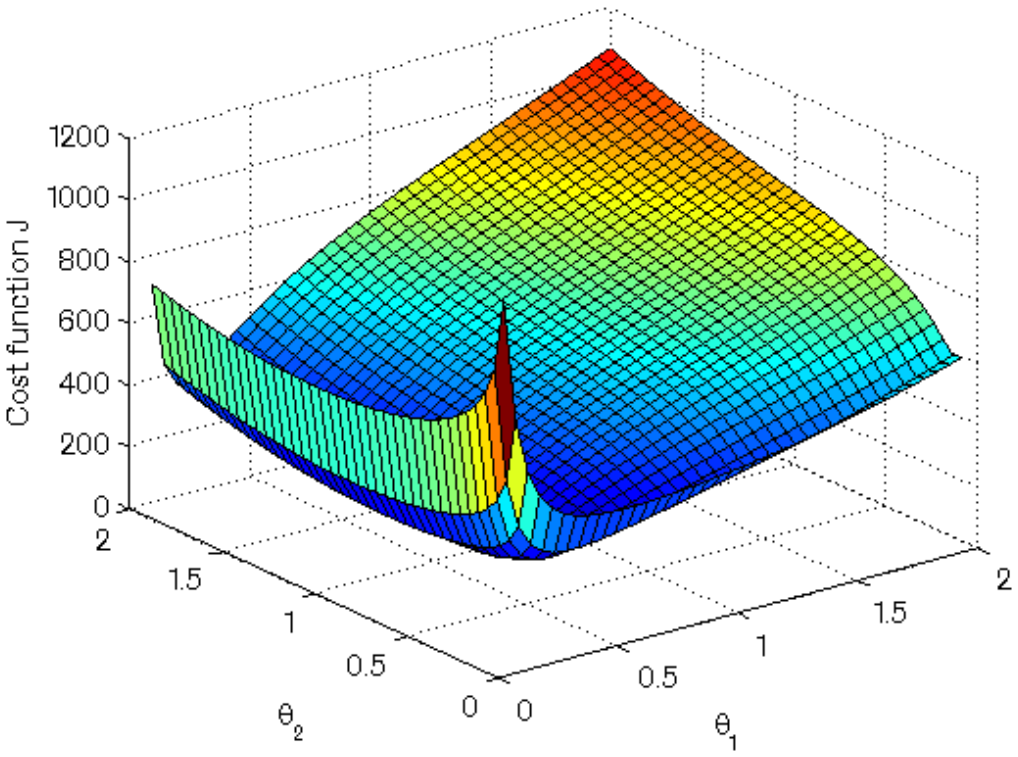

Figure 3-22 Cost function surface for 75% damage

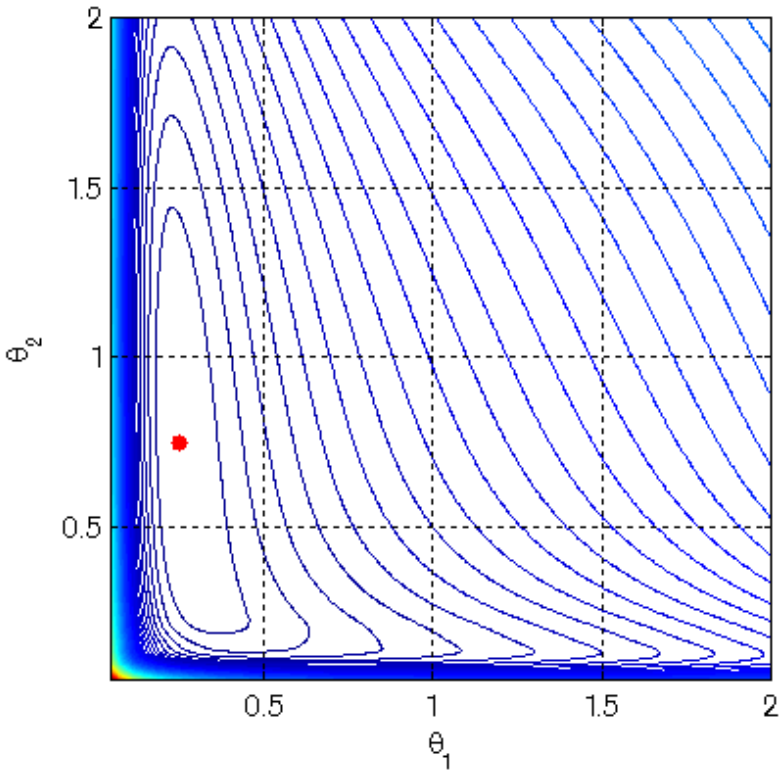

Figure 3-23 Cost function contour plot for 75% damage

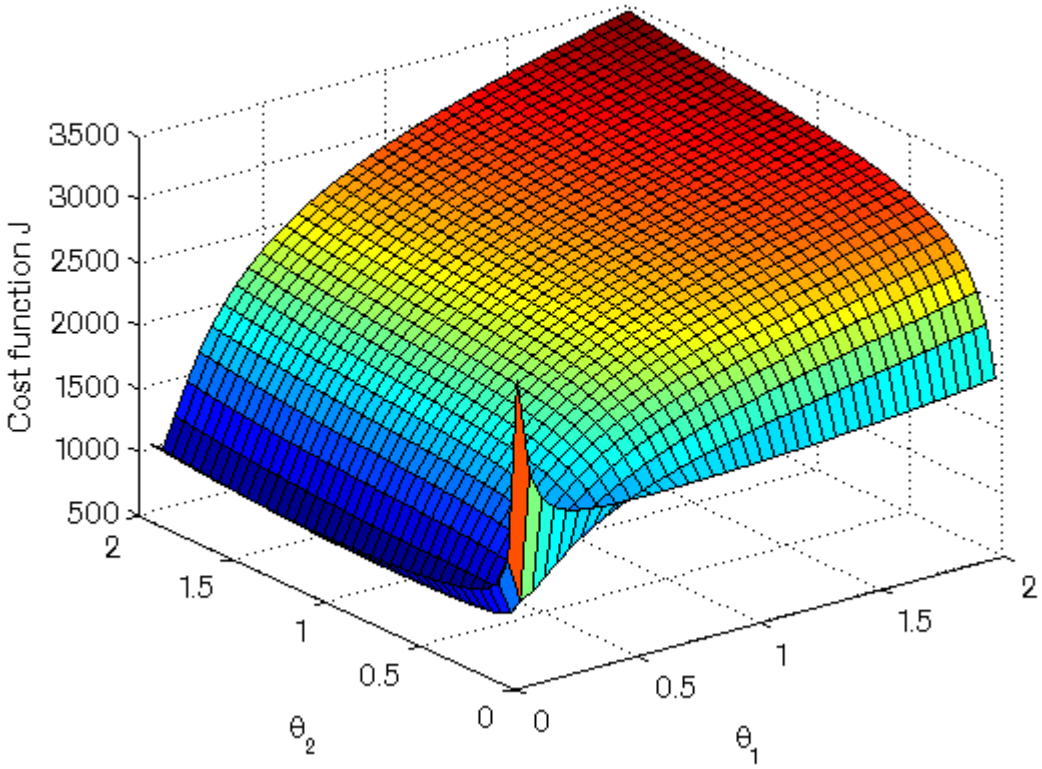

Figure 3-24 Cost function surface for 100% damage

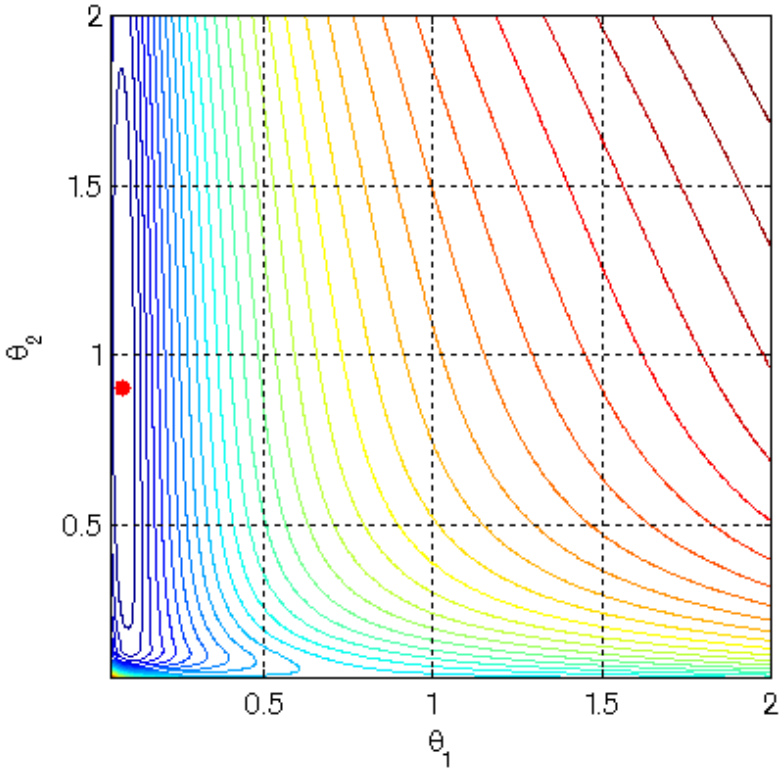

Figure 3-25 Cost function contour plot for 100% damage

The estimation method was repeated for the 200 undamaged states and 4 x 100 damaged state observations. For each observation the cost surface is calculated and the minimum found, to give the parameter estimate. These estimates, grouped by damage state, are illustrated in Figure 3-26.

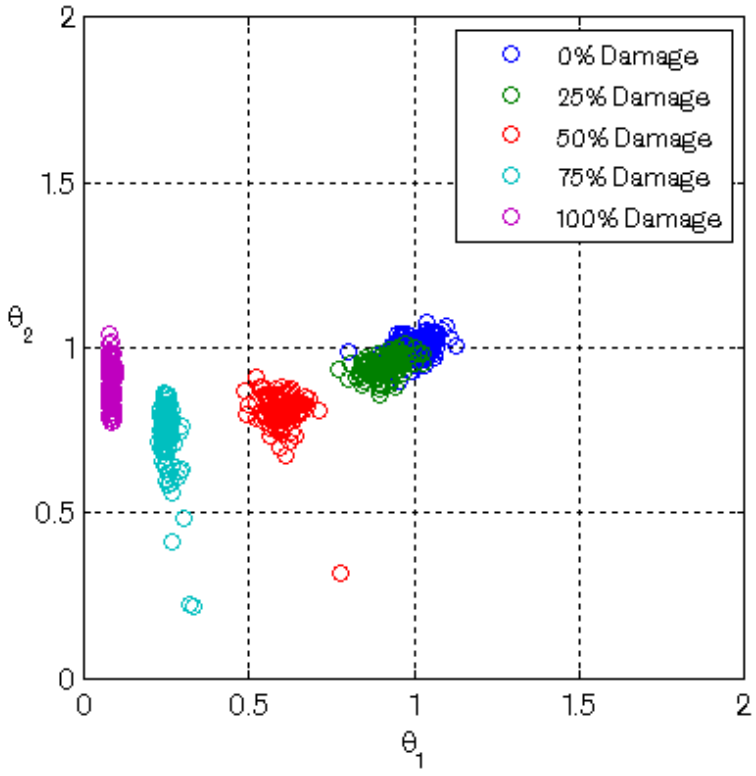

Figure 3-26 Parameter predictions- noisy data

3.6.1 Damage location outcomes

Analysing parameter estimates to arrive at damage location diagnoses is not a clear cut process. The outcome of the updating process will in general be a modified parameter vector showing changes to several or all of the updating parameters. Additional assumptions must be applied in arriving at a final diagnosis of the state of the structure; a typical assumption is that if damage has occurred, it will have done so at a single location only. The parameter displaying the largest charge is consequently diagnosed as being the location of damage. Applying this interpretation to the mean of the damage states given in Figure 3-18 to Figure 3-25, the method is shown to have worked well in locating damage to parameter θ_1. For each damage state a greater deviation in the estimate of θ_2 than in θ_1 was returned.

The results of updating using data that includes random measurement variability are presented in Figure 3-26. For the distributed data, a greater degree of confidence may be attributed to parameter predictions that display a low degree of variability. Considering parameter θ_1, some separation was observed between the undamaged and 25% damaged states, and excellent separation was observed for the 50%, 75% and 100% damage states. Some deviation was observed for the parameter θ_2, although with a lower degree of confidence attributable according to the definition above.

In summary, damage location based on model-updating has been successfully demonstrated for the experimental structure. The results overall are broadly indicative of damage occurring at the location of parameter θ_1 as opposed to θ_2, in agreement with the true damage location. It should be noted that as damage identification has been pursued only for location, the relationship between the parameter values and the extent of the experimentally introduced damage has not been established. As such, the values of the parameter estimates should not be interpreted as giving a quantitative indication of the actual damage extent. It should also be noted that as only two parameters were considered, the location task undertaken may be viewed as a somewhat simple. The handling of larger parameter sets would require a more careful application of regularisation methods.

3.7 Sensitivity to systematic variability

The model updating approach treats damage identification as a problem of parameter estimation. Parameter estimation methods are constrained to consider only the specified parameter set, and may give inaccurate results when the features used for correlation are sensitive to factors outside this set. A pertinent example is the known sensitivity of measured structural responses to environmental and operational changes. This sensitivity presents a substantial challenge for damage identification methods, particularly as the sensitivity of the responses to such effects may be in the order of or greater than their sensitivity to damage. A second example is in the sensitivity of numerical model response predictions to modelling choices, for example the specification of mesh density, the simplifications made in modelling complex regions of the structure such as joints, and the parameterisation of damage.

In this section, the challenges presented by such *systematic* variabilities are illustrated for the wingbox structure. Two cases are considered:

- Sensitivity to measurement variability, such as that caused be environmental and operational effects. The specific form of variability introduced is a mass modification.
- Sensitivity to model-form discrepancy. The analysis presented in the previous section is repeated using the predictions of a model with modified boundary conditions.

For both cases, the analysis of the previous section was repeated while maintaining the same parameter set. This serves to illustrate the inaccuracies that may arise by attempting to reconcile discrepancies caused by the sensitivity to factors outside the parameter set. The damage identification outcomes are compared to those observed in the previous section and conclusions drawn.

3.7.1 Sensitivity to systematic measurement variability

Due to equipment limitations, it was not possible to investigate environmental effects (temperature, humidity etc) for the wingbox structure. However, it was possible to investigate the effect of mass modification as an example of the problems presented by environmental and operational factors. While somewhat abstract in the context of the wingbox, mass variability is a pertinent factor in civil and aerospace applications. In aerospace applications this variability may arise due to fuel usage and payload changes; in civil applications it may arise, for example, from vehicle loading on a bridge deck or rainwater loading on a building.

Structural modification data were acquired from the structure prior to the introduction of saw-cut damage. The modification made was the addition of a small mass in the vicinity of the damage location marked in Figure 3-1. Bolts were attached to each side of the stringer at the specified location using ceramic glue, on which additional weights were mounted as shown in Figure 3-27. Data were recorded for masses of 25g, 50g, 75g and 100g (including bolts and fasteners). The acquisition sequence was identical to that for the saw-cut damaged data, with 100 x two-average samples taken for each state.

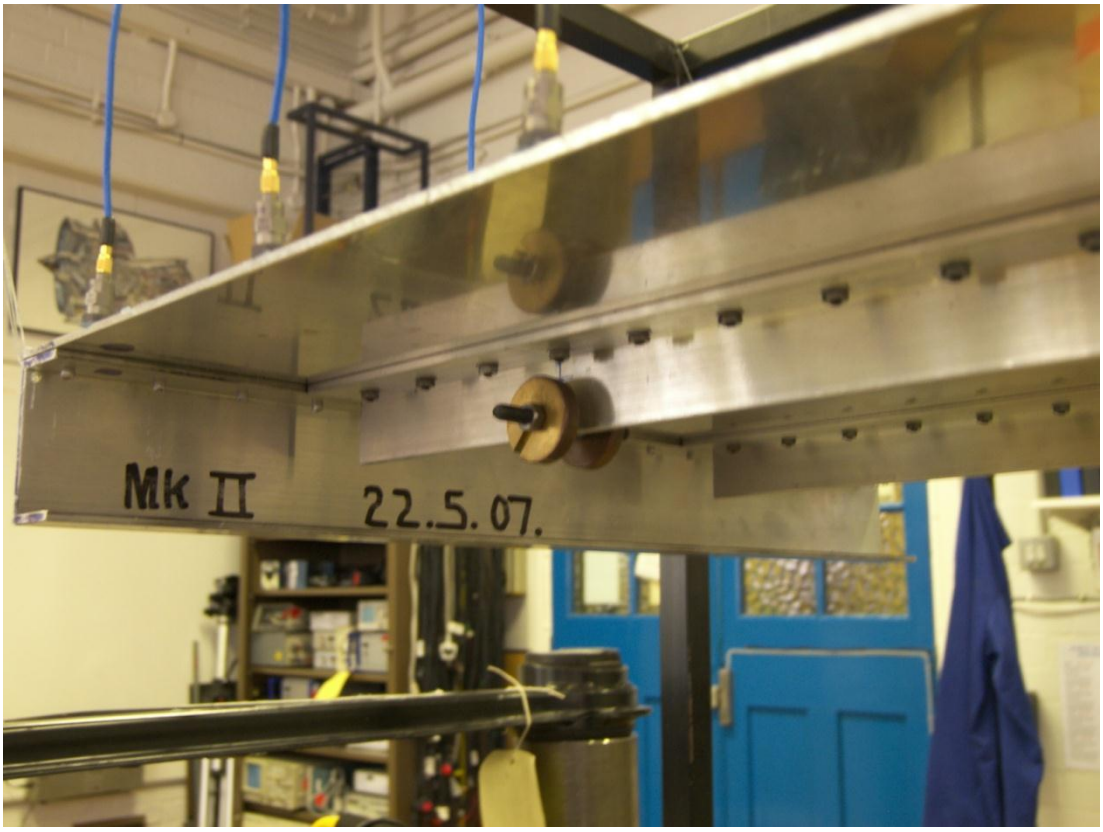

Figure 3-27 Mass addition (100g)

Natural frequencies were again estimated through curve fitting of the data recorded at sensor location S9. The experimentally measured responses are summarised in Table 3-3.

Mode No	ω_m (mean ± 1 std) [Hz]				
	0%	25g	50g	75g	100g
E2	70.05 ±0.25	70.05 ±0.27	70.01 ±0.21	70.00 ±0.20	69.94 ±0.23
E3	104.39 ±0.29	104.36 ±0.35	104.23 ±0.23	104.15 ±0.30	104.04 ±0.32
E4	116.16 ±0.33	116.09 ±0.33	116.12 ±0.38	116.19 ±0.37	116.20 ±0.37
E5	186.07 ±0.28	185.77 ±0.24	185.50 ±0.28	185.17 ±0.24	184.76 ±0.25
E6	197.32 ±0.27	196.91 ±0.36	196.65 ±0.31	196.28 ±0.25	195.85 ±0.31
E7	236.32 ±0.52	235.95 ±0.52	235.38 ±0.64	235.04 ±0.64	234.67 ±0.64
E9	294.02 ±0.38	293.74 ±0.31	293.46 ±0.32	292.99 ±0.31	292.40 ±0.35
E11	346.8 ±0.32	346.15 ±0.40	345.56 ±0.37	345.04 ±0.37	344.62 ±0.38
E12	391.83 ±0.74	390.60 ±1.47	389.47 ±0.77	388.09 ±0.71	385.63 ±0.66
E15	457.44 ±0.41	454.76 ±0.40	452.51 ±0.45	450.72 ±0.86	449.09 ±0.68

Table 3-3 Mean and standard deviations of experimentally identified modes: mass addition

3.7.2 Mass addition outcomes

The updating process was repeated, maintaining the parameter set $\theta = [\theta_1\ \theta_2]^{\mathrm{T}}$. The damage location outcomes for the mass addition case are given in Figure 3-28.

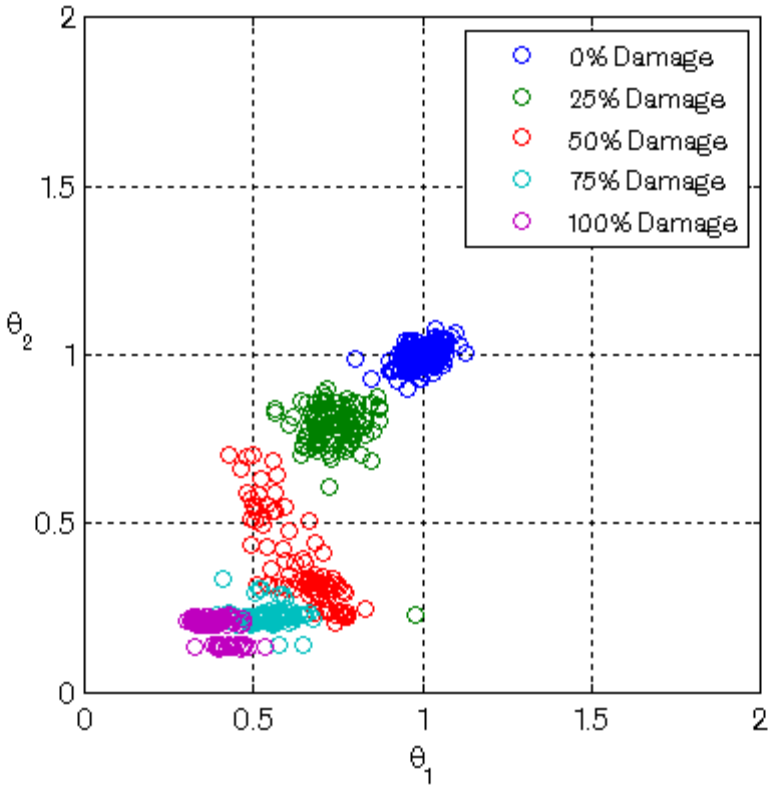

Figure 3-28 Parameter predictions: mass addition

It is observed from Figure 3-28 that the stiffness parameters θ_1 and θ_2 are both sensitive to the responses changes that were in truth caused by mass addition. Given the outcomes in Figure 3-28, it is possible that damage may erroneously be diagnosed as having occurred at one or both locations, particularly for the 75g and 100g mass addition states. In addition, it is observed that a reasonably high level of confidence (as defined in the previous section) would likely be ascribed to such a conclusion. As an aside, the bi-modality in parameter θ_2 observed for the 75g and 100g states is explained as an artefact of interpolation difficulties encountered near the lower extremes of the parameter range.

It is concluded that the estimation method employed coped poorly in the mass addition case, with mass addition being indicative of the kind of systematic deviations in the data that may be caused by sensitivity to factors outside the parameter set, such as operational and environmental effects.

The investigation of the robustness of updating approaches to other environmental and operational factors that are representative of those experienced in application, for example arising from temperature effects, is left as a topic for further work. As a

final comment, the similarity in the nature of the response changes that occur has led to the development of methods for simulating damage using mass modification, allowing the selection of features that may be indicative of damage without the requirement that damage be introduced into the structure [39].

3.7.3 Sensitivity to model-form variability

The aim of this section is to investigate sensitivity to model-form discrepancy. The model updating approach is closely tied to the numerical model adopted. The outputs of the numerical model are dependent upon the decisions made in the modelling process. By necessity, simplifications are made in this process. For simple regions of the structure, for example the individual top-plate and stiffening elements from which the wingbox is constructed, modelling decisions will relate to the element type and mesh density. The representation of complex regions of the structure, for example jointed connections, presents a more challenging set of modelling choices which may encompass simplifications such as generic elements [20, 40-42]. Large discrepancies in the model-predicted responses may arise from the choices made for such regions. For damage identification studies, the form of the damage model represents a further important decision. Parameterisation of damage is discussed in the previous chapter. Finally, it should be noted that it is not usually possible to 'parameterise' such decisions in a meaningful way, hence even if the importance of such decisions is recognised it is not generally possible to incorporate them into a parameter estimation process.

By adopting an updating strategy based on correlation of response deviations as opposed to two-stage correlation as detailed in Chapter 2, first-order modelling effects may be removed. The objective here is to illustrate the sensitivity to model-form discrepancies that remains. The approach taken is to compare the damage location outcomes for the initial model $\mathcal{M}_1$, with those for a similar model $\mathcal{M}_2$ which has been subjected to a minor modification.

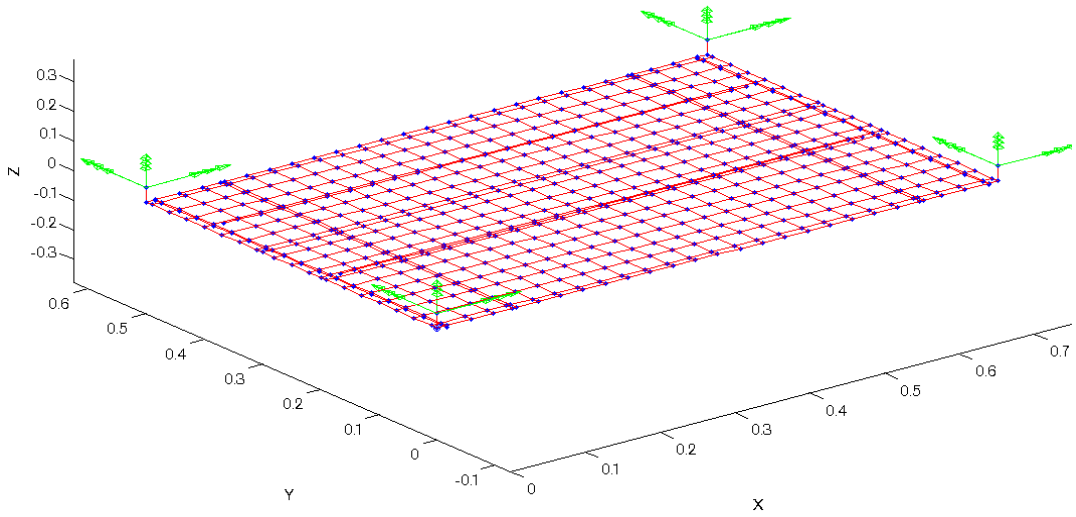

Figure 3-29 Undeformed finite element model $\mathcal{M}_2$ with modified boundary conditions

The modification adopted relates to the representation of boundary conditions. The structure was initially modelled with free-free boundary conditions. Here, a closer approximation of the experimental boundary conditions is made. They are represented as four two-node spring elements, fixed at one end and attached to the corners of the structure at the other, as shown in Figure 3-29. The stiffness of the strings is approximated from the physical springs. This represents a relatively minor modification of the structure, with no change to the modelling of damage. The analysis presented in the previous section is repeated using the predictions of the modified model.

3.7.4 Modified boundary condition outcomes

The damage location outcomes for the modified model are given in Figure 3-30. To allow comparison, the results for the initial model are repeated in Figure 3-31.

It is observed that there is little discrepancy between the sets of results for the 0%, 25% and 50% damage cases. However, a greater degree of discrepancy is observed for the 75% and 100% damage case. Of particular note is the bimodality that occurs for the 75% damage case using the modified model. This may represent an example of the non-uniqueness to which inverse approaches can be prone.

These discrepancies arise despite the response surface only being perturbed to a very small degree by the change in model-form. This may be ascribed to the shallow topography of the cost function surface in the vicinity of the solution. The slight bimodality observed for the 75% damage state indicates that the change in the

topography that resulting from the use of the modified model allied with a small degree of random measurement noise is sufficient to cause a marked shift in the local minimum of the cost surface.

Attention is turned briefly to the effect of the modification introduced on the outputs of the two models. The responses of the initial model $\boldsymbol{\omega}(\theta)_{\mathcal{M}_1}$ and modified model $\boldsymbol{\omega}(\theta)_{\mathcal{M}_2}$ for the undamaged state are given in Table 3-4. The scale of the absolute error between the modified and initial model responses is shown to be relatively large in comparison to the experimentally observed deviations due to damage given in Table 3-2.

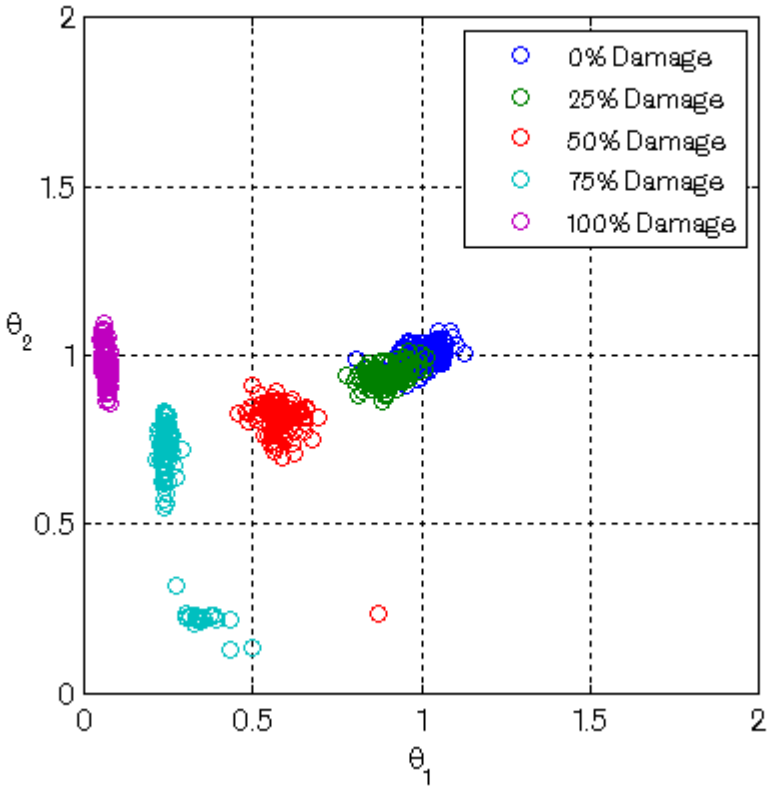

Figure 3-30 Parameter predictions: model $\mathcal{M}_2$ with modified boundary conditions

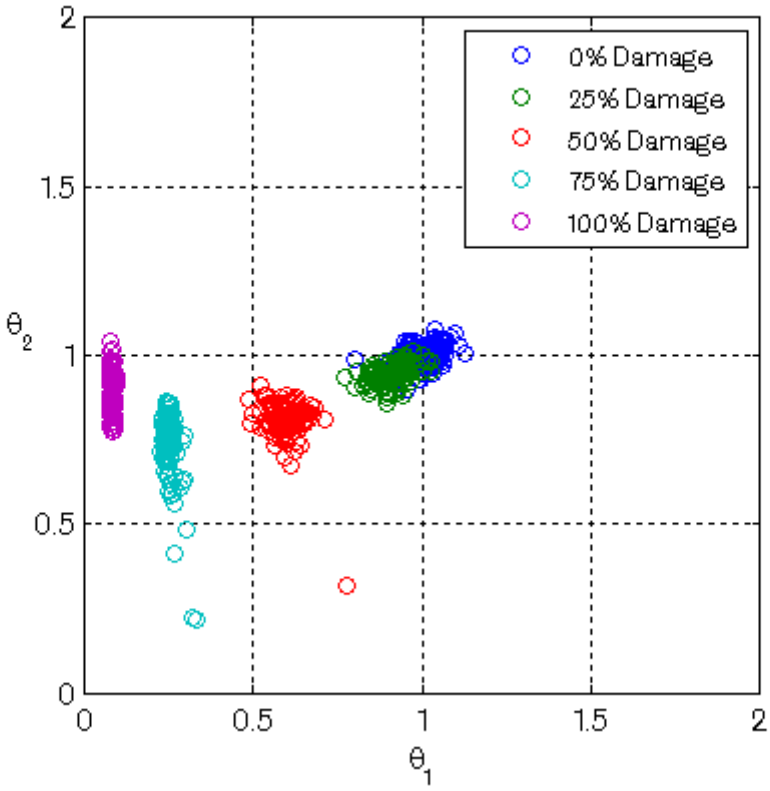

Figure 3-31 Parameter predictions: model $\mathcal{M}_1$ with free-free boundary conditions

Mode No	Predicted frequency [Hz]		Frequency deviation $\Delta\omega$ [Hz]
	$\boldsymbol{\omega}(\theta^{Undam})_{\mathcal{M}_1}$	$\boldsymbol{\omega}(\theta^{Undam})_{\mathcal{M}_2}$	
A8	64.44	69.53	5.09
A9	101.65	104.20	2.55
A10	105.65	108.76	3.10
A13	180.08	181.65	1.58
A14	190.73	190.85	0.12
A16	219.84	221.77	1.93
A17	251.81	252.45	0.64
A22	343.58	343.75	0.17
A24	366.79	368.48	1.69
A27	436.93	437.48	0.56
Mean			1.74

Table 3-4 Response differences due to model-form variability- undamaged state

The scale of the errors in the responses due to model discrepancy motivates the use of response deviations, as opposed to the responses themselves, for damage identification. The mean absolute error between the response predictions of the initial model $\boldsymbol{\omega}(\theta)_{\mathcal{M}_1}$ and modified model $\boldsymbol{\omega}(\theta)_{\mathcal{M}_2}$ as a function of damage extent is given in Figure 3-32. The mean absolute error between the deviations $\delta\boldsymbol{\omega}(\theta)_{\mathcal{M}_1}$ and $\delta\boldsymbol{\omega}(\theta)_{\mathcal{M}_2}$ that are used for updating from equation (3.4) are also given.

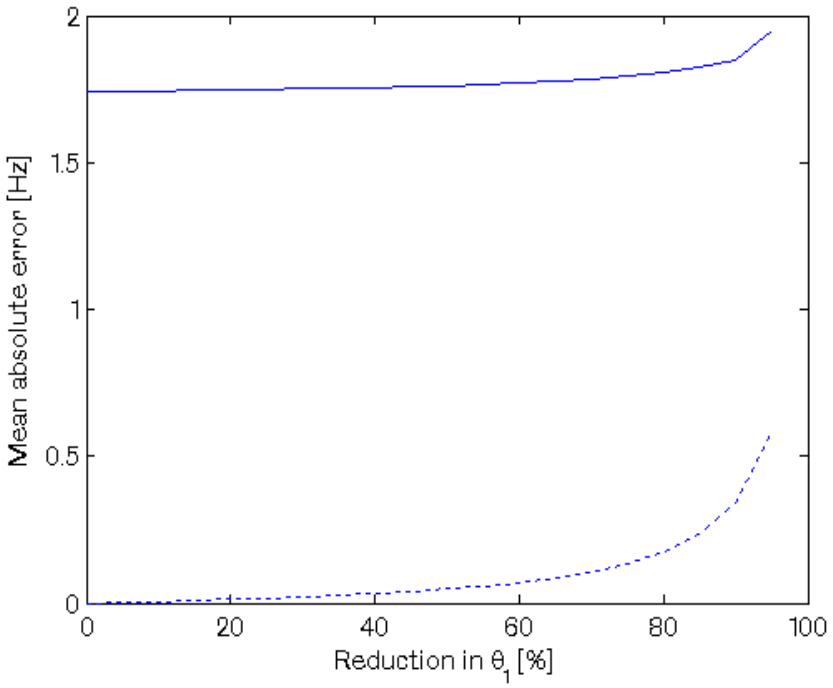

Figure 3-32 Mean error between responses $\boldsymbol{\omega}(\theta)_{\mathcal{M}_1}$ and $\boldsymbol{\omega}(\theta)_{\mathcal{M}_2}$ (solid line) and deviations $\delta\boldsymbol{\omega}(\theta)_{\mathcal{M}_1}$ and $\delta\boldsymbol{\omega}(\theta)_{\mathcal{M}_2}$ (dashed line) vs. stiffness reduction in parameter θ_1

From Figure 3-32, it is apparent that the mean error between the deviations $\delta\boldsymbol{\omega}(\theta)_{\mathcal{M}_1}$ and $\delta\boldsymbol{\omega}(\theta)_{\mathcal{M}_2}$ is substantially lower than the corresponding error between the responses $\boldsymbol{\omega}(\theta)_{\mathcal{M}_1}$ and $\boldsymbol{\omega}(\theta)_{\mathcal{M}_2}$. This is taken as providing justification for adopting an updating approach based upon achieving correlation of response deviations, with first-order model-form effects being removed. The error in the deviations for the undamaged state is zero, and remains low until the stiffness of parameter θ_1 has been reduced by over 60%. However, at high damage extents this error rises rapidly, in agreement with statements in the literature [31] that the effect of model discrepancies grows as the extent of damage increases.

In summary, the outcomes of damage identification using model-updating are shown to be sensitive to modelling choices, and this sensitivity increases at higher damage extents. The changes observed in the damage location outcomes were relatively small, even at higher damage levels. However, the model discrepancy leading to the change in the estimates introduced was also comparatively minor. What is perhaps crucial is that the changes did not relate to the damage model. It would be expected that changes to the damage model would have a far greater effect on damage identification outcomes, although assessing such changes would in general be complicated by the lack of equivalence between different parameterisations of damage.

3.8 Discussion

Damage location using numerical model updating was investigated using an experimental structure. Location was restricted to two parameters. It was found that the model-updating approach gave results that were in agreement with the actual damage introduced, although interpretation of the parameter estimates required the application of a degree of engineering judgement. Sensitivity of the damage location outcomes to systematic measurement variability, such as may arise from environmental and operational factors, and to model-form discrepancy were also investigated. It was found that the damage location outcomes were not robust to these sources of variability.

The following pros and cons of the method have been highlighted through the case study and from the literature.

Pros

- No damage state data used
- For the case investigated, a reasonable level of robustness to random measurement noise was observed
- First-order model-form errors can be eliminated through minimising damage residuals, rather than adopting a two step approach
- A simplified representation of damage (stiffness loss) was sufficient for level 2 damage identification (i.e. location)

Cons

- Purely parametric estimation has no mechanism for handling non-stationary effects affecting the measured data, such as might be associated with environmental and operational factors. Similarly, there is no mechanism for handling model-form errors beyond elimination of first-order effects
- Parameter estimation is prioritised over feature selection. Feature selection is driven by the availability of correlated features as opposed to their ability to identify damage. The natural frequency responses are relatively insensitive to the damage parameters.
- Interpretation of damage location outcomes is not trivial, even for the two-parameter case considered.

3.8.1 Further issues

The size of the parameter set is limited by issues of rank deficiency and ill-conditioning, contributed to by the number of candidate parameters typically being greater than the number of measured features available for updating. A reduction of the candidate set must be sought and/or additional constraints applied However, care must be taken in how such constraints are applied. Regularisation in the form of constraints to keep parameter values similar - as commonly applied when the updating objective is improved model correlation - is not appropriate for the damage identification task where only a single parameter or small subset of parameters is expected to vary. Adapting the regularisation term to promote the identification of a small number of parameter changes is challenging and, if applied

wrongly, would have the potential to mislead. Higher levels of damage identification, for example to allow assessment, may require more complex parameterisations than the stiffness reduction adopted in the present chapter. This may increase the number of parameters that are candidates for updating, presenting further problems for regularisation. Methods for subset selection to reduce the size of the parameter set are presented in [16] and [43].

The outcomes of parameter estimation are highly sensitive to the choice of regularisation parameter λ. In this case study, a value of λ that worked well for illustrating other issues was adopted *a posteriori*. Setting regularisation parameters in a principled manner represents a challenging task, the subtleties of which are not covered here; useful references on the topic include [29, 32, 44].

Finally, the computational expense of the method is considered. For any damage identification method, the computational effort may be separated into that required for development and that for diagnosis. For the model updating-based method, the development costs are associated with the development and validation of a single FE model, and, while potentially challenging, would not be expected to be onerous. However, the updating process must be executed for every observation for which damage diagnosis is sought. The computational cost associated with each diagnosis is thus not negligible and may be limiting in some applications, for example if real-time diagnosis is sought.

3.8.2 Conclusions

Damage location using numerical model updating is a mature topic of research whose difficulties are well known from the technical literature. Attention has been drawn in particular to:

- The unsuitability of parameter estimation-based methods for handling data that displays sensitivity to factors outside the parameter set. These factors include sensitivity to model-form discrepancy and systematic measurement variability in addition to non-represented parametric factors.
- Issues of ill-conditioning and rank deficiency. A major limitation of the model-updating approach is the restriction on the number of parameters considered at one time.

- Difficulties in ascribing confidence to the damage identification outcomes. In addition to the computational cost of running multiple parameter estimation cycles, the requirement to set regularisation parameters in a principled manner at each iterative step makes the quantification of uncertainty for model updating extremely challenging.

The conclusions above motivate the investigation of an alternative approach to model-based damage identification in Chapter 4.

4 Forward Prediction of Law-Based Models for SHM

4.1 Introduction

Law-based models may be applied to the SHM problem using one of two approaches: the model updating (or inverse) approach; and the forward approach. The focus of the inverse approach is to infer the parameters of a law-based model using measured data. The observed data is 'fed back' through the model and damage inferred from consequent parametric changes. Conversely, the focus of the forward approach is to make law-based predictions with which to inform the inference problem. The inference problem itself may consequently be handled using other means, for example a statistical (or data-based) model that is better-suited to handling uncertain problems. The law-based model serves as a proxy for the experimental structure, allowing predictions to be made for states that cannot be observed experimentally.

Law-based approaches to SHM have largely focused on inverse methods for solving the damage identification task In Chapters 2 and 3, methods based upon the updating of physically-representative numerical models were introduced and demonstrated. While a degree of success was achieved, several drawbacks were also identified. Principal among these were the need to solve a typically ill-conditioned inverse problem with the resulting need for careful regularisation; and the need for the *a priori* identification of damage sensitive parameters. Additionally, it is apparent that these methods may not be well-suited to dealing with inherent modelling and experimental uncertainties. The effect of these uncertainties may be hugely detrimental to the accuracy of the identification task.

The alternative means of employing a law-based model is in the forward mode. Forward approaches for the purposes of damage identification have received relatively little attention in the SHM literature in comparison to inverse approaches. However, advances in the fields of probabilistic model validation suggest the possibility of replacing the solution of an inverse problem with the solution of a sequence of stochastic, forward problems [45]. The information gained from the law-based model may be applied at various levels. Key ways in which predictive modelling may be applied for SHM include:

1) To inform feature selection.
2) To generate data to inform statistical inference.

The interrogation of even a relatively coarse model may help identify features that are suitable for damage identification. If a sufficiently 'trusted' model exists, the use of the model to generate predictions for training a statistical model may be considered. Ultimately, the statistical model will be required to make diagnoses regarding the state of the experimental structure using newly-acquired data. The application of probabilistic approaches may be a step towards being able to return a level of confidence in the predictions.

In this chapter, the law-based portion of a possible two-stage forward approach is investigated. A framework for SHM is presented whereby the outputs of a forward, probabilistic modelling task are used to inform the development of a discriminative statistical model. Practical tools from the field of probabilistic verification and validation (V&V) are introduced for the purposes of aiding the predictive stage. It is found that in addition to offering a sound methodology for developing physically representative models, several of these tools may be adapted further to offer benefits specific to SHM. The second stage of the approach (the inference task) is outside the scope of this chapter. However, data-driven methods for handling the inference task are considered in Chapters 5-7.

The layout of the chapter is as follows. In Section 4.2 a framework for considering the probabilistic verification and validation of mathematical models is introduced. In Section 4.3, this framework is expanded with particular reference to the SHM modelling task, and tools that may be employed for the realisation of such a framework are introduced. This is followed by two case studies. In Case Study 1, it is proposed that by considering both analytical and experimental uncertainties during

model development, features that are sensitive to damage, while robust in the face of inherent uncertainties, may be identified. Case Study 2 provides an example of uncertainty propagation using samples drawn using a space-filling method. The calibration of simplified damage models is also considered. The subject of both case studies is the wingbox structure introduced in Chapter 3. Conclusions are summarised at the end of the chapter.

4.2 Introduction to probabilistic validation and verification

The broad aim of V&V is to lend credibility to developed numerical models and assess confidence in their predictive ability [46]. In this section a conceptual basis for considering the V&V task is given, prior to the introduction in Section 4.3 of practical methods for executing this task. A general outline for model verification and validation is presented in Figure 4-1. This outline is an interpretation of the framework developed by the ASME committee on V&V in computational solid mechanics and summarised in ref [47].

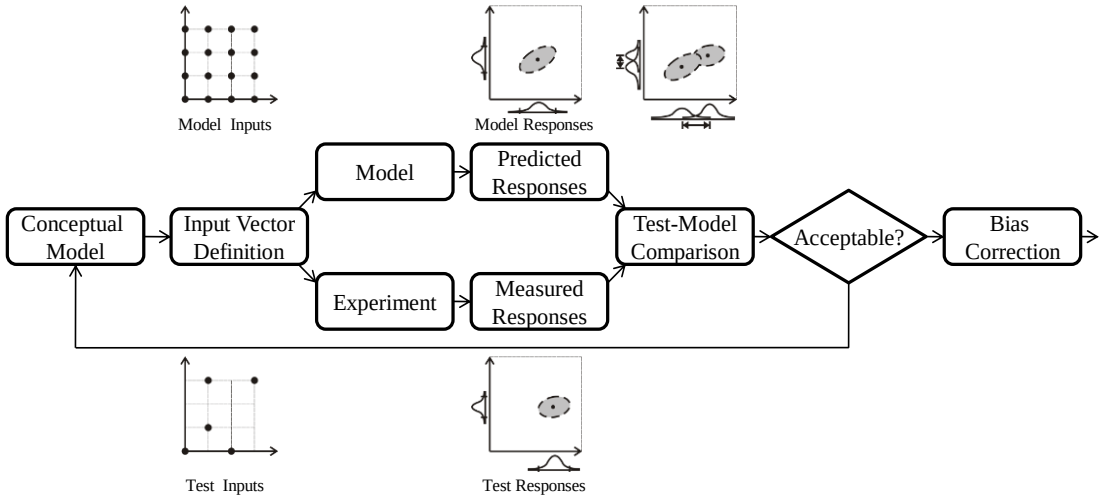

Figure 4-1 Conceptual outline of the model verification and validation task

It is perhaps helpful at this point to introduce the notion of input spaces and response spaces. Points in the input space describe the various states of the system; points in the response space describe the structure's resulting behaviour at each of the input states. Each dimension of the input space corresponds to an observable input variable: in a general sense, perhaps a temperature, a bolt pre-load value or a flow rate. The range of possible response space variables verges on the limitless, including static responses, dynamic responses, and the many response variables that

may be derived from them. In the SHM context, candidate input states might relate to operational variables (mass of fuel in tanks, swing bridge open or closed) or environmental variables (ambient temperature, wind level): however the input domain of most interest, and that which cannot typically be explored experimentally, is that relating to damage.

The ideal scenario for training a statistical model with which to discriminate damage would perhaps be to explore the entire input domain of interest at some suitably high resolution, via experiment. This is the data-based approach to SHM, and the infeasibility of introducing damage into real structures has been highlighted. Even within the domain of 'normality', experimental evaluation of all points in the input space will not be feasible. This sparseness of points in the input domain is represented in Figure 4-2(a). The challenge for the predictive modelling task is to probe those regions of the input space that are of interest, but which are not accessible via testing for some reason (economical, technical or otherwise, Figure 4-2(c)). In principle, predictions may be made for all outstanding points of interest in the input space (Figure 4-2(b)). In particular, the absence of experimentally-derived damaged state data may be tackled using predictive modelling to fill gaps in the knowledge of the input domain. Note that, despite the inference of Figure 4-2 (b), there is no assumption that the model need be evaluated at every point of the discretised input space. A more efficient methodology may be to evaluate the model at a carefully selected subset of the input points, and from these to infer the responses of the model at intervening points. Sampling methods are returned to in Section 4.3.5.

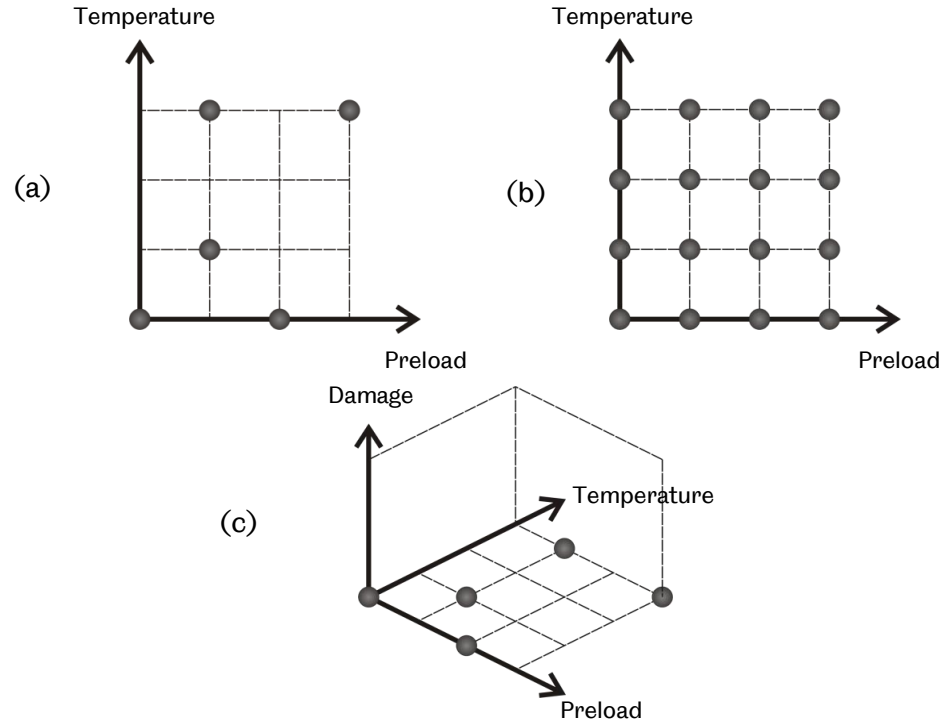

Figure 4-2 Illustrations of (a) the experimental input space; (b) the model input space; and (c) the experimental input space, illustrating the non-measurable 'damage' domain

4.2.1 From deterministic to probabilistic prediction

Any computer model can be regarded as a function $f(\boldsymbol{x})$ of its inputs $\boldsymbol{x}$. Although this function is deterministic and based upon known mathematical relationships, it is often sufficiently complex as to be mathematically intractable. So from a practical point of view $f(\boldsymbol{x})$ may be regarded as an unknown function, since the output of the model is unknown until the model is run. A similar argument may be applied to an experiment. The output of an experiment may be considered to be a function of its inputs $g(\boldsymbol{x})$. Given the inherent modelling and experimental uncertainties, both $f(\boldsymbol{x})$ and $g(\boldsymbol{x})$ may be treated as probabilistic functions of their inputs.

The implication of applying non-deterministic functions $f(\boldsymbol{x})$ and $g(\boldsymbol{x})$ is that the elements of the outputs will be non-deterministic. The uncertainty in the inputs and outputs may be expressed in many forms: parametric probability distributions (e.g. a normal distribution, with mean and standard deviation specified for each element); upper and lower intervals; upper and lower confidence bounds at a specified confidence level etc. Several of these are discussed in the context of structural dynamics problems in [46]. A graphical illustration of two responses distributions is given in Figure 4-3(a) and Figure 4-3(b), as predicted by a model and as measured experimentally. The variance of the distributions is symptomatic of the *aleatoric*

uncertainties - those pertaining to random occurrences, for example the random measurement variability associated with the experimentally-derived responses. Comparing model and experimental response (Figure 4-3 (c)), the bias between the means of the distributions is symptomatic of *epistemic* uncertainties – those pertaining to a lack of knowledge, related for example to model-form errors arising from discretisation and modelling decisions, and systematic measurement errors. The responses thus contain a quantified statement of belief in both the pattern of changes in the included responses arising from changes in the input state, and in the robustness of the responses to modelling and experimental error, both aleatoric and epistemic. Deciding whether to accept or reject the model on the basis of the test-model comparison is arguably the crux of the model validation task. The decision is ultimately subjective, but statistical metrics are available to reduce this subjectivity [48].

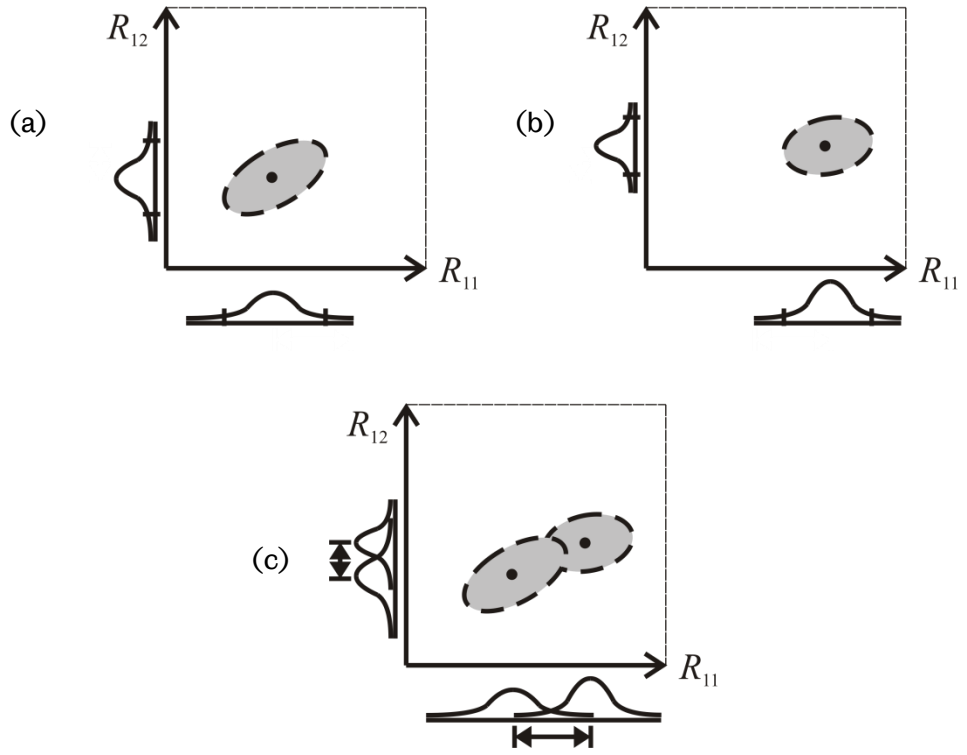

Figure 4-3 Hypothetical depiction of two responses observed at a given input point (a) as predicted by a model; and (b) as measured experimentally. The responses are compared in (c).

If the model is accepted within the defined thresholds, the probabilistic model response matrices and their associated input states may be passed forward to the feature selection and statistical modelling stage. As an intermediate step, bias

correction may be employed to account for any residual bias between test and model responses [49]. The task for the statistical model is to estimate the input state vector for newly presented observations of the experimental structure responses, ideally with an associated degree of confidence. A grey-box model may be profitably employed for this task i.e. a model for which physical laws for the input-output relationships are partially known *a priori*, as opposed to unknown (black box) or fully specified (white box). These themes are not considered further in this chapter, but are returned to in Chapter 8.

On a point of generality, no restriction is placed upon the model employed for prediction. FE models validated against experimental data have particular applicability to vibration-based SHM. However, model selection is dependent upon the problem at hand. Analytical models are available for predicting delamination of composite panels, for example, although they can be complex to implement – the model proposed in [50] requires 85 equations to be implemented. A combination of predictions from an analytical model *and* a numerical model may be considered better still - more credence may be given to model predictions if there is a high level of agreement between multiple models. The counter consideration is developmental cost.

The following section is concerned with introducing principled methods for predictive modelling with specific reference to SHM. Tools from the field of probabilistic model verification and validation are advocated for this task. Applications of these tools for sensitivity analysis and uncertainty propagation are illustrated via case studies in Sections 4.4 and 4.5.

4.3 Validation and verification for SHM

In Figure 4-4 a possible 'SHM-specific' realisation of the model validation process outlined in Figure 4-1 is presented. In this section each step of the proposed framework is expanded, with particular reference to methods available from verification and validation. Specific adaptations of these methods for the purposes of SHM are highlighted.

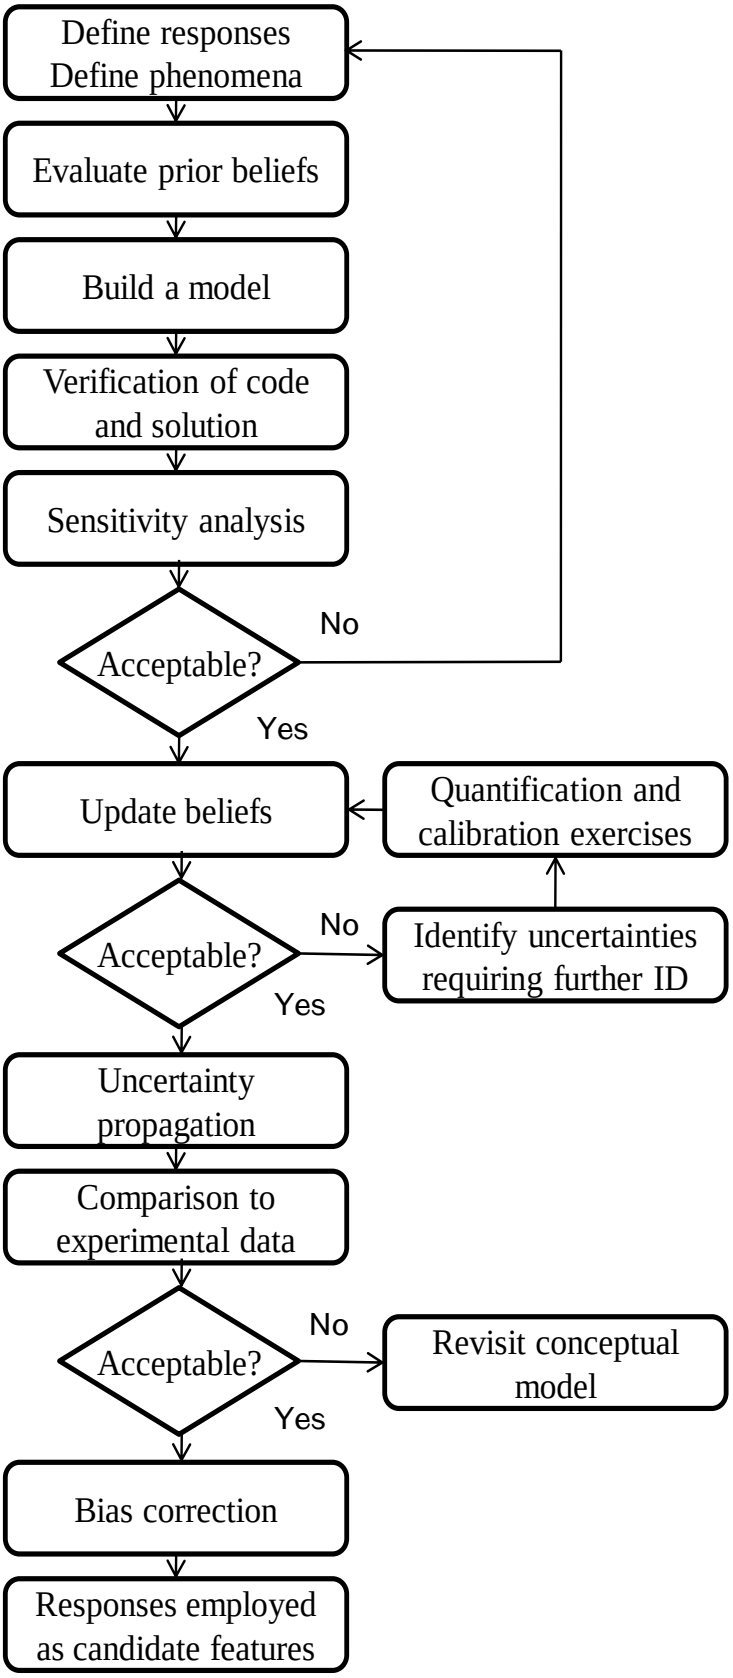

Figure 4-4 A model verification and validation framework for SHM

4.3.1 Define responses and model uncertainties, and evaluate prior beliefs

An extensive set of both responses and model uncertainties may be defined. The ability to specify a large range of responses, and be able to make an assessment of their suitability as indicators of damage *a priori*, is a major benefit of the forward approach, and requires the method of damage introduction to be included as an input at the sensitivity analysis stage. The balancing factor is computational cost. The computational efficiency of the sensitivity analysis employed for making this assessment alleviates this concern to a large extent for all but the most computationally-demanding models.

The evaluation of prior beliefs with regard to the inputs requires two factors to be stated for each input. First, an estimate is made of the 'impact' of uncertainty in each input on the responses. This will be specific to the chosen responses. Secondly, prior beliefs regarding the uncertainty related to each input are stated. Evaluation of each factor may initially be made on the basis of expert opinion, ideally encompassing past experience of both testing and numerical modelling of similar problems. Both factors may at this stage be stated either qualitatively or quantitatively. A decision as to the treatment of each source of uncertainty is made on the basis of the beliefs expressed: if the degree of uncertainty and the expected effect of a particular phenomena are believed to be low, a decision may be made to fix the value of the phenomena; if either are high or unknown, sensitivity analysis and uncertainty quantification will be employed to update these initial beliefs. A principled approach to this task may make use of a table such as that given in Figure 4-5. This is an example of what is variously referred to as a Phenomena Identification and Ranking Table (PIRT) or Input Uncertainty (IU) Map. This map will continually be returned to and revised during the course of the validation process as more information is acquired. Ultimately, the table should contain all the quantified information required to inform the uncertainty propagation task. For SHM tasks, the specification of the damage model is always likely to be among the inputs identified for further attention.

Input		Effect	Uncertainty	Current Status
Material Properties	Mass Densities: all components	5	Low uncertainty	Validate component models
	Elastic Moduli: all components	5	Low uncertainty	Validate component models
	Poisson's Ratios: all components	3	Low uncertainty	Validate component models
Geometric Properties	All components	5	Low uncertainty	Validate component models
Numerical Parameters	Mesh Density: all components	3	Unspecified	Verify component models
	BCs: Free-Free suspension	4	Unspecified	To be specified
Environmental Conditions	Temperature	3	Unspecified	Fixed - Assumed controllable
	Humidity	1	Unspecified	Fixed - Assumed controllable
Bolted Joint Properties	Choice of joint model	4	Unspecified, epistemic	To be specified
Riveted Joint Properties	Choice of joint model	4	Unspecified, epistemic	To be specified
Damage Model Properties	Choice of damage model	5	Unspecified, epistemic	To be specified

Figure 4-5 Example of a Phenomena Identification and Ranking Table

4.3.2 Verification of code and solution

The outputs of discrete computational codes are only approximations of the underlying, continuous physical equations. Verification activities may be separated into two aspects – *code verification* and *solution verification*. Code verification covers assessment of whether the computational code employed is executed correctly, without programming bugs, significant round-off errors etc. This is not considered at length here - it might be assumed that the commercial FE codes typically employed for model-based SHM have been subjected to a sufficiently rigorous degree of quality assurance.

Of more direct interest is solution verification, an assessment of the degree of error arising from the discretisation of the underlying, continuous physical equations. The degree of numerical discretisation Δx, be it a grid size or time step or similar, comprises a core modelling decision. For any discretised model, there exists a domain $\Delta x_{min} \leq \Delta x \leq \Delta x_{max}$ where the overall numerical error is dominated by truncation errors due to discretisation. This domain defines the *regime of asymptotic convergence*. For finer discretisation $\Delta x \leq \Delta x_{min}$ round-off errors begin to dominate the numerical error; coarser discretisations $\Delta x_{max} \leq \Delta x$ are not valid for the numerical solution of the continuous equations. Solution verification allows

the convergence of a numerical solution $y(\Delta x)$ to an 'exact' value y^{exact} to be assessed. This allows the numerical error of a model prediction at a given level of discretisation to be predicted [51] and a compromise reached between solution accuracy and computational cost.

The solution error is given by,

$$Solution\ error = \|y^{exact} - y(\Delta x)\| \tag{4.6}$$

In order to evaluate the solution error, an estimate of the exact solution y^{exact} must be made. An estimate can made by considering the observed rate of convergence of the model response $y(\Delta x)$ for varying mesh sizes Δx. Consider a model with three possible element sizes: fine (Δx_f), medium (Δx_m) and coarse (Δx_c).The refinement ratio R between each mesh size is given by,

$$R = \frac{\Delta x_c}{\Delta x_m} = \frac{\Delta x_m}{\Delta x_f} = \cdots \tag{4.7}$$

Note that while it is not an absolute requirement that R be kept constant between mesh refinements, the mathematical solution is simpler for constant R. Assuming monotonic convergence and that the response $y(\Delta x)$ is scalar, the rate of convergence p of the observed responses is,

$$p = \frac{\log\left(\frac{y(\Delta x_m) - y(\Delta x_c)}{y(\Delta x_f) - y(\Delta x_m)}\right)}{\log(R)} \tag{4.8}$$

With this estimate of the rate of convergence made, an estimate of the exact but unknown solution may be made using the Richardson extrapolation [52],

$$y^{exact} = \frac{R^p y(\Delta x_f) - y(\Delta x_m)}{R^p - 1} \tag{4.9}$$

This allows an estimate of the solution error for models with varying mesh density to be calculated.

If the model is sufficiently complex for computational cost to be a concern, mesh density (and thus numerical accuracy) may be traded for computational efficiency. A good general approach might be to adopt a mesh density at the coarser extreme of the regime of convergent behaviour. Models in this region will be computationally

efficient, but with the knowledge that this efficiency may be traded for increased accuracy in a broadly predictable fashion if required.

An example of asymptotic convergence for the 11th mode of a 2.91mm gauge, 750mm x 500mm Aluminium plate, with free-free boundary conditions, is given in Figure 4-6. The response is evaluated for element sizes from 125mm x 125mm (requiring 24 elements) to 6.25mm x 6.25mm (requiring 9600 elements). The resulting response predictions are given in Table 4-1.

Model label	Δx_{xxc}	Δx_{xc}	Δx_{c}	Δx_{mc}	Δx_{m}	Δx_{mf}	Δx_{f}	Δx_{xf}	Δx_{xxf}
Mesh size [mm]	125.0	62.5	50.0	31.3	25.0	15.7	12.5	10.0	6.3
No elements	24	96	150	384	600	1536	2400	3750	9600
Response [Hz]	189.91	186.42	186.10	185.78	185.71	185.64	185.62	185.61	185.58

Table 4-1 Mesh sizes and model responses for the 11th mode of a rectangular plate

Applying equation (4.9) to the results in Table 4-1, the rate of convergence of the solution may be estimated to be $p = 1.78$. The exact solution is estimated to be $y^{exact} = 185.59$Hz. The solution error results are illustrated in Figure 4-6. Asymptotic convergence of the solution is clearly apparent for much of the investigated range, breaking down at a mesh size of around 10mm =10^{-2} m. The plot is augmented with the time taken to execute the model on a desktop PC (1.8 GHz dual core processor, 3Gb RAM).

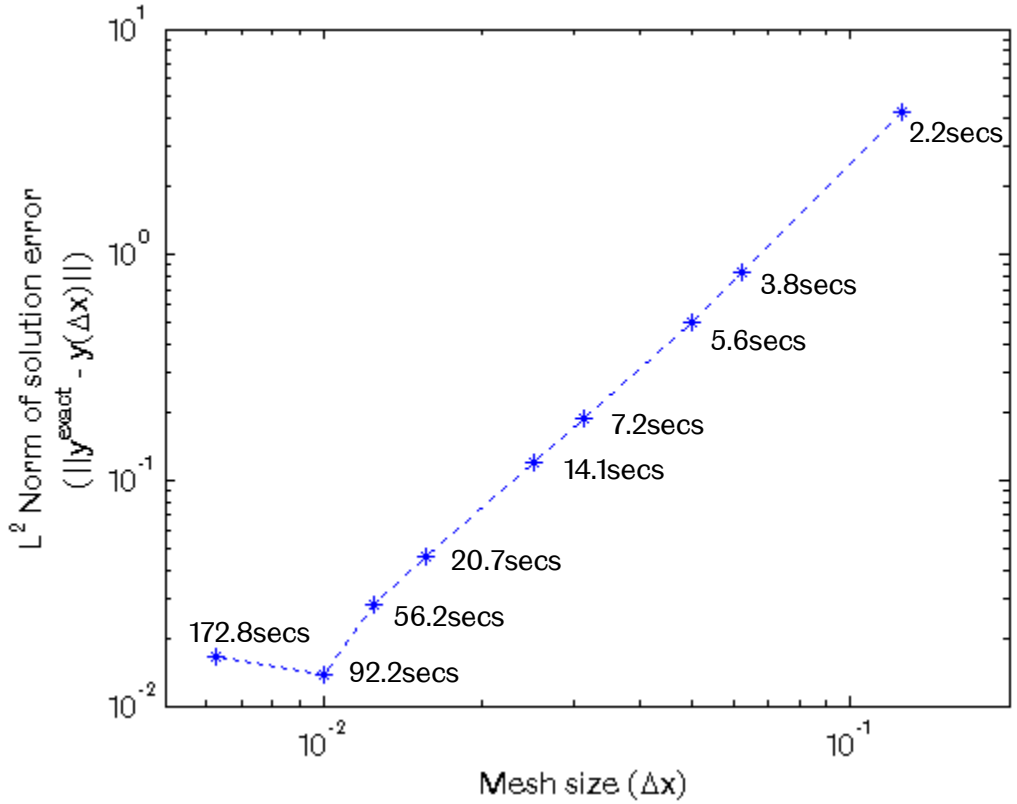

Figure 4-6 Example of convergence of solution for a thin metallic plate

With this information in hand, an informed decision on the mesh density appropriate to the SHM modelling task can be made. In early model development, the cost of execution may be prioritised over the solution accuracy, for example to allow low-cost sensitivity analysis. At later stages of development the accuracy of solution may take precedence, and the mesh density can be selected accordingly.

4.3.3 Sensitivity analysis

Sensitivity analysis allows an early assessment of the suitability of the identified responses as candidate features for damage identification. This is a key SHM-specific adaptation of the general V&V framework. By including input points relating to damage, in addition to prior beliefs relating to model uncertainties, an assessment both of the discriminatory performance of the response set and of the robustness of individual responses to modelling uncertainty may be made.

If, at this stage, it is found that the sensitivity of the chosen responses to damage is unacceptably low relative to the influence of model uncertainty *or* if it is found that the response set chosen does not allow sufficient separability of the damage states of interest, a decision may be made to review the responses chosen. The risk of

developing a sophisticated model that is nevertheless incapable of distinguishing between damage classes may thus be greatly reduced, if not eliminated entirely.

At the other end of the scale, if it is found that the sensitivity of the defined responses to damage is far greater than their corresponding sensitivity to other sources of uncertainty *and* that the responses allow acceptable separation of the input states, there may be little justification for further investment in model development, and the uncertainty quantification step introduced below may be bypassed. In practice, it would be expected that the sensitivity of responses and separability of input states would lie somewhere between these extremes. The outcomes of the sensitivity analysis would in this case be fed back into the PIRT, updating beliefs regarding the effect of uncertainties. The updated PIRT is then used to inform decisions on the best use of subsequent computational and experimental effort.

Practical tools for conducting sensitivity analysis include: effects screening, local sensitivity analysis and global sensitivity analysis - an introduction to these methods may be found in ref [53]. An application of Bayesian meta-modelling methodology applied to global sensitivity analysis of large, nonlinear FE models is presented by Becker *et al* [54].

An example of the application of effects screening for the wingbox is given here in the first case study (Section 4.5).

4.3.4 Quantification and Calibration exercises

The uncertainties expressed in the PIRT may be quantified through additional measurement and testing. The testing method to be employed will be heavily dependent upon the parameter being evaluated. It is important to note that the method of testing is not restricted to that for which predictions will ultimately be made and that in many cases modal data will not be the most appropriate for calibration. As a trivial example, directly measurable properties such as geometric parameters are most easily and accurately quantified through direct measurement.

Where the properties are not directly observable, they may be inferred using inverse estimation methods. The challenges faced when implementing inverse approaches for the purposes of SHM have been covered in some depth in Chapters 2 and 3. There are a number of distinctions between the application of inverse methods

when applied as introduced in Chapter 2 and for the purposes of quantification and calibration. The emphasis may be placed upon the identification of the properties of particular substructures, and characterised using an appropriately small parameter set. Doing so may alleviate problems of ill-conditioning in many instances. There is also a distinction in the testing methods and features that may be employed for calibration/quantification. Whereas the methods introduced in Chapter 2 are traditionally reliant upon modal characteristics, the use of whichever method is most appropriate to the given identification task (given resource constraints) is advocated. Research into alternative methods for quantifying unknown structural properties is ongoing, with numerous application-specific methods having been proposed [55]. Finally, in Chapter 2 deterministic approaches were introduced as typical of updating applications. In this chapter emphasis is placed upon probabilistic approaches, although in cases where the degree of uncertainty associated with a particular parameter is deemed sufficiently 'low', a decision may be made to pursue a deterministic approach.

The outcome of the test or measurement should be expressed in a form consistent with the method of uncertainty propagation that will subsequently be employed, for example as a distribution or an interval, and fed back into the PIRT. Testing may also be used for the purposes of model refinement, either through the validation of basic model assumptions, or the calibration of unknown model parameters. Model-updating methods may be employed for the calibration task: where updating is required, the inverse problem to be tackled should be structured so as to be as well-posed as possible.

The selection and calibration of the damage model, for example by destructively assessing individual components in isolation, is a key way in which the general modelling framework may be tailored for SHM. Deciding upon what level of model complexity is 'sufficient' for representing damage in the numerical model, and then estimating *a priori* the parameters of the damage model, is a challenging task in model-based approaches to SHM. An example is given in ref [19], where several approaches for estimating the stiffness reduction arising from the introduction of an open crack in the absence of any experimental damage data are discussed. If it is possible to conduct individual coupon or component tests for given damage scenarios, a more informed decision on the adequacy of the damage model may be reached. Direct calibration of parameters of the selected model, on the basis of

experiment, may replace the difficult and uncertain *a priori* parameter estimation task, allowing a primary source of epistemic modelling uncertainty to be tackled.

4.3.5 Uncertainty propagation

Uncertainty propagation through sampling is accurate and adaptable, allowing the calculation of probability densities, intervals, confidence intervals etc for the responses, but is the most computationally-intensive aspect of the process. Numerous choices exist for the propagation method and they are only briefly introduced here. The basic premise of uncertainty propagation is to draw samples of the uncertain phenomena, run the model for each of these sampled points, and thus form distributions for each of the response parameters. The choice of statistical sampling method has a substantial bearing on the numerical cost of uncertainty propagation. Random sampling (e.g. Monte Carlo (MC)) methods are a popular starting point largely for their simplicity. As the name suggests, sample points from the uncertain parameter distributions are drawn at random using an appropriate sampling distribution, for example a normal or uniform distribution over the parameter range. This can lead to many runs of the model being required in order to achieve an acceptable resolution for the response distributions, particularly at the extremes of the parameter distribution. Space-filling methods (e.g. Latin hypercube sampling (LHS)) may offer equivalent accuracy in fewer runs, by distributing the sample points over the parameter space in a stratified manner. Many further methods build on this: a good introduction to further sampling methodology may be found in ref [56]. Numerically efficient uncertainty propagation may be pursued by fitting an emulator to the model. Among the options available to this end is Gaussian process emulation; a useful and accessible tutorial on Gaussian process emulation is provided in [57]. Aspects of another emulation option, response surface methodology, were applied in Chapter 3. The use of model emulation is returned to in the discussion of future work in Chapter 9.

4.3.6 Comparison to experiment: validation metrics and bias correction

At the validation stage, a difficult and ultimately subjective question is posed: does the model provide an acceptable physical representation of reality? Validation metrics allow quantitative comparison of measured and predicted responses for given points in the input space, providing evidence to inform this decision. An introduction to multivariate metrics taking into account uncertainty is given by

Rebba and Mahadevan [48]. The complication for SHM is that the region of the input space that is of most interest (that pertaining to damage) typically cannot be explored experimentally, and thus predictions at points lying in this region cannot be validated. The measurable input space is instead restricted to the realm of 'normality', and whatever undamaged cases are available therein: returning to the example postulated in Figure 4-2(a), perhaps dynamic tests of the undamaged structure across a range of recorded temperatures, and with a number of levels of bolt pre-loading. Two possible scenarios for model validation for the SHM task are suggested.

The first scenario is customary in many forms of model validation, statistical, physical, or otherwise. A selection of the available experimental results are 'held back' as a *validation set*, and the remainder used as a *training set* in the model development phase detailed above. For SHM, the validation set would comprise a subset of the gathered undamaged state results. It is perhaps instructive to consider bounds on the degree of belief held in model predictions in terms of interpolation and extrapolation. It would be hoped that if the model is capable of representing those points that were made available during development accurately, it should be able to predict responses for points lying between them only marginally less accurately, with an additional, bounded, degree of uncertainty. Depending on the relative locations of the training and validation points in the input space (and it is likely that those points at the extremes of the input space will have been incorporated into the training set rather than the validation set, in order to increase the range of validity of the model to the greatest extent possible), this will typically be a process of interpolation. This form of validation predominantly allows assessment of confidence in the model's fidelity-to-data in the region of the input space relating to 'normality' (i.e. the undamaged states). It is accepted that bounds on the prediction uncertainty will grow rapidly upon extrapolation into regions outside 'normality'.

A second scenario would be to modify the structure in order to probe areas of the input space outside the realm of normality and, if possible, to explore those areas that are representative of damage. Ideally, this would be done by gathering data from the structure in its damaged state: of course, it is the absence of this very data that is the primary motivation for using predictive modelling. In the atypical case that damage data is available for some or all of the damage states of interest, the data

may be treated as described above. An alternative is to evaluate the model in further, experimentally-realisable regions of the input space. An example of this would be through attaching a mass to the experimental structure in some non-destructive fashion, and running the model with an analogous mass at the corresponding location. This form of structural modification has been employed in the past in order to check whether an FE model remains physically-representative following an updating exercise. However, mass addition may be of additional interest in predictive modelling for SHM, as the input space explored may be related to that which would be probed through introducing damage. Indeed, mass addition has been investigated as a means of simulating damage in data-driven approaches to SHM [58, 59]. The possibility of probing regions of the input space away from normality raises an interesting possibility for validating SHM-specific models. Here, interest lies largely in gaining knowledge of *changes* in structural responses arising from different input states rather than in the normal state responses themselves. If it is possible to quantify changes in response due to damage-like structural modifications both numerically and experimentally, then validation may be performed based on the model's ability to predict these changes. From an SHM perspective, far greater credence may be ascribed to a model that is shown to be capable of predicting such changes to an acceptable level, over a model that is validated only in its 'normal' state.

Following both scenarios, it would be parsimonious to consider rerunning the modelling exercise using a new training set comprising both the initial training set *and* the validation set. This avoids any waste of expensive experimental data- particularly in the special case of damaged state data, where great benefit would be derived from reincorporating this data to refine the model. If the agreement between the test and physical model responses is found to be acceptable, the model predictions are adopted as candidate features for the statistical model. Prior to training the statistical model, bias correction may be employed in order to account for any remaining response bias. If the agreement is deemed unacceptable, the model and the experimental method may be revised.

The application of validation and verification tools to the SHM task is illustrated via two case studies. The first case study investigates the use of sensitivity analysis to aid the initial assessment of features and to guide subsequent model development. The second illustrates uncertainty propagation when applied to damaged-state response

predictions. The subject used for both case studies is the wingbox structure introduced in Chapter 3.

4.4 Case study 1: Identification of robust, damage sensitive features

4.4.1 Approach

This case study provides an example of how a numerical model may be interrogated in order to inform feature selection. The achievement of reliable model-based predictions for SHM is complicated by the presence of both modelling and experimental variability, motivating the investigation of methods that take these uncertainties into account. Of particular concern are the effects on the model predictions of model-form errors arising from modelling decisions.

Taking finite element (FE) modelling as a pertinent example, these epistemic uncertainties arise from factors including the choice of element types; selection of mesh densities; *a priori* specification of parametric values; and decisions on modelling simplifications, including the representation of complex substructures such as joints. The number of modelling decisions made can become large, even for relatively simple structures. It is therefore desirable to be able to assess the effect of these modelling decisions on the quality of model predictions at an early stage in model development. *Statistical effects analysis* is a computationally efficient tool for this process, and has found application in the assessment of factors that dominate variability in the model responses [60]. In the case study presented here, statistical effects analysis is used to assess the relative impacts of both modelling uncertainties and damage on candidate features for SHM, allowing an assessment of the robustness of these candidate features in the face of epistemic modelling uncertainties.

The case study illustrates some of the procedures that may be undertaken early in the model development process. First, an initial assessment is made of the separability of the identified damage states in the response space for the model with nominally-valued parameters. The suitability of the predefined response set may thus be assessed: if it is found that it is not possible to distinguish between damage states, consideration can be given to expanding the response set prior to investing in further model development. Separability of response states in this deterministic

scenario may be considered the minimum requirement for the SHM system, with experimental and modelling uncertainties serving to confound this separability. An assessment of the sensitivity of the responses to damage is also made: in the absence of information about other sources of variability, the default strategy might be to assume that those features that are most sensitive to damage will be the most 'robust' to these sources of variability. As more is learnt about the nature of the uncertainties, this assumption will be superseded.

Secondly, statistical effects analysis is introduced for evaluating the effect of uncertain modelling parameters on the variance of the absolute responses of the structure. While primarily intended as a means of introducing the statistical effects analysis method prior to applying it to an SHM-specific problem, the outcomes are in themselves of interest in model development. If closer agreement between numerical and experimental responses is sought, parameters that contribute greatly to the variance of the responses will be candidates for calibration. Statistical effects analysis allows dominant parameter effects to be assessed on a global scale, serving as an alternative to local sensitivity analysis. While derivative-based local sensitivity analysis is a popular tool for identifying dominant parameters in linear models, it is found to be inappropriate when the model inputs are uncertain and the linearity of the model is unknown [61]: an example highlighting the drawbacks of the method when applied to nonlinear structures is given in ref [60]. The use of designs of computer experiments (DOEs) to reduce the required number of model runs is highlighted in this section.

Finally, statistical effects analysis is extended to evaluate the contribution of both uncertain factors and damage to the variance of the model outputs, allowing an assessment of the robustness of outputs in the face of uncertainty and the identification of factors requiring further consideration. The factors considered are a selection of model parameters that represent sources of epistemic uncertainty arising from modelling decisions. In the case that the model parameters are shown to contribute substantially to the output variance, modelling and testing efforts may be directed at quantifying the 'true' values of these parameters in order to improve the model predictions. Conversely, if it is found that the effect of the values assigned to particular model parameters on the responses is minor in comparison to the effect of introducing damage, there may be little justification for expending further resources on efforts to quantify these parameters. In the context of this study, the

responses are considered 'robust' to epistemic modelling uncertainty. Findings are discussed and summarised at the end of the chapter.

4.4.2 Numerical model

The numerical model used for illustration is the wingbox structure introduced in Chapter 3. The model, including the investigated damage locations, is illustrated in Figure 4-7.

Recall that the fasteners are meshed using two-node spring elements. The stiffness matrix of these elements may be user-defined, allowing the six longitudinal, torsional, shear and bending stiffness of the spring elements to be decoupled and independently specified. The choice of values for these stiffening parameters is a challenging task, and they represent a potentially large source of modelling uncertainty. Nominal values were estimated based upon the geometries and materials of the fasteners of the laboratory structure.

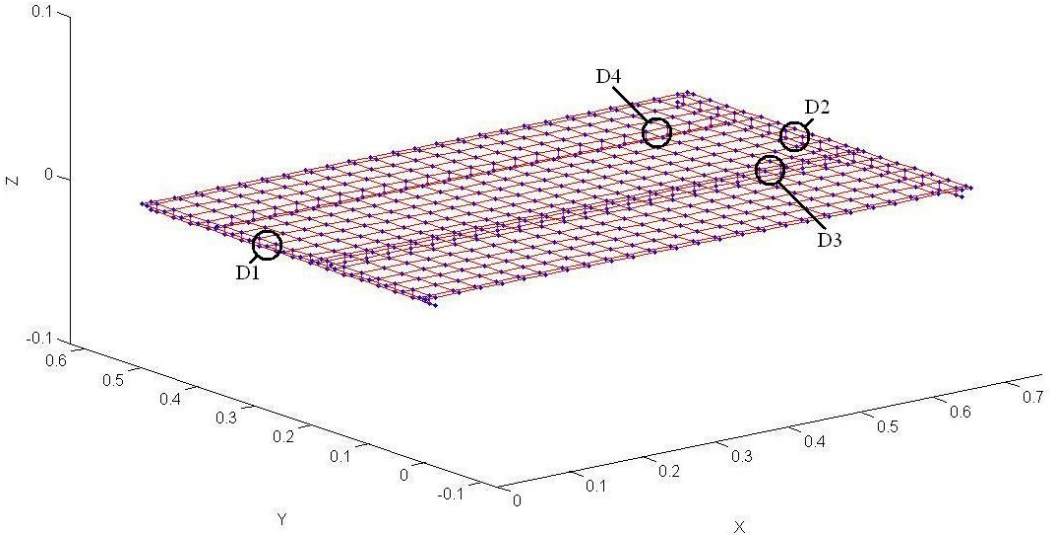

Figure 4-7 FE model of the wingbox structure, with investigated damage locations indicated

4.4.3 Specification of model uncertainties

The effect of a selection of factors on the model outputs was investigated. An advantage of employing effect screening is that an extensive group of factors may be investigated with relatively little computational expense. In this case, 15 factors were specified in total, comprising:

- The six defining stiffness parameters of the rivet elements ($K_{x,Bolt}$, $K_{xy,\ Bolt}$, $K_{xz,\ Bolt}$, $K_{xx,\ Bolt}$, $K_{zz,\ Bolt}$ and $K_{yy,\ Bolt}$)

- The six defining stiffness parameters of the bolt elements ($K_{x,Bolt}$, $K_{xy,\ Bolt}$, $K_{xz,\ Bolt}$, $K_{xx,\ Bolt}$, $K_{zz,\ Bolt}$ and $K_{yy,\ Bolt}$)
- The Young's modulus of the plate component (E_{Plate}), rib components (E_{Rib}) and stringer components ($E_{Stringer}$).

The nominal values of these factors are given in Table 4-2. Two further levels are specified for each factor, corresponding to the maximum and minimum bounds of 'belief' in the parameter values and this defining an interval for each factor. For the factors pertaining to Young's modulus these intervals are set to be relatively narrow, at ±2% of the nominal value. The values of the stiffening elements represent a somewhat broader region of lack-of-knowledge, and this is reflected in the range of values specified for bounds of belief. These are set at ±50% relative to the nominal value. Note that the application of sensitivity analysis pursued here is *non-probabilistic* – the distribution of the model outputs is not sought. A survey of non-probabilistic approaches to propagation in structural dynamics is given in [62]; application of non- probabilistic approaches to nonlinear problems is discussed in [63].

Parameter label	Parameter description	Parameter Value			Unit
		Min	Nominal	Max	
θ_1	E_{Plate}	67.6	69.0	70.4	GPa
θ_2	E_{Rib}	67.6	69.0	70.4	GPa
θ_3	$E_{Stringer}$	67.6	69.0	70.4	GPa
θ_4	$K_{x,Bolt}$	$1.3x10^8$	$2.5x10^8$	$3.8x10^8$	N/m
θ_5	$K_{xy,\ Bolt}$	$4.8x10^6$	$9.6x10^7$	$1.4x10^8$	N/m
θ_6	$K_{xz,\ Bolt}$	$4.8x10^6$	$9.6x10^7$	$1.4x10^8$	N/m
θ_7	$K_{xx,\ Bolt}$	54.5	109.0	163.5	N/m
θ_8	$K_{yy,\ Bolt}$	69.5	139.0	208.5	N/m
θ_9	$K_{zz,\ Bolt}$	69.5	139.0	208.5	N/m
θ_{10}	$K_{x,Rivet}$	$4.1x10^7$	$8.2x10^7$	$1.2x10^8$	N/m
θ_{11}	$K_{xy,Rivet}$	$1.6x10^7$	$3.2x10^7$	$4.8x10^7$	N/m
θ_{12}	$K_{xz,Rivet}$	$1.6x10^7$	$3.2x10^7$	$4.8x10^7$	N/m
θ_{13}	$K_{xx,Rivet}$	18.0	36.0	54.0	N/m
θ_{14}	$K_{yy,Rivet}$	23.0	46.0	69.0	N/m
θ_{15}	$K_{zz,Rivet}$	23.0	46.0	69.0	N/m

Table 4-2 Nominal, maximum and minimum values of factors identified for the FE model

Damage was investigated at single locations on each of the 4 stiffening elements, as indicated in Figure 4-7. Damage is introduced as a reduction in the Young's modulus of a single element. It is acknowledged that this is a hugely simplified representation

of real damage, and that problems are encountered in calibrating a parametric value (the Young's modulus of an element) to the actual level of damage in the structure (for example, the depth of an incipient crack) [19]. Methods of introducing parameterised damage into FE models form a body of research in their own right, and are not covered here. For the present study it is sufficient that the dimensions of the input space are representative of damage in some meaningful fashion, and changes in stiffness parameters are deemed suitable for illustrating the method.

The responses employed for identifying dominant parameters are the 10 natural frequencies of the structure $\boldsymbol{\omega}^{Undam}$, as identified in Chapter 3 and described in Table 3-1.

The responses used for assessing the sensitivity and separability of the damage states, and the robustness of the candidates features, are normalised frequency shifts associated with the same set of ten natural frequencies,

$$\delta\boldsymbol{\omega} = diag\left(\boldsymbol{\omega}^{Undam}\right)^{-1}\left(\boldsymbol{\omega}^{Dam} - \boldsymbol{\omega}^{Undam}\right) \tag{3.5}$$

Where $\boldsymbol{\omega}^{Dam}$ and $\boldsymbol{\omega}^{Undam}$ are the values of the natural frequency vector in the current and the undamaged state respectively

4.4.4 Damage sensitivity and separability of damage states

A key motivation for the use of forward, predictive modelling-based methods for SHM is that input-output relationships may be evaluated prior to training the damage identification algorithm, simply by running the model in its damaged states. The objectives of this section are, for the initial model with nominally-valued parameters, to assess:

- the sensitivity of individual features, and
- the separability of the states in the multivariate response space

A deterministic representation of the response states is readily achievable by running the model once per input state. Assessment of the separability of the data would ideally be performed using the statistical approach that will ultimately be employed for damage identification. In this case, principal component analysis (PCA) was employed to provide a graphical visualisation of the separability of the damage

states. Responses were generated for single-site damage at each of the locations indicated in Figure 4-7.

The separability outcomes are given in Figure 4-8 and Table 4-3. Four damage extents were assessed for each location, corresponding to 10%, 20%, 30% and 40% reductions in the Young's modulus of the identified element. A quantified statement of the separability of the damage states in the response space would be preferable to the graphic illustration of Figure 4-8, but is not pursued here. From Figure 4-8 it might be concluded that:

1) In the absence of uncertainty, damage at the four locations is separable using the specified statistical method
2) The degree of separability appears to increase with the extent of damage

It might also be noted that the degree of separability is rather greater for some damage locations than for others- note in particular the example of damage at locations D1 and D2. The concern raised is that separability may diminish once uncertainty is taken into account, and may prompt consideration of *probabilistic* forward modelling to allow quantification and propagation of uncertainties, and thus the provision of richer information for the damage identification task.

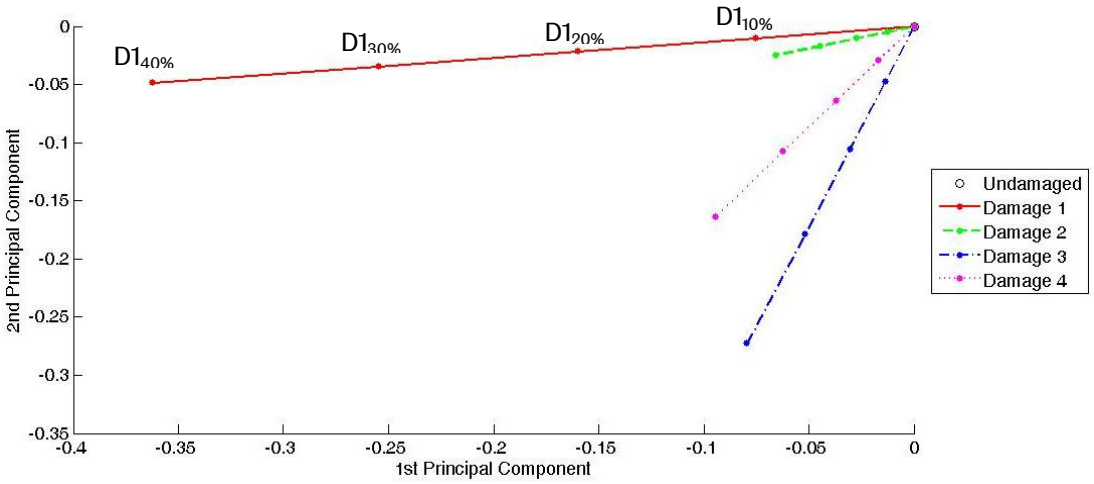

Figure 4-8 PCA analysis of the structure predictions, illustrating separability of the damage states

Damage state	Response value [%]									
	$\delta\omega_1$	$\delta\omega_2$	$\delta\omega_3$	$\delta\omega_4$	$\delta\omega_5$	$\delta\omega_6$	$\delta\omega_7$	$\delta\omega_8$	$\delta\omega_9$	$\delta\omega_{10}$
Undamaged	0	0	0	0	0	0	0	0	0	0
$D1_{20\%}$	-0.020	-0.018	-0.056	-0.002	-0.012	-0.021	-0.132	-0.010	-0.060	-0.035
$D2_{20\%}$	-0.019	-0.002	-0.002	-0.003	0.000	-0.004	-0.002	-0.001	-0.092	-0.006
$D3_{20\%}$	-0.026	-0.028	-0.014	-0.008	-0.043	-0.002	-0.001	-0.004	-0.032	-0.093
$D4_{20\%}$	-0.026	0.000	-0.053	-0.012	-0.020	-0.004	-0.001	-0.049	-0.038	-0.040

Table 4-3 Values of $\delta\omega$ for 5 states of the nominal structure

The results of this analysis for five of the input states (undamaged and 20% damage at each of the four locations in turn) are given in Figure 4-3. It is noted that the sensitivity of the responses to damage is 'small' – the largest element of $\delta\boldsymbol{\omega}$ for the damage states investigated at this level corresponds to a 0.132% reduction in the specified natural frequency. In practice, deviations of this magnitude may be swamped by variability caused by temperature deviations, for example. This is useful information for assessing the suitability of the selected response set prior to further development of the model and, subsequently, the damage identification algorithm.

As noted in the introduction to the case study, when the extent of the confounding modelling and measurement uncertainties are unknown, the best assumption that can be made is that responses that are sensitive to damage are the most likely to be robust to the unknown uncertainties. This is somewhat unsatisfactory as the basis of a damage identification methodology, and further prompts use of forward modelling methods in order to better understand the sources and effects of uncertainty.

4.4.5 Statistical effects analysis for identifying dominant parameters

Statistical effects analysis allows evaluation of the extent to which the variability of the model outputs is explained by variability in the model inputs, either individually (main effects) or in combination (interaction effects). A standard application of statistical effects analysis in model development is in identifying dominant structural parameters. Dominant parameters may be candidates for subsequent *calibration*, in order to reduce bias between model and experimental outputs, and/or *quantification*, if random variation of parametric values of the real structure is expected. In this section the model inputs are the factors identified in Table 4-2, and the outputs considered are selected natural frequencies of the structure, $\omega_{1\text{-}10}$. Effect analysis is used to evaluate sources of variability in the outputs of the undamaged model. Only the results of the main effects are illustrated below.

First, a DOE is selected specifying input parameter sets at which the model will be run. Options for the DOE include full-factorial designs, partial factorial designs and Taguchi orthogonal designs. A full-factorial design for this problem, investigating only the maximum and minimum factor values, would require 2^{15}= 32768 runs of the model, which is clearly prohibitive. While the model used for illustration here required only around 12 seconds per run on the desktop PC used, the two-level full-factorial analysis would require around 4.5 days to execute what is intended to be a cost-efficient step in the model development process. In the analysis below, a Taguchi orthogonal design was used [64], allowing consideration of the main and linear interaction effects of 15 factors at two levels in 16 runs without substantial loss of accuracy compared to the full-factorial approach. The design used is given in Table 4-4. The minimum factor value given in Table 4-2 is coded as Level 1 and the maximum factor value is coded as Level 2.

Run	Factor values														
	θ_1	θ_2	θ_3	θ_4	θ_5	θ_6	θ_7	θ_8	θ_9	θ_{10}	θ_{11}	θ_{12}	θ_{13}	θ_{14}	θ_{15}
1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
2	1	1	1	1	1	1	1	2	2	2	2	2	2	2	2
3	1	1	1	2	2	2	2	1	1	1	1	2	2	2	2
4	1	1	1	2	2	2	2	2	2	2	2	1	1	1	1
5	1	2	2	1	1	2	2	1	1	2	2	1	1	2	2
6	1	2	2	1	1	2	2	2	2	1	1	2	2	1	1
7	1	2	2	2	2	1	1	1	1	2	2	2	2	1	1
8	1	2	2	2	2	1	1	2	2	1	1	1	1	2	2
9	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2
10	2	1	2	1	2	1	2	2	1	2	1	2	1	2	1
11	2	1	2	2	1	2	1	1	2	1	2	2	1	2	1
12	2	1	2	2	1	2	1	2	1	2	1	1	2	1	2
13	2	2	1	1	2	2	1	1	2	2	1	1	2	2	1
14	2	2	1	1	2	2	1	2	1	1	2	2	1	1	2
15	2	2	1	2	1	1	2	1	2	2	1	2	1	1	2
16	2	2	1	2	1	1	2	2	1	1	2	1	2	2	1

Table 4-4 Taguchi orthogonal design L16

The global contribution of input parameters to variability in the model outputs is evaluated by applying statistical tests to the data generated by the model. Effect analysis was performed here using the R^2 (or R-square) statistic [65] which, for each output, estimates the ratio between the variance attributable to a given effect and the total variance observed. The R^2 statistic for each main or interaction effect is given by,

$$R^2 = 1 - \frac{\sum_{l=1}^{N_{level}} \sum_{j=1}^{N_{data}^{(l)}} \left(y_j^{(l)} - \bar{y}^{(l)}\right)^2}{\sum_{j=1}^{N_{data}^{(l)}} \left(y_j - \bar{y}\right)^2} \tag{4.6}$$

where y_j is the jth output of interest, N_{level} is the total number of levels at which the output of interest was evaluated, and $N_{data}^{(l)}$ is the total number of data points available at the given level of the output of interest. The value of the R^2 statistic lies between 0 and 1, with values close to 1 indicating that the effect contributes substantially to variance in the given output. The results of the R^2 effect analysis conducted on the 15 factors and 10 outputs of the wingbox problem are presented in Figure 4-9.

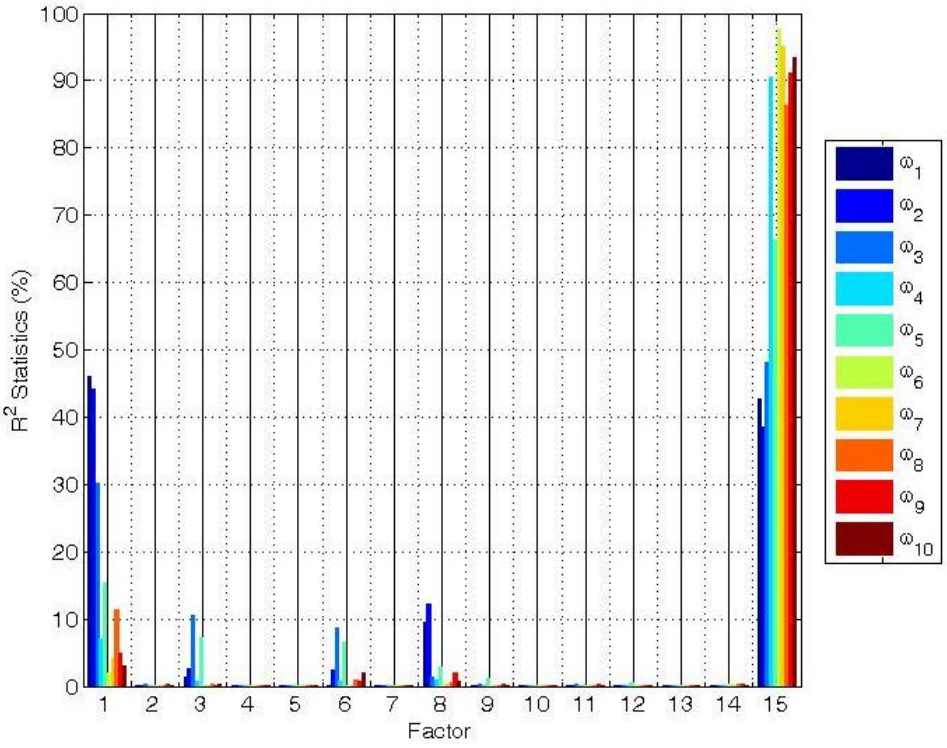

Figure 4-9 R^2 analysis of main effects on ω values of the undamaged structure

From Figure 4-9 it is observed that variability in the outputs is dominated by θ_{15}, the bending stiffness of the rivet elements about the local z-axis - this axis corresponds to the global x-axis in Figure 4-7. The Young's Modulus of the plate θ_1 was identified as a further substantial source of variability for several of the outputs. Three other factors (θ_3, θ_6 and θ_8) were identified as more minor sources of variability. The impact of the ten remaining factors was considered negligible. Parameters to which the responses are sensitive would be candidates for calibration and/or uncertainty

quantification in further model development; there is little justification for further assessment of parameters to which the responses are globally insensitive, and a decision would likely be made to omit them from further uncertainty analysis.

4.4.6 Statistical effects analysis for assessing robustness of features for SHM

Statistical effects analysis may be extended to assess the robustness of model-generated response features to epistemic uncertainties arising from the choice of model parameters. Knowledge of the importance of these uncertainties is of great value in both initial assessment of the suitability of the feature set for the purposes of SHM, and more generally in guiding further development of the model. If it is found that variance in the response features is dominated by damage and that modelling choice uncertainty contributes little, there may be little justification for expending further effort on the calibration and/or quantification of these uncertainties. Conversely, factors that are identified as contributing substantially to variance of the response set may be prioritised for calibration and/or quantification. If further quantification is infeasible or undesirable, a decision may be made to prioritise those responses that demonstrate robustness to the identified uncertainties and discard those that do not.

The approach taken is to simultaneously consider both the uncertain factors and damage states using the R^2-based effect analysis of the previous section. The 16-run DOE given in Table 4-4 is implemented once per damage state. The responses may thus be evaluated at each point in the DOE. In the illustrative example, the 5 input states investigated are the undamaged state and 20% damage at each of the four damage locations in turn, requiring a total of 80 model runs. The requirement to run the DOE once per damage state leads to a linear scaling of the computational requirements with the number of investigated states. This scaling of computational expense will be a consideration for practical application of the method. However, it is suggested that the improved understanding of the model that arises through the analysis should outweigh the computational cost incurred, allowing better direction of subsequent modelling and testing efforts.

The response set investigated comprised the set of 10 normalised changes in natural frequency $\delta\boldsymbol{\omega}$ whose separability was previously assessed. From the arguments presented in Chapter 3 it is expected that the relative responses $\delta\boldsymbol{\omega}$ will

(a)

(b)

(c)

(d)

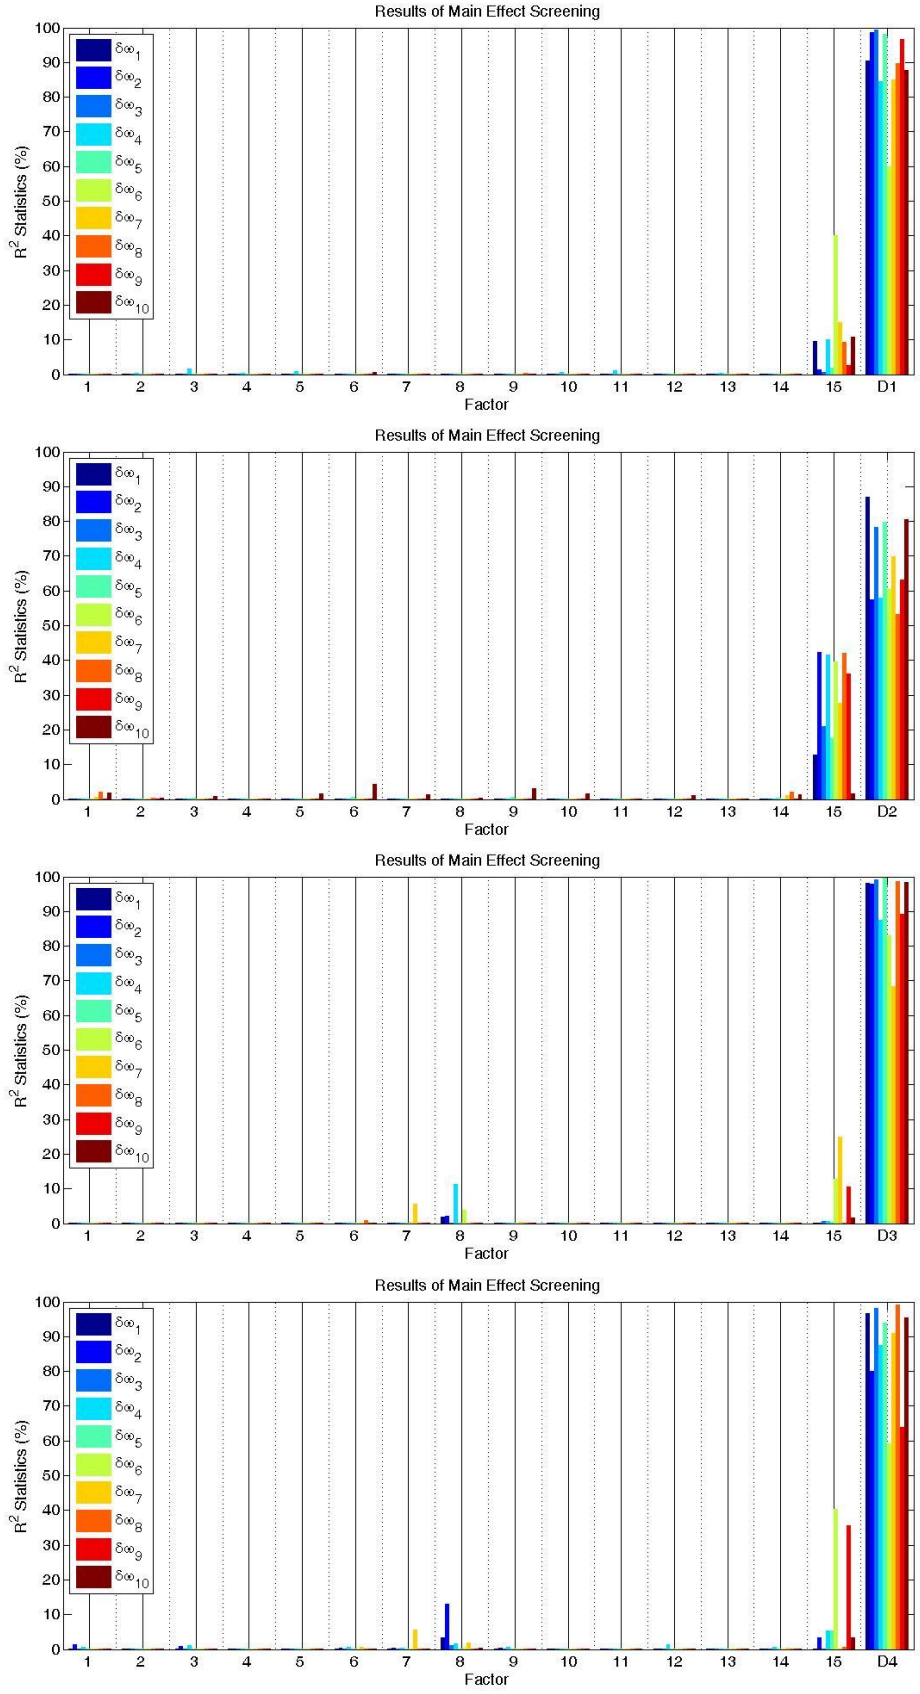

Figure 4-10 R^2 analysis of main effects on $\delta\boldsymbol{\omega}$ for damage states (a) D1, (b) D2, (c) D3, (d) D4

demonstrate greater robustness to model-form error than the absolute responses $\boldsymbol{\omega}$.

The quantity of information produced by sensitivity analyses tends to be extensive and only selected results are presented, intended to best illustrate the most important effects in each damage state. Each state is considered in turn and only the main effects of the 15 uncertain factors and damage are evaluated. Results are presented in Figure 4-10.

It is observed that the predicted responses for all four damage states are dominated by the introduction of damage. However, many of the response predictions also displayed sensitivity to factor θ_{15}, particularly those of damage state D2. Some of the responses at damage states D3 and D4 display a small degree of sensitivity to factors θ_7 and θ_8. The response predictions are otherwise found to be relatively insensitive to the model parameters chosen. The implication is that, in the context of the current study, the response predictions may be taken to be robust to the main effects of the investigated model-form variability. The outcome of this analysis would likely be a decision to further investigate factor θ_{15}, the bending stiffness of the bolt elements about the global x-axis.

4.4.7 Analysis

In this case study three pragmatic steps for early-stage evaluation of predictive models for SHM have been illustrated for the wingbox structure. A visual evaluation of the separability of damage states was made using the 'nominal' model; dominant parameters in explaining the variability of absolute outputs of the model in its undamaged model were identified; and the 'robustness' of a candidate SHM feature set to the effects of modelling uncertainties was evaluated. The latter two steps demonstrated differing applications of statistical effects analysis. It is suggested from the case study provided that effect analysis can offer a computationally inexpensive means of gathering highly informative data for guiding further model development.

For the case study used for illustration, a key finding was that despite a gross lack of knowledge expressed with regard to a selection of uncertain model inputs (±50% relative to the nominal value for stiffness components of the fastening elements), and the low level of damage specified (a 20% reduction in Young's modulus of a

single element), the feature set was shown to offer a surprising degree of 'robustness' to epistemic modelling uncertainty. One factor in particular was identified as requiring further quantification. The same factor had previously been identified as contributing substantially to output variability for the undamaged structure.

In the example problem, consideration is given only to main effects. Extension to quadratic (and higher) order interaction effects is possible through appropriate DOE selection, and should further improve understanding of the model outputs. Model-form errors were limited in the case study to substantial uncertainties in parametric values. Demonstrating robustness to variability arising from other modelling decisions (for example simplifications and element choice) would lend a substantial amount of credibility to the model predictions.

Finally, a relatively homogeneous set of model outputs was assessed in the case study (normalised natural frequencies). A benefit of the effect analysis approach is that a broad range of both inputs and outputs may be evaluated, Given that many model outputs will have potential for use as features for SHM, the method may be useful in assessing the relative merits of candidate features. An alternative extension of effect analysis is in model-based assessment of the robustness of features to other sources of uncertainty, with environmental effects being a particularly important example in civil and aerospace applications of SHM.

4.5 Case study 2: Propagation of uncertainty

Uncertainty propagation may be applied for a number of purposes. SHM-specific objectives suggested in this chapter include: to allow the confidence held in the model to be assessed; to inform feature selection; and, potentially, to generate data with which to train a damage identification algorithm. The aim of this case study is to demonstrate uncertainty propagation principles using the wingbox structure. The objective of probabilistic validation is to build an impression of the distribution of the outputs of the model in the face of uncertain inputs in order to answer the question:

- Does the developed numerical model of the wingbox accurately represent 'reality'?

In this case study this is devolved into two constituent questions:

- Does the deterministic model accurately represent reality?
- How would uncertainty in some of the model inputs affect its response predictions?

The problem adopted for demonstration is one of damage assessment, with damage occurring at a single location and varying depth. The approach employed is described below.

4.5.1 Approach

The subject of the case study is the wingbox model developed in Chapter 3. To answer the questions posed the objectives are: to produce natural frequency predictions $\boldsymbol{\omega}$ for the 25%, 50%, 75% and 100% damage states; and to compare the predicted outcomes to the corresponding experimentally observed values.

The procedure undertaken in doing so is as follows. First, the damage model is calibrated. Recall that the wingbox model introduced in Chapter 3 featured a highly-simplified parameterisation of damage, the reduction of Young's modulus values for parameters at the location of damage. In order for predictions to be made for particular damage depths, a relationship between damage depth and stiffness must be established. In this instance, the damage parameter is calibrated using the predictions of a second model which features a more complex parameterisation of damage. In principle, calibration could be undertaken in a similar fashion using experimental data. For example, destructive testing of components or sub-

structures should allow an accurate estimate of the damage depth relationship to be made. The calibration undertaken here may thus be considered as indicative of how damage parameter calibration using experimental data may be conducted.

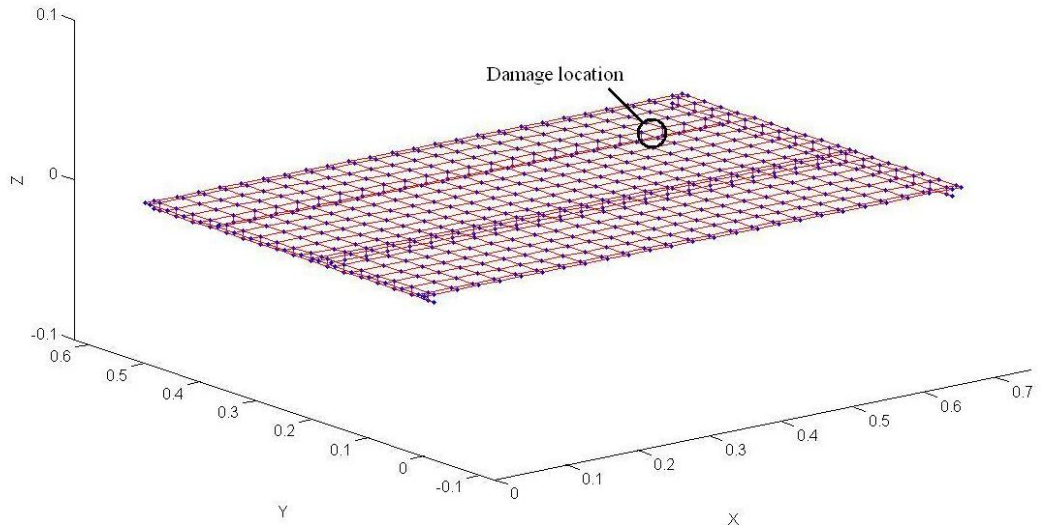

Figure 4-11 Wingbox model $\mathcal{M}_1$, with the location of damage indicated

Next, uncertain parameters are defined; their statistical distributions are specified and samples from the resulting parameter space are drawn. In this study, two parameters are defined as 'uncertain'. The level of uncertainty is expressed by specifying an interval for each parameter and assuming a uniform distribution between the interval bounds. The parameters were in this case selected simply to demonstrate the uncertainty propagation approach, rather than on the basis of any initial screening; they are the thickness of the wingbox top-sheet and the bending stiffness of the rivets used for attaching the rib components. Sampling is performed using a *Latin hypercube* approach.

Finally, propagation is undertaken and the distributed response predictions are compared to experiment. Response predictions are produced by executing the model for each point specified in the Latin hypercube design. Visual comparison of the experimentally-measured and model-predicted responses informs discussion of the two questions posed.

4.5.2 Damage parameter calibration

In the model of the wingbox used, damage is parameterised as a reduction in element stiffness at a single location. Given that the objective of the case study is damage assessment, a relationship between the parameter values and the true

damage extent must be established. Options for estimating damage parameter values were introduced in Section 4.4.4 and [14]. They include approaches from fracture mechanics and direct calibration to destructive testing results. The approach adopted here is to calibrate the responses of the simple model, which shall be referred to as model $\mathcal{M}_1$, to those of a second, more complex numerical model referred to as $\mathcal{M}_2$. The damage parameter is referred to as θ_D. While a valid approach in itself, this also serves a proxy to the case when destructive testing results are available and calibration to these results is sought.

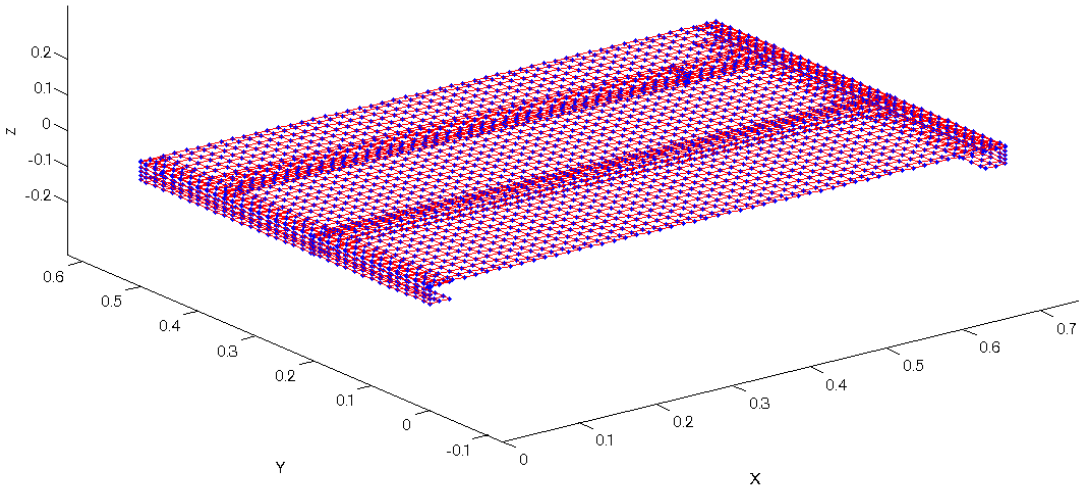

Figure 4-12 Finite Element model $\mathcal{M}_2$

The model $\mathcal{M}_2$ is illustrated in Figure 4-12. The ribs and stringers of the structure are represented using shell elements in this model, while the element mesh is defined such that damage at the 25%, 50%, 75% and 100% extents can be represented through the removal of entire elements (shown in Figure 4-13). The model comprises 3208 shell elements and 80 beam elements. The numerical cost of executing model $\mathcal{M}_2$ is subsequently significantly greater than that for executing model $\mathcal{M}_1$, which contains 560 shell elements and 164 beam elements. The representation of joints and (free-free) boundary conditions is the same for models $\mathcal{M}_1$ and $\mathcal{M}_2$.

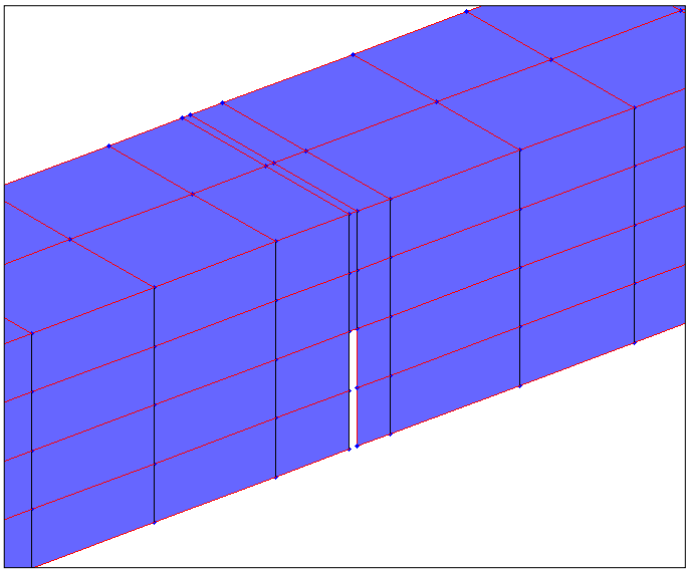

Figure 4-13 Parameterisation of damage to the stringer component of model $\mathcal{M}_2$ through element removal

The aim of the calibration exercise is to find values of the damage parameter θ_D for each damage state. Calibration of the damage parameter of model $\mathcal{M}_1$ was carried out by finding the minimum of a simple least-squares cost function for each damage state. The cost function used was,

$$J(\theta_D) = \left\| \delta\boldsymbol{\omega}(\theta)_{\mathcal{M}_1} - \delta\boldsymbol{\omega}(\theta)_{\mathcal{M}_2} \right\|_2^2 \tag{4.7}$$

where $\delta\boldsymbol{\omega}(\theta)_{\mathcal{M}_1}$ is the vector of normalised response change predictions from model $\mathcal{M}_1$,

$$\delta\boldsymbol{\omega}(\theta)_{\mathcal{M}_1} = diag\left(\boldsymbol{\omega}\left(\theta^{Undam}\right)_{\mathcal{M}_1}\right)^{-1}\left(\boldsymbol{\omega}(\theta)_{\mathcal{M}_1} - \boldsymbol{\omega}\left(\theta^{Undam}\right)_{\mathcal{M}_1}\right) \tag{4.8}$$

and $\delta\boldsymbol{\omega}(\theta)_{\mathcal{M}_2}$ is the vector of normalised response change predictions from model $\mathcal{M}_2$,

$$\delta\boldsymbol{\omega}(\theta)_{\mathcal{M}_2} = diag\left(\boldsymbol{\omega}\left(\theta^{Undam}\right)_{\mathcal{M}_2}\right)^{-1}\left(\boldsymbol{\omega}(\theta)_{\mathcal{M}_2} - \boldsymbol{\omega}\left(\theta^{Undam}\right)_{\mathcal{M}_2}\right) \tag{4.9}$$

for element removals corresponding to the 25%, 50%, 75% and 100% extents. The MAC was used to ensure mode matching was preserved between the damage states.

The cost function $J(\theta_D)$ was evaluated across a range of parameter values for each damage extent. The model $\mathcal{M}_2$ was executed for element removal corresponding to

the 25%, 50%, 75% and 100% damage extents. The responses $\boldsymbol{\omega}(\theta)_{\mathcal{M}_1}$ of model $\mathcal{M}_1$ were evaluated across a range of parameter values from 0.05 to 1 times the initial value, with cubic spline interpolation used to increase the resolution of the response vector predictions with damage. The results of evaluating the cost function for each damage state are presented in Figure 4-14. Parameter values corresponding to the cost function minimum are adopted for each damage state.

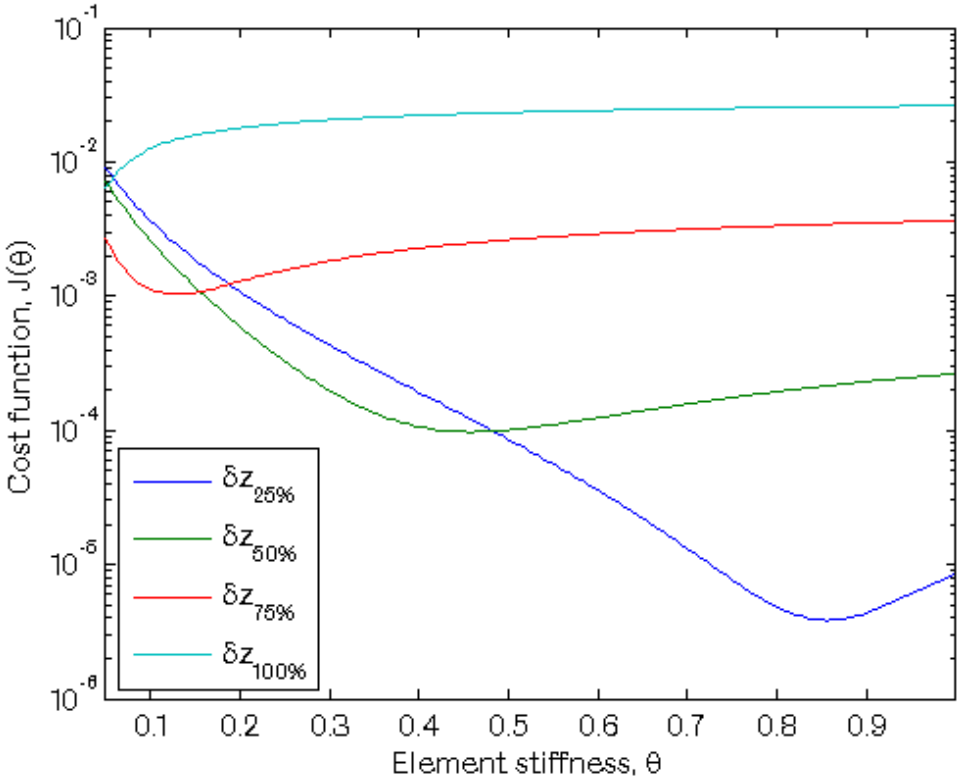

Figure 4-14 Cost function evaluation for calibration of the damage parameter θ_D

Due to the differing scales of the cost function plots for each damage state, a log-scaled y-axis is used to more clearly illustrate the minima. Clear minima are observed for the 25%, 50% and 75% damage states. If a physically realisable minimum exists for the 100% damage state, it lies outside of the evaluated parameter range. The value at the lower limit of the parameter range is used for this state. The values of the damage parameter θ_D - the damage parameter values of the 'simple' model $\mathcal{M}_1$based on calibration to the 'complex' model $\mathcal{M}_2$ - that are adopted for each state are given in Table 4-5.

Damage extent	Parameter value θ_D	Corresponding Young's modulus of element [GPa]
Undamaged	1	69.0
25%	0.855	59.0
50%	0.460	31.7
75%	0.135	9.3
100%	0.050	3.5

Table 4-5 Calibrated damage parameter values

4.5.3 Define responses and uncertainties

With the values of the damage parameter for each damage state having been established, it is possible to execute the model once for each state and make deterministic predictions of how the responses will change with damage. By making these predictions and comparing them to the experimental results a first response to the question of whether the model accurately represents 'reality' may be made.

However, suppose that for some reason there is a degree of uncertainty with regard to the precise thickness of the top-plate to use in the model. Suppose also (and perhaps more feasibly) that the bending stiffness of the rivets attaching the C-channel rib elements to the top-plate is not known with confidence. How would these uncertainties affect the response predictions? This is an example of the type of question that may be tackled via uncertainty propagation. In a practical implementation of uncertainty propagation, the factors to be included should have been properly-identified via a sensitivity analysis such as that in Case Study 1, and wherever feasible, their statistical distributions should have been quantified and included in a schema such as the PIRT table. In this case, the factors have been chosen solely for the purposes of illustration. These factors and their assumed bounds are summarised in Table 4-6.

Parameter label	Parameter description	Parameter Value			Unit
		Min	Nominal	Max	
θ_1	t_{plate}	2.85	2.91	2.95	mm
θ_2	$K_{xy,Rivet}$, $K_{xz,Rivet}$	$2.2\text{x}10^7$	$3.2\text{x}10^7$	$4.2\text{x}10^7$	N/m

Table 4-6 Uncertain parameters and distributions

4.5.4 Sampling methodology

The factors in Table 4-6 define a two-dimensional parameter space. Samples were drawn from this space by creating a space-filling Latin Hypercube design (LHD). Latin Hypercube sampling (LHS) is a constrained, randomised method for drawing samples from a design space, and is based on a Latin square. Further information on LHS and other methods may be found in [66]. An study of the properties of LHS for uncertainty propagation, including comparison to other techniques, is given in [67]. For the case study, a *maximin* design was used, which seeks to maximise the minimum distance between any two sample points. The resulting design allows the parameter space to be adequately represented in far fewer runs than would be possible using fully randomised sampling.

The implementation of UP pursued here is as follows. 100 design points were drawn to represent the two factors identified in Table 4-6 using a maximin LHD. The design was created using the input generator provided in the sensitivity analysis package GEM-SA [68]. The resulting set of design points are represented in scatterplot form in Figure 4-15.

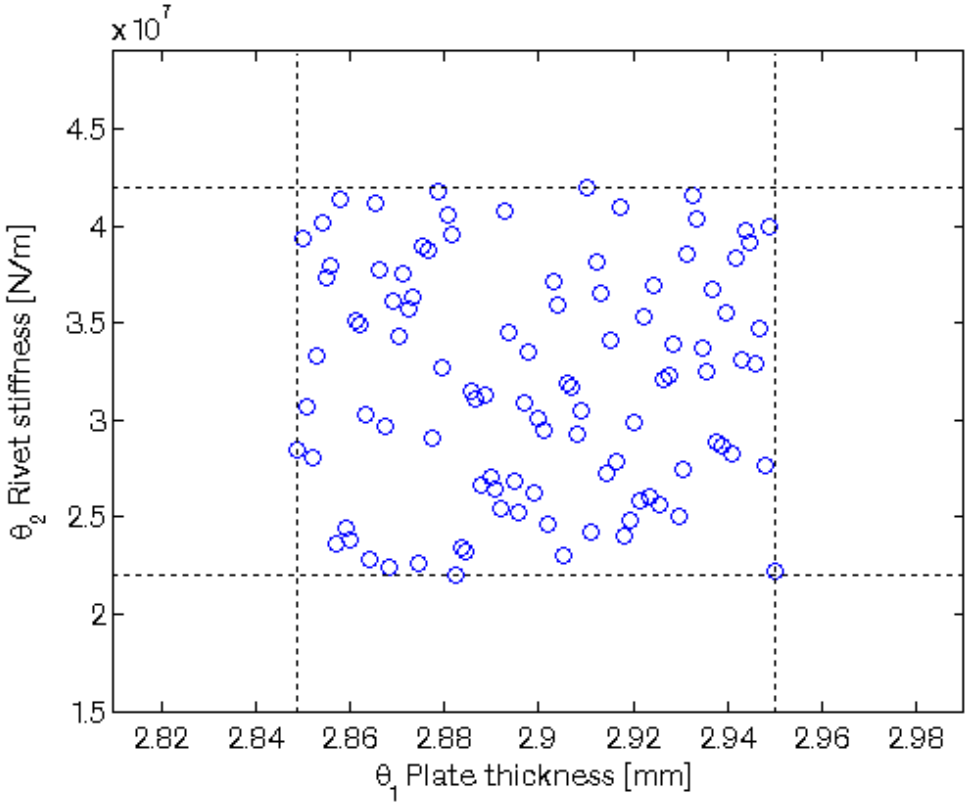

Figure 4-15 Two-dimensional Latin hypercube design used for uncertainty propagation

For each of the five states of the structure identified in Table 4-5 the model was executed for each of these 100 design points. This resulted in 500 runs in total. The model is also run in its 'nominal' state, with the factors identified in Table 4-6 remaining unchanged. The decision to take 100 samples per state is an arbitrary one; alternative sampling methodologies are considered in the discussion at the end of the case study.

4.5.5 Results

The results of the uncertainty propagation exercise are presented in Figure 4-16 to Figure 4-25. For each of the ten mode pairs previously identified (see Table 3-1) the experimentally observed mode is plotted in blue, and the model predictions in red. All observations or predictions at each damage state are included, with summary statistics (mean ± 2 standard deviations) overlaid upon the data for both cases. For the numerical prediction, the predictions of the model when run with the nominal values from Table 4-6 are overlain on the plot in black. The same scaling is used for the experimental and model prediction plots.

Note that the predictions are presented in unmodified form, with no adjustment made to account for the observed test-FE discrepancy. This could be accounted for through bias correction: basic, additive bias correction based on the undamaged state values may be applied by adding a factor $\left(\boldsymbol{\omega}_{Test}^{Undam} - \boldsymbol{\omega}_{FE}^{Undam}\right)$ to all the model-predicted values. As the comparison made here is purely graphical and the pattern of change is readily apparent, bias correction was omitted for simplicity.

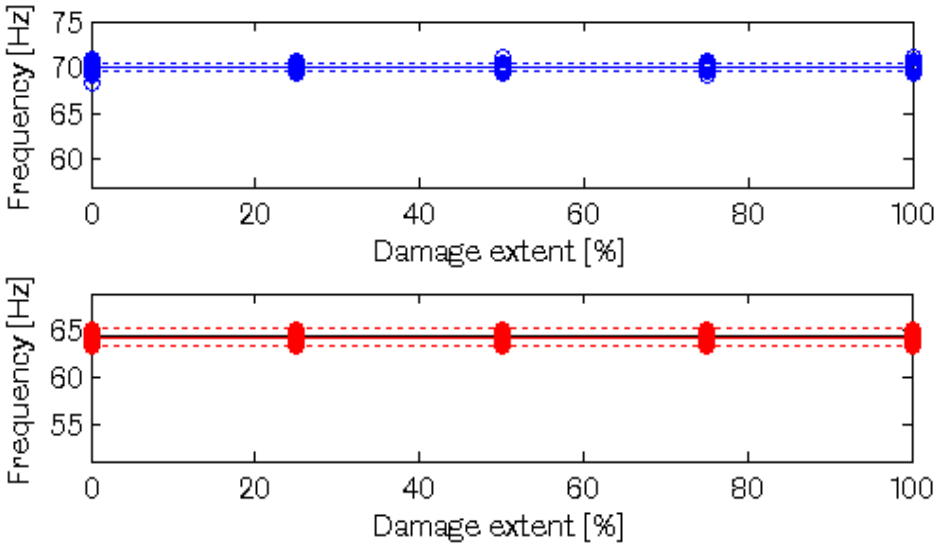

Figure 4-16 Natural frequencies, experimental mode E2 and analytical mode A8

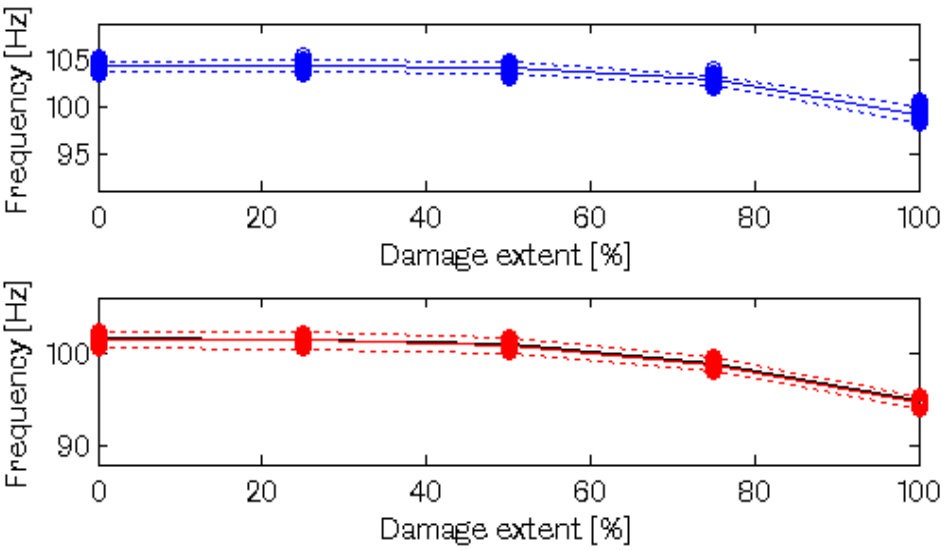

Figure 4-17 Natural frequencies, experimental mode E3 and analytical mode A9

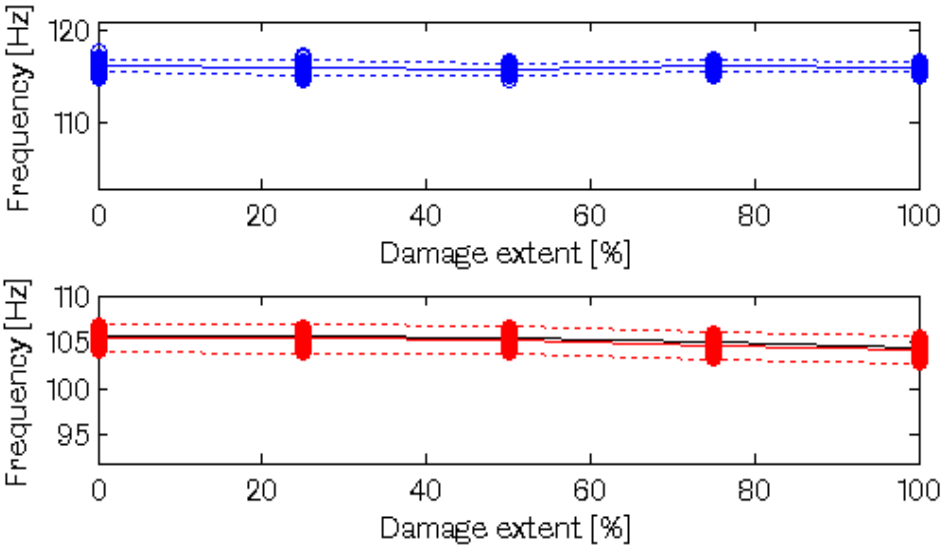

Figure 4-18 Natural frequencies, experimental mode E4 and analytical mode A10

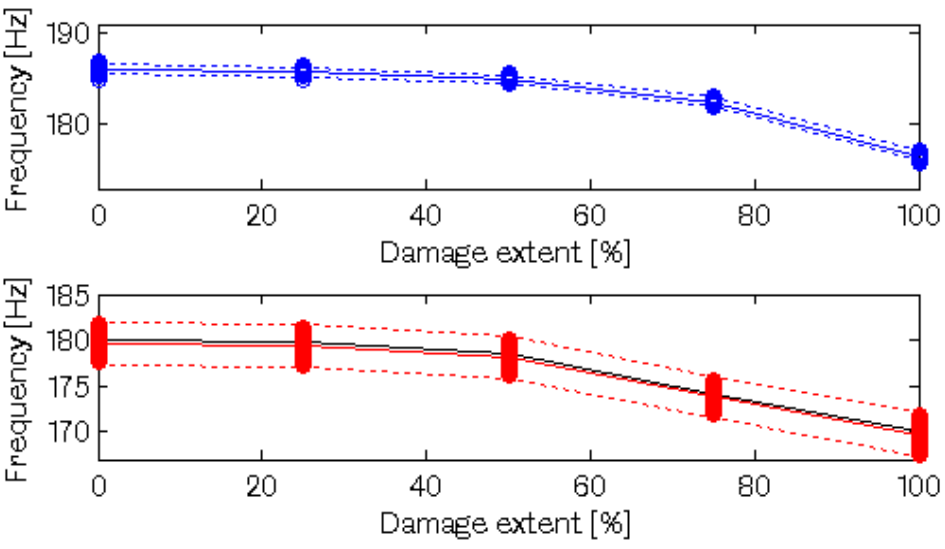

Figure 4-19 Natural frequencies, experimental mode E5 and analytical mode A13

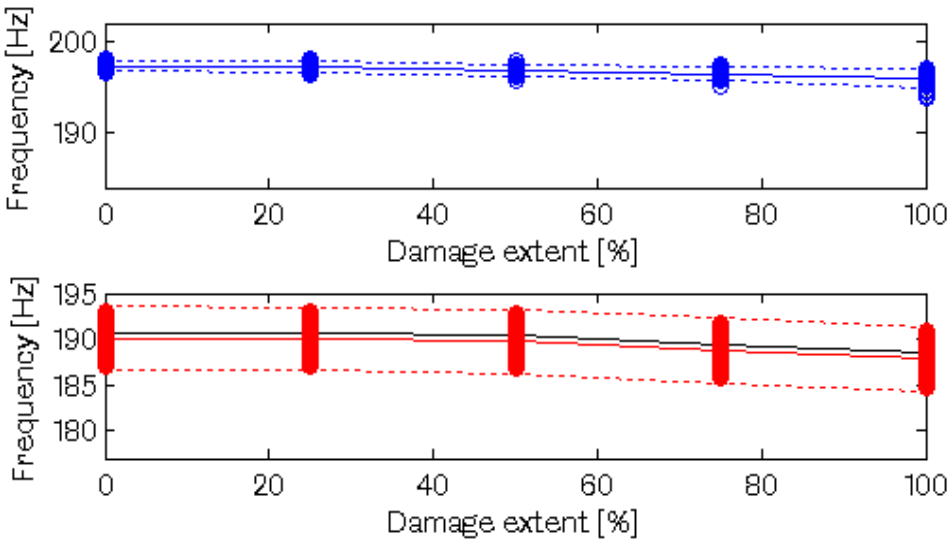

Figure 4-20 Natural frequencies, experimental mode E6 and analytical mode A14

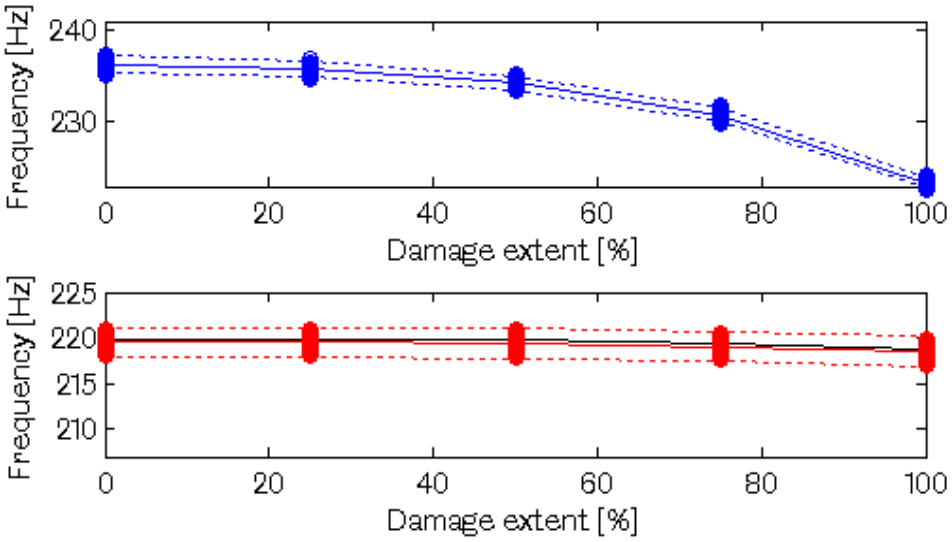

Figure 4-21 Natural frequencies, experimental mode E7 and analytical mode A16

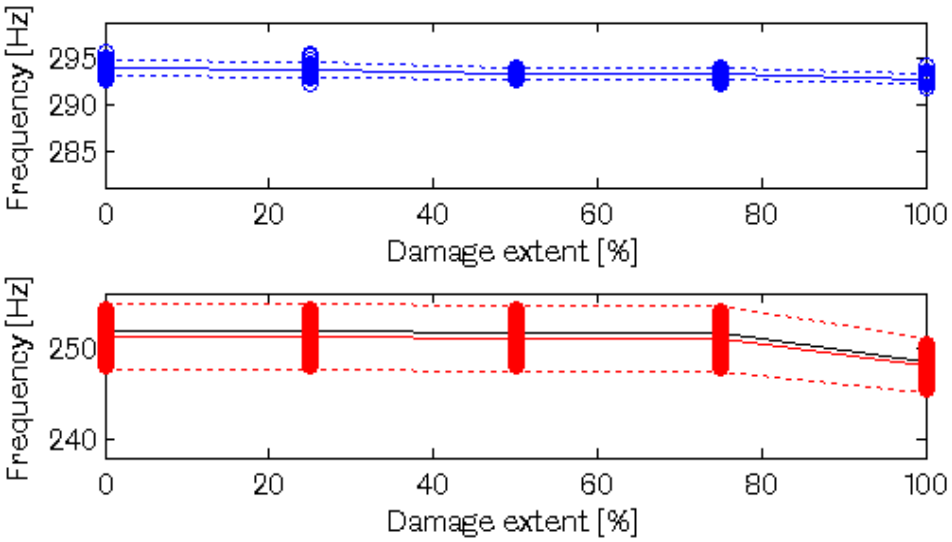

Figure 4-22 Natural frequencies, experimental mode E9 and analytical mode A17

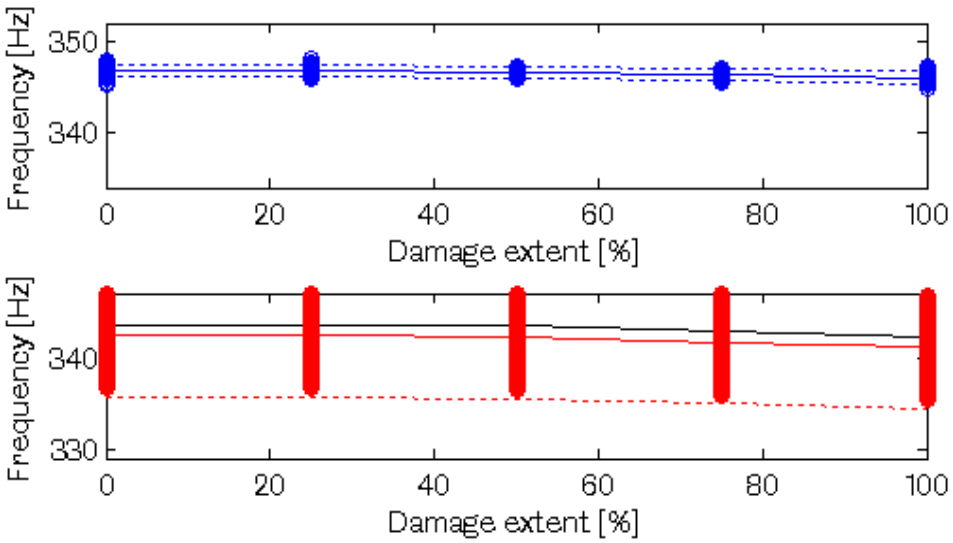

Figure 4-23 Natural frequencies, experimental mode E11 and analytical mode A22

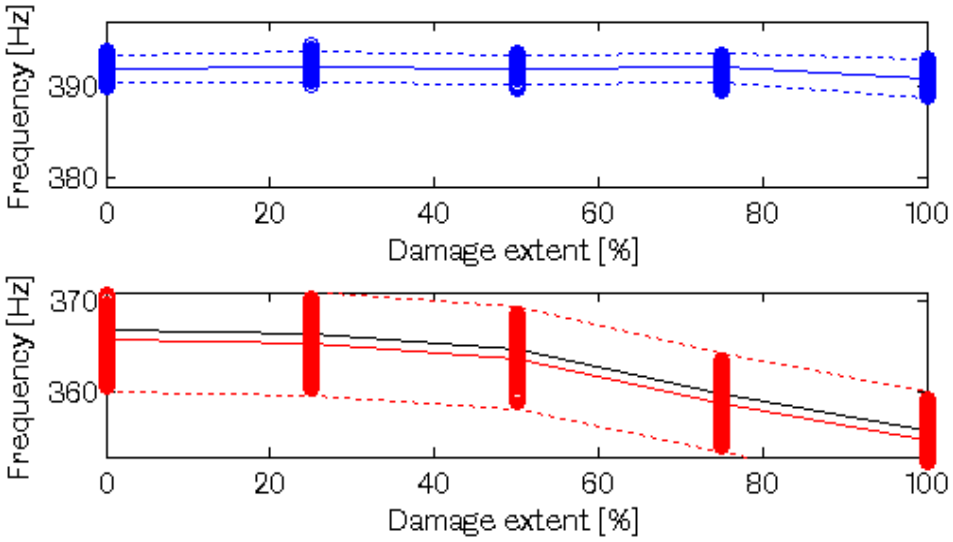

Figure 4-24 Natural frequencies, experimental mode E12 and analytical mode A24

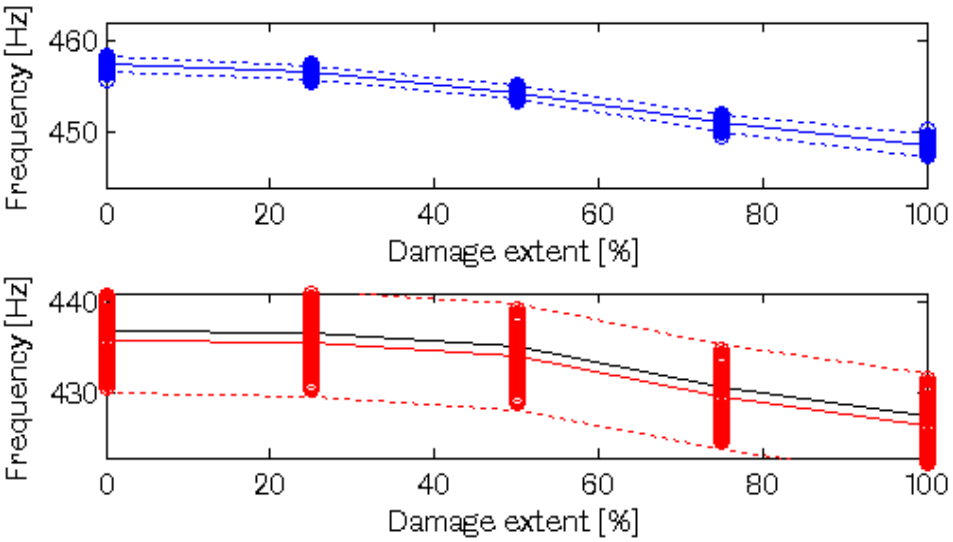

Figure 4-25 Natural frequencies, experimental mode E15 and analytical mode A27

4.5.6 Analysis

The two questions posed were: does the developed numerical model of the wingbox accurately represent reality, and how would uncertainty in certain of the model inputs affect its response predictions? These questions are considered in turn on the basis of the above plots.

The question of representing reality is considered first, focusing on the outputs of the nominally-valued model (in black). From an initial visual comparison of the experimental and numerical outputs in Figure 4-16 to Figure 4-25 it is apparent that the success of the model in predicting the pattern of change with damage has been somewhat mixed. The model appears to have correctly identified the insensitivity of experimental modes E2, E4, E11 and the low sensitivity of mode E6. For mode E9, a sharp drop in response value for the 100% damage-depth case was predicted by the model but did not transpire in the experimental data; otherwise predictions for this mode appeared to correctly identify the mode as insensitive to damage. For the damage-sensitive features, changes resulting from damage in modes E3, E5 and E15 appear to have been predicted correctly. The application of further validation metrics would be required in order to judge the accuracy of the predictions for each state in more detail, but the visual comparison is nevertheless encouraging.

However, the predictions made for modes E7 and E12 cause concern. Mode E7 demonstrated a marked sensitivity to damage in practice, but the model failed to predict this. Conversely, mode E12 was predicted to be highly sensitive to damage, but this was not observed experimentally. Two possible explanations present themselves: either the model is flawed to the extent that it is invalid for this combination of damage states and modes; or, more prosaically, the wrong modes are being compared. Other candidate mode-pairs are considered in Figure 4-26 to Figure 4-28 for experimental mode E7, and Figure 4-29 to Figure 4-31 for mode E12. For both, there is another candidate analytical mode for which similar modeshape and frequency correlation values were observed to those of the originally adopted modes. However, the frequency deviation with damage for analytical modes A16 and A26 is a much more accurate reflection of the observed behaviour of experimental modes E7 and E12 respectively.

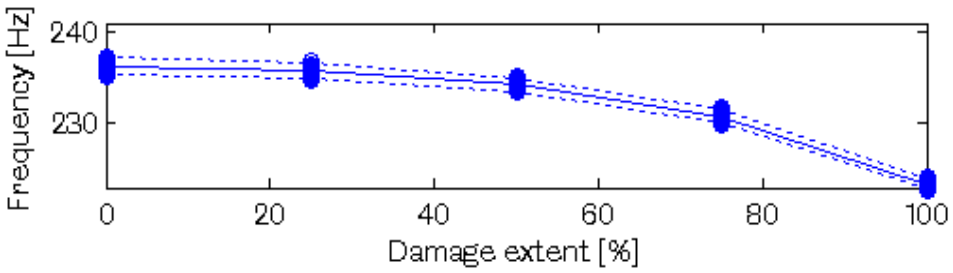

Figure 4-26 Natural frequencies, experimental mode E7

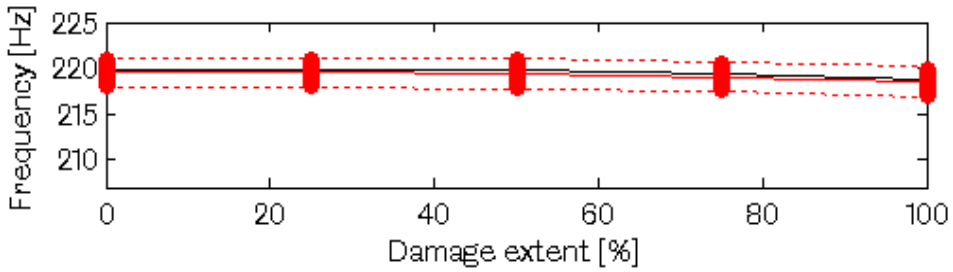

Figure 4-27 Natural frequencies, analytical mode A16: MAC 32.72, $\delta\omega$=-7.50%

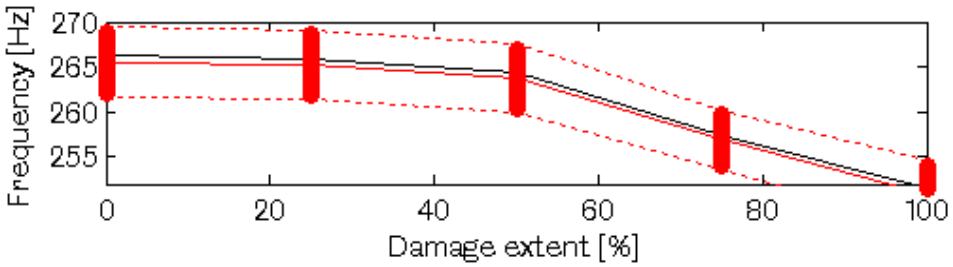

Figure 4-28 Natural frequencies, analytical mode A18: MAC 30.00, $\delta\omega$=+12.64%

The conclusion drawn is that in both cases it is likely that the wrong modes are being compared. This serves to draw attention to the challenge of validation/mode-matching based up modeshape correlation and frequency differences alone. However, the fact that the error in mode-matching was revealed through a form of structural modification (the introduction of damage) is indicative of a further tool that may be applied to the mode-matching problem. Non-destructive structural modifications such as mass addition may be a useful means of adding information during the mode-matching and validation exercises.

NOTE. The original mode-matching presented in this case study was that employed in the model updating case study in Chapter 3. The implication is that the errors in mode-matching will have adversely affected the damage location outcomes, However, this is not viewed as being unfairly pejorative to the model updating method, as mode-matching based on undamaged state data only is considered to be standard practice in updating-based approaches to damage identification, as it

might be in predictive approaches in cases where no damaged sate data is available. Additionally, no direct comparison between methods is sought or made.

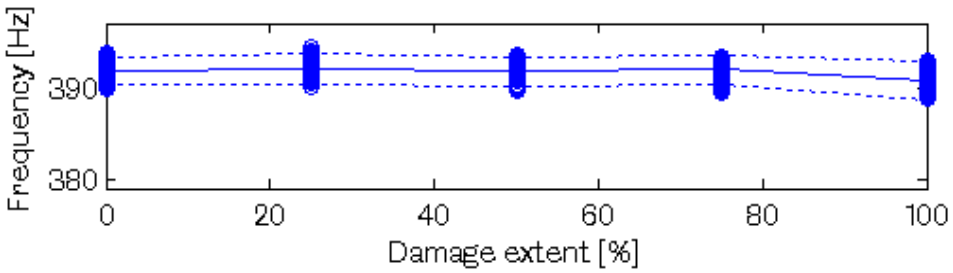

Figure 4-29 Natural frequencies, experimental mode E12

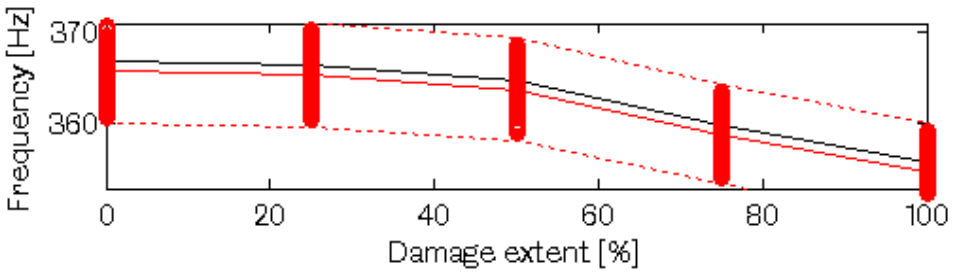

Figure 4-30 Natural frequencies, analytical mode A24: MAC 60.06, $\delta\omega$=-6.83%

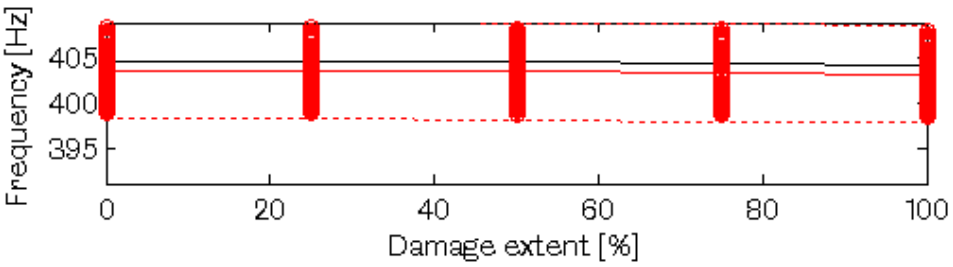

Figure 4-31 Natural frequencies, analytical mode A26: MAC 36.00, $\delta\omega$=+3.28%

Next, the effect of input uncertainty is considered. It is immediately apparent from considering the distribution of the predicted data that certain of the effect of the investigated uncertainties on certain of the modes is far greater than on others. The answer to the question 'how does uncertainty in the model inputs affect its response predictions?' may thus be very tersely stated as 'substantially, in some cases'. To achieve a more precise understanding of how inputs (or combinations of inputs) explain variance in the responses sensitivity analysis may be applied. However, a more useful answer may be arrived at by applying quantified validation metrics. Here, only visual comparison was employed. The next step would be to apply statistical tests to ascertain whether the model is trusted, using the generated response predictions and the experimentally-observed response distributions.

As a final comment, note that there is no claim that the best use of the computational 'budget' has been made. While a Latin hypercube design was used for

illustration, an outcome similar to that used for here for analysis (effectively finding the max/min of the returned predictions) may have been achieved in fewer runs using other uncertainty analysis options, for example an interval approach.

The decision to specify five levels of the damage parameter θ_D was taken primarily to allow direct visual comparison with the experimental results at the corresponding damage levels. In order to cover the full domain of model realisations using a parsimonious number of runs, it would in general be preferable to allow the damage parameter θ_D to take any valid value (i.e. in the range [0 1]) and to include it as a factor in the Latin hypercube design. This also leads towards the idea of meta-modelling. By specifying appropriate points from the domain of possible models, it is possible to construct a meta-model that acts as a fast-running surrogate for the numerical model. The advantage of doing so is that the relationship between model inputs and outputs may be estimated from a limited number of runs. This opens up the possibility of pursuing global sensitivity analysis of numerical models in a parsimonious number of runs [57]. Further, it opens up the possibility of being able to infer inputs from observed outputs in a manner akin to data-based approaches. This is of great interest in damage identification where damage parameters may be adopted as inputs, and meta-modelling is returned to as an area for future work in Chapter 9.

4.6 Discussion

In this chapter probabilistic, forward approaches to law-based modelling for the purposes of informing statistical damage identification have been presented. The approach has been presented as an extension of an existing framework for the verification and validation of numerical models, with several adaptations of the general V&V approach that offer potential for SHM having been introduced. Aspects of the approach have been illustrated via two case studies. The ideas covered in this chapter are somewhat diffuse and, by necessity, in part speculative. However, it appears that merging forward identification with concepts from probabilistic verification and validation has the potential to be a powerful alternative to inverse modelling approaches (as introduced in Chapter 2 and 3) and purely data-driven damage identification (as will be covered in Chapters 5-7).

The presented case studies illustrate some of the ways in which V&V tools may be applied to an SHM problem, and serve to highlight both the pros and cons of adopting a forward approach.

The first case study concerned the application of *statistical effects analysis* at the sensitivity analysis stage in order to assess the effectiveness and robustness of the selected feature set. From the first case study, it was concluded that the application of sensitivity analysis at an early stage in model development, in the presence of both damage and sources of modelling uncertainty, may be a useful tool for informing decisions both on the strategy for further model development, and on the ability of the selected response set to robustly separate the damage states of interest at the statistical discrimination stage. The ability to interrogate the model to inform decisions is an important aspect of the forward approach, and the broader application of sensitivity analysis methods for this purpose is strongly advocated.

The second case study gave an example of uncertainty propagation for the wingbox structure, making use of Latin hypercube sampling. The case study included the calibration of the damage model in order to allow predictions to be made for damage assessment at a single location. Uncertain parameters were defined; their statistical distributions specified; and samples from the resulting parameter space wee drawn using a Latin hypercube approach. It was demonstrated that by propagating uncertainty through the model and comparing the distributed response predictions to those from experiment, a more informative answer may be provided to questions of validation than if deterministic prediction alone is pursued.

Despite the somewhat broad nature of the work covered, several initial conclusions may be drawn. The first among these is that a general framework for tackling the probabilistic, forward modelling task already exists under the umbrella of computational model verification and validation, and that many of the tools developed for V&V appear amenable to the model-based SHM task. This is an agreeable finding, as adopting a forward, probabilistic framework may have numerous positive implications for model-based SHM in areas such as model development, model selection, experimental design, and feature evaluation. In one of the preliminary case studies included in this chapter it was shown that a forward approach may be employed to assess both the robustness of model predictions to modelling uncertainty (both epistemic and aleatoric), and the sensitivity of these

predictions to damage. It is hoped that through further investigation other aspects of the forward approach to SHM may be realised through adaptation of existing tools.

A potential drawback of the proposed approach is computational expense, although this is largely confined to the development stage. Once the statistical model has been trained, minimal computational cost should be associated with classifying newly presented data. The potential for reducing the required computational effort through, for example, meta-modelling is one area highlighted for further research.

The overall conclusion reached is that forward modelling approaches may offer a useful alternative to inverse approaches for the SHM task and are deserving of further investigation. In parallel, the other major area for future work is in tackling the second part of damage identification task: making use of information from the law-based model in the development of a discriminative statistical model. This has not been considered in this chapter, but the lack of work in this area may be considered the major hurdle to be overcome in applying a forward approach.

A final comment to be made relates to the existence of a degree of crossover between the forward and inverse model-based approaches, to the extent that parametric updating was an important part of the second case study. However, several aspects of the forward approach that are distinct from the inverse approach have also been highlighted. The commonalities and distinctions between the approaches are returned to in Chapter 8. Also considered in Chapter 8 are the class of damage identification algorithms commonly referred to as *data-driven approaches*. These so-named approaches are the subject of Chapters 5-7.

5 A Data-Driven Approach: Multiple Damage Location

5.1 Introduction

In Chapters 2 to 4, attention has been given to law-driven approaches to the damage identification task. In Chapters 5 to 7, the focus is turned to data-driven approaches to this problem. The data-driven approach is illustrated through consideration of a challenging and unresolved problem in SHM: the correct identification of damage simultaneously occurring at multiple locations. The steps involved in constructing a data-driven classifier for the purposes of multiple damage identification are incrementally developed over the course of the next three chapters. The approach and discussion developed in Chapters 5 to 7 allow the data-driven approach to be exemplified via a practical example. The similarities and distinctions between this and the law-based approach are expanded upon in Chapter 8.

The layout of the chapter is as follows. First, the problem of multiple damage location is introduced and previous studies in this field are reviewed. Next, an overview of the approach developed in Chapters 5 to 7 is given. Finally, the first stage in this approach - experimental data acquisition - is described. Particular attention is paid to test sequencing in order that the effects of systematic and random test variability are well-represented in the datasets used for training and testing.

5.2 Multiple damage location

It is more or less accepted that damage identification at Level 2 (damage location) via statistical pattern recognition requires the adoption of a supervised learning approach, thus requiring the gathering of data from the structure in its damaged state. In the particular case that damage is believed to occur at only one of a finite

discrete set of n locations, this would require the data to be gathered for each of the corresponding n damage states. In the more general case of damage occurring concurrently at more than one location, the naïve approach might be to gather damage data for all combinations of damage location. The number of states for which data would be required in order to cover all combinations thus grows exponentially with the number of locations n. It has been stated earlier in this thesis that damaged state data even for single site damage is unlikely to be available, and it would appear unwise to rely on the availability of damage data for all possible combinations of damage location in the general case. It is apparent that the lack-of-data problem, a major obstacle in the diagnosis of single-site damage, is a potentially critical issue in multiple-site damage location.

5.2.1 Previous studies

The multi-site damage identification task represents an important and challenging problem in SHM but has received little dedicated attention in the literature. Where multiple damage identification has been addressed, it has typically been done so with the aid of a law-based model. In [69] a study on multiple damage location and assessment in a cantilever beam is presented. The solution procedure adopted is to form residuals between features drawn from the structure and the predictions of a representative FE model. The particular approach taken was to identify three "fundamental functions" from the low-frequency modal characteristics of the structure. These fundamental functions are the natural frequencies, modal curvature and mass-normalised modeshape residuals between the simulated and the damaged structure. An objective function is formed from the weighted sum of these residuals. Multi-site damage location and severity assessment was performed through optimisation of this function.

The approach is applied to both a simulated and an experimental structure. For the experimental case, 50-average data were employed, with the intention of averaging out uncorrelated experimental variability. Notably, the effect of boundary conditions at the clamped end of the beam was also taken into account, with 10 repetitions of the test performed for each damage state. It was found that the variability arising from the clamping condition was low (the relative standard deviation of the identified natural frequencies due to the boundary condition was in all cases below 0.004% of the mean value). Other sources of variability were not explicitly taken into consideration.

A similar approach is presented in [70], with the damage location assurance criterion (DLAC) method extended to the multiple damage case. The DLAC is a cross-correlation function between measured and theoretical frequency changes resulting from damage, given by,

$$DLAC(i) = \frac{\|\delta \boldsymbol{z}_m \delta \boldsymbol{z}_i\|^2}{\|\delta \boldsymbol{z}_m\|^2 \|\delta \boldsymbol{z}_i\|^2} \tag{5.2}$$

where $\delta \boldsymbol{z}_m$ is the vector of measured frequency changes and $\delta \boldsymbol{z}_i$ is the vector of frequency changes due to damage at location i as predicted by the model**.** The location returning the greatest value is returned as the predicted location of damage. This approach relies upon having a database of frequency changes associated with each damage state, supplied by a law-based (numerical or analytical) model. The computational cost associated with executing the numerical model for all combinations of damage location is deemed restrictive for multiple damage location problems.

The approach is instead extended to a Multiple-Damage Location Assurance Criterion (MDLAC). Instead of executing a model for each damage location, the first-order sensitivity of the global stiffness matrix is estimated analytically and only one model execution is required. The sensitivities calculated for each single-site damage location form the basis of the MDLAC,

$$MDLAC(\delta \boldsymbol{D}) = \frac{\|\delta \boldsymbol{z}_m \delta \boldsymbol{z}(\delta \boldsymbol{D})\|^2}{\|\delta \boldsymbol{z}_m\|^2 \|\delta \boldsymbol{z}(\delta \boldsymbol{D})\|^2} \tag{5.3}$$

where $\delta \boldsymbol{D}$ is a vector of stiffness reductions for all included elements, and $\delta \boldsymbol{z}(\delta \boldsymbol{D})$ is the estimated vector of parameter changes due to $\delta \boldsymbol{D}$ as calculated from first-order sensitivities. Prediction of damage requires the damage vector $\delta \boldsymbol{D}$ returning the highest MDLAC value, to be found, and involves the solution of an optimisation problem. The approach is demonstrated for three numerical examples, each being truss-type structures. It was shown to work well for multiple damage location in these scenarios. The effect of modelling errors was not investigated, although the removal of first-order modelling errors is highlighted by the authors. Similarly the effect of measurement error was not assessed; however the authors highlight previous work indicating that the DLAC approach is tolerant of low-level experimental variability.

An incremental development of the MDLAC-based approach is presented in [71]. There, a two-stage method to identify structural damage locations and extents is pursued using evidence theory and a micro-search genetic algorithm. At the first stage, location of damage is pursued using Dempster-Shafer evidence theory, allowing fusion of the results of frequency MDLAC and modeshape MDLAC outcomes. Once the damage has been located, the damage extent is estimated. This is pursued by forming an objective function that maximises both the frequency and modeshape MDLAC values and seeking its solution by applying a micro-search GA. The approach was demonstrated for a simulated cantilever beam and compared to the standard MDLAC approach. The evidence theory approach was found to be more effective for damage location than the frequency and modeshape MDLAC approaches alone, and the micro-search GA offered computational advantages over alternative search methods due to a reduction in the size of the search space.

In [72] it is noted that few multiple damage identification algorithms focus on the handling of experimental variability, such as may be associated with measurement noise and material uncertainty. A probabilistic approach to two-site damage identification was developed. The features adopted were modal curvatures, which display both sensitivity to damage and a high level of susceptibility to measurement noise. A two-stage process of damage location and assessment of extent is again applied. The method relies on a database of damage responses generated by a model. However, Latin hypercube sampling is employed in generating the response database, allowing a drastic reduction in the number of model runs required compared to full-factorial sampling. The adoption of an efficient sampling method allows a probabilistic assessment of the extent of damage at the identified locations in the presence of experimental variability to be achieved. The approach is demonstrated experimentally for a beam structure with free-free boundary conditions.

The problem of multiple damage location from a data-driven perspective is not extensively considered in the published literature. In the studies that are available, it has generally been accepted that a representative physical model is required in order to generate a database of response changes due to damage. It is noted that the structures investigated in the reviewed literature have typically exhibited a low level of complexity.

The approach pursued in this section of the thesis is intended to allow diagnosis of multiple-site damage without resort to a law-based model or full damage state data. Only a greatly reduced subset of the full dataset corresponding to single-site damage is used. The method is demonstrated for a complex decommissioned structure.

5.3 Approach

The aim of the approach developed in Chapters 5-7 is to dramatically reduce the number of states for which training data is required. This is pursued by constructing a data-driven classifier that is capable of distinguishing occurrences of single-site damage, and evaluating the success of this classifier when presented with multiple-site damage data.

The study is experimental in nature, with the structure used being the wing of a Piper Tomahawk trainer aircraft. The features used are discordancy measures, and Support Vector Machines (SVMs) are employed for classification. Three objectives must be met in order to achieve this aim:

1) Data acquisition
2) Data reduction, feature selection and evaluation
3) Training, validation and testing of statistical classifiers

These objectives form the basis of Chapters 5 to 7 respectively. They are summarised below.

5.3.1 Objective 1: Data acquisition

In the present chapter the wing structure and the procedure for data acquisition are described. The wing structure features five removable inspection panels and the removal of these panels is adopted as a proxy for the introduction of damage. Particular attention is paid to test sequencing in order that the effects of systematic and random test variability are represented in the datasets used for training and testing. The raw data recorded are FRF spectra from a network of 15 sensors.

Two sequences of tests are conducted, resulting in two sets of data being produced. Dataset A contains 1000 observations of the structure in its 'normal' state, and 200 observations of the structure for each of the single-site panel removal states. Dataset A is primarily used as *training* data. It is the only data employed for feature selection and for training statistical classifiers. Dataset B provides the *testing* data

used to evaluate the selected features, and ultimately for evaluating the performance of the developed classifier. Dataset B contains 2300 normal, 1000 single-site damage and 1300 multi-site damage state observations.

5.3.2 Objective 2: Feature selection and evaluation

The objective pursued in Chapter 6 is to reduce the high-dimensional response data into a low-dimensional feature set that is suitable for training statistical classifiers, and to evaluate the individual performance of the chosen feature set.

The recorded FRF data are first converted into acceleration transmissibility spectra. Multivariate discordancy measures derived from sections of these spectra are adopted as candidate features. Features are selected on the basis of Dataset A, and evaluated using the observations of Dataset B. The hypothesis of this chapter is that it may be possible to find features that offer a high degree of discrimination between multiple-site damage states, despite only single-site damage states being available for feature selection.

5.3.3 Objective 3: Statistical classification

The final objective, detailed in Chapter 7, is the development and evaluation of statistical classifiers using the selected feature set. Support Vector Classification is employed for the classification task.

First, an SVM-based classifier is trained using training and validation sets drawn from Dataset B, and evaluated using a testing set also drawn from Dataset B. This allows the validity of the SVM approach to be assessed. Secondly, a further classifier is trained using data drawn from Dataset A. and its performance evaluated when applied to a testing set drawn from Dataset B, which contains multi-site damage data. The ability of the classifier not only to identify damage states on which it has not been trained, but also its ability to generalise between tests may thus be evaluated.

5.4 Experimental data acquisition

5.4.1 The wing structure

The experimental structure considered in this chapter is an Aluminium aircraft wing, shown *in situ* in the laboratory in Figure 5-1. The wing is mounted in a cantilevered fashion on a substantial, sand-filled steel frame. The sensor network and data

acquisition equipment may also be observed. The wing is in fact mounted upside-down in order to allow access to the inspection panels mounted on the underside of the wing. The reasons for this are explained later.

It is immediately apparent that the structure under investigation is of far greater complexity than the wingbox structure investigated using model-based approaches in Chapters 3 and 4. It includes various complicating features, including stiffening elements (both accessible and inaccessible), inspection panels, riveted and welded connections, as well as auxiliary structures such as aileron mounting points. There is some minor damage apparent associated with either the in-service use or post-service handling of the structure. The boundary conditions represent a further complicating factor. Achieving a truly rigid mounting is non-trivial, and while the dynamics of the frame have not been explicitly tested it can be assumed that the objective of providing a truly rigid mounting was not fully met. Accurately modelling the rigid connection presents a more challenging task than modelling the lightly sprung boundary conditions of the wingbox.

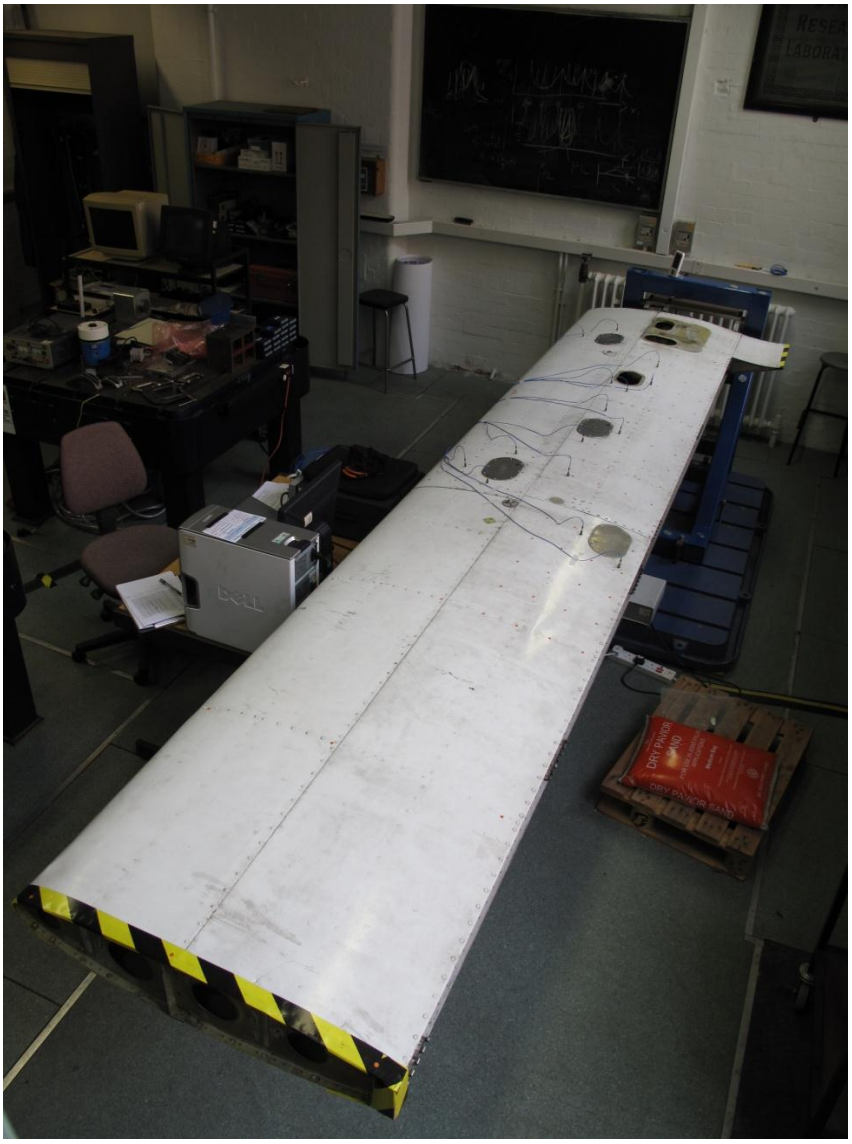

Figure 5-1 Piper Tomahawk trainer aircraft wing

Producing an FE model capable of making accurate predictions of dynamic behaviour for such a structure would appear a somewhat unenviable task, and it is when considering more complex structures that the possibility of tackling damage identification without recourse to physics-based models becomes most desirable. These themes are returned to in the synthesis discussion of Chapter 8.

5.4.2 Sensor placement

15 PCB 353B16 piezoelectric accelerometers were mounted on the upper (as mounted) surface of the wing using ceramic cement. The location of the sensors, inspection panels and sub-surface stiffening elements (dotted lines) are shown schematically in Figure 5-2. The sensors were placed in an *ad hoc* fashion on the basis of previous experience, with no formalised sensor placement optimisation

undertaken. The sensors were placed so as to form transmissibility 'paths' between sensor pairs - these paths are indicated in Figure 5-2 and specified in full in Chapter 6. The sensors are denoted S1 to S15. Placement directly above stiffening elements was avoided as it was believed that such locations would offer poor observation of localised changes in structural flexibility.

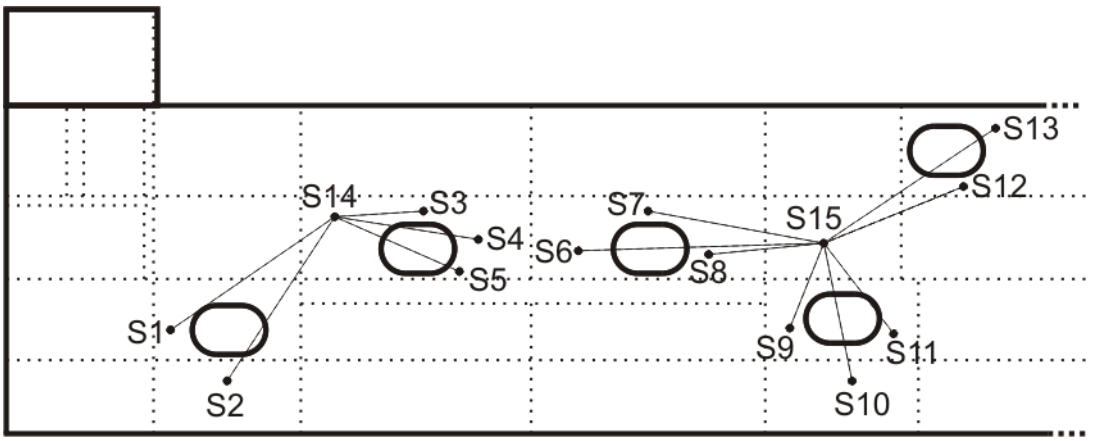

Figure 5-2 Schematic sensor placement diagram

5.4.3 Equipment

Experimental data acquisition was performed using a DIFA SCADAS III unit controlled by LMS software running on a desktop PC. All measurements were recorded within a frequency range of 0-2048 Hz with a resolution of 0.5 Hz. The structure was excited with a band-limited white Gaussian signal using a Gearing and Watson amplifier and shaker mounted beneath the wing. Both the real and imaginary parts of the accelerance FRFs were recorded at 15 response locations using single-axis accelerometers. Five-average samples were recorded in all cases as this was found to offer a good compromise between noise reduction and acquisition time.

5.4.4 Damage introduction

In order to introduce damage in a repeatable, and in some sense realistic, way advantage was taken of the presence of inspection panels on the underside of the wing, with the wing being mounted upside-down to enable access. Five such panels were employed in the present study, all of the same dimension and orientation. The panel dimensions are given in Figure 5-3.

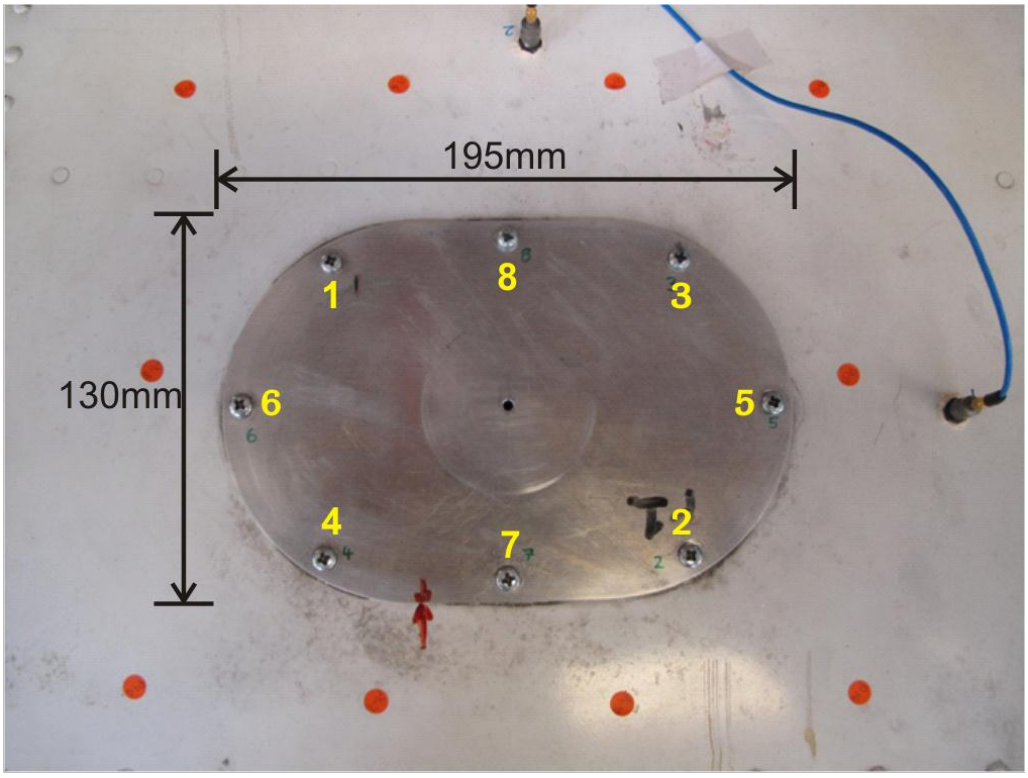

Figure 5-3 Inspection panel P1, with extremal dimensions and attachment sequence

Through removal of these panels, a 'gross' level of damage could be introduced at each panel location. This form of damage state is denoted in the study by the letters DO, followed by the numbers of the panels removed e.g. DO_1_5 refers to the removal of panels 1 and 5. The undamaged case, with all five intact panels in place, is referred to as N (normal condition). In this chapter the terms 'panel off' and 'damage' are used interchangeably. The normal state (all panels attached) and the damage state DO_2 (removal of panel 2 only) for panel 2 are shown in Figure 5-4.

The advantages of using the removal of a panel as a proxy for damage are that:

1. It is non-destructive;
2. it is repeatable; and,
3. the primary effect of panel removal (a localised reduction in structural stiffness) is expected to be similar to the effect of introducing gross damage at the same location

The disadvantage is that the repeatability is not perfect. During preliminary studies it was found that the removal and reattachment of the panel led to substantial variability in the FRF observations. This and other sources of variability were considered in greater detail prior to the specification of the test sequence.

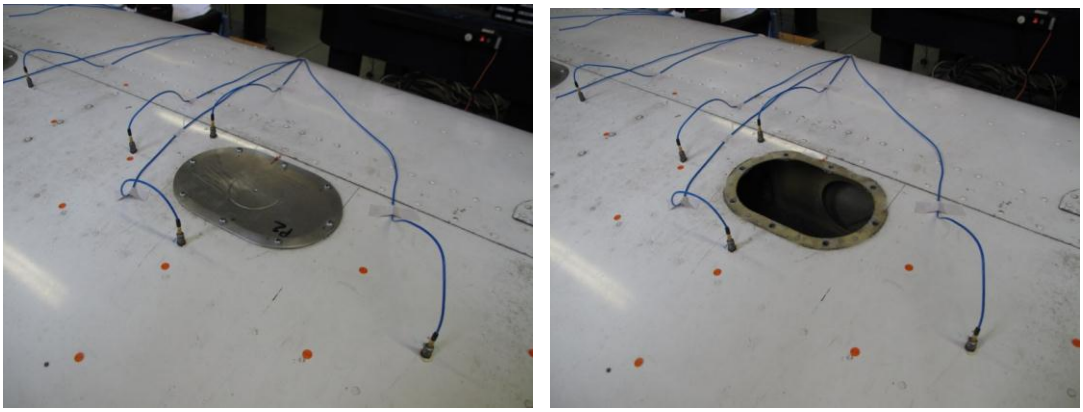

Figure 5-4 Inspection panel P2 in normal condition (left) and damaged condition (right)

5.5 Sources of measurement variability

Consistent with the aims of the *statistical pattern recognition* (SPR) Paradigm [6], consideration is given to potential sources of variability at the test planning stage. The ultimate aim is to develop a classifier that will generalise well to newly-presented data. By identifying and allowing for sources of variability at an early stage the prospects of achieving good generalisation at a later stage may be increased. With regard to variability, the objectives of the current study are:

- To reduce variability to whatever extent is possible.
- To ensure the remaining variability is appropriately represented in the dataset.

Systematic variability due to panel boundary conditions, operational variability and environmental effects are considered. This is in addition to random variation between runs (i.e. 'noise'). Although not pursued here, the process of identifying and quantifying the effects of experimental variability could be formalised using a form such as the PIRT table introduced in Chapter 4.

5.5.1 Boundary condition variability

The inspection panels are attached using eight screws. It has been observed in previous studies [73, 74], that the measured FRFs may display marked sensitivity to the boundary condition variability that arises from the removal and reattachment of the panels. The degree of variability observed can be problematic when attempting to construct a sensitive, robust classifier to identify damage states. In the experimental portion of the study undertaken for this chapter, emphasis was placed upon gathering an improved representation of states that would be used for training

and validating the developed classifier. To this end, systematic randomisation is applied in the development of the training set test sequence in order to incorporate the effect of panel removal and replacement in the data used for classifier training.

In order to reduce the random variability associated with the panel boundary conditions, a torque-controlled electric screwdriver was used for removing and replacing the panels. The order in which the bolts were replaced was also taken into account, with a consistent sequence of attachment maintained throughout the test. The bolt numbering (in yellow) in Figure 5-3 indicates the order of bolt reattachment. These efforts to reduce random variability were consistently applied for both the training and testing sets. In practical applications, this level of consistency would not be expected. The training dataset should be structured so as to capture the true level of variability as fully as possible, in order that a classifier that generalises well to new data may be built.

5.5.2 Operational variability

The variable loadings exerted upon aerospace structures when in use are numerous and may have marked effect on the dynamic responses of the structure. They include passenger and cargo payloads, fuel loading and external excitation loadings. External excitation is provided by wind effects in addition to flight events such as takeoff and landing. These loadings are often difficult to accurately quantify, and are in general non-stationary during the flight window. The effects of operational variability on response measurements has been investigated in the SHM literature, both for civil [75, 76] and mechanical [77] applications.

For the laboratory-based experiment the effects of operational variability are assumed negligible. The only potentially variable loading exerted on the structure is the excitation provided by the shaker. The forcing amplitude is considered to be a controlled variable (although no feedback loop was employed to maintain the forcing amplitude) and was kept constant throughout testing. The measured data, FRFs, may be considered to represent a transfer function between inputs and outputs of the system. As such the FRF measurements should be not display significant sensitivity to small fluctuations in the amplitude of excitation provided that the structure does not exhibit significant nonlinearity. A linearity test was not conducted in this case. The FRFs are, however, converted into transmissibility spectra during feature extraction. Transmissibility spectra represent a transfer

function between acceleration responses at a pair of points. One of the motivations for employing transmissibility data is that the effect of excitation variability may be further mitigated [73]. It is concluded that the effects of operational variability should be very small for the wing structure.

5.5.3 Environmental effects

It has been established that structural responses may display sensitivity to environmental factors such as temperature and humidity. Such effects are a major confounding factor in applying SHM to real-world structures, where changes in response due to temperature may be of a similar order to changes caused by substantial damage. Humidity has also been identified as an important factor for composite structures.

For the metallic, laboratory-based wing under investigation, it was not expected that environmental factors would have a significant effect upon the responses of the structure. Additionally, the degree of temperature variation was expected to be very low - on the order of several degrees Celsius. However, it was seen as good practice to record temperatures during the test, if not for direct investigation for the current study then in support of any future testing that may be carried out on the structure. Electronic thermometers with a nominal sensitivity of $\pm 0.1°C$ were placed at either end of the wing and temperatures recorded at the beginning of each test run. Humidity data were not recorded.

5.6 Test sequencing

Two tests were conducted, resulting in two data sets being available with which to develop and test classifiers. A test sequence was developed for each test to reflect the testing objectives, and taking into consideration the sources of measurement variability identified. Randomisation and blocking were applied in the specification of the test sequences in order to account for the effects of measurement noise and panel boundary condition variability.

Dataset A comprises 1000 normal state observations and 1000 single-site damage state observations. The primary objective for the dataset was to allowing training and validation of a classifier using normal condition and single–site damage condition data only, but which is capable of generalising to the multi-site damage identification task using data from a separate test. This is a demanding objective, and

particular attention was paid to full randomisation of the panel boundary condition. 'Randomisation', as used here, refers to the removal and replacement of panels to ensure that the latent boundary condition variation (that which is present despite the use of a torque controlled screwdriver and care over the order of screw tightening)be represented in the dataset.

Dataset B comprised 2300 normal, 1000 single-site damage and 1300 multi-site damage state observations. The primary objective of this test was to provide data with which to test the classifier developed using the single-site data of Dataset A. The secondary objective was to allow development of 'best possible' classifier to be trained and validated using both single-site and multi-site damage data. As the primary objective of the dataset is to provide a 'testing set', and as a far greater number of structural states were to be included, a marginally less stringent to randomisation of boundary conditions was taken than for Dataset A. Randomisation was limited to the removal and replacement of individual panels, rather than full randomisation of the panel boundary conditions.

For both tests, a large number of response measurements were taken with the aim of capturing random measurement variability. For Dataset A, repetition was applied to further enhance the dataset. In developing the test sequences, a balance is sought such that both random and systematic variabilities may be well represented in the datasets without exceeding a reasonable timescale for testing.

5.6.1 Dataset A

The test sequence used for gathering Dataset A is given in Table 5-1. The test sequence is presented as a hierarchy. The test consists of 10 blocks, each comprising a normal condition run and a single-site damage condition run. Each run contains 100 observations of five-averaged response data. The use of blocks of testing allowed two levels of boundary condition randomisation to be introduced.

In order to evaluate the variability arising from single-panel removal within each block, the panel of interest was removed and replaced after every 20 observations for the normal condition runs. Within each block only the panel of interest is removed - the remaining four panels remain in place. The intention is that the sensitivity of the responses to removal of a single panel may be assessed, and that a training set is gathered that is more fully representative of observations of

undamaged and damaged states that may be encountered. For the panel-off case, where there are no single-plate boundary condition effects to consider, a straight run of 100 observations was recorded. Between blocks, all five panels were removed and replaced. This introduced full randomisation of the boundary conditions. In order to add further information, blocks were executed for each of the five panels, and then the sequence repeated.

The normal condition runs, which included the removal and replacement of one panel after every set of 20 observations, took approximately 45 minutes each to complete. Approximately 40 minutes were required for each damaged condition run. Removing and replacing all panels between runs took approximately five minutes. Each block thus took approximately 90 minutes to complete, with a total time requirement of around 15 hours. Testing was conducted over a five day period. The temperatures recorded at either end of the wing are denoted T1 and T2. The observed values of T1 and T2 at the start of each run are included in Table 5-1. The testing was carried out by the author.

Block no	Run no	Run Label	No of Observations	T1 [°C]	T2 [°C]	Date	Time
1	1	N1	5 x 20	21.1	21.5	16/07/09	11:21
	2	DO1_R1	100	21.5	21.9	16/07/09	14:49
			Replace all panels				
2	3	N2	5 x 20	21.5	21.9	16/07/09	15:58
	4	DO2_R1	100	21.5	21.7	16/07/09	17:29
			Replace all panels				
3	5	N3	5 x 20	19.9	19.9	17/07/09	08:16
	6	DO3_R1	100	20.2	20.3	17/07/09	09:34
			Replace all panels				
4	7	N4	5 x 20	20.3	20.5	17/07/09	10:34
	8	DO4_R1	100	20.5	20.9	17/07/09	12:12
			Replace all panels				
5	9	N5	5 x 20	20.5	20.8	17/07/09	13:44
	10	DO5_R1	100	20.4	20.8	17/07/09	15:09
			Replace all panels				
6	11	N6	5 x 20	19.8	19.2	20/07/09	09:50
	12	DO1_R2	100	20.1	20.2	20/07/09	11:16
			Replace all panels				
7	13	N7	5 x 20	20.3	20.6	20/07/09	13:27
	14	DO2_R2	100	20.6	20.9	20/07/09	15:09
			Replace all panels				
8	15	N8	5 x 20	20.3	20.4	22/07/09	10:28
	16	DO3_R2	100	20.5	20.9	22/07/09	11:44
			Replace all panels				
9	17	N9	5 x 20	20.9	21.2	23/07/09	14:16
	18	DO4_R2	100	20.9	21.2	23/07/09	15:40
			Replace all panels				
10	19	N10	5 x 20	20.3	20.5	23/07/09	11:19
	20	DO5_R2	100	20.6	20.8	23/07/09	12:42

Table 5-1 Dataset A test sequence

5.6.2 Dataset B

The test sequence used for gathering Dataset B is summarised in Table 5-2 and Table 5-3. The dataset was developed for use in several studies. It contains a larger number of damage states (1 normal condition and 5 single-site damage states, plus 26 multi-site damage states) plus a greater number of repetitions for each state with various degrees of randomisation, particularly of the normal condition.

The test was conducted in blocks comprising a number of normal condition runs and a number of runs from one of the damaged state conditions. For each of the single-site damage states, each block contained 20 normal condition runs alternated with 20 damaged condition runs. For the multiple-site damage states, each run contained five normal and five damaged condition runs. Each run contains 10 observations of five-averaged data. Within each block only the panel (or panels) of interest are removed and replaced. For example in the first block, Run 1 compromised 10 observations of the normal condition N1. Panel 1 was then removed

and Run 2, comprising 10 observations of the damaged condition DO1, was recorded. The panel was then replaced and Run 3, a further 10 observations of N1, was performed. Note that the format of Table 5-2 and Table 5-3 differs from that of Table 5-1.

Systematic randomisation of boundary conditions through removal and replacement of all panels between blocks was not pursued, although gradual randomisation was introduced as panels were removed and replaced during the test sequence. This is of some importance to any conclusions drawn with regard to the approach. The primary objective of Dataset B is to provide a training set to allow the evaluation of the performance of the developed classifiers when presented with new data comprising normal, single-site and multi-site data. It is not intended to be the basis of an exhaustive examination of how the developed classifiers perform across the full domain of possible observations, in the presence of variability. Any conclusions drawn must thus include the caveat that they are valid only for the data with which the classifiers were tested; and caution should be employed in drawing general conclusions on the basis of a limited testing dataset.

Blocks 1 to 5 required approximately 200 minutes each to complete. Blocks 6 to 31 involved fewer observations but the removal or replacement of between two and five panels after every 10 observations. On average, approximately 40 minutes were required for each of these 'multiple-damage' blocks. The total testing time was approximately 34 hours and was distributed over 11 days of testing. The minimum and maximum ranges of temperatures T1 and T2 within each block are given. The testing was carried out by Dr Graeme Manson.

Block no	Run no	Run Label	No of Observations	T1 [°C]	T2 [°C]	Dates
1	1-40	N1	20x10	min 21.0	min 21.0	29/06/09-
		DO1	20x10	max 21.7	max 21.7	30/06/09
2	41-80	N2	20x10	min 21.5	min 22.0	30/06/09-
		DO2	20x10	max 22.6	max 23.1	01/07/09
3	81-120	N3	20x10	min 22.6	min 23.4	01/07/09-
		DO3	20x10	max 23.5	max 25.1	01/07/09
4	121-160	N4	20x10	min 23.2	min 22.6	02/07/09-
		DO4	20x10	max 24.2	max 29.3	02/07/09
5	161-200	N5	20x10	min 22.3	min 22.3	02/07/09-
		DO5	20x10	max 24.2	max 24.8	03/07/09

Table 5-2 Dataset B test sequence: part 1

Block no	Run no	Run Label	No of Observations	T1 [°C]	T2 [°C]	Dates
6	201-210	N6 DO1_2	5x10 5x10	min 22.7 max 22.9	min 23.0 max23.0	03/07/09
7	211-220	N7 DO1_3	5x10 5x10	min 22.9 max 23.0	min 23.1 max 23.1	03/07/09
8	212-230	N8 DO1_4	5x10 5x10	min 22.8 max 23.0	min 23.1 max 23.2	03/07/09
9	231-240	N9 DO1_5	5x10 5x10	min 22.9 max 23.1	min 23.2 max 23.2	03/07/09
10	241-250	N10 DO2_3	5x10 5x10	min 21.9 max 22.0	min 22.1 max 22.2	06/07/09
11	251-260	N11 DO2_4	5x10 5x10	min 22.1 max 22.3	min 22.3 max 22.5	06/07/09
12	261-270	N12 DO2_5	5x10 5x10	min 21.3 max 21.5	min 21.7 max 21.8	06/07/09
13	271-280	N13 DO3_4	5x10 5x10	min 21.3 max 21.5	min 21.7 max 21.8	07/07/09
14	281-290	N14 DO3_5	5x10 5x10	min 21.3 max 21.4	min 21.5 max 21.7	07/07/09
15	291-300	N15 DO4_5	5x10 5x10	min 21.3 max 21.5	min 21.5 max 21.7	07/07/09
16	301-310	N16 DO1_2_3	5x10 5x10	min 20.3 max 20.4	min 20.1 max 20.8	09/07/09
17	311-320	N17 DO1_2_4	5x10 5x10	min 20.3 max 20.5	min 20.7 max 20.9	09/07/09
18	321-330	N18 DO1_2_5	5x10 5x10	min 20.5 max 20.6	min 20.7 max 20.9	09/07/09
19	331-340	N19 DO1_3_4	5x10 5x10	min 20.7 max 20.8	min 20.9 max 21.1	09/07/09
20	301-310	N20 DO1_3_5	5x10 5x10	min 20.7 max 20.8	min 21.1 max 21.1	09/07/09
21	311-320	N21 DO1_4_5	5x10 5x10	min 19.1 max 19.4	min 19.2 max 19.5	10/07/09
22	321-330	N22 DO2_3_4	5x10 5x10	min 20.0 max 20.3	min 20.4 max 20.7	10/07/09
23	331-340	N23 DO2_3_5	5x10 5x10	min 20.3 max 20.4	min 20.7 max 20.8	10/07/09
24	341-350	N24 DO2_4_5	5x10 5x10	min 20.4 max 20.6	min 20.8 max 20.9	10/07/09
25	351-360	N25 DO3_4_5	5x10 5x10	min 19.6 max 20.6	min 19.4 max 20.0	13/07/09
26	361-370	N26 DO1_2_3_4	5x10 5x10	min 20.8 max 20.9	min 21.1 max 21.2	13/07/09
27	371-380	N27 DO1_2_3_5	5x10 5x10	min 20.9 max 21.1	min 21.3 max 21.5	13/07/09
28	381-390	N28 DO1_2_4_5	5x10 5x10	min 20.8 max 21.1	min 21.2 max 21.5	13/07/09
29	391-400	N29 DO1_3_4_5	5x10 5x10	min 20.9 max 20.9	min 21.1 max 21.2	13/07/09
30	401-410	N30 DO2_3_4_5	5x10 5x10	min 20.8 max 21.1	min 21.3 max 21.5	14/07/09
31	411-420	N31 DO1_2_3_4_5	5x10 5x10	min 21.1 max 21.2	min 21.5 max 21.7	14/07/09

Table 5-3 Dataset B test sequence: part 2

5.7 Discussion

In this chapter the motivation behind the multiple-damage location problem has been introduced, with reference made to previous studies. The approach to be investigated in the following chapters - based upon data-driven statistical pattern recognition - has been introduced, and the data acquisition portion of the approach presented. The investigated structure and the experimental setup have been introduced, and the recording of the two data sets used in the remainder of the study has been described.

A great deal of thought has been given to the issue of test sequencing. Decisions made with regard to the test sequence may have a substantial bearing on the ability of the developed classifier to handle experimental uncertainty, and ultimately its ability to generalise to new data. Care should be taken in order to produce a training dataset that is representative of the conditions under which classification will be sought. For the structure under study, it was felt that the boundary condition variability associated with the removal and replacement of inspection panels was the most significant source of random variation in the response spectra. Accordingly, care was taken in order to reduce this variability to as great an extent as possible by using a torque-controlled screwdriver and raking care over screw tightening, and to account for the outstanding variability in the recorded data through randomisation. The attention paid to test sequencing represents a development from previous studies in this area [73, 74, 78]. There appears to be scope for further development in this regard. The application of sensitivity analysis methods in the design of experiments appears to be a logical step in optimising the SHM data acquisition process.

The main outcome of the chapter is the gathering of two datasets for the training and testing of statistical damage classifiers. In Chapter 6, a reduced feature set is selected from the recorded data sets. In Chapter 7 these features are used to train classifiers for the multiple damage identification task.

6 A Data-Driven Approach: Feature Selection and Evaluation

6.1 Introduction

In Chapter 5, an approach was proposed to alleviate the problem of the explosion in the number of damage states that occurs when dealing with multi-site damage location problems. The approach involves the construction and evaluation of a statistical classifier tasked with identifying multiple-site damage despite being trained with single-site damage data only. The experimental structure was introduced and the first stage in the development of the classifier - the acquisition of two datasets dubbed A and B with which to train and test the classifier - was described.

The aim in this chapter is to describe the reduction of the high-dimensional data recorded in Chapter 5 into a set of low-dimensional features suitable for the training of statistical classifiers in Chapter 7, and to evaluate the individual performance of the chosen feature set. The objectives of the feature selection and evaluation task may be summarised in several steps.

- Extraction of a candidate feature set from raw experimental data.
- Selection of a low-dimensional feature set containing features that are sensitive to *single-site* damage, using a training set drawn from Dataset A.
- Evaluation of these features for the identification of *single-site* damage, using a testing set drawn from Dataset B.

- Evaluation of these features for the identification of *multi-site* damage, using a testing set drawn from Dataset B.

Each of these steps is covered in turn. The hypothesis of this chapter is that it may be possible to find features that offer a good level of discrimination between multiple-site damage states, despite only single-site damage states being available for feature selection.

6.2 Feature extraction

Damage leads to the structural response of the structure deviating from that observed when it is in its initial, undamaged condition. The structural response features considered in this study are transmissibility spectra. By comparing examples of undamaged and damaged spectra, regions of the spectra that are sensitive to particular damage states may be identified. These regions form the basis of features that may subsequently be used to train a statistical damage classifier. In the interests of developing a statistical classifier, it is desirable that the feature set used is of low dimension. Achieving a suitably concise feature set requires further condensation of the identified spectral region. In this study the additional data reduction is performed using a discordancy measure - the *Mahalanobis squared-distance* (MSD) - between the newly-presented (and therefore potentially-damaged) state and the previously-recorded undamaged state. The resulting features are the discordancy values associated with damage sensitive regions of the transmissibility spectra. In Chapter 7, classifiers are trained using observations of the resulting concise set of damage sensitive features.

The data recorded in Chapter 5 comprised 15 accelerance FRFs, each containing 4097 spectral lines. The procedure undertaken for reducing this high-dimensional data to a low-dimensional set of features in this study is as follows. The FRF data are first converted into transmissibility spectra, which have been shown to perform well in similar damage identification scenarios [73, 74, 78-80]. In this study, log magnitude transmissibility spectra are employed, which exhibit further useful properties when employed for feature selection, as discussed below. Next, regions of the transmissibility spectra that are sensitive to damage at a particular location are selected. In this instance, feature selection is performed manually on the basis of a principled set of objectives and aided by visual tools. Finally, the identified low-

dimensional spectral ranges are reduced to a scalar feature value through the application of the MSD measure.

These steps are presented below. First, the transmissibility paths adopted for the study are specified, and the MSD discordancy measure is introduced. In the following section, the procedure used for the selection of damage-sensitive features from the wing data is described.

6.2.1 Extraction of transmissibility data

The raw test data comprised real and imaginary components of the accelerance FRFs between the 15 response locations and the single force input from the shaker. Prior to further processing, this data was converted to transmissibility spectra using the relationship,

$$T_{ij}(\omega) = \frac{H_{ik}(\omega)}{H_{jk}(\omega)} \tag{6.1}$$

where $H_{ik}(\omega)$ is the FRF between input location k and response location i, $H_{jk}(\omega)$ is the FRF between input location k and response location j, and $T_{ij}(\omega)$ is the resulting transmissibility spectra between response location i and response location j. A useful step taken here is to take the log of the resulting transmissibility spectra in order to take advantage of some useful properties of log relationships. One advantage of doing so is that the log scale aids visualisation during feature selection, with 'peaks' in the transmissibility spectra typically returning values several orders of magnitude greater than those for intervening regions of the spectra. Additionally, note that,

$$\log(T_{ij}) = \log\left(\frac{H_{ik}}{H_{jk}}\right) = -\log\left(\frac{H_{jk}}{H_{ik}}\right) = -\log(T_{ji}) \tag{6.2}$$

The practical benefit of this simple property is that a single logged spectra contains information on the magnitude of the transmissibilities 'in both directions', with only the sign of the spectra affected by the choice of numerator and denominator. Spectra corresponding to different damage states may thus be considered using a single plot, rather than one for each 'direction'. Changing the selection of numerator and denominator would have the effect of 'flipping' the spectra about the y=0 axis, with the information contained in the spectra otherwise unaffected.

6.2.2 Specification of experimental transmissibility paths

Log magnitude transmissibility spectra were generated for 13 transmissibility paths. The locations of the 5 inspection panels are illustrated schematically in Figure 6-1, with the positions of the sub-surface stiffening elements shown as dotted lines.

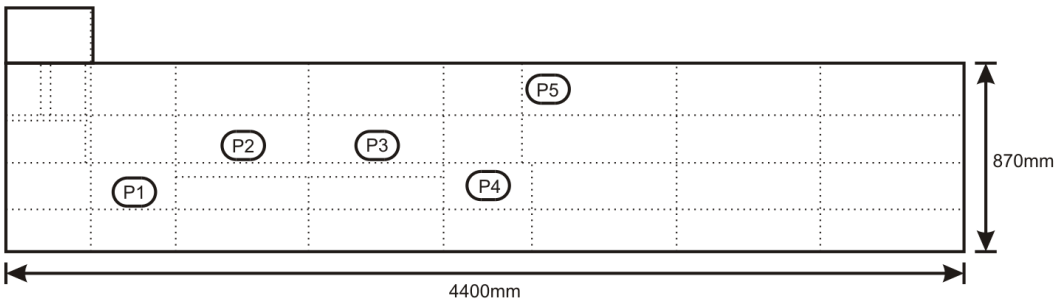

Figure 6-1 Schematic of wing inspection panels and stiffening elements

In total, 15 accelerometers were mounted on the upper surface of the wing. The locations of the sensors, denoted S1 to S15, are shown schematically in Figure 6-2. Placement directly above stiffening elements was avoided as it was believed that such locations would offer poor observation of localised changes in structural flexibility.

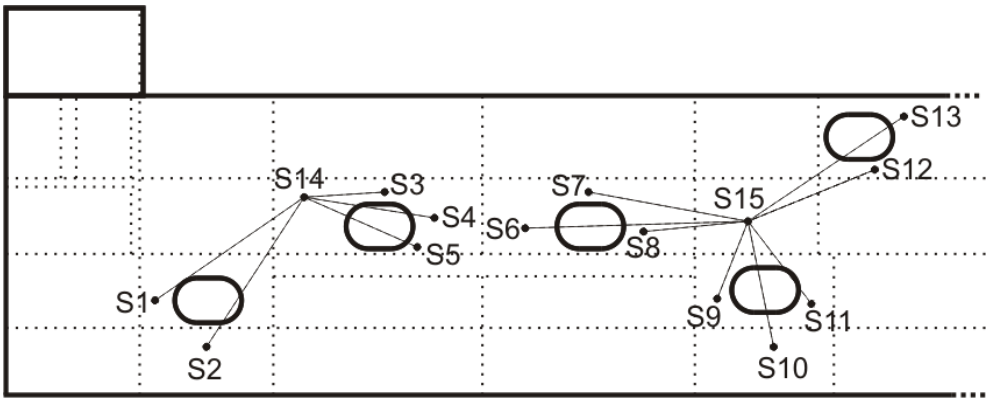

Figure 6-2 Schematic of sensor locations

The sensors were placed so as to form transmissibility 'paths' between sensor pairs. The sensors are arranged into two groups, the first covering panels P1 and P2 and the second covering panels P3, P4 and P5. Each group contains a 'reference' sensor which is common to all transmissibility paths within the group and a number of 'response' sensors. For the first group the reference sensor is S14 and the response sensors are S1-S5; for the second group the reference sensor is S15 and the response sensors are S6 to S13. Transmissibility paths are formed between pairs of response and reference sensors. The 13 resulting pairings are specified in Table 6-1.

Panel Label	Path Label	Reference Transducer	Response Transducer
P1	T1	S14	S1
	T2	S14	S2
P2	T3	S14	S3
	T4	S14	S4
	T5	S14	S5
P3	T6	S15	S6
	T7	S15	S7
	T8	S15	S8
P4	T9	S15	S9
	T10	S15	S10
	T11	S15	S11
P5	T12	S15	S12
	T13	S15	S13

Table 6-1 Specification of transmissibility paths

The paths are indicated schematically in Figure 6-3. Each panel has either two or three transmissibility paths associated with it. It is expected that the effects of removing a panel may observed in all spectra to some extent, but that the effect will be most apparent in the spectra associated with that panel.

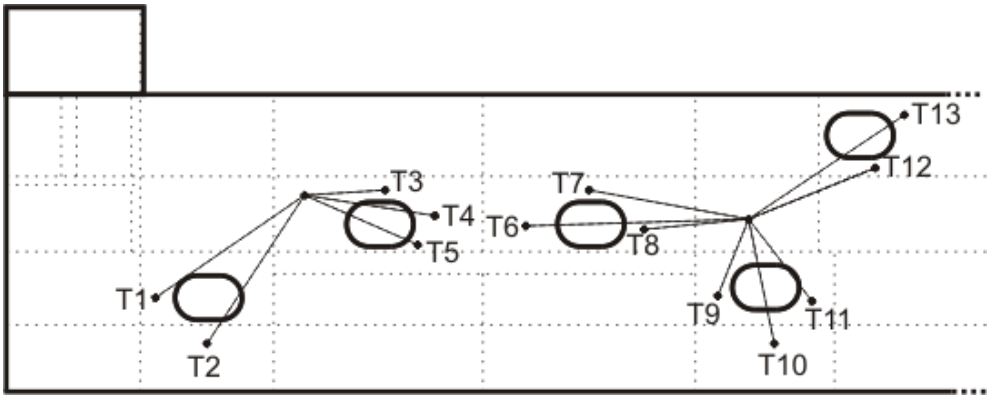

Figure 6-3 Schematic of transmissibility paths

The resulting transmissibility data is very high-dimensional. Each observation comprises 13 spectra, each of which contains 4097 spectral lines, giving a 53261-dimensional dataset. In order to alleviate the 'curse of dimensionality' [81], some degree of dimensionality reduction must be pursued prior to applying pattern recognition methods. In this study dimensionality reduction is realised as the selection of damage-sensitive regions of the transmissibility spectra, and the reduction of these multivariate regions to univariate features using a discordancy measure. First, changes in the spectra as a result of damage are considered.

6.2.3 Sensitivity of transmissibility spectra to damage

An example of the sensitivity of transmissibility spectra to panel removal is given in Figure 6-4. The figure shows a short section of the transmissibility spectrum T1 for each of the structural states recorded in Dataset A. The spectra demonstrate a clear distinction between the removal of panel P1 and the other damage states. A window capturing this behaviour may be used to form a damage-discriminative feature.

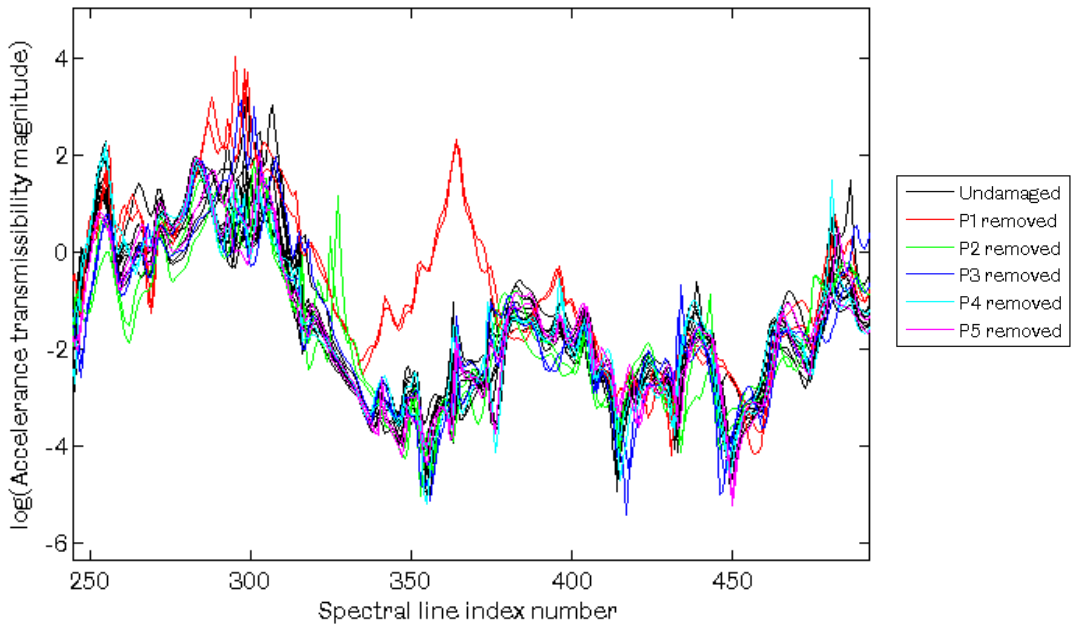

Figure 6-4 Section of transmissibility spectrum T1 for Dataset A

Note that in creating Figure 6-4, averaging has been applied in order to reduce the number of observations that require plotting. For the purposes of visualisation during feature selection, selective averaging was carried out. It was found that the factor explaining the largest proportion of variance was the removal and replacement of all panels. Output noise provided a smaller contribution to the variability of the spectra. By taking the average of each run specified in the test sequence (given in Table 5-1) the variability associated with removal and replacement of all the panels between repetitions could be retained, and other factors 'averaged out'. This proved useful for the purposes of visualisation as the number of spectra that required plotting was reduced from 2000 observations to a single, averaged spectrum for each of the 20 runs.

Selection of damage-sensitive windows of the spectra is the first stage in reducing the dimension of the recorded data. Further dimensionality reduction is achieved

through outlier analysis using a discordancy measure. The discordancy measure employed is the Mahalanobis squared-distance, introduced below.

6.2.4 Mahalanobis squared-distance (MSD) measure

The MSD is a multivariate extension of the univariate discordancy measure. Discordancy measures allow deviations from normality to be quantitatively evaluated.

A brief summary of the technique is given here: the technique is described in [23] and validated for an experimental structure in [73, 74, 79]. For a multivariate data set consisting of *n* observations in *p* variables, the MSD may be used to give a measure of the discordancy of any given observation. The scalar discordancy value D_ξ is given by,

$$D_\xi = \left(\boldsymbol{x}_\xi - \overline{\boldsymbol{x}}\right)^{\mathrm{T}} [\boldsymbol{\Sigma}]^{-1} \left(\boldsymbol{x}_\xi - \overline{\boldsymbol{x}}\right) \tag{6.3}$$

where $\boldsymbol{x}_\xi$ is the potential outlier, $\overline{\boldsymbol{x}}$ is the mean of the sample observations and $\boldsymbol{\Sigma}$ is the sample covariance matrix. The MSD approach offers a measure of the extent by which a sample differs from the population regardless of the manner in which it differs.

6.3 Feature selection from single-site training data

Feature selection from a spectrum can be done in several ways. It may be performed manually, on the basis of engineering judgement; algorithmically, primarily by applying some form of combinatorial optimisation; or a combination thereof. The three approaches are illustrated in [78]. There, outlier analysis was performed on transmissibility spectra from an experimental structure subjected to damage. The aim was to select a set of nine features as inputs to a *multi-layer perceptron* (MLP) classifier. First a manual approach was taken to select 44 candidate features from the full spectral dataset, and from these a further reduction to the required nine features was made on the basis of engineering judgement. Secondly, a Genetic Algorithm (GA) was employed to select an 'optimal' subset of nine features from the manually selected set of 44: an example of a manual/algorithmic approach. Finally, a GA was run to select nine features from the full spectral dataset, with no prior selection of feature ranges. The broad conclusion was that the manual/algorithmic approach worked 'best' for the application presented, followed by the fully

algorithmic approach and the fully manual approach. The great benefit of the fully algorithmic approach is in removing the requirement for an exhaustive manual search for candidate features

In this chapter, the objective is to select features that are robustly indicative of the removal of panels using observations from Dataset A. Objectively optimal feature selection is not pursued in this instance, although this is discussed as an area of future work. The feature selection method is instead performed manually. This process is guided by a series of considerations, and aided by appropriate visualisations. In total, 10 features were selected for each of the five panels. The performance of these features when applied to Dataset B is evaluated in the Section 6.4.

6.3.1 Discordancy plots

The aim of the feature selection step is to identify features that discriminate between the normal condition and the removal of a single panel, for each of the five panels in turn.

For the present work, a predominantly manual feature selection approach was employed, guided by evaluation of the discordancy values of all possible 20-spectral-line windows of the spectra. As each spectrum contains 4097 lines, 4078 such windows existed for each of the 13 spectra, resulting in a total of 53014 *candidate features*.

Discordancy values for each of these features were evaluated using the mean-averaged data recorded for each of the 20 runs in Dataset A. Plotting the discordancy values across the feature set for each of the structural states allows the sensitivity of the features to be visualised. An example for a limited range of the candidate feature set is given in Figure 6-5. The two repetitions of removal of panel P1 are shown in red. The remaining 18 damage states (10 normal condition runs and eight other single panel removal runs) are shown in green.

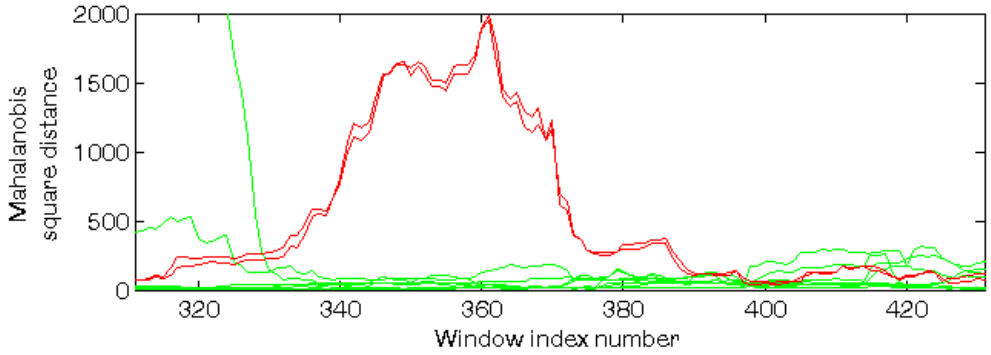

Figure 6-5 Discordancy plot for Window Index Numbers 311 to 431

From Figure 6-5, it is apparent that the feature displaying the greatest sensitivity to the removal of the panel is identified by the window with the index number 361, which corresponds to spectral lines 361-380. However, sensitivity to damage alone is an insufficient criterion for selecting features that are robustly indicative of damage. A more wide-ranging set of criteria were instead applied.

6.3.2 Feature selection criteria

Four criteria were considered when assessing the 'strength' of a candidate feature.

1) The feature should be sensitive to the removal of the specified panel.
2) In the interests of promoting robustness there should be consistency in the feature values between the two available repetitions.
3) Preferably, the feature should be insensitive to the removal of panels other than that specified. The success of the classifier is predicated upon the discriminative ability of the combined feature set; however, is expected that if features that are capable of discriminating between damage classes individually can be identified, the discriminative ability of the combined set will also be increased.
4) Preferably, there should be degree of consistency in the above considerations between the selected feature and those in the immediate vicinity. The feature selection process is guided by a desire to attribute some degree of confidence to the performance of the feature when presented with new data. It would appear, at the very least intuitively, that more confidence may be placed in a feature that is one of several contiguous features that perform well, rather than a feature that performs conspicuously better than its neighbours.

6.3.3 Feature selection plots

An extended list of candidate features was identified using the MSD plots, with between 15 and 20 features identified for each panel. Each of these features was manually assessed in order to reduce the feature set to a final size of 10 features per panel. Inspection was aided by employing visualisations of the form given in Figure 6-6.

The top plot gives an expanded view of MSD values for the identified feature and for those in its vicinity. The values plotted in red are the discordancy values for the window when the panel of interest is removed (Panel 1 in the example plot in Figure 6-6); the values in green are for the normal conditions and removals of the other four panels individually (i.e. Panels 2-5 in the example). The dashed vertical line denotes the selected feature itself. The *window index number* corresponds to the index of the first spectral line in the 20-line window used to form the feature.

The lower plot displays the portion of the transmissibility spectra corresponding to the selected window for Dataset A. The window itself is denoted as lying between the dashed vertical lines. In common with the upper plot, the spectra in red correspond to single-site removal of the panel of interest, and the spectra in green to the single-site removal of the other panels. The averaged normal condition spectra are plotted in black.

The suitability of the candidate features was manually evaluated on the basis of the selection criteria specified above, and a final feature set defined.

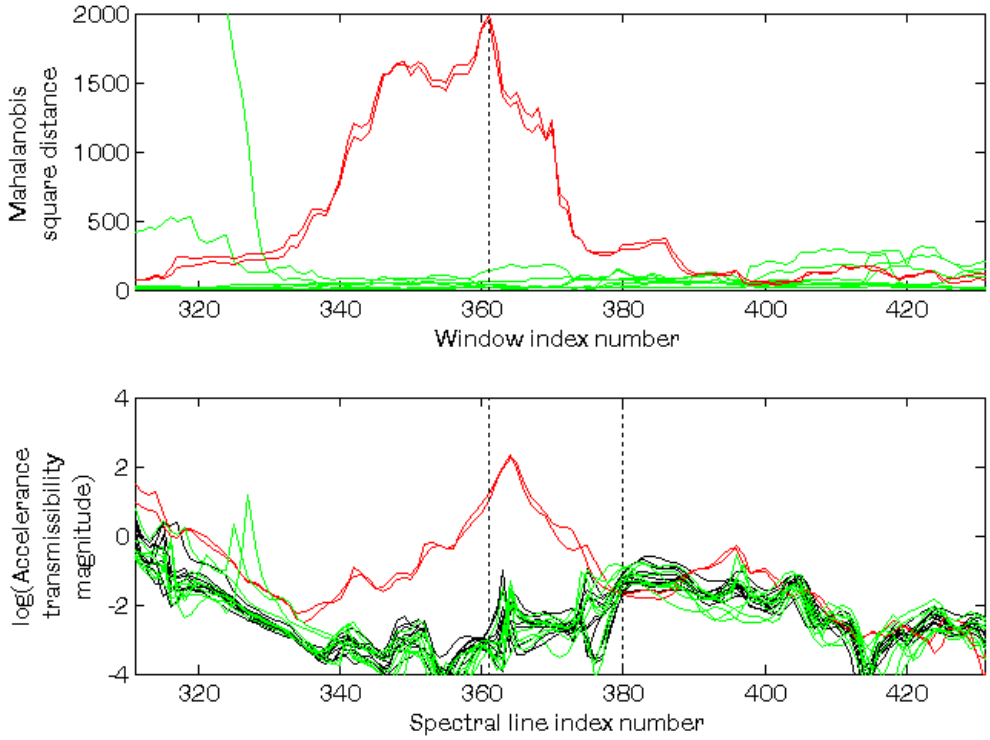

Figure 6-6 Feature selection plots for Window Index Number 361

Returning to the transmissibility plots after initially assessing the features on the basis solely of discordancy values allows the information available for making a final decision to be extended. In several cases, considering the spectra led to some adjustment of the originally identified feature. One example of this is where a greater level of apparent robustness may be achieved through a small adjustment of the window. A second example is where the window returning the highest MSD values is found to contain multiple characteristic structures, such as 'peaks' in the spectra. In such cases, it may be desirable for the window to span a single peak well, rather than partially span two or more such peaks.

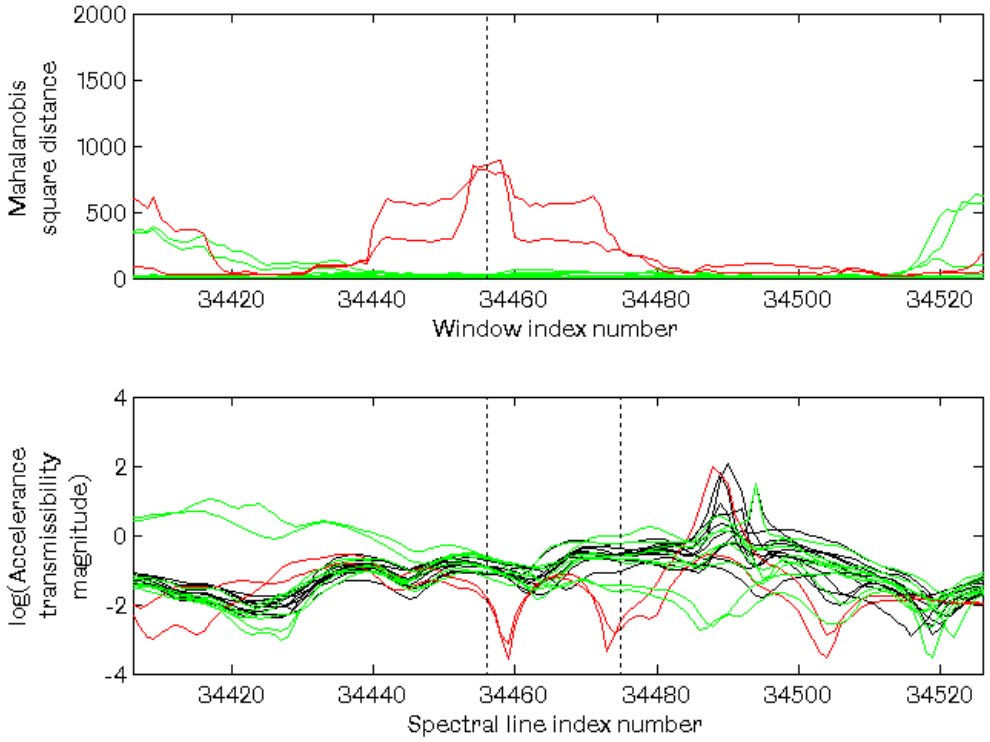

Figure 6-7 Feature selection plots for Window Index Number 34456

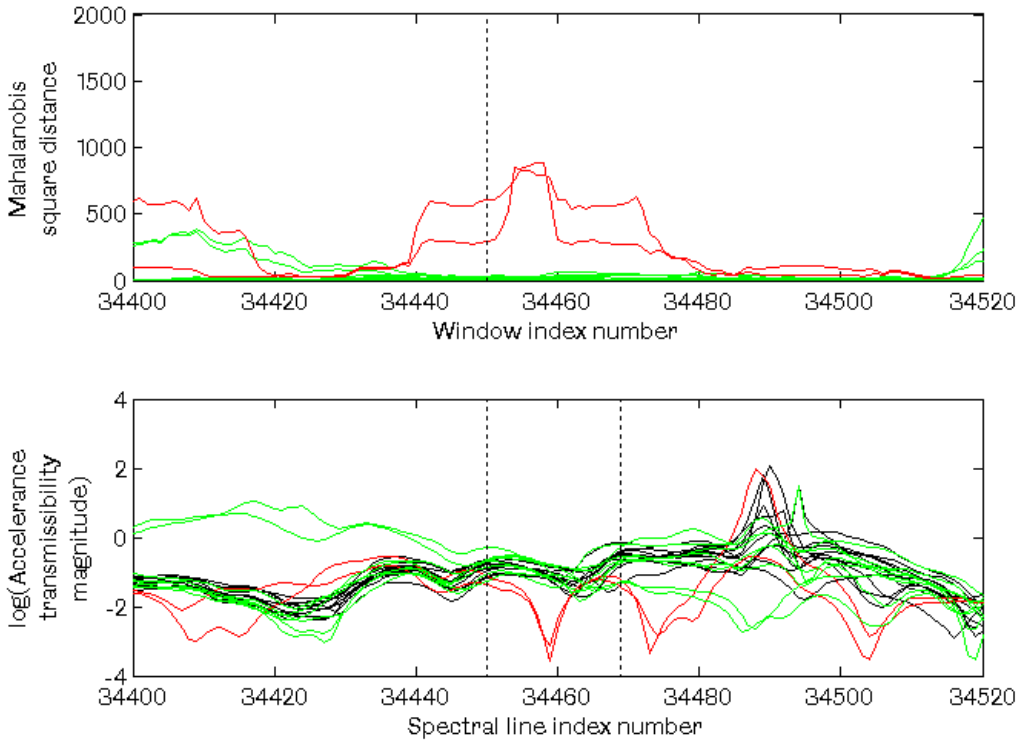

Figure 6-8 Feature selection plots for Window Index Number 34450

As an example, consider the feature identified in Figure 6-7. The original feature window contained two peaks. In the interests of promoting robustness in the feature set the window illustrated in Figure 6-8 was preferred to that shown in Figure 6-7 Feature selection plots for Window Index Number 34456 Despite a lower discordancy value than that for the initially identified feature, plus a lower level of consistency between the two observed repetitions, the feature shown in Figure 6-8 was adopted due to the presence of only one characteristic in the defined window.

6.3.4 Feature selection results

Through the process of assisted manual selection detailed above, the feature set was reduced from a candidate set of 53014 features to a final set of 50 features. The results of the feature selection exercise using single-site training data are given in Table 6-2.

For completeness, the discordancy values for each averaged repetition (named R1 and R2) are included. Discordancy values were a primary consideration for feature selection, and the values returned in these columns allow the reader some insight into the relative sensitivity of the features to damage on the allotted panel, and to the consistency in values between repetitions.

Panel Label	Path Label	Feature Label	Window Index No	Frequency Range [Hz]		Discordancy Value	
				min	max	R1	R2
P1	T1	F1	361	180.0	189.5	1984.7	1951.5
	T1	F2	567	283.0	292.5	1065.5	1098.2
	T1	F3	673	336.0	345.5	2472.4	1106.8
	T1	F4	774	386.5	396.0	1739.8	4076.8
	T1	F5	1018	508.5	518.0	2556.6	1570.4
	T1	F6	2199	1099.0	1108.5	850.6	331.4
	T2	F7	4365	133.5	143.0	1191.1	1202.0
	T2	F8	4548	225.0	234.5	640.7	756.2
	T2	F9	5133	517.5	527.0	2954.1	1539.1
	T2	F10	5384	643.0	652.5	520.0	590.9
P2	T3	F11	8399	102.0	111.5	37061.4	46374.4
	T3	F12	9519	159.5	169.0	9200.0	13665.8
	T3	F13	8514	662.0	671.5	3985.7	4589.1
	T4	F14	13282	495.0	504.5	7862.0	2111.5
	T4	F15	14118	913.0	922.5	12697.1	6185.8
	T4	F16	14545	1126.5	1136.0	7217.6	16148.7
	T4	F17	15065	1386.5	1396.0	6618.7	4428.5
	T5	F18	16928	269.5	279.0	13707.2	20261.0
	T5	F19	19163	494.5	504.0	2733.9	1876.9
	T5	F20	17378	1387.0	1396.5	26096.6	7578.5
P3	T6	F21	20662	88.0	97.5	1696.9	2131.5
	T6	F22	20903	208.5	218.0	600.5	897.6
	T6	F23	22377	945.5	955.0	2255.0	1785.5
	T6	F24	22487	1000.5	1010.0	2908.9	1290.0
	T7	F25	25018	217.5	227.0	2866.9	2199.5
	T7	F26	26464	940.5	950.0	1829.5	1586.1
	T8	F27	29385	352.5	362.0	2287.7	2553.0
	T8	F28	30011	665.5	675.0	847.5	482.3
	T8	F29	30960	1140.0	1149.5	1422.8	2402.9
	T8	F30	31255	1287.5	1297.0	3091.4	1107.0
P4	T9	F31	33228	225.5	235.0	997.2	900.5
	T9	F32	33397	310.0	319.5	2457.0	2926.0
	T9	F33	34450	836.5	846.0	291.9	603.7
	T10	F34	38390	758.0	767.5	709.9	502.2
	T10	F35	38881	1003.5	1013.0	591.5	614.4
	T10	F36	39173	1149.5	1159.0	916.1	448.3
	T11	F37	41350	189.5	199.0	1468.2	969.4
	T11	F38	41833	431.0	440.5	1079.0	1868.8
	T11	F39	43152	1090.5	1100.0	1615.3	2566.3
	T11	F40	43472	1250.5	1260.0	915.1	1736.2
P5	T12	F41	45787	359.5	369.0	1151.4	1041.6
	T12	F42	47288	513.0	522.5	174.3	228.8
	T12	F43	47328	1110.0	1119.5	715.7	570.4
	T12	F44	46094	1130.0	1139.5	322.4	601.6
	T13	F45	49512	1701.5	1711.0	298.3	109.3
	T13	F46	49824	173.5	183.0	565.7	377.4
	T13	F47	49903	308.0	317.5	306.7	465.8
	T13	F48	51428	329.5	339.0	384.0	350.7
	T13	F49	49781	369.0	378.5	562.8	1539.2
	T13	F50	48471	1131.5	1141.0	444.7	831.8

Table 6-2 Features identified using Dataset A

6.4 Feature evaluation: single-site and multi-site damage

In this section, the discriminatory performance of features selected using Dataset A are evaluated when these are applied to Dataset B. The objectives are two-fold. The first objective is to verify that the selected features are, individually, capable of discriminating between *normal* and *single-site* panel removals for a previously unobserved dataset. It is expected that the features should be capable of making this distinction with a level of accuracy approaching that achieved for the training set.

The second objective is to evaluate the performance of the features in discriminating between the previously unobserved *normal* and *multi-site* panel removal data in Dataset B. The hypothesis of the present chapter is that a good degree of discrimination might be achieved in the multi-site case, despite no multi-site data being used for feature selection or eventual training. In the present chapter the performance of the individual features is evaluated, with the hope that this performance may approach that achieved for single-site cases. In Chapter 7 the requirement will be that the degree of discrimination of the combined feature set should be sufficient to permit the training of a family of classifiers that can robustly separate normal and damage states.

Evaluation is performed in incremental steps. Initially, a visual assessment of the transmissibility spectra in the vicinity of the selected features is made. This allows comment to be made on the general 'behaviour' of the spectra when single- and multi-site damage is introduced into the structure. Secondly, the discordancy values associated with each feature are evaluated for all 4600 observations contained in Dataset B. Finally, a summary statistic is introduced to allow the discriminative capability of each feature to be concisely stated, both for single-site and multi-site panel removal observations.

6.4.1 Visual evaluation of selected features

An initial assessment of the performance of the features selected in Table 6-2, when applied to the test set, was made through a qualitative examination of the spectra and through the calculation of discordancy values for every observation in the test set. First, spectral plots were produced to allow the behaviour of the transmissibility spectra in the region of the selected features to be assessed. In order to reduce the

number of spectra to be visualised, averaging was applied in a manner consistent with that previously applied for feature selection. For each of the 62 runs (31 normal condition, five single-panel removals and 26 multiple-panel removals) named in the test sequence in Table 5-2, one mean-averaged spectra was calculated and included.

Secondly, the feature values returned for each of the 4600 observations in the testing set were quantitatively evaluated. The discordancy value D_ξ for each feature and for each observation is given by,

$$D_\xi = \left(\boldsymbol{x}_\xi - \overline{\boldsymbol{x}}_A\right)^{\mathrm{T}} [\boldsymbol{\Sigma}_A]^{-1} \left(\boldsymbol{x}_\xi - \overline{\boldsymbol{x}}_A\right) \tag{6.4}$$

where $\boldsymbol{x}_\xi$ is the vector of newly-presented spectral data from Dataset B, $\overline{\boldsymbol{x}}_A$ is the mean vector of the normal condition data from Dataset A and $[\boldsymbol{\Sigma}_A]$ is the covariance matrix of the normal condition data from Dataset A. Note that while this requires a large number of discordancy evaluations to be made (4600 observations x 50 features in this case, requiring a total of 230000 evaluations), comparative computational efficiency is maintained as the most computationally costly term - the inversion of the covariance matrix - need only be performed once per feature.

The transmissibility and discordancy outcomes for two of the 50 features are presented in Figure 6-9 and Figure 6-10. For each feature, four plots are given. The first plot is the discordancy value data used for selecting features from Dataset A. Removal of the panel of interest is shown in red and the removal of the other four panels are shown in green. The second plot shows the portion of the transmissibility spectra corresponding to the chosen feature for each of the runs in Dataset A. Single-site removal of the panel of interest is denoted in red; single-site removal of the other panels is shown in green; and normal condition data is given in black.

The third plot shows transmissibility spectra from Dataset B for the chosen feature window. The colour scheme used is as for the second plot. Solid red lines denote that only the panel of interest is removed (i.e. *single-site* states) and dashed lines denote that the panel is one of several removed (i.e. *multi-site* states). Spectra in green correspond to damage states for which the panel of interest is not removed, with solid lines again corresponding to single-site states and dashed lines corresponding to multi-site states.

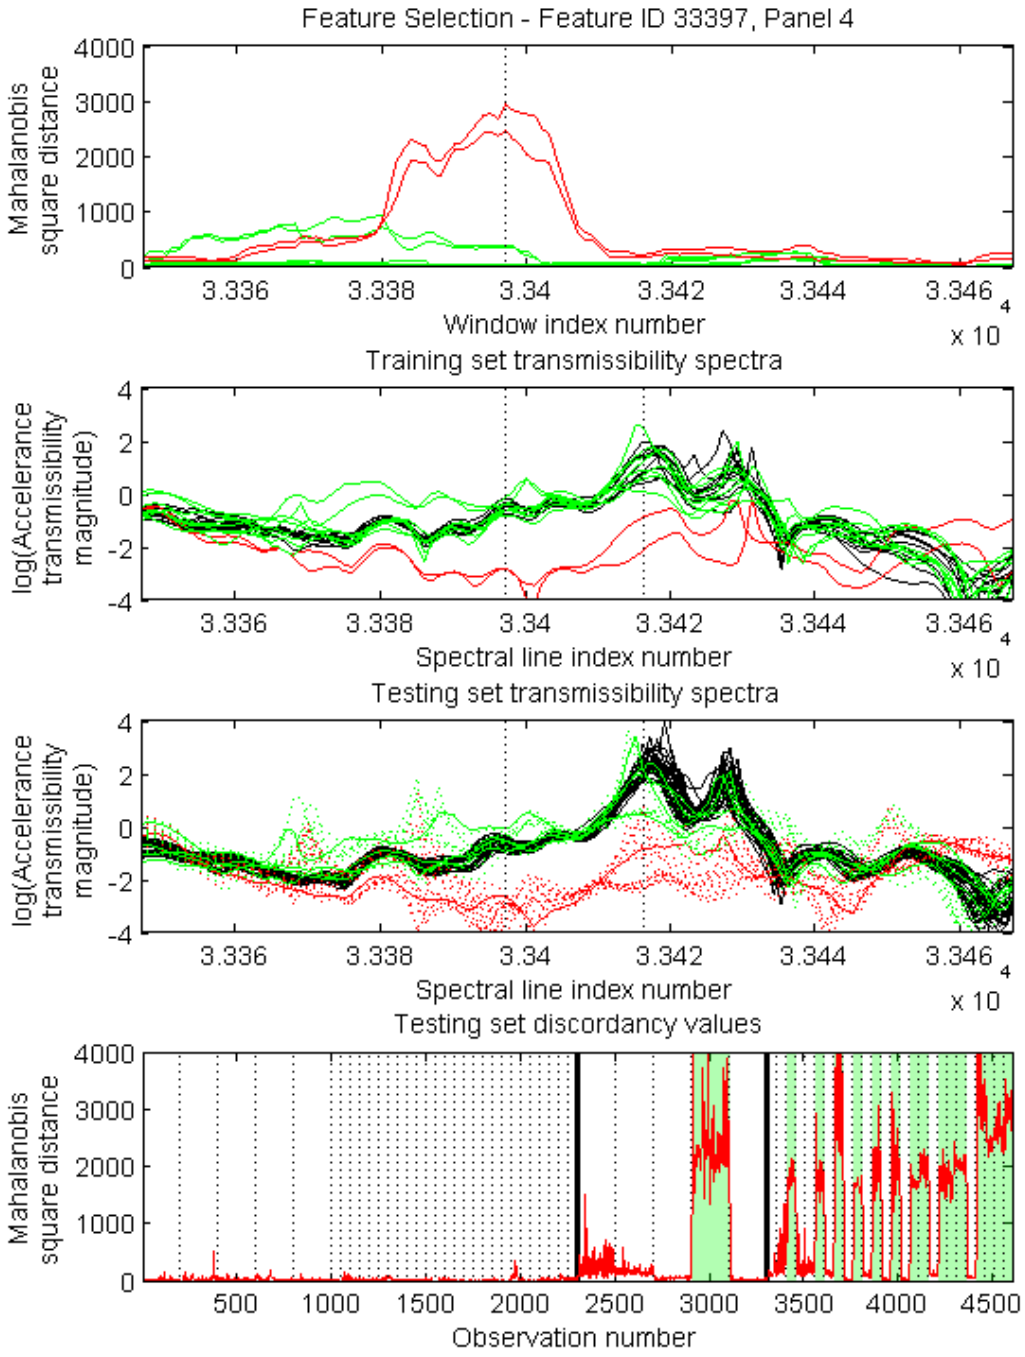

Figure 6-9 Feature evaluation plots for feature F32 (Window Index Number 33397)

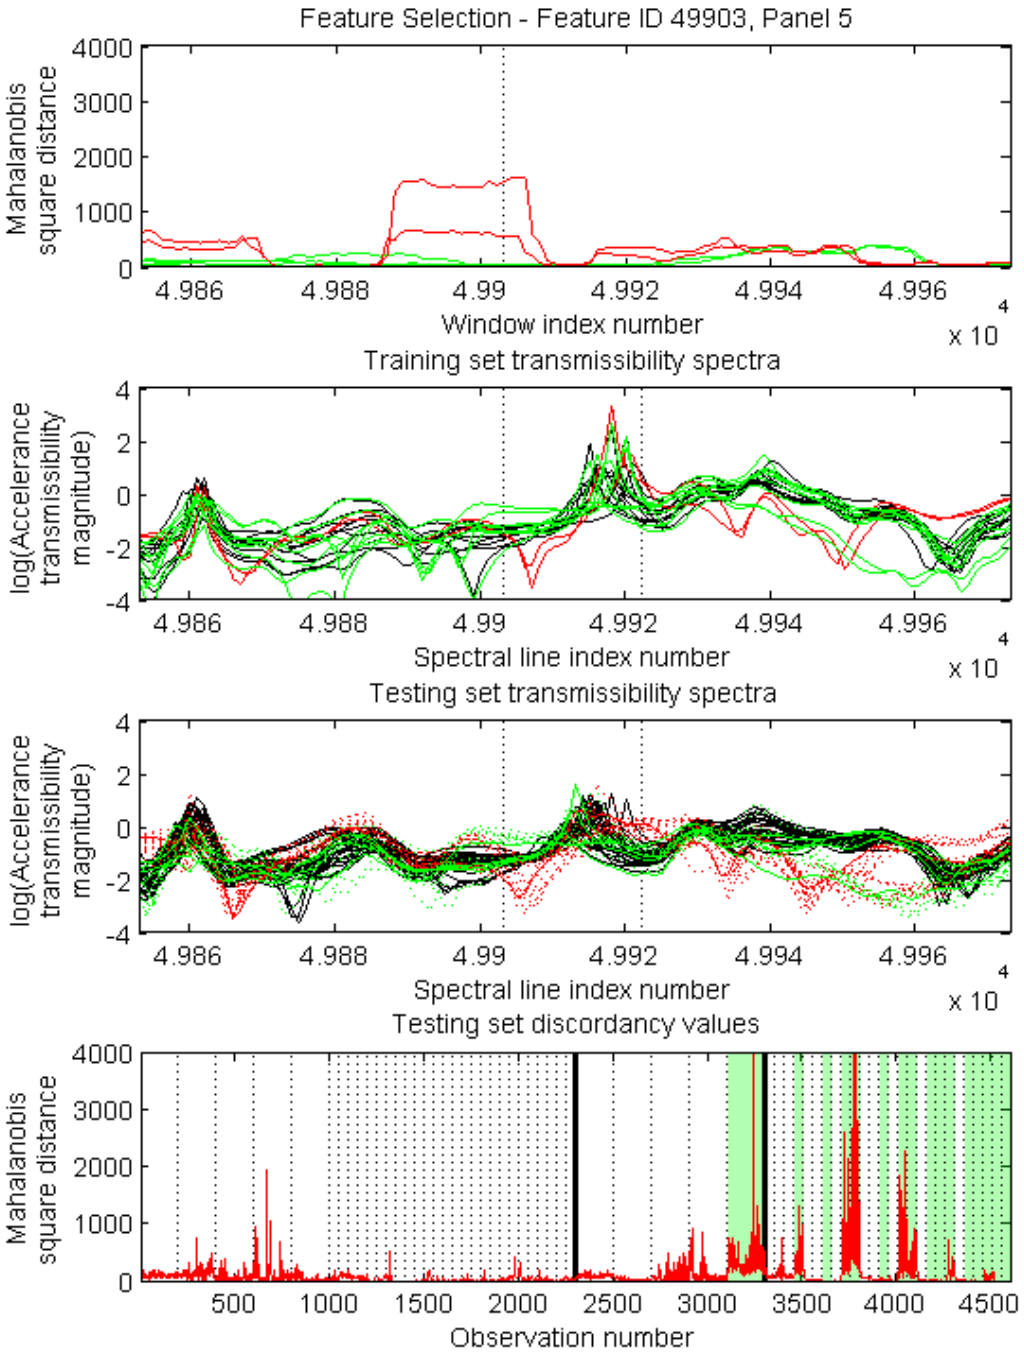

Figure 6-10 Feature evaluation plots for feature F47 (Window Index Number 49903)

The final plot gives discordancy values for the feature, evaluated for all observations of Dataset B. The plot is divided according the state of the structure using dashed vertical lines. The states corresponding to each group of observations are given in Table 6-3: in summary, observations 1-2300 are for the normal condition, 2301-3300 correspond to single-panel removal and 3301-4600 correspond to multiple-panel removal. These groups are further delineated using solid vertical lines. Damage states that include removal of the specified panel are highlighted in green.

Observation No	State	Observation No	State
1-200	N1	2301-2500	DO1
201-400	N2	2501-2700	DO2
401-600	N3	2701-2900	DO3
601-800	N4	2901-3100	DO4
801-1000	N5	3101-3300	DO5
1001-1050	N6	3301-3350	DO1_2
1051-1100	N7	3351-3400	DO1_3
1101-1150	N8	3401-3450	DO1_4
1151-1200	N9	3451-3500	DO1_5
1201-1250	N10	3501-3550	DO2_3
1251-1300	N11	3551-3600	DO2_4
1301-1350	N12	3601-3650	DO2_5
1351-1400	N13	3651-3700	DO3_4
1401-1450	N14	3701-3750	DO3_5
1451-1500	N15	3751-3800	DO4_5
1501-1550	N16	3801-3850	DO1_2_3
1551-1600	N17	3851-3900	DO1_2_4
1601-1650	N18	3901-3950	DO1_2_5
1651-1700	N19	3951-4000	DO1_3_4
1701-1750	N20	4001-4050	DO1_3_5
1751-1800	N21	4051-4100	DO1_4_5
1801-1850	N22	4101-4150	DO2_3_4
1851-1900	N23	4151-4200	DO2_3_5
1901-1950	N24	4201-4250	DO2_4_5
1951-2000	N25	4251-4300	DO3_4_5
2001-2050	N26	4301-4350	DO1_2_3_4
2051-2100	N27	4351-4400	DO1_2_3_5
2101-2150	N28	4401-4450	DO1_2_4_5
2151-2200	N29	4451-4500	DO1_3_4_5
2201-2250	N30	4501-4550	DO2_3_4_5
2251-2300	N31	4551-4600	DO1_2_3_4_5

Table 6-3 Observation numbers corresponding to structural states

Figure 6-9 illustrates feature F32, which performed well in identifying both single-site and multiple-site removals involving panel P3. The second plot illustrates the clear distinction between the removal of panel P3 (in red) and the other states included in Dataset A, and this distinction lead to the selection of the feature. The third plot illustrates the same feature window for the transmissibility spectra of

Dataset B. It is observed that the spectra behave in a very similar fashion to that found for Dataset A, both for the removal of panel P3 alone (shown as a solid line) and where panel P3 was one of several removed (shown as dashed lines). The clear distinction between the removal of panel P3 and other states is maintained. This is reflected in the discordancy values illustrated in the fourth plot. The feature 'fires' strongly for the observations of single panel removal, and similarly strongly for each of the states in which it was one of several panels removed (highlighted with a green background). This very encouraging level of performance was observed for the vast majority of the 50 identified features, and supports the stated hypothesis.

In Figure 6-10 a feature that performed less well is illustrated. While the featured 'fired' to a reasonable degree for the single-site removals, it was unable to distinguish all of the multiple site damage removals. It was in general found to be more challenging to identify features that performed well for panel P5. However, the reason for feature F47 performing less well is likely to be attributable to the nature of the spectra captured within the identified feature window. The spectra comprise two 'peaks', only one of which demonstrates a marked shift as a result of damage. This was apparent to some degree from the observations of Dataset A in the second plot, but is more strongly evident in the third plot for the observations of Dataset B.

The visual assessment of the selected feature windows suggests that there is similarity in the behaviour of spectra due to single-site and multi-site panel removal over some portions of the frequency ranges, and that this behaviour is reflected in the returned feature values. However, the visual assessment of a large number of features is a somewhat onerous task. A concise method for presenting the quantitative evaluation results for the selected features is instead introduced. It should be made clear that the multiple-site damage data was not used for feature selection.

6.4.2 Quantitative evaluation of selected features

The evaluation process generates a large amount of graphical data, and a summary statistic was sought with which to evaluate the discriminatory effectiveness of the selected features. Receiver Operating Characteristic (ROC) curves offer such a measure. The ROC curve is introduced below, and the area under the ROC curve (AUC) is adopted as a summary statistic for quantifying the discriminative ability of the selected features.

Receiver Operating Characteristic (ROC) curves allow a simple benefit (true-positive rate) vs. cost (false-positive rate) analysis for two-class data. A brief introduction to ROC curves and their application to the feature evaluation problem are given here. A more complete introduction to ROC analysis is given in [82].

The two-class classification problem consists of correctly assigning a class label to a newly-presented observation. If the class labels are 'Positive' and 'Negative', the four possible outcomes are *True Positive*, *True Negative*, *False Positive* and *False Negative*. The criteria leading to the four outcomes are summarised in Figure 6-11.

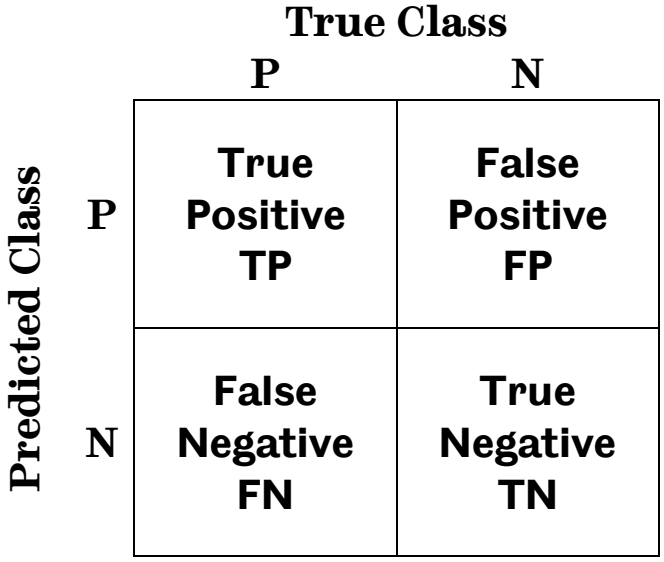

Figure 6-11 Two-class confusion matrix

Separation of the data into two classes necessarily involves the setting of a threshold, with observations returning values above the threshold being assigned to one class (e.g. Positive) and observations returning values below the threshold being assigned to the other (e.g. Negative). The success of the resulting classifier in accurately separating the data may be evaluated using two operating characteristics: the *true positive rate* (TPR) and the *false positive rate* (FPR). It is desirable for the classifier to return a high rate of true positives and a low rate of false positives.

The true positive rate is given by,

$$TPR = \frac{TP}{TP + FN} \tag{6.4}$$

where TP is the number of *true positives* (observations that come from the positive set and are correctly classified as positive) and FN is the number of *false negatives* (observations that come from the positive set but are incorrectly classified as negative). The true positive rate is alternatively referred to as the *sensitivity*, and gives the rate at which observations from the positive class are correctly classified as positive.

Similarly, the false positive rate is given by,

$$FPR = \frac{FP}{FP + TN} \tag{6.5}$$

where FP is the number of *false positives* (observations that come from the negative set but are incorrectly classified as positive) and TN is the number of *true negatives* (observations that come from the negative set and are correctly classified as negative). The false positive may be thought of as the residual of the *specificity*, the rate at which observations from the negative class are correctly classified as negative.

Both characteristics take values in the range [0 1]. A perfect classifier would return a true positive rate of 1 and a false positive rate of 0.

The Receiver Operating Characteristic gives a comparison between the TPR and FPR operating characteristics for a given threshold level. Evaluating the ROC value for a range of possible threshold values enables the plotting of an ROC curve. An example of such a curve is given in Figure 6-12. The ROC curve may be employed to:

1. Evaluate the discriminative ability of the data.
2. Select a level for the threshold that is appropriate to the problem.

It is for the first of these functions that it is applied here. As an example, the curve in Figure 6-12 offers an insight into the discriminatory performance of the feature F46, intended to detect the removal of panel P5. A ‘Negative’ class label implies the panel has not been removed and a ‘Positive’ class label implies that it has. The data used in is from the multiple-site damage case. The dashed line corresponds to a classifier offering no discrimination. The further the curve corresponding to a classifier lies above this line, the greater its discriminatory performance. In this case, the discriminatory performance of this feature for the data presented may qualitatively

be described as 'good'. A quantitative evaluation of the level of discrimination comes from considering the area under the ROC curve.

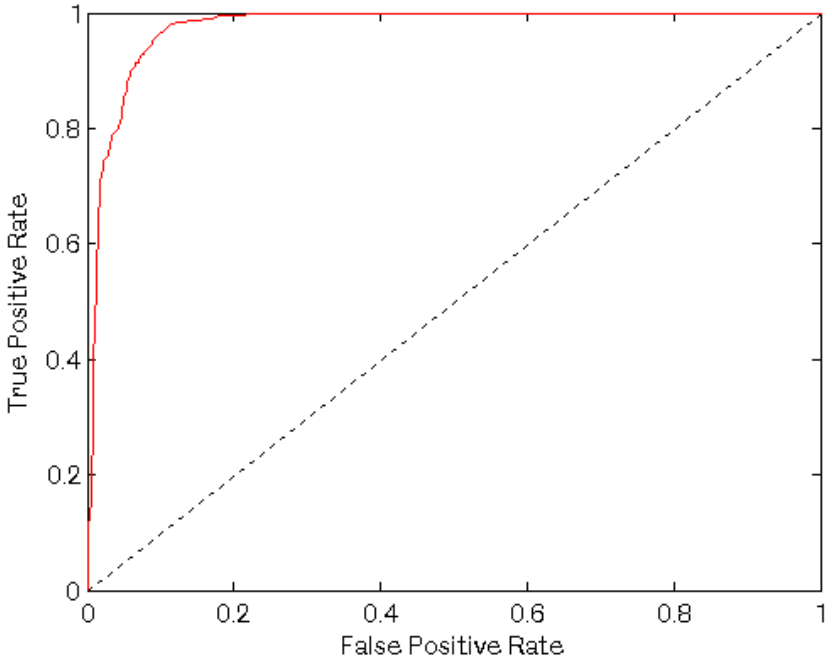

Figure 6-12 Example of an ROC curve for feature F46

The area under the ROC curve (AUC) is a simple means of further summarising classifier performance. As the name suggests the AUC is given by the area lying below the plotted Receiver Operating Characteristic curve (see Figure 6-13).

The AUC characteristic takes a value in the range [0 1], where a perfect classifier would return an AUC value of 1. The AUC has a useful interpretation as the probability that the classifier will rank a randomly selected positive instance higher than a randomly chosen negative instance [82]. This AUC is in fact the same quantity that is estimated when calculating the Wilcoxon statistic [83]. For the example of feature F46 given in Figure 6-13, the AUC value when discriminating between panel-on and panel-off data for Dataset B was 0.976. This is taken to represent an excellent level of discrimination between classes.

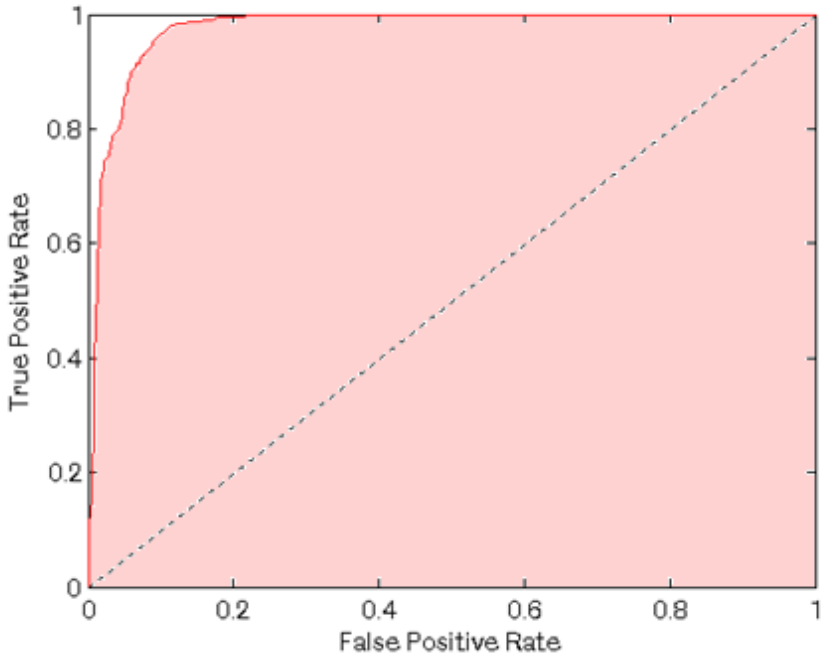

Figure 6-13 Area under the ROC curve for feature F46

For the data considered, the two classes are 'specified panel removed' and 'specified panel not removed'. The AUC is adopted as a summary statistic for comparing the classification performance of the individual features in separating normal condition data first from single-site data and then from multi-site data.

For each of the 50 features identified using Dataset A, two ROC curves were calculated. The first curve illustrates the ability of the feature to separate normal data from single-site damage data for Dataset B. The second illustrates the ability of the feature to separate normal data from multi-site damage data for Dataset B. The area under each of the ROC curves was calculated to give an AUC value. The desired outcome is that the individual features would display a good degree of separability, evidenced by an AUC value of close to 1, for both cases.

6.4.3 Feature evaluation results

The AUC values calculated for the discrimination of single-site and multi-site damage for each of the five panels are presented in Table 6-4 to Table 6-8. The outcomes are summarised in Table 6-9.

Feature ID	Area Under Curve		
	Single-Site	Multi-Site	Difference
F1	0.994	0.989	-0.005
F2	1.000	1.000	0.000
F3	0.896	0.922	0.026
F4	1.000	1.000	0.000
F5	0.913	0.901	-0.011
F6	0.902	0.893	-0.009
F7	1.000	1.000	0.000
F8	0.999	0.998	-0.001
F9	1.000	1.000	0.000
F10	0.998	0.998	0.000
Mean	0.970	0.970	

Table 6-4 Area under ROC curve for Panel P1 features, Dataset B

Feature ID	Area Under Curve		
	Single-Site	Multi-Site	Difference
F11	1.000	1.000	0.000
F12	1.000	1.000	0.000
F13	1.000	1.000	0.000
F14	1.000	1.000	0.000
F15	1.000	1.000	0.000
F16	0.994	1.000	0.006
F17	1.000	0.999	0.000
F18	1.000	1.000	0.000
F19	0.999	1.000	0.000
F20	0.993	1.000	0.007
Mean	0.999	1.000	

Table 6-5 Area under ROC curve for Panel P2 features, Dataset B

Feature ID	Area Under Curve		
	Single-Site	Multi-Site	Difference
F21	0.998	0.999	0.001
F22	0.986	0.984	-0.002
F23	0.996	0.997	0.002
F24	1.000	0.983	-0.017
F25	1.000	1.000	0.000
F26	0.999	0.998	-0.001
F27	1.000	1.000	0.000
F28	0.947	0.983	0.035
F29	0.985	0.991	0.006
F30	1.000	0.999	-0.001
Mean	0.991	0.993	

Table 6-6 Area under ROC curve for Panel P3 features, Dataset B

Feature ID	Area Under Curve		
	Single-Site	Multi-Site	Difference
F31	1.000	1.000	0.000
F32	1.000	1.000	0.000
F33	0.979	0.988	0.010
F34	0.991	0.996	0.005
F35	0.981	1.000	0.019
F36	0.990	0.999	0.009
F37	0.999	0.999	0.000
F38	0.994	0.997	0.003
F39	1.000	1.000	0.000
F40	0.996	0.995	-0.002
Mean	0.993	0.997	

Table 6-7 Area under ROC curve for Panel P4 features, Dataset B

Feature ID	Area Under Curve		
	Single-Site	Multi-Site	Difference
F41	0.999	0.998	-0.001
F42	0.846	0.760	-0.086
F43	0.966	0.983	0.017
F44	0.687	0.794	0.107
F45	0.973	0.919	-0.054
F46	0.972	0.986	0.014
F47	0.991	0.777	-0.214
F48	0.975	0.976	0.001
F49	0.975	0.537	-0.438
F50	0.946	0.941	-0.004
Mean	0.933	0.867	

Table 6-8 Area under ROC curve for Panel P5 features, Dataset B

Panel	Mean Area Under Curve	
	Single-Site	Multi-Site
P1	0.970	0.970
P2	0.999	1.000
P3	0.991	0.993
P4	0.993	0.997
P5	0.933	0.867
Overall	0.977	0.966

Table 6-9 Summary of areas under ROC curves, Dataset B

6.4.4 Analysis of feature evaluation results

The outcome of applying the features selected using the training dataset to the testing dataset may be summarised thus:

- A mean AUC value of 0.977 across all features was achieved for the single-site panel removal.
- A mean AUC value of 0.966 across all features was achieved for the multi-site panel removal.
- Perfect classification was achieved by 18 of the individual features for single-site panel removal.
- Perfect classification was achieved by 19 of the individual features for multi-site panel removal.

Some have worked exceptionally well, others less so. Some of the features 'fired' unexpectedly when presented with normal condition observations. A small degree of inter-test variability between Datasets A and B is observed through comparison of normal condition spectra for these features. This finding serves to reiterate the importance of gathering a training set that is truly representative of the conditions that may be encountered.

It is also found that the removal of some of the panels is distinctly more easily detected than the removal of others, and that while some transmissibilities provide a rich source of features, others offer little discrimination between states. It is notable that transmissibility paths comprising response sensors that are separated from the damage location by a stringer appear to be less successful than those for which the reference sensor is 'close' to the damage. This leads to the suggestion that it may be beneficial to position sensors to be proximal to damage. Of greater practical relevance is the further suggestion that for stiffened-panel structures such as the aircraft wing considered, it may be beneficial to place at least one sensor within each 'panel' ('panel' in this sense being a region of the wing top-sheet bounded on all sides by stiffening elements) that is to be monitored for damage. While these suggestions are perhaps somewhat unsurprising in hindsight, this is nevertheless a useful demonstration of the importance of sensor placement, and gives some insight into how an extended network of sensors for detecting top sheet damage in similar structures may be formed. This suggestion will be returned to in the conclusions of Chapter 7.

A final point relates to the use of Mahalanobis squared-distance measures for feature reduction. The MSD approach offers a measure of discordancy: that is, a measure of to what extent a sample differs from the population mean regardless of the manner in which it differs. This is an attractive property of features for novelty detection, where the only training data are representations of 'normality' and the aim of the classifier is to detect deviations from this normality on the basis of feature values. Novelty detection is thus a two-class problem: either the structure has strayed from normality, which may be indicative of damage; or it has not. It does not matter where the structure has moved to in the feature space: all that matters is that it is no longer in the region of feature space deemed 'normal'. For multi-class problems, the new location of the structure in the feature space becomes more important. The classifier seeks to separate classes, hopefully in a robust fashion, on the basis of feature values. As such, it is desirable that the feature set allow the robust separation of classes in feature space.

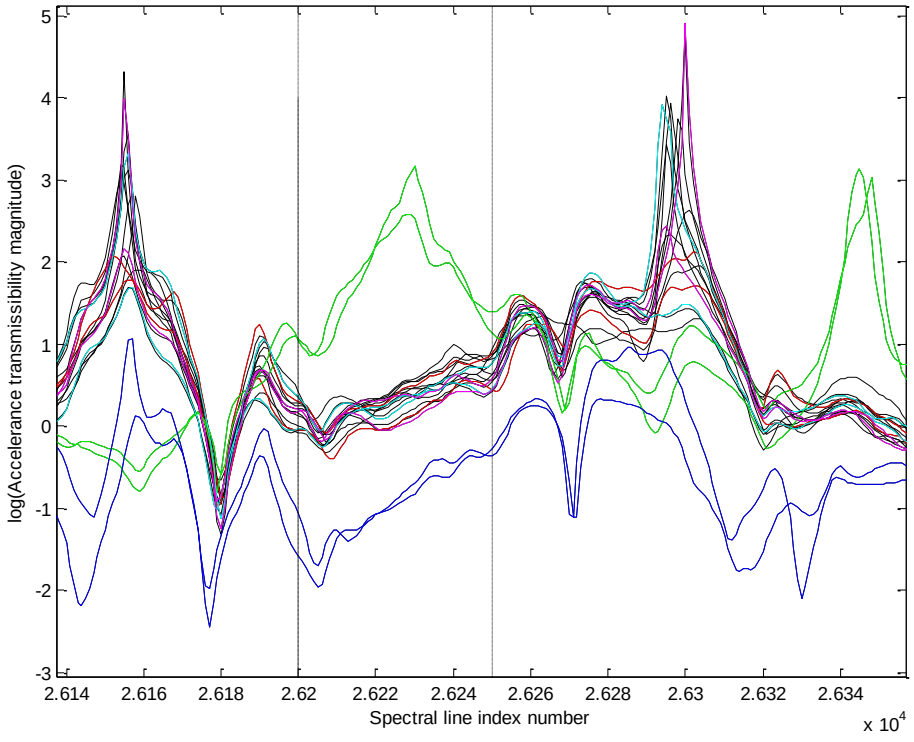

Figure 6-14 Section of transmissibility spectra T7

As an example of a drawback of the MSD approach, consider the example in Figure 6-14. Here, a short section of transmissibility spectra T7 is presented. The normal condition data is coded in black, and single-site removal of panels P1, P2, P3, P4 and

P5 in red, green, blue, cyan and magenta respectively. From a visual inspection it is apparent that this portion of the spectrum is sensitive to the removals of panel P2 (in blue) and panel P3 (in green). It would further appear relatively straightforward to construct some feature capable of separating the data into three classes: no panels removed; P2 removed; and P3 removed. However, a typical application of the MSD approach would be to select a window from the spectrum and calculate discordancy values for the three classes. In this case, the discordancy values returned for the removals of panels P2 and P3 are very similar, despite the clear difference in the nature of the respective spectra. A useful piece of information has thus been lost in reducing the feature dimension.

This serves as a reminder that the dimensionality reduction offered by the MSD approach does not come for free, and it may be useful in the context of multi-class identification to seek a measure which retains information on the 'direction' in which the data moves as structural changes occur.

6.5 Discussion

Overall, it was found that the individual features selected on the basis of the training Dataset A (comprising normal and single-site panel removal data) performed very well when presented with single-site panel removal data from a previously unobserved dataset. This result, while expected, serves to validate the test sequencing introduced in Chapter 5.

Of primary interest, however, is the finding that the features also performed very well when presented with multi-site panel removal data, despite no multi-site data being used for feature selection. The multi-site performance of the classifiers averaged over the 50 selected features is very close to that when presented with single-site data, and is evidenced both through visual inspection and through applying quantitative measures. The conclusion drawn is that the study supports the hypothesis of the chapter: for the structure states investigated, it has been possible to find features that offer a high degree of discrimination between multiple-site damage states despite only the single-site damage states having been observed. The significance of this result is the suggestion that the problem of the explosion in the number of damage states that must be observed in order to build a classifier capable of identifying multiple-site damage may be circumvented in some circumstances. It is, of course, not possible to draw general conclusions on the basis

of one case study, and further investigation into the validity of the approach is warranted.

It should be remembered that while evaluating features, the primary concern is that they should 'fire' when damage occurs at the specified location and not fire when the structure is undamaged. It is of only secondary concern that the feature should not fire when damage is not present at the specified location, but is present at other locations on the structure. If such features are found, it is intuitively easier to interpret the patterns of firing features. However, the application of statistical pattern recognition methods is specifically intended to reveal underlying patterns within the observed data that would at the very least be challenging for a human observer to interpret

Having identified features that are individually capable of discriminating between damage states to a good degree, focus moves to the application of statistical pattern recognition in order to employ the features for damage classification. The expectation is that combining the features in an appropriate fashion should offer a further improvement in discriminatory performance. In the following chapter, the features identified form the basis of statistical classification using support vector machines.

7 A Data-Driven Approach: Support Vector Classification

7.1 Introduction

The preceding two chapters have been concerned with different stages of the statistical pattern recognition (SPR) paradigm. In Chapter 5, the experimental structure and test programme were described, including the interrogation of previously-acquired data when specifying the test schedule. Chapter 6 focused on feature selection. The final stage of the SPR paradigm, addressed in this chapter, is the statistical modelling of the selected features. Multiple-class support vector machines (SVMs) are investigated for this purpose.

The application of SVM methods has received relatively little attention in the damage identification literature in comparison to that received by other pattern recognition approaches, notably neural network and nearest neighbour formulations. SVMs have, however, been demonstrated to possess several properties that suggest they may be well-suited to the damage identification task. They have been shown to be competitive with other methods when applied to real engineering datasets [26], and to generalise well from the small datasets usually encountered in damage identification problems. The SVM is also able to support different classes of discriminant function (for example linear, polynomial or radial basis functions) without requiring substantial modification of the basic learning algorithm. Finally, the SVM may be considered to be a universal approximator. This capability means that an SVM can represent any function to arbitrary accuracy, provided no limit is placed upon the number of free weights used. The advantage for the damage identification task is that the full 'toolbox' of discriminant functions is in principle available when training a damage classifier.

7.2 Objectives

The first hypothesis addressed in this chapter is that support vector machines may be an effective option for the classification of multi-site damage data, where observations of each class are available for training. The second hypothesis, in common with the preceding two chapters, is that a classifier may be trained using observations of normal and single-site damage data only; and that when applied to a testing set that includes multi-site damage observations, this classifier may offer a level of classification accuracy approaching that of a classifier trained using observations of all damage states. The motivation for the second hypothesis is that the number of structural states for which training data is required may be greatly reduced. Two objectives are undertaken with the aim of addressing these hypotheses.

The first objective of the current chapter is to train an SVM-based classifier using normal, single-site and multi-site damage data; and to evaluate its performance on a testing set gathered during the same experiment also containing observations of normality, single-site damage and multi-site damage. In this manner the efficacy of the SVM approach may be demonstrated.

The second objective is to train an SVM-based classifier using normal and single-site damage state data from a separate test, and to evaluate its performance when applied to the same multi-site testing set. The task faced by this SVM classifier is challenging: not only is the classifier being asked to identify damage states on which it has not been trained, but it is expected to generalise between tests.

The structure of the chapter is as follows. The theoretical background to the SVM algorithm is given in Sections 7.3 and 7.4. In section 7.5 the approach used for the multi-class damage identification task is described. In Sections 7.6 and 7.7 the performances of the classifier systems are compared, findings discussed, and conclusions drawn

7.3 Statistical learning theory

Support vector machines are a product of the field of statistical learning theory (SLT) [84]. Prior to discussion of the SVM, the principles underpinning SLT are first briefly described. The theoretical basis provided in this section is primarily drawn

from [26]. The theory is presented in much greater detail in numerous textbooks [85-87], with excellent tutorials provided in both [88] and [89]. The intention here is not to provide a comprehensive theoretical background for methods of pattern recognition, but rather to provide an introduction to SVM theory at a depth sufficient to illustrate the motivation and application of the method.

7.3.1 Empirical risk minimisation

The aim of statistical learning theory is to establish a relationship between two sets of data through the application of a learning algorithm. The algorithm seeks to establish a mapping between a d-dimensional input space, denoted $\boldsymbol{x}$, and a one-dimensional output space, denoted y, on the basis of a set of training data $(\boldsymbol{x}_k, y_k)$ where $k = 1, \dots, N$ with N being the number of observations in the training set. The relationship is modelled on the basis of the training set as,

$$y = f(\boldsymbol{x}, \boldsymbol{w}) \tag{7.1}$$

where $\boldsymbol{w}$ is the vector of free parameters in the model. The task for the learning algorithm is to set the values of $\boldsymbol{w}$ so as to maximise the fit of the model on the basis of a set of training data, and for this model to generalise well to new data drawn from the same distribution. The nature of the model is defined by the nature of the output space: for regression problems, y will is a continuous variable; for classification problems, y is a class label. It should be noted that the definition of the output space as one-dimensional does not result in a loss of generality: problems with multivariate outputs can often be reduced to a set of independent, univariate problems.

One measure of model fit is in terms of the mean error in the function, termed the *actual risk*:

$$R(\boldsymbol{w}) = \int \boldsymbol{dxdy}\, |y - f(\boldsymbol{x}, \boldsymbol{w})| p(\boldsymbol{x}, y) \tag{7.2}$$

The hope is that the model that is ultimately selected should minimise this risk. Unfortunately, as the joint density between the inputs and outputs $p(\boldsymbol{x}, y)$ is generally unknown, the value of the actual risk cannot be computed. An approximation to the actual risk, termed the *empirical risk*, may instead be calculated:

$$R_{emp}(\boldsymbol{w}) = \sum_{i=1}^{N} |y_i - f(\boldsymbol{x}_i, \boldsymbol{w})| \tag{7.3}$$

Empirical risk minimisation (ERM) offers a potential means of approximating the best fitting solution, provided that,

$$\lim_{N\to\infty} R_{emp}(\boldsymbol{w}) = R(\boldsymbol{w}) \tag{7.4}$$

which is satisfied through the theory of large numbers in an intuitive fashion, and that,

$$\lim_{N\to\infty} \min_{\boldsymbol{w}\in\mathcal{H}_n} \left(R_{emp}(\boldsymbol{w})\right) = \min_{\boldsymbol{w}\in\mathcal{H}_n} \left(R(\boldsymbol{w})\right) \tag{7.5}$$

within a space of hypothesised models $\mathcal{H}_n$. The criterion that the minima of the actual and empirical risk should converge in the limit will be satisfied only for suitably restricted regions of $\mathcal{H}_n$. In addition to these considerations, the scalar value R_{emp} offers no information about whether a model selected in this manner will generalise well to other datasets drawn from the same joint distribution $p(\boldsymbol{x}, y)$.

A basic principle of Statistical Learning Theory is that a bound on the actual risk that holds with a probability $1 - \delta$ may be defined. This bound is of the form,

$$R(\boldsymbol{w}) \leq R_{emp}(\boldsymbol{w}) + \Phi(h, \delta, N) \tag{7.6}$$

where the confidence interval Φ is given by,

$$\Phi(h, \delta, N) = \sqrt{\frac{h \ln\left(\frac{2N}{h} + 1\right) - \ln\left(\frac{\delta}{4}\right)}{N}} \tag{7.7}$$

In addition to being a function of δ and the number of training points N, Φ is also dependent on h, the *Vapnik-Chervonenkis* (VC) *dimension*. Minimisation of the empirical risk plus this bound forms the basis of *Structural Risk Minimisation*. These two linked concepts are introduced below.

7.3.2 VC dimension

The VC dimension offers a measure of the capacity of a statistical classification algorithm $f(\boldsymbol{x}, y)$, and scales with the complexity of the model basis used for fitting f. It is defined as being the number of elements of the largest set of points that a

given algorithm can completely 'shatter' – that is, separate perfectly in all configurations. For illustration, consider the example of three data points in two-dimensions in Figure 7-1. If each data point can be assigned a class label of either +1 or -1, there are 2^3=8 possible configurations of the points. If linear indicator functions are adopted as the classification algorithm f it is possible to shatter the data points in all 8 configurations.

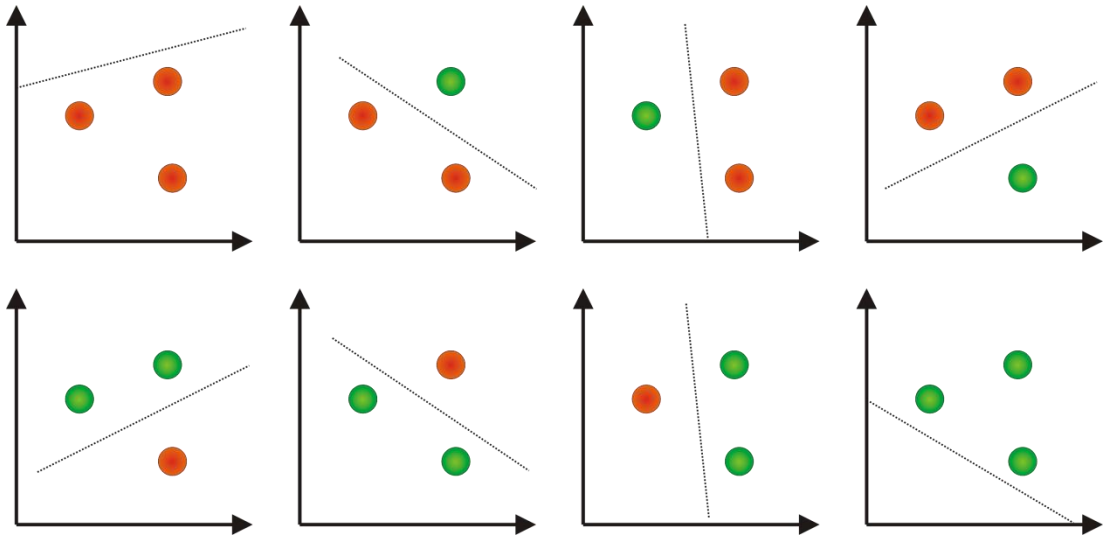

Figure 7-1 Shattering of the 2^n combinations of n=3 data points, using a linear classification algorithm

However, the same is not true for four data points. Figure 7-2 gives an illustration of one configuration in which a single linear indicator function is incapable of shattering the data points. A more complex algorithm, such as that shown in the Figure 7-2, would be capable of appropriately shattering the data.

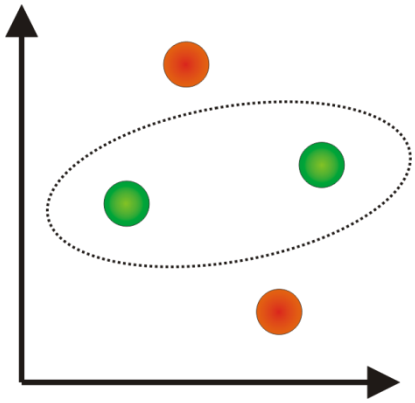

Figure 7-2 Shattering of n=3 data points by a more complex algorithm. Shattering using a linear classification algorithm is not possible

The VC dimension of the linear indicator function is thus three. The result extends to a VC dimension equal to n+1 for any non-collinear n-dimensional space. In general, the VC dimension h_i of a given algorithm is usually somewhat difficult to compute, although an estimate of the upper bound on h_i will suffice instead for the purposes of structural risk minimisation.

7.3.3 Structural risk minimisation

Structural risk minimisation (SRM) consists of solving the minimisation problem for the empirical risk plus the confidence bound,

$$min\left(R_{emp}(\boldsymbol{w}) + \Phi(h, \delta, N)\right) \tag{7.8}$$

As the size of the confidence bound Φ scales with the VC dimension h as shown in equation 7.7, minimisation of the bound penalises model complexity.

The concept is to generate a group of models (alternatively termed 'structures') with different bases. Each model has a VC dimension h_i and each spans a hypothesis space S_i such that,

$$h_i \leq h_{i+1} \tag{7.9}$$

leading to,

$$S_i \subset S_{i+1} \tag{7.10}$$

A graphical illustration of a hierarchy of models is given in fig Figure 7-3. The value of h_i is computed (or, where appropriate, estimated) for each model, and the model that minimises the sum of the empirical risk and the confidence bound (equation 7.8) is selected. An interesting discussion of the comparative ability of SRM and ERM formulations to generalise is given in [88].

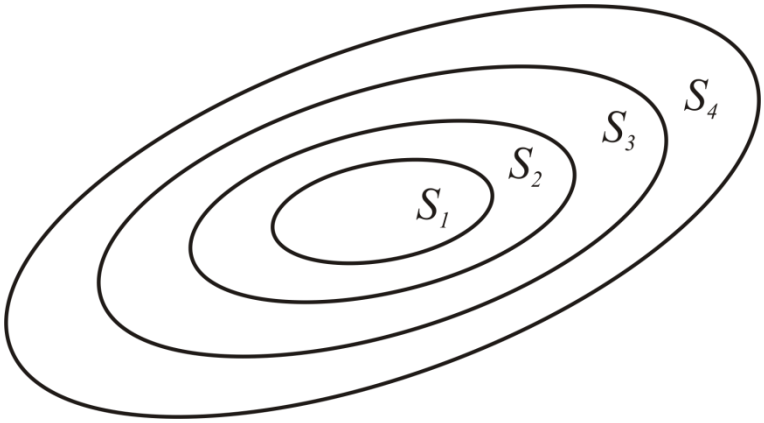

Figure 7-3 Hierarchy of models, parameterised by hypothesis space

7.4 Classification via support vector machines

Consider the two class dataset presented in Figure 7-4. The goal of the classification task is to induce a function on the basis of this training dataset that will generalise well to the separation of as yet unseen data. This function takes the form of a hyperplane, with data lying on one side of the hyperplane assigned to one class, and data on the other side to another. As shown in Figure 7-4, even if linear classifiers alone are considered there may be many hyperplanes that separate the two classes. In order to select a hyperplane that generalises well, one might intuitively select a hyperplane that lays 'somewhere in the middle', as far as possible from each class.

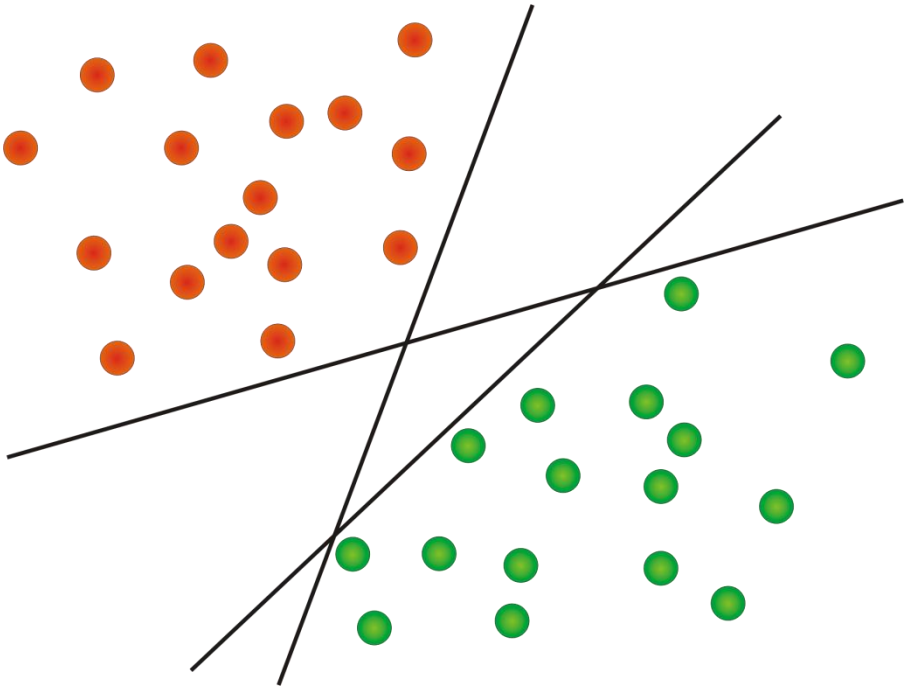

Figure 7-4 Example two-class data, with arbitrary separating hyperplanes

Support vector classification offers a formal means of achieving this intuitively correct result. The principle underpinning the approach is to minimise the structural risk by selecting the hyperplane that maximises the margin between itself and the data in each class.

7.4.1 Linearly separable case

A separating hyperplane is defined,

$$D(\boldsymbol{x}) = \langle \boldsymbol{w}, \boldsymbol{x} \rangle = 0 \tag{7.11}$$

Note that in defining (7.11), a constant term is required in order for the offset of the hyperplane from the origin to be set. Here, this is achieved by including by a 0^{th} component, held at unity, in the input vector $\boldsymbol{x}$. The offset may thus be specified as the 0^{th} component of the vector of free weights $\boldsymbol{w}$. This represents something of a simplification of the issue of constant term - further discussion may be found in [89]. Each data point $\boldsymbol{x}_k$ in the training set is assigned to a class on the basis of the of the separating condition given by,

$$\begin{aligned} \boldsymbol{x}_k \in C_1 &\Longrightarrow D(\boldsymbol{x}_k) = \langle \boldsymbol{w}, \boldsymbol{x}_k \rangle \geq 1 \\ \boldsymbol{x}_k \in C_2 &\Longrightarrow D(\boldsymbol{x}_k) = \langle \boldsymbol{w}, \boldsymbol{x}_k \rangle \leq -1 \end{aligned} \tag{7.12}$$

where C_1 and C_2 are classes, with the respective class labels (1) and (-1). This may be expressed concisely as,

$$D(\boldsymbol{x}_k) = y_k \langle \boldsymbol{w}, \boldsymbol{x}_k \rangle \geq 1 \tag{7.13}$$

where y_k is the class label. The Euclidean distance of each data point from the separating hyperplane is given by,

$$\frac{D(\boldsymbol{x}_k)}{\|\boldsymbol{w}\|} \tag{7.14}$$

An interval that contains the separating hyperplane $D(\boldsymbol{x})$ but does not contain any of the data points $\boldsymbol{x}_k$ is defined as a *margin*. τ is a margin if,

$$y_k \frac{D(\boldsymbol{x}_k)}{\|\boldsymbol{w}\|} \geq \tau \tag{7.15}$$

is satisfied for all k. The optimal margin for the hypothetical dataset is illustrated in Figure 7-5. Note that a margin will exist only if the data are separable - the non-separable case is discussed below.

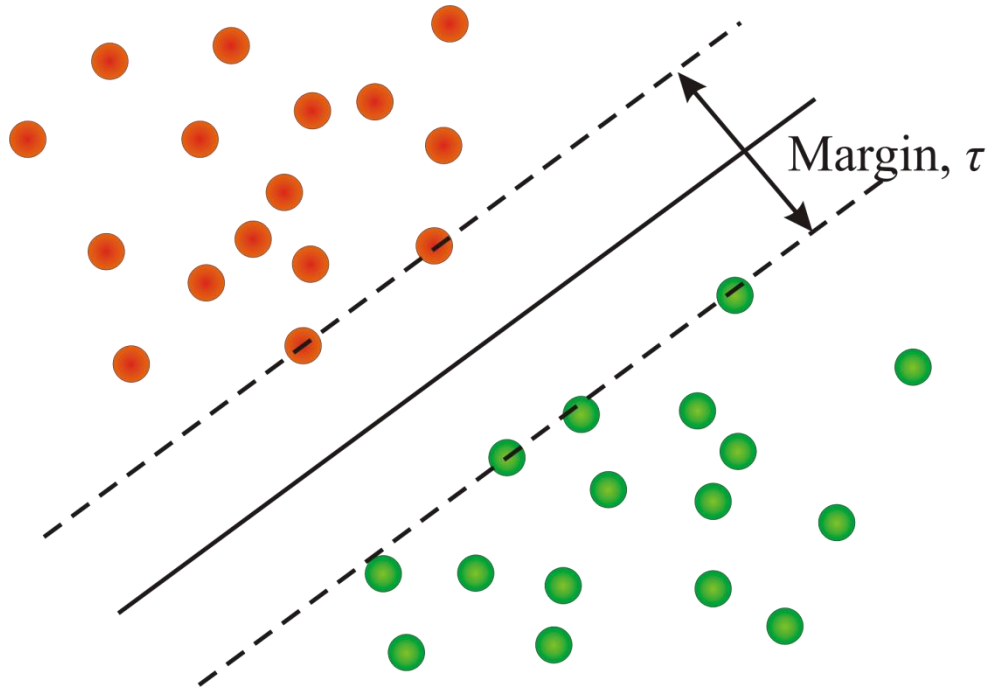

Figure 7-5 Optimal separating hyperplane

Note also that the hyperplane has until this point been arbitrarily parameterised. The scaling of the parameters $\boldsymbol{w}$ may usefully be fixed by,

$$\|\boldsymbol{w}\|\tau = 1 \tag{7.16}$$

With this parameterisation the definition of the margin (7.15) becomes equivalent to the definition of the separating condition (7.13).

The objective of the Support Vector Machine is to minimise the structural risk through maximisation of the separating margin, τ. It is apparent from equation 7.16 that this is equivalent to minimisation of the parameter norm $\|\boldsymbol{w}\|$. The minimisation is constrained by the separating condition (7.13). A convex quadratic programming problem arises from defining an appropriate objective function,

$$Q(\boldsymbol{w}) = \frac{1}{2}\|\boldsymbol{w}\|^2 - \sum_{i=1}^{N} \alpha_i y_i (\langle \boldsymbol{w}, \boldsymbol{x}_i \rangle - 1) \tag{7.17}$$

with Lagrange multipliers α_i. The dual formulation of the problem is formed through use of the *Kuhn-Tucker (KT) conditions*, with the model parameters $\boldsymbol{w}$ expressed n terms of the multipliers α_i. The dual formulation asserts that at the optimum,

$$\frac{\partial Q}{\partial \boldsymbol{w}} = 0 \Longrightarrow \boldsymbol{w} = \sum_{i=1}^{N} \alpha_i y_i \boldsymbol{x}_i \tag{7.18}$$

and that a multiplier α_i may only be non-zero if the constraint,

$$\alpha_i y_i (\langle \boldsymbol{w}, \boldsymbol{x}_i \rangle - 1) = 0 \tag{7.19}$$

is satisfied. It is found that the data points $\boldsymbol{x}_i$ that satisfy this constraint lie at the limit of the margin (see Figure 7-6). These data points are the *support vectors*. An observation from Figure 7-6 is that in the linearly separable case the data points acting as support vectors may comprise only a small subset of the training data. The hyperplane calculated by the SVM would be exactly the same if training were repeated with all non-support vectors removed.

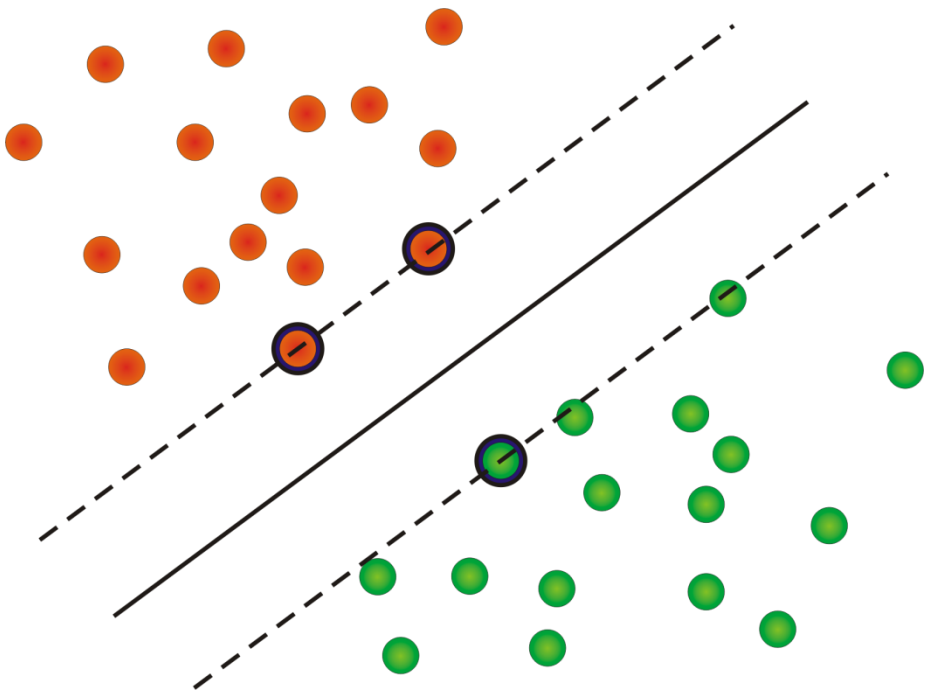

Figure 7-6 Optimal separating hyperplane, with support vectors highlighted

Construction of the dual formulation of the minimisation problem is achieved by substituting the relationship given in equation 7.18 into the objective function specified in equation 7.17. The resulting objective function is,

$$Q(\boldsymbol{w}) = \sum_{i=1}^{N} \alpha_i - \frac{1}{2} \sum_{i=1}^{N} \sum_{j=1}^{N} \alpha_i \alpha_j y_i y_j \langle \boldsymbol{x}_i, \boldsymbol{x}_j \rangle \tag{7.20}$$

and is to be maximised with respect to the multipliers α_i, subject to constraints,

$$\begin{gathered} \sum_{i=1}^{N} \alpha_i y_i = 0 \\ \alpha_i \geq 0, i = 1, \ldots, N \end{gathered} \tag{7.21}$$

Quadratic programming may be employed for the optimisation task. With an optimal vector of multipliers thus identified, the optimal separating hyperplane for the linearly separable case may be expressed through substitution of (7.18) into (7.11) as,

$$D(\boldsymbol{x}) = \sum_{i=1}^{N} \alpha_i y_i \langle \boldsymbol{x}_i, \boldsymbol{x} \rangle \tag{7.22}$$

up to a constant offset (or *bias*) term. For convenience, the bias term is omitted here.

7.4.2 Linearly non-separable case

As previously intimated, some modification of the method is required for the general case that the two classes are not linearly separable. Two options present themselves. First, a more complex form for the hyperplane may be assumed: extension of the approach to include nonlinear boundaries will be briefly returned to at the end of the present section. The second option is the inclusion of non-negative *slack variables* ξ_i such that,

$$D(\boldsymbol{x}_k) = y_k \langle \boldsymbol{w}, \boldsymbol{x}_k \rangle \geq 1 - \xi_k \tag{7.23}$$

The slack variables ξ_i are a measure of the accepted rate of misclassification error. The objective of the optimisation is now to minimise the rate of misclassification in addition to minimising the VC dimension of the classifier. This is reflected in the

objective function, which is modified to include an additional penalty term accounting for the slack variables ξ_i,

$$Q(\boldsymbol{w}) = \frac{1}{2}\|\boldsymbol{w}\|^2 + C\sum_{i=1}^{N}\xi_i - \sum_{i=1}^{N}\alpha_i y_i(\langle \boldsymbol{w}, \boldsymbol{x}_i\rangle - 1 + \xi_i) - \sum_{i=1}^{N}\mu_i\xi_i \tag{7.24}$$

The dual formulation becomes,

$$Q(\boldsymbol{w}) = \sum_{i=1}^{N}\alpha_i - \frac{1}{2}\sum_{i=1}^{N}\sum_{j=1}^{N}\alpha_i\alpha_j y_i y_j\langle \boldsymbol{x}_i, \boldsymbol{x}_j\rangle \tag{7.25}$$

with constraints,

$$\sum_{i=1}^{N}\alpha_i y_i = 0$$
$$C \geq \alpha_i \geq 0, i = 1, \ldots, N \tag{7.26}$$

The dual formulation, and thus the expression for the hyperplane returned as upon its solution, is identical to that for the linearly separable case with the exception that the bounds of the Lagrange multipliers have been modified. This modification leads to a *soft boundary* classifier, as opposed to the *hard boundary* classifier introduced for the linearly-separable case. C serves as a form of regularisation parameter and trades off the margin size against the training error rate. C should best be chosen based upon knowledge of the noise in the data. It is consistent with the aims of structural risk minimisation to seek a classifier that is robust to the choice of C over a range of values.

Until this point only linear hyperplanes have been considered. Nonlinear hyperplanes may also be employed, and in this chapter use will be made of *radial basis function* (RBF) kernels. The theoretical basis of nonlinear hyperplanes is, however, not included here. Further details of this and other aspects of single=class SVMs can be found in [85-87].

7.4.3 Multi-class support vector machines

Support vector machines are designed, without loss of generality, for the resolution of data into two classes. Extension to multi-class problems is a topic of ongoing research, with approaches typically falling into two categories- those which seek to consider all classes at once using a 'one-against-all' classifier, and those that seek to construct a system of binary 'one-against-one' classifiers [90]. The strategy employed in this chapter is to train an ensemble of 'one-against-one' SVMs. This decision is informed in part by the computational cost of the approaches.

The computational cost involved in the application of support vector classification is largely attributable to the optimisation of solutions based upon a matrix H. H is a $(N \times N)$ square matrix, N being the number of observations employed for training- the iterative Matlab QP solver is used for performing the optimisation step in the current work. At each iteration step, multiplication and inversion of the matrix is required in order to compute the search direction. The computational complexity of this operation is of the order $O(N^3)$. Computational cost considerations thus preclude the use of very large numbers of observations for training SVMs. As the number of observations required for training a 'one-against-all' classifier will typically be higher than that for a 'one-against-one' classifier in order to adequately represent the dataset, 'one-against-one' classifiers may be preferred on the grounds of computational cost. The structure of the classifier system is described in greater detail in the following section.

7.5 Approach

The hypotheses to be tested are that:

1) SVMs may be an effective option for the classification of multi-site damage data, where observations of each class are available for training, and
2) that a classifier trained using observations of normal and single-site damage data may offer a level of classification accuracy approaching that of a classifier trained using observations of all damage states. when applied to a testing set that includes multi-site damage observations

To address these hypotheses two classifiers are constructed and tested. Classifier 1 is trained using data from the structure in all its damage states, including observations of multi-site damage. Classifier 2 is trained using data from the

structure only in its normal and single-site damaged states. The objective for both classifiers is to achieve a high rate of correct classification when presented with a *testing set* of previously unseen data, which contains observations from the structure in single-site, multiple-site and undamaged states.

The same specification is used for both classifiers, and they differ only in the data used for training and validation. The separation of data into training, validation and testing sets is considered below, preceded by specification of the classifier structure. The methodology employed for setting hyperparameter values (the *validation step*) is also given.

7.5.1 Classifier specification

The classifiers were created using the MATLAB Support Vector Machine Toolbox [91]. The acquisition of the data used for training, validating and testing the classifiers has been detailed in Chapter 5 and is recalled below. The 50 features employed are log discordancy values derived from transmissibility data. The selection of the feature set is described in Chapter 6 and the chosen features are detailed in Table 6-2. The classification architecture is based upon the dichotomous Support Vector Machine (SVM) extended to the multi-class problem.

Each 'classifier' in fact comprises an ensemble of 6 binary SVM classifiers. The first SVM (labelled SVM1) seeks to separate damage-state data from normal-state data. Each observation is classified as either 'undamaged' (coded as +1) or 'damaged' (coded as -1). SVM1 thus acts as a damage detection step.

Five further dichotomous SVMs (labelled SVM2-6) seek to indicate whether removal has occurred for each of the five panels in turn. Each SVM seeks to class an individual location as 'undamaged' or 'damaged'- SVM2 seeks to classify panel 1 as on or off, SVM3 relates to panel 2 etc. The undamaged (panel on) and damaged (panel off) classes are again coded +1 and -1 respectively. The classes used are summarised in Table 7-1.

SVM	Features employed	Positive (+1) Class	Negative (-1) Class
SVM1	F1-F50	All Undamaged Data	All Damaged Data
SVM2	F1-F10	Damaged, P1 not removed	Damaged, P1 removed
SVM3	F11-F20	Damaged, P2 not removed	Damaged, P2 removed
SVM4	F21-F30	Damaged, P3 not removed	Damaged, P3 removed
SVM5	F31-F40	Damaged, P4 not removed	Damaged, P4 removed
SVM6	F41-F50	Damaged, P5 not removed	Damaged, P5 removed

Table 7-1 Classifier specification: classes and features

The features used in training, validating and testing each SVM are also given in Table 7-1. During initial development it was found that the performance of the individual SVMs 2-6 was substantially diminished if normal condition data was included in the training set. As such, training of these classifiers was conducted using a training set comprising damaged state data only. The positive (+1) class thus comprises all damage states that did not involve the removal of the specified panel. The negative (-1) class comprises all those damage states that did involve the removal of the specified panel, either in isolation or in combination with other panel removals.

It was also found that greatly improved results were achieved for SVMs 2-6 when only the 10 features selected for each panel were used for classification, instead of making use of all 50 features. SVM1 is unique among the 6 SVMs employed in being trained using all 50 features identified in Chapter 6.

Radial basis kernels were employed in the discriminant function. Given the ability of the SVM to act as a universal approximator, it is in principle possible to incorporate the selection of the discriminant function into the validation process, with the possibility of developing an improved classifier being countered by the increased computational cost associated with the validation phase. This is left as an extension to the present work.

The data presented to the classifier are log discordancy values, as detailed in Chapter 6. Normalisation of the data is recommended to aid the conditioning of the optimisation problem [81]. In this instance, each feature was normalised to the interval [0 1], with 1 being the maximum value of the feature observed in dataset A. A further option, not pursued here, is normalisation to unit variance through division by the feature variance.

7.5.2 From classifier outputs to decisions

The output of the family of classifiers is a six element coded sequence, reflecting the predictions of each SVM. For example:

- [1, 1, 1, 1, 1, 1] indicates that the structure is undamaged
- [-1, -1, 1, 1, -1, 1] indicates that damage has occurred to the structure, and that the damage is identified as occurring at locations 1 and 4.

Note that contradictory sequences are possible. For example:

- [1, -1, -1, 1, 1, 1] indicates that damage has not occurred to the structure, but damage is nevertheless identified at two of the five locations.
- [-1, 1, 1, 1, 1, 1] indicates that damage has occurred to the structure, but damage is not identified at any of the five locations monitored.

A decision must be made as to the interpretation of such sequences. In this chapter, the interpretation of the contradictory sequences above is as follows:

- [1, -1, -1, 1, 1, 1] - precedence is given to the prediction of the detection classifier SVM1. The structure is classified as undamaged.
- [-1, 1, 1, 1, 1, 1] - such instances are labelled 'Unclassified'

The full set of decisions based upon the SVM outputs is given in Table 7-2.

SVM Prediction						Decision
SVM1	SVM2	SVM3	SVM4	SVM5	SVM6	
1			Any			No panels removed
-1	-1	1	1	1	1	Panel 1 removed
-1	1	-1	1	1	1	Panel 2 removed
-1	1	1	-1	1	1	Panel 3 removed
-1	1	1	1	-1	1	Panel 4 removed
-1	1	1	1	1	-1	Panel 5 removed
-1	-1	-1	1	1	1	Panels 1 & 2 removed
-1	-1	1	-1	1	1	Panels 1 & 3 removed
-1	-1	1	1	-1	1	Panels 1 & 4 removed
-1	-1	1	1	1	-1	Panels 1 & 5 removed
-1	1	-1	-1	1	1	Panels 2 & 3 removed
-1	1	-1	1	-1	1	Panels 2 & 4 removed
-1	1	-1	1	1	-1	Panels 2 & 5 removed
-1	1	1	-1	-1	1	Panels 3 & 4 removed
-1	1	1	-1	1	-1	Panels 3 & 5 removed
-1	1	1	1	-1	-1	Panels 4 & 5 removed
-1	-1	-1	-1	1	1	Panels 1,2 & 3 removed
-1	-1	-1	1	-1	1	Panels 1,2 & 4 removed
-1	-1	-1	1	1	-1	Panels 1,2 & 5 removed
-1	-1	1	-1	-1	1	Panels 1,3 & 4 removed
-1	-1	1	-1	1	-1	Panels 1,3 & 5 removed
-1	-1	1	1	-1	-1	Panels 1,4 & 5 removed
-1	1	-1	-1	-1	1	Panels 2,3 & 4 removed
-1	1	-1	-1	1	-1	Panels 2,3 & 5 removed
-1	1	-1	1	-1	-1	Panels 2,4 & 5 removed
-1	1	1	-1	-1	-1	Panels 3,4 & 5 removed
-1	-1	-1	-1	-1	1	Panels 1,2,3 & 4 removed
-1	-1	-1	-1	1	-1	Panels 1,2,3 & 5 removed
-1	-1	-1	1	-1	-1	Panels 1,2,4 & 5 removed
-1	-1	1	-1	-1	-1	Panels 1,3,4 & 5 removed
-1	1	-1	-1	-1	-1	Panels 2,3,4 & 5 removed
-1	-1	-1	-1	-1	-1	Panels 1,2,3,4 & 5 removed
-1	1	1	1	1	1	Unclassified

Table 7-2 Decision rules

Classifier 1 and Classifier 2 differ only in the data that is used in their development: all other factors (classifier structure, validation procedure, features employed etc) are kept the same. With the structure of the classifiers specified, attention is turned to the data used to train, validate and test the classifiers.

7.5.3 Training, validation and testing sets

Two experimental datasets are used for developing and testing the classifiers. The first, referred to as dataset A, comprises 1000 normal state and 1000 single-site damage state observations of the wing structure, gathered during a test conducted on 16th-22nd July 2009. The second dataset, referred to as dataset B, comprised 2300

normal, 1000 single-site damage and 1300 multi-site damage state observations, from a test conducted on 29th June-14th July 2009. The datasets are illustrated pictorially in Figure 7-7. The brief description of the datasets given here is intended to serve as a reminder of their content, with the data acquisition described in detail in Chapter 5.

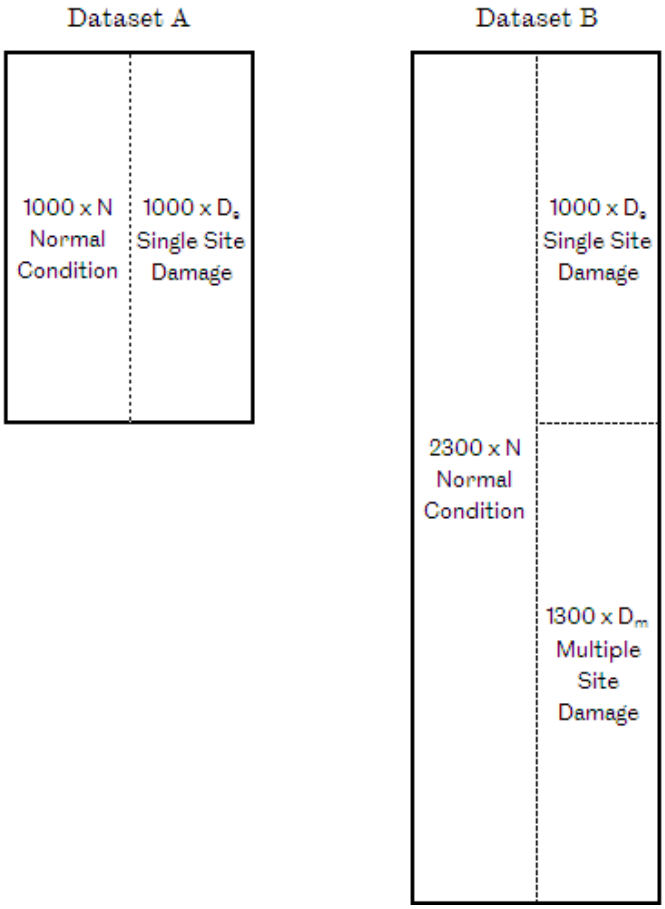

Figure 7-7 Datasets A and B

Conducting supervised learning in a principled fashion necessitates separating the data into three, non-overlapping sets: the *testing set*, *validation set* and *training set*. Each serves a purpose in the development and testing of the classifier.

- The *training* set is used to set the values of the classifier parameters
- The *validation* set is used to set the values of the classifier hyperparameters (discussed in Section 7.5.6)
- The *testing* set is used to verify that the developed classifier works for an independent set of observations.

Datasets A and B are each separated into training, validation and testing sets, denoted A_{TRN}, $A_{VAL}$and A_{TST} for Dataset A and B_{TRN}, $B_{VAL}$and B_{TST} for Dataset B. The distribution of observations in each set is summarised in Figure 7-8 Separation of the datasets into training, validation and testing sets, where N refers to normal state observations; D_S refers to single-site damage observations; and D_M refers to multi-site damage observations. A consideration in separating the data was that the training and validation sets drawn from Datasets A and B be the same size so as not to favour either Classifier 1 (trained using Dataset B) or Classifier 2 (trained using Dataset A). Recall that Dataset A alone was used for feature selection.

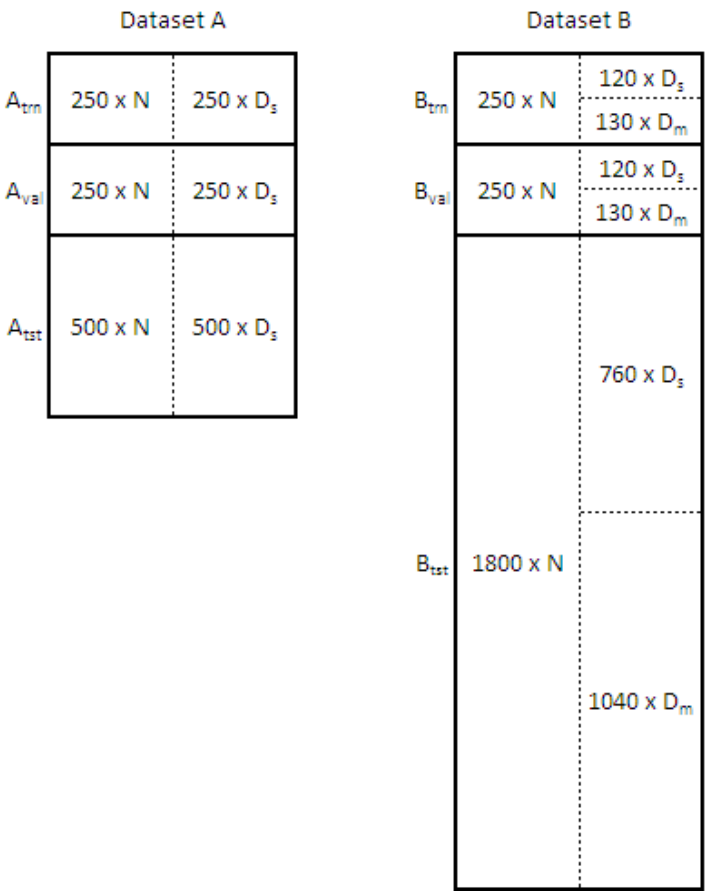

Figure 7-8 Separation of the datasets into training, validation and testing sets

The data used for each classifier in order to meet the stated objectives are given below.

7.5.4 Classifier 1

The first hypothesis of this chapter is that SVMs may be an effective option for the classification of multi-site damage data, where observations of each class are available for training. This hypothesis is addressed using Classifier 1.

Classifier 1 is trained using the training and validation sets B_{TRN} and B_{VAL}. These sets include observations of the structure in its normal, single-site damaged and multiple-site damaged states. The ability of the developed classifier to correctly identify damage is evaluated using the testing set B_{TST}.

The validity of the SVM classification approach for the presented multiple damage identification problem may thus be assessed. The performance of Classifier 1 is also taken as being a 'best case' benchmark, given the features selected and the modelling decisions made. It is against this benchmark that the performance of Classifier 2 may be measured.

7.5.5 Classifier 2

The second hypothesis of this chapter is that a classifier trained using observations of normal and single-site damage data may offer a level of classification accuracy approaching that of a classifier trained using observations of all damage states when applied to a testing set that includes multi-site damage observations. This hypothesis is addressed using Classifier 2.

Classifier 2 is developed using the training and validation sets A_{TRN} and A_{VAL}, which contain observations of normal and single-site damage only. The classifier is again evaluated using testing set B_{TST}, which contains observations of all damage states. The performance of a classifier trained in the absence of observations of multi-site damage data when presented with observations from all damage states may thus be assessed. The performance of this classifier is compared to that of the 'best case' performance given by Classifier 1. As an intermediate step, the classifier is tested using the testing set A_{TST} to check that the classifier can successfully locate single site damage.

In addition to the restriction to data from a greatly reduced number of damage states this data was drawn from a separate sequence of tests, the classifier is trained using data from one sequence of tests, and tested using data acquired from a second sequence of tests second dataset conducted several days earlier. This is a somewhat severe test of the capacity of developed classifier to generalise to newly presented data that contains not only unseen states, but which may be subject to some degree of inter-test variability. It is common in applications of statistical

pattern recognition to real data for a single dataset to be used, and for interleaved data to be used for training, validation and testing.

7.5.6 Validation step

For the classifiers employed in this chapter, two hyperparameters must be set:

- C, the misclassification tolerance parameter, and;
- α, the radial basis kernel width.

The validation step allows an informed decision as to the most appropriate values of the hyperparameters to be made in order to meet a defined set of objectives. An example objective might be to maximise the probability of correct classification when the classifier is applied to the validation set. It was found that for the classifiers developed in this chapter, the probability of correct classification (P(correct)) was insufficient to discriminate between hyperparameters, as an extended region of the hyperparameter space returned a perfect (i.e. 100% correct) classification. A set of objectives are instead defined.

1. Maximisation of probability of correct classification $P(Correct)_{VAL}$
2. Maximisation of the margin of separation, for the case where $P(Correct)_{VAL}=1$ for more than one combination of hyperparameters
3. Minimisation of the number of support vectors (maintaining (no. of support vectors)>0)

The expectation is that maximising the margin of separation between classes will improve the performance of the classification boundary under generalisation. However, the complexity of the decision boundary should be considered alongside the margin of separation that is achieved. The complexity of the decision boundary is reflected in the number of support vectors utilised in defining the boundary (related to the VC dimension).Where several classifiers offer the same level of prediction, the simplest is the best - the popularly applied principle of Occam's Razor – and good generalisation to previously unseen datasets will typically be promoted through selecting the simplest possible classifier that meets objectives (1) an (2). It is possible for a support vector machine to be constructed with no support. This is undesirable, and leads to the additional specification in objective (3) that the number of support vectors should be non-zero.

Generally, formalisation of the objectives as a quantifiable objective function to be maximised or minimised as appropriate would be preferred, with the subsequent requirement that the suitability of the developed objective function be assessed in order to verify that there is no pathology in the parameter selection process. An additional layer of complexity is thus added. In this case, the objectives were instead interpreted 'by eye', with the final decision on the hyperparameter values made manually. The validation step applied to the training of SVM1 is presented as an example of the approach applied.

An initial assessment over a broad range of values allowed an approximate range for the hyperparameter values to be established. The values used at the validation step for all 6 SVMs were:

- $C_i = 10^i$, where $i = \{0,1, \dots ,8\}$
- $\alpha_j = \{0.2,0.4, \dots ,1\}$

The SVM was trained for each of the resulting 45 hyperparameter pairs (C_i, α_j) using the testing set A_{TST}. The SVM resulting from each repetition of training was applied to the validation set A_{VAL}. The probability of correct classification for SVM1 with different combinations of the hyperparameters when applied to the validation set is shown in Table 7-3.

		Kernel width, α				
		0.2	0.4	0.6	0.8	1.0
Misclassification tolerance parameter, C	10^0	1.00	1.00	1.00	1.00	1.00
	10^1	1.00	1.00	1.00	1.00	1.00
	10^2	1.00	1.00	1.00	1.00	1.00
	10^3	1.00	1.00	1.00	1.00	1.00
	10^4	1.00	1.00	1.00	1.00	1.00
	10^5	1.00	1.00	1.00	1.00	1.00
	10^6	1.00	1.00	1.00	1.00	1.00
	10^7	1.00	1.00	1.00	1.00	1.00
	10^8	1.00	1.00	1.00	1.00	1.00

Table 7-3 Probability of correct classification for SVM1 for the validation set A_{VAL}

It is apparent that the probability of correct classification alone offers little evidence for choosing one pair of hyperparameter values over another. As a result, the margin of separation and the number of support vectors used is evaluated. The margins of

separation achieved for the defined hyperparameter pairs are presented in Table 7-4.

		Kernel width, α				
		0.2	0.4	0.6	0.8	1.0
Misclassification tolerance parameter, C	10^0	0.167	0.433	0.551	0.525	0.454
	10^1	0.167	0.433	0.551	0.525	0.454
	10^2	0.167	0.433	0.551	0.525	0.454
	10^3	0.167	0.433	0.551	0.525	0.454
	10^4	0.167	0.433	0.551	0.525	0.454
	10^5	0.167	0.433	0.551	0.525	0.454
	10^6	0.167	0.433	0.551	0.525	0.454
	10^7	0.167	0.433	0.551	0.525	0.454
	10^8	0.167	0.433	0.551	0.525	0.454

Table 7-4 Margin of separation for SVM1 for the validation set A_{val}

The margins are sensitive in this case to the choice of kernel width, α. As the objective is to maximise the margin of separation, a value of α = 0.6 is adopted. The margin size does not, however, allow for any discrimination between candidate values for the misclassification tolerance parameter, C. The number of support vectors required for the SVM does, however, allow some degree of discrimination. The results for SVM1 are presented in Table 7-5. The objective is to select the least complex SVM, assumed to be that which requires the lowest non-zero set of support vectors. On this basis, given the value already determined for α, a value of C = 10^5 would be selected.

		Kernel width, α				
		0.2	0.4	0.6	0.8	1.0
Misclassification tolerance parameter, C	10^0	406	131	89	67	54
	10^1	406	131	89	65	48
	10^2	406	131	89	65	48
	10^3	405	131	88	65	48
	10^4	403	129	85	61	45
	10^5	363	77	44	42	34
	10^6	0	0	0	1	6
	10^7	0	0	0	0	0
	10^8	0	0	0	0	0

Table 7-5 Number of support vectors for SVM1 for the validation set A_{val}

This validation step was repeated for each of the 6 SVMs from which the classifiers are comprised. The hyperparameters adopted are included in the results section below.

7.6 Results

7.6.1 Classifier 1: Multi-site damage

The aim of the first case study is to develop a 'best possible' classifier using the training and validation datasets $\mathrm{B}_{\mathrm{TRN}}$ and $\mathrm{B}_{\mathrm{VAL}}$, and testing the classifier on the test set $\mathrm{B}_{\mathrm{TST}}$. The validity of the SVM classification approach for the multiple damage identification problem may thus be evaluated. The classifier is the 'best possible' given the features selected in Chapter 6 and the modelling decisions specified in this chapter.

The classifier is trained using dataset $\mathrm{B}_{\mathrm{TRN}}$, which comprises 250 normal condition observations, 24 observations for each of the 5 single-site damage states and 5 observations for each of the 26 multiple-site damage states. The model hyperparameters are set using the validation set $\mathrm{B}_{\mathrm{VAL}}$. The hyperparameter values adopted are presented in Table 7-6.

SVM	Misclassification tolerance, C	Kernel width, α
SVM1	10^5	0.8
SVM2	10^5	0.4
SVM3	10^5	0.4
SVM4	10^5	0.4
SVM5	10^5	0.4
SVM6	10^5	0.2

Table 7-6 Hyperparameter values for Classifier 1

The classifier is tested on the testing dataset $\mathrm{B}_{\mathrm{TST}}$, which contains 1800 normal condition observations, 152 observations for each of the 5 single-site damage states and 40 observations for each of the 26 multiple-site damage states. The classification results are presented in confusion matrix form in Figure 7-9. Shading is used to delineate the normal, single-site and multiple-site damage states.

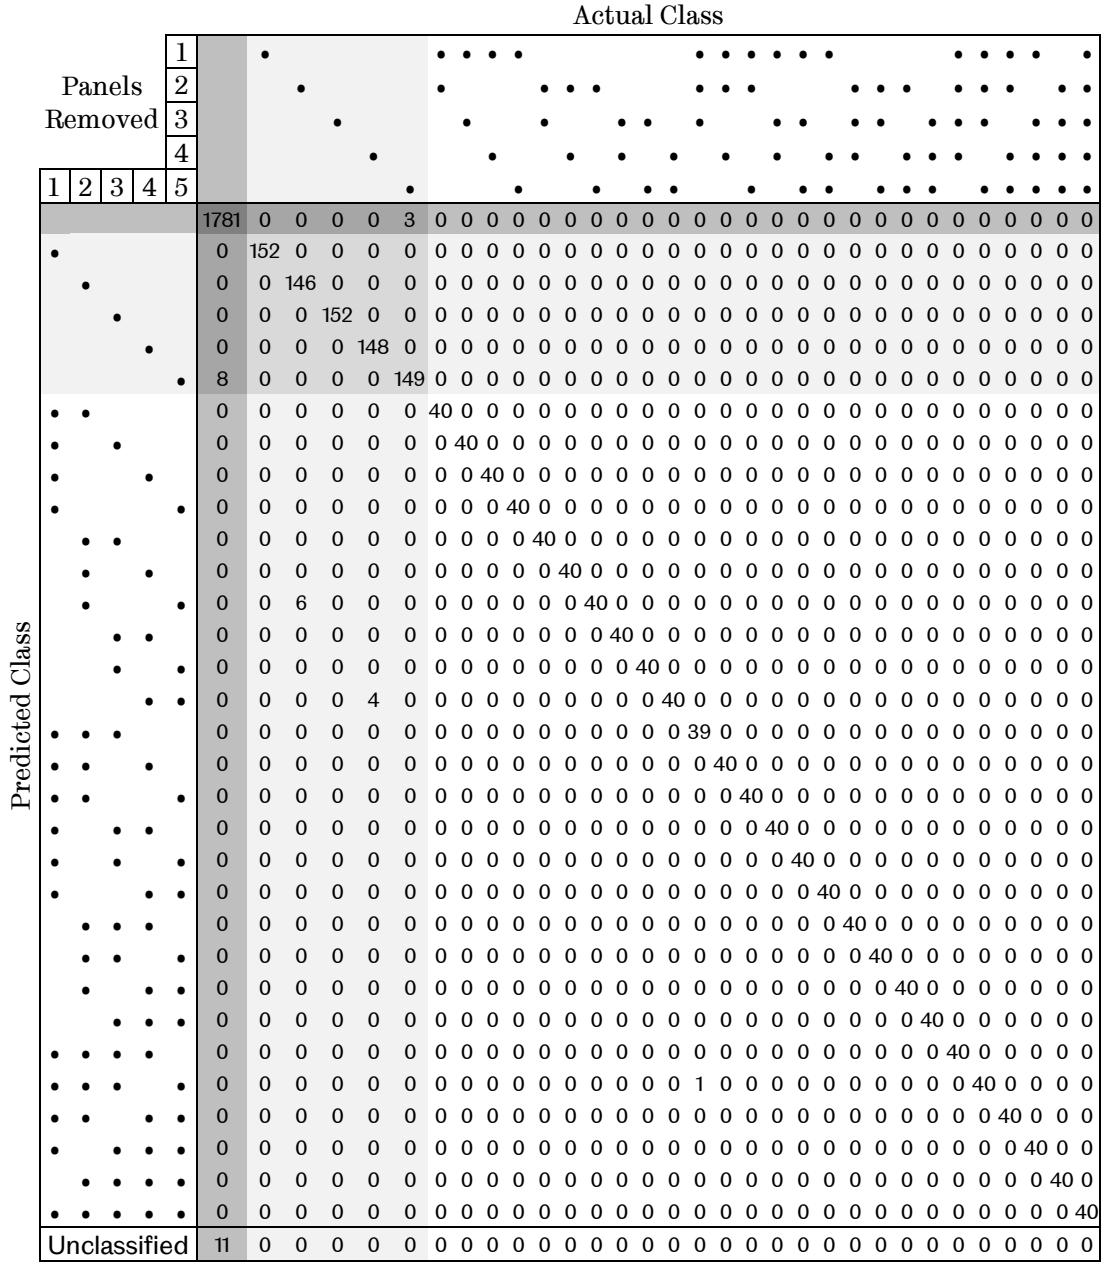

Figure 7-9 Confusion matrix for Classifier 1 applied to testing set B_{TST}

The results for each group of the damage states are summarised in Figure 7-8.

Condition	Correctly Classified	Incorrectly Classified	Unclassified
Undamaged	98.94%	0.50%	0.61%
Single-site damage	98.29%	1.71%	0.00%
Multi-site damage	99.90%	0.10%	0.00%

Table 7-7 Summary of results for Classifier 1 applied to testing set B_{TST}

The overall classification results are excellent. For two of the five single-site damage classes, a perfect (100% correct) rate of classification was achieved. The results are even more encouraging for the multiple-site damage classes, with a perfect rate of classification for 25 of the 26 classes.

There was in fact a single misclassification from a total of 1040 observations of multiple-site damage data, with the removal of panels [1 2 3] misclassified as the removal of [1 2 3 5]. It is worth noting that every misclassification involved panel 5. For the single-site damage, six observations of panel [2] removal were classified as removal of panels [2 5]; four observations of panel [4] removal were classified as removal of panels [4 5]; and three observations of panel [5] removal were classified as undamaged. Of the 1800 undamaged state observations, eight were classified as the removal of panel [5].It was found in Chapter 6 that the selection of features for panel P5 was more challenging than for panels P1-P4. This difficulty in identifying 'strong' features for panel P5 manifests itself as an adverse effect, albeit limited in severity, on the performance of the classifier in relation to the panel. This highlights the critical dependence of the classifier performance on the quality of features that can be identified. This may be of concern in the identification of smaller damage extents than the panel removal case investigated in Chapters 5 to 7.

In summary, a 'best possible' classifier has been developed and gave an excellent level of performance in correctly classifying testing set data. This gives confidence that the application of SVMs developed in this chapter is valid for the identification of multiple-site damage. Attention now moves to whether such a classifier can be developed in the absence of training data for the multiple-site damage cases.

7.6.2 Classifier 2: Single-site damage

Unlike for the previous classifier, an intermediate, independent testing set $\mathrm{A_{TST}}$ was available for testing Classifier 2. This proves useful as an initial check of the developed classifiers ability to identify single-site damage using data gathered during the same test sequence.

The classifier is trained using dataset $\mathrm{A_{TRN}}$, which comprises 250 normal condition observations and 50 observations for each of the 5 single-site damage states. The same set of 50 features is employed as for Classifier 1. The model hyperparameters

are set using the validation set A_{VAL}. The hyperparameter values adopted are presented in Table 7-8.

SVM	Misclassification tolerance, C	Kernel width, α
SVM1	10^5	0.6
SVM2	10^5	0.6
SVM3	10^5	0.6
SVM4	10^5	0.4
SVM5	10^5	0.6
SVM6	10^5	0.6

Table 7-8 Hyperparameter values for Classifier 2

The classifier is initially tested using the set A_{TST}. The expectation is that the classifier will achieve a high level of correct classification for this testing set. The results are presented as a confusion matrix in Figure 7-10.

Panel Removed	Actual Class: 1	2	3	4	5
1	500	0	0	0	0	0
2	0	100	0	0	0	0
3	0	0	100	0	0	0
4	0	0	0	100	0	0
5	0	0	0	0	100	0
Unclassified	0	0	0	0	0	100
Predicted Class	0	0	0	0	0	0

Figure 7-10 Confusion matrix for Classifier 2 applied to testing set A_{TST}

As foreseen, the classifier performs well on the testing set A_{TST}. The classifier is observed to achieve 100% correct classification for the 500 normal and 500 single-site damage state observations presented.

7.6.3 Classifier 2: Multi-site damage

The true test of the classifier comes when it is applied to the testing set B_{TST}, which contains multi-site damage data. in addition to previously unseen observations of the structure in its undamaged and single-site damaged states. An additional confounding factor is that the data comes from a separate testing programme to that used for training: the classifier is being asked to generalise across tests, handling any inter-test variability that may have arisen.

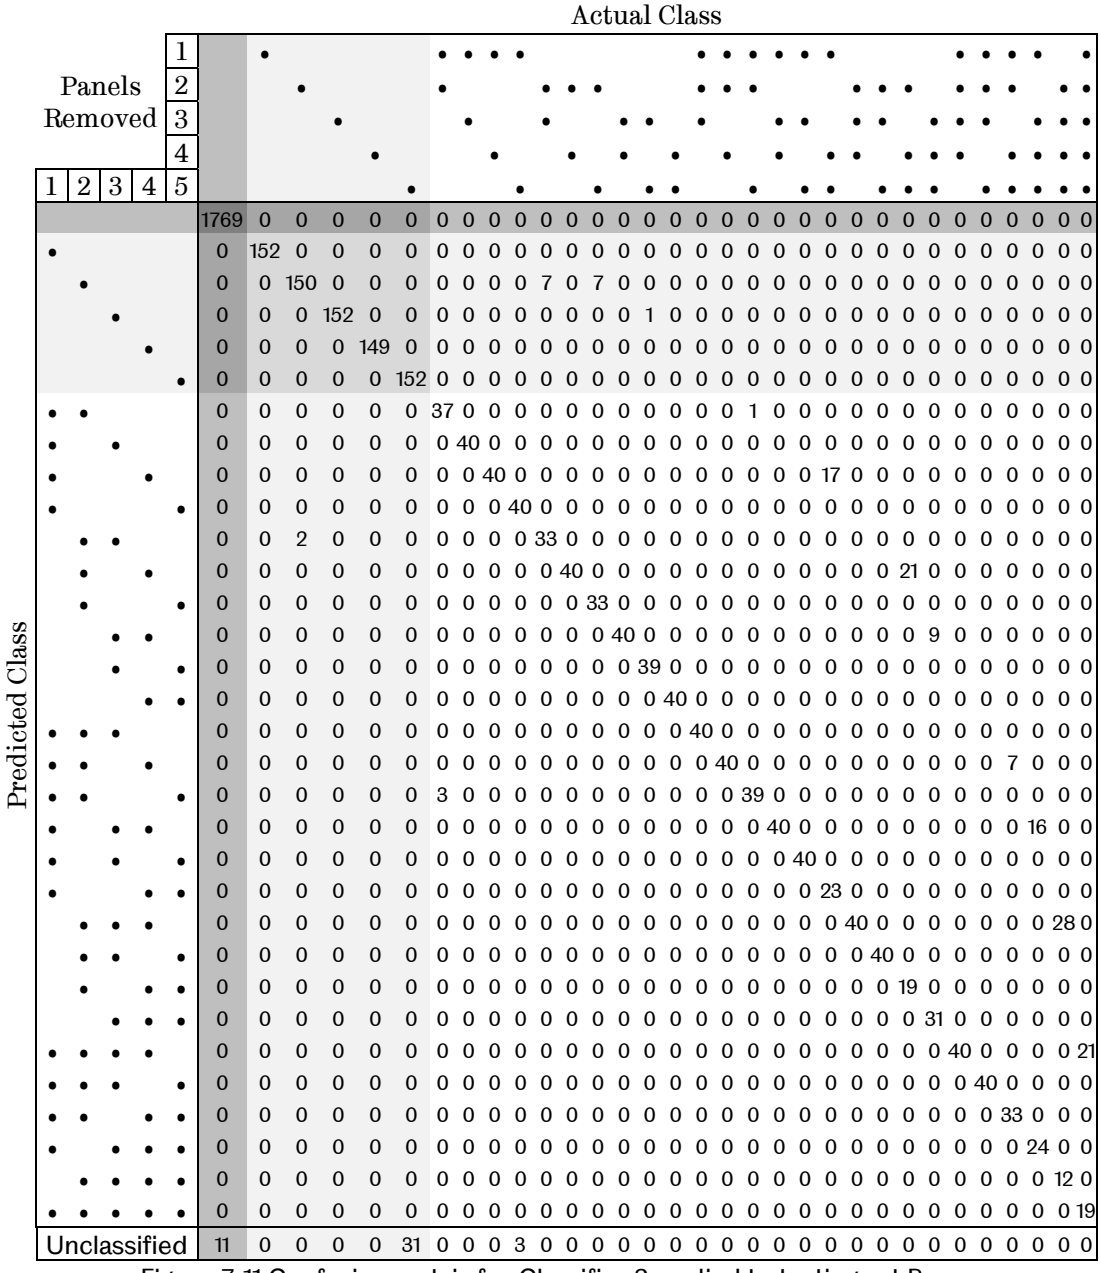

Figure 7-11 Confusion matrix for Classifier 2 applied to testing set B_{TST}

The full classification results are presented in confusion matrix form in Figure 7-11. The results are summarised, in terms of the probability of perfect correct classification for the three sub-categories of structural states in Table 7-9.

Condition	Correctly Classified	Incorrectly Classified	Unclassified
Undamaged	98.28%	0.00%	1.72%
Single-site damage	98.34%	0.26%	0.39%
Multi-site damage	86.73%	13.27%	0.00%

Table 7-9 Summary of results for Classifier 2 applied to testing set B_{TST}

The classification outcomes presented in Figure 7-11 and Table 7-9 are highly encouraging. For the structure in its normal state, there were no false indications of damage, and 1769 of the 1800 observations tested were correctly classified as undamaged. The remaining 31 observations were returned as 'unclassified', with SVM1 indicating damage but SVMs2-6 not indicating that damage had occurred on panels 1-5. This is taken as an indication that the classifier failed to generalise fully between tests, and may be attributable to the training set not fully representing the variability that may arise between tests. A second observation is that while the objective of SVM1 was damage identification, the 50 features used were selected based essentially upon damage location criteria. The criterion for selecting the features was that they should be indicative of damage on individual panels, rather than indicative of damage regardless of source.

The classifier performed exceptionally well for the single-site damage states, with only 5 non-correct classifications out of 760 observations. Of these, three were returned as 'unclassified' with SVM1 indicating damage but SVM5 failing to classify that panel 4 had been removed. A further two observations were misclassified, with the classifier indicating the removals of panels 2 and 3 where in fact only panel 2 had been removed. This was one of only two instances in which the classifier falsely indicated damage at any location. The other was for the removal of panels 1 and 2, for which the classifier indicated that panel 5 had also been removed.

In total, there were 128 observations (out of a total of 1040 observations) for which the classifier trained solely on single-site data failed to perfectly identify multi-site

damage states. Of these, 125 observations missed the removal of one panel but were otherwise correct. 118 of these observations were due to removal of panel 5 having been missed; the remaining seven were due to panel 2 being missed. The suggestion, as initially raised in Chapter 6, is that the feature set identified for indicating the removal of panel 5 is comparatively 'weaker' than those for the remaining four panels. This will warrant further investigation if the approach is to be applied for less severe damage scenarios for which less discriminatory feature sets will be available. Alternative methods of feature selection and normalisation may be sought in such cases.

Overall, the success of the classifier when applied to multi-site data is encouraging. It appears that given a suitably discriminatory feature set, the SVM approach is capable of achieving an exceptionally high level of correct classification for single-site damage, and a good level of classification for the much more challenging task of identifying multi-site damage.

7.7 Discussion

In this chapter SVMs are first shown to be an applicable as the basis of multi-class classifiers for tackling the multiple damage location task. A family of binary support vector machines has been shown to be capable of tackling a multi-class damage identification problem, with excellent (near perfect) results when damaged-state data for all the structural states are incorporated in the training set. The major contribution of the chapter, however, is that it is shown that a classifier trained using normal condition and single-site damage data only may be capable of identifying the presence of multi-site damage with a high degree of accuracy. This hugely reduces the number of states for which damaged condition data is required.

The case investigated is challenging. However, the removal of inspection panels represents a relatively gross level of damage, and the damage locations are coarsely dispersed. A question that remains is how well the method would cope with more challenging scenarios? This concern is promoted by the difficulty experienced in achieving perfect diagnosis for panel P5. The features selected for identifying this panel were somewhat weaker than those found for other panels. It is likely that the true efficacy of the method will only become apparent when considering more challenging cases. It is also foreseen that modelling choices will play a greater role as the discriminative capacity of the feature set decreases, and that arriving at correct

decisions with regard to the model structure and kernel choice, for example, will subsequently be of greater significance in more marginal cases. Comparison of options for multi-class identification and further investigation of validation methods are among the areas that warrant further study.

A final point to consider is the existence or otherwise of an underlying physical explanation for the success of the approach. The approach relies upon the presence of features that are sensitive to the removal of one panel in isolation and which also 'fire' when that panel is one of several removed, but little regard has been given to *why* that should be. An possible explanation arises through considering the structure of the wing. The investigated section of the wing is made up of a thin Aluminium topsheet mounted on a framework of more substantial stiffening elements. This results in the topsheet being divided into many rectangular regions of flexible material, bounded on each side by more rigid elements. A simple suggestion that may be made on the basis of engineering insight, is that the coupling between these regions may be rather weak, with the effect that stiffness changes on individual regions would be expected to be most obvious in the response measurements taken from that region and less so in responses taken from other locations. It is conceivable that this may lead to a degree of independence between regions, which would, intuitively, aid the ability of features selected using single-site damage data to generalise to multiple-site damage states. While this was not specifically considered during sensor placement it is notable from Figure 6-2 that 11 of the 13 'response' sensors were placed within the same region as the panel they were employed to monitor.

The conclusion drawn is that while the approach pursued in this section of the thesis doesn't require a law-based model for reference, its success still appears reliant on having some insight into the physical underpinnings of the observed phenomena. The ways in which law-based and data-based modelling methods may support one another are considered in the synthesis of Chapter 8.

8 Synthesis of Approaches

8.1 Introduction

The objective of the current chapter is to draw together the approaches presented and discussed throughout the thesis. This is pursued through discussion of the relative strengths and weaknesses of law-driven and data-driven approaches to the diagnostic element of Structural Health Monitoring (SHM), with the intention of indicating where they are best used in practice. Given the degree of commonality that has been made apparent, the chapter will also discuss how the two approaches can support each other in the development of best practice and some speculation will be made as to how the approaches might be combined in order to exploit the strengths of both.

Current approaches to the diagnostic problem which is central to the field of SHM are usually based on two main possibilities: an inverse problem formulation and a machine learning approach. The first of these approaches, often called the model-based or law-based approach is usually applied by constructing a physics-based model of the structure of interest (e.g. a Finite Element (FE) model) and correlating it with experimental data. Once the model is established, it can be used in a monitoring phase by periodically updating the parameters of the model, usually by linear-algebraic methods. The nature of the problem means that the linear-algebraic formulation is often ill-posed and requires careful regularisation [14, 15]. This is the approach pursued and discussed in Chapters 2 and 3. The alternative approach to diagnostics in SHM, often called the data-driven approach, also involves the construction of a model, but this model is usually statistical. The model is established by means of machine learning or pattern recognition and may involve the use of classifiers or novelty (outlier) detectors [22, 92]. An example of a data-driven approach applied to a challenging problem has been presented in Chapter 5-7. It

must be recognised that the problem of implementing a credible SHM strategy in any real-world context is much more wide-ranging than the choice of a diagnostic methodology. The broader aspects of SHM are however, not discussed here; the reader may consult [93] and [94] for more background; reference [10] must be considered as the current definitive guide to the subject.

Both of the approaches discussed above have substantial support in the literature of SHM; however, they arguably have different strengths and weaknesses, which potentially make their domains of application problem-dependent. The methods also show a degree of commonality which is sometimes overlooked. In the first case, as observed above, both approaches can be said to be model-based. The distinction is in the type of model. If one classifies models into white, grey and black-box models according to their degree of *a priori* physical content; one would observe that the inverse problem approach seeks to establish a white-box model, while the machine learning approach uses a grey or black-box model. The advantage of the former is precisely that it exploits any available physical knowledge of the system of interest; the advantage of the latter is that it automatically accommodates any uncertainty in the specification of the system or structure. Another common aspect of the approaches is that they require measured features from the structure. In the inverse-problem approach, features are generally needed in order to update the *a priori* physical model in order to bring the normal condition model into better accord with reality. Features are also required for the damage identification task: these may or may not be the same features that were employed for updating the normal condition model. The machine learning approach uses measured features in order to form the statistical model and to reduce the dimensionality of the problem as far as possible. In both cases, the selected features must be sensitive to the damage. For the inverse problem this is a requirement for a non-trivial update; for the machine learning approach, the features are essentially everything and must reflect any information about the damage. Feature selection, is therefore an issue for all SHM methodologies, but is arguably most discussed in the machine learning context.

8.2 Classification of models

The discussion later in the chapter will require the specification of a taxonomy which allows one to distinguish between classes of models. The two most common means of distinguishing model types are covered by the following.

8.2.1 Data-driven and law-driven models

A good reference on this classification of models is [95]. A convenient way of expressing the differences between the two types of model is by means of 'bullet points'.

Law-driven models

- Based on accepted laws attributed to the system – 'physical' in nature
- Suited to prediction, potentially for unobserved system states
- May be used to inform critical data acquisition decisions
- Typically highly-parameterised
- Generally do not accommodate uncertainty

Data-driven models

- Based on observed input/output relationships (supervised learning) or outputs only (unsupervised learning) – 'statistical' in nature
- Suited to recovering inputs from observed outputs
- May be parsimoniously parameterised
- Naturally accommodate uncertainty

This is a useful picture; however, one should regard these statements as representing 'extreme' viewpoints; things are seldom black and white. Consider the suitability of the two paradigms for 'prediction' purposes. It is well-known that data–driven models like neural networks should only be used in situations where the input data do not depart dramatically from those used during the 'training' of the model [81]; in other words, such models can only be used to *interpolate* with any real confidence. In contrast, one would imagine that an appropriately derived law-driven or physics-based model could be used to *extrapolate* i.e. make predictions in situations removed from the current SHM context; physics should after all be universal. However, one should bear in mind that law-driven models may be the best

that physics has to offer, yet still be subject to restriction; for example, the Navier-Stokes equation is restricted to situations where one has confidence in a continuum assumption. Further, when one describes the Navier-Stokes equation as a physics-based model one is misrepresenting the situation in terms of prediction capability as the equation has no known analytical solutions; prediction is only possible by passing to numerical solution methods which are themselves subject to limitations and represent a different class of model to the original partial differential equation. Finally, one can observe that a perfect specification of a deterministic model still does not guarantee prediction accuracy; this is one of the hard lessons learned from the discovery of deterministic chaos [96].

One can also consider the question of 'parameterisation' of the models. Although it is stated above that law-driven models may be 'highly-parameterised'; this need not be the case. If one considers a large FE model, there may be very many parameters indeed; in principle, each individual element could have independent material parameters. In fact, only a small number of 'substructure' material constants will be needed. Of the huge number of coordinates which specify the mesh geometry, only the subset which fixes the geometry of the structure of interest matter in any real sense. 'Internal' nodal coordinates can be varied more or less with impunity (as long as one does not violate aspect ratio constraints etc). It is this fact that only a very small subset of the parameters 'matters' which allows FE model updating to be computationally feasible for model improvement. For SHM the additional requirement is that the model should be capable of representing the structure in all damage states of interest. Meeting this requirement while avoiding an explosion in the number of required parameters (and subsequent issues of ill-conditioning) is perhaps one of the greater challenges faced in the application of model updating-based methods.

The parallel statement that data-driven models can be 'parsimoniously-parameterised' is by no means a general rule either; neural networks may need hundreds of internal parameters in order to capture the input-output behaviour of a system with appropriate fidelity; this observation impacts considerably on the applicability of data-driven models, as it imposes severe constraints on the amount of data which must be available for 'training'.

8.2.2 Black, white and grey box models

"Here they come, every colour of the rainbow: black, white, brown" [97]

It is informative to consider the nature of the information upon which a model is based. A white box model is built solely upon the best possible understanding of the underlying physics of the system. For a purely white model, no assumptions or approximations regarding model structure would be made, and the 'whiteness' of a given model reflects the depth to which a complete physical understanding is pursued. In practice, limits on our understanding of the physical universe preclude the existence of pure white box models; the discussion on the Navier-Stokes equation above makes this point clear. One should also note that both the Navier-Stokes equation itself and the discretisation usually necessitated for solution are often brought together under the term 'white box.

Conversely, a black box model seeks to describe the behaviour of a system with no reference to its internal structure or 'physics'. Such models are instead built solely around observed input-output behaviour. This does not always present practical limitations as the universal approximation capability of many machine learning model paradigms means that a black-box model can in principle capture input-output behaviour perfectly; however, as discussed above, for predictions with such models it is not usually safe to stray too far away from the situations in which the training data were generated.

A grey-box model is a compromise between these two extremes i.e. a model for which the physics dictating the input-output relationships are partially understood *a priori*, but which allows for the inclusion of approximations from empirical observations. The relative 'whiteness' or 'blackness' of the model may be viewed as dependent upon the number and quality of the assumptions made in specifying the physical understanding of the system, and the degree to which the model relies upon approximations made from the observed data. A good example here may be in specifying an initially linear FE model for a structure. If experiment were to make it clear that the actual structure was in fact nonlinear, one could add individual nonlinear elements and calibrate their coefficients using experimental observations. Even if the nonlinearities are non-polynomial in nature, they will always be approximated effectively by an appropriate high-order polynomial; 'model'

nonlinearity such as this is therefore non-physical and converts the initial white FE model to a grey box.

8.3 The modelling task for damage identification

The purpose of this section is to discuss the modelling capabilities of the two main approaches: law-based and data-based. The 'true' structure may be regarded as a function (or functional) of its inputs. This function specifies the mapping of inputs to responses for all states of the structure, both damaged and undamaged. For damage identification, the modelling task is to specify a model capable of approximating the true structure across all states of interest.

The discussion is illustrated through consideration of a space of functions or functionals encompassing all possible model structures. The 'true' structure is denoted as *T* in the following figures. Where model updating has been employed to reduce the residual distance between model and structure, the 'optimal' model is denoted as *O* in the figures. In the interests of visualisation, the abstract concept of an infinite-dimensional functional space is presented in two dimensions.

8.3.1 Physics-based modelling

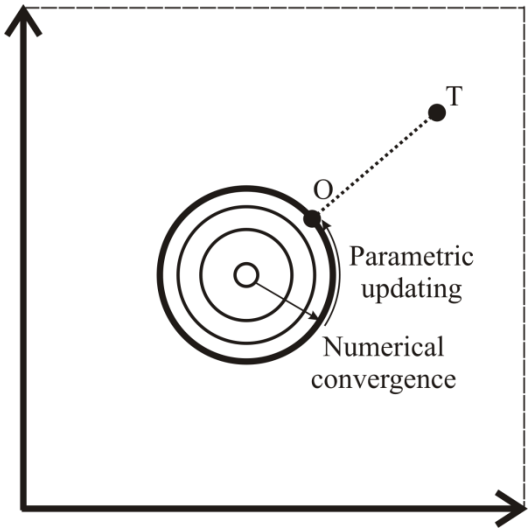

Figure 8-1 Functional space portrait of a linear physics-based model

Suppose one begins by considering the family of linear physics-based models as depicted in Figure 8-1. This class of models, by nature of its restrictions, forms a subset or spans a subspace of all the possible functional forms for a model. The particular subspace spanned is dictated by the model-form selected and the values assigned to its parameters. The extent of this subspace (the modelling 'power' or

capability) grows with the dimension of the model – as more degrees of freedom are added to, say, a finite element discretisation of the structure, so the number of possible models expands. One can specify a complexity parameter: in this case the number of elements is meaningful. So for linear FE models, one has a complexity parameter: $N_{phys=}\{N_{elements}\}$.

The creation of a numerical model necessarily involves the discretisation of continuous physical laws, manifested in the specification of time-steps or element sizes. The discretisation error arising from this process limits the capability of the model to accurately reflect the physics of the true system. As the step or element-size is decreased (and thus the order of the model increased), convergence of the subspace to some outer limit or boundary would be expected. Discretisation of continuous equations and the concept of a regime of asymptotic convergence of predictions are covered in Chapter 4.

Within the spanned subspace, the 'optimal' model is that which minimises the distance (in terms of the geometry of the function space) between the model and that of the 'true' structure. This optimisation may be guided by direct measurement of parameter values and/or calibration of parameters through updating. Where parametric updating methods are employed for this purpose, the optimisation process will typically be constrained in order to ensure that the model parameters remain physically representative. e.g. for FE updating, a principled approach to choosing the 'closest' model within the class could be based on least-squares minimisation of residual errors; possible constraints on the minimisation might be the requirement to preserve the sparsity structure of the physical matrices during updating [29]. If the true system falls within the explanatory range of the model class, the updating procedure will bring the distance down to zero. If the true system is without the model class as depicted in Figure 8-1, one can only hope to get as close as possible.

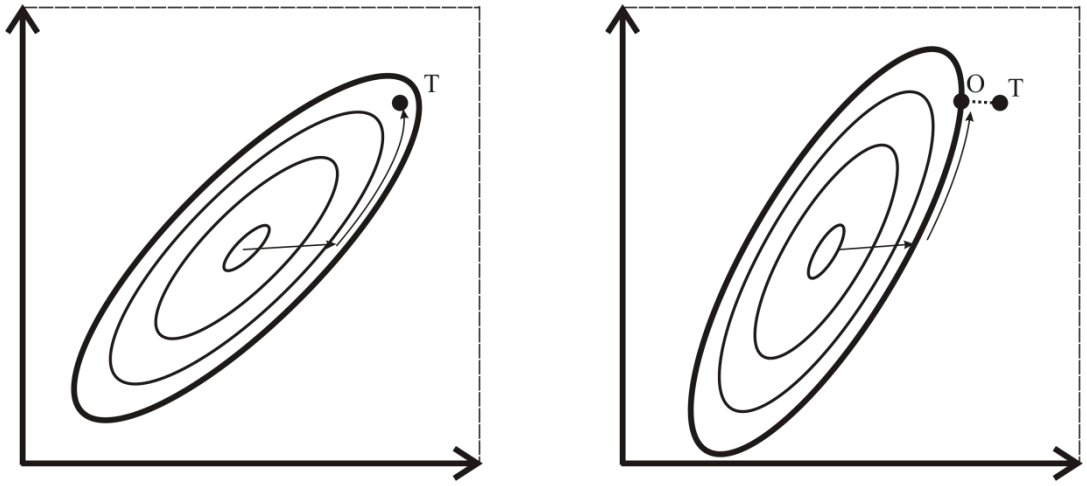

Figure 8-2 Functional space portrait of a nonlinear physics-based model where
(a) the included nonlinearities accurately reflect those of the true structure, and
(b) there are discrepancies between the included nonlinearities and those of the true structure

One can now progress to, say, nonlinear FE models. Such models will have at least the explanatory power of linear FE models, but will clearly span a greater volume of the function space. Staying within the dictates of law-based modelling, adding finite numbers of specific nonlinearities will mean the linear FE class is only extended in certain directions in the function space. In general, this may well mean that the extended class still does not include the true system as depicted in Figure 8-2(b); however, it may be that it does, as in Figure 8-2(a). The complexity parameters for the class must also specify the number of added 'nonlinearities': $N_{phys=}\{N_{elements}, N_{non-linearities}\}$. Any given physics-based class, e.g. nonlinear FE models, will still be subject to restrictions on the explanatory power even if the number of elements grows without bound. This example represents the case discussed earlier where addition of extra terms increases the explanatory power of a white box model but converts it into a grey box in the process. Concentrating on a linear FE model with 'added' nonlinearity is extremely relevant for SHM as many damage types will convert an initially linear structure into a nonlinear one. Further, the exact form of the nonlinearity relevant to a fatigue crack in a metal or a delamination in a composite laminate may not be precisely known.

8.3.2 Data-based modelling

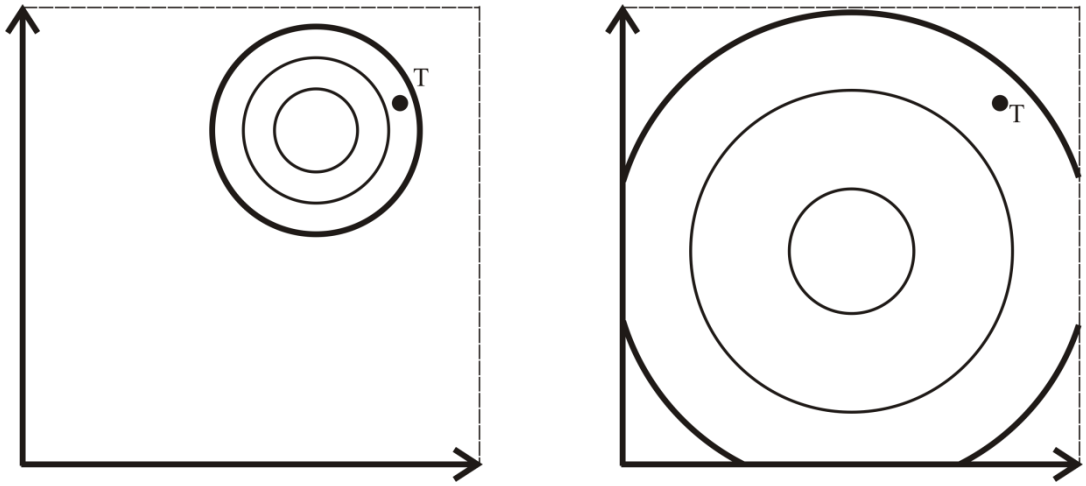

Figure 8-3 Functional space portrait of a data-based approach where:
a) the structure of the data-based model is 'close' to the true structure, and
b) the structure of the data-based model is not 'close' to the true structure

As in the case of physics-based models, one can regard a given data-based paradigm as spanning a subspace of possible functions/functionals. However, there is a critical difference. Many classes of data-based models can be proved capable of acting as universal approximators; these classes include: multi-layer perceptron neural networks, radial-basis function networks [81] and Support Vector Machines [85] (as employed in Chapter 7). The universal approximation capability means that as the complexity parameter (number of free weights etc.) increases without bound, any function/functional can be represented arbitrarily accurately. In terms of Figure 8-3, this means that as the number of parameters increases, any point in the space is reachable and one should always be able to encompass the 'true' point T for some complexity parameter: $N_{data=}\{N_{free\ params}\}$. Unfortunately, the greater the number of parameters, the greater the amount of training data that is required and this is the critical problem in the context of SHM. The 'true' function in an SHM context is often a classifier which maps data to a class label specifying the state of health of a given structure. If the classifier is a dual class (i.e. healthy/damaged) novelty detector, it may only need to be trained on examples of data from the healthy structure; this is called unsupervised learning. If the classifier is more refined and is required to indicate location or severity of damage, training data from the damage states will be needed; the problem is now one of supervised learning. In an engineering context, this presents a great problem as it may not be economically feasible to generate fault data exemplars from high-value structures. Training of a

complex model with many parameters may demand a great deal of training data in a situation where acquiring any damage state data at all will be a formidable problem.

8.3.3 Hybrid approach

The 'hybrid' approach presents a potential solution to the problem just discussed. One can use a physics-based approach to establish a model which explains as much as possible of the function/functional behaviour as possible. The situation, as depicted in Figure 8-4(a), will be that the initial physics-based model will establish a point *O*, close to the truth. From this point, a data-based adjunct will be able to bridge the gap between *O* and *T* because of the universal approximation capability of the data-based approach. The advantage here is that the data-based component will potentially explain the 'residual' between *O* and *T* with a need for fewer parameters (than a full data-based model) and will therefore require less training data.

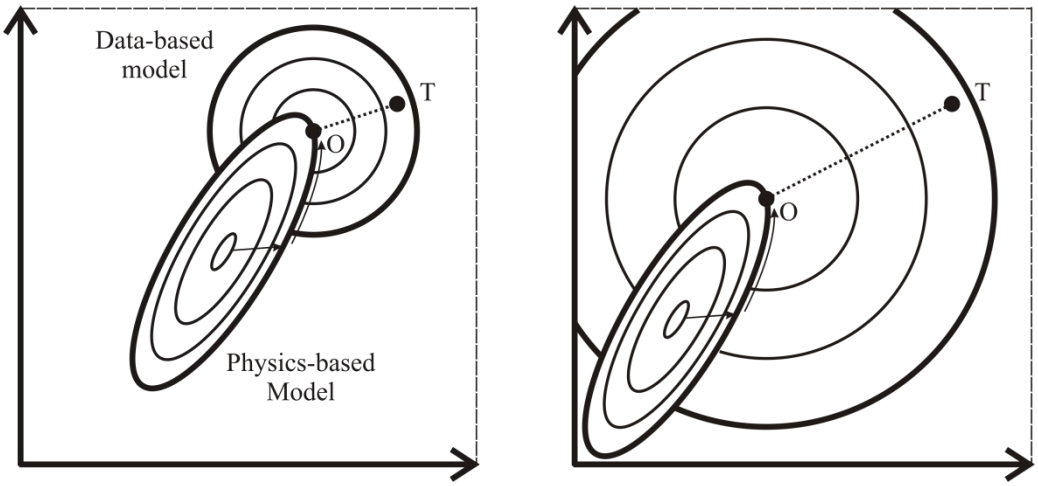

Figure 8-4 Functional space portrait of a hybrid physics- and data-based approach where
a) the physics-based model is 'close' to the true structure, and
b) the physics-based model is not 'close' to the true structure

The complexity measures in this case are $N_{phys=}\{N_{elements}, N_{non-linearities}\}$ plus $N_{data=}\{N_{free\ params}\}$. In the situation shown in Figure 8-4(b), the initial physics-based model has less explanatory power than that in Figure 8-4(a) and as a result the parameter count for the data-based component will be higher. The initial stages of a possible hybrid approach are discussed in Chapter 4.

8.4 Application of approaches

Each method may be considered as comprising a training phase and a monitoring phase. The training phase for the data-based approach is precisely as described in all

the relevant texts in machine learning or pattern recognition. The monitoring phase comprises the application of the learnt 'rules' to newly-presented data from the possibly damaged structure. In terms of the law-based approach, the 'training phase' means the initial building of the physics-based model on the basis of all available prior information, followed by updating and validation exercises using experimental observations to bring the model further into accord with reality. The monitoring phase for a physics-based approach will typically comprise further updating steps on the basis of newly-presented data from the structure.

For the hybrid approach, the training phase is considered in the discussion below to involve the development of the 'optimal' law-based model, and the training of a data-based model using the (bias-corrected) predictions of this law-based model. The monitoring phase is once again the application of the data-based model to new data. It should be noted that the objectives of the law-based modelling task are thus somewhat different for the hybrid approach and the purely physics-based approach. For the hybrid approach it is advantageous for the reliance on deterministic updating to be superseded by application of the rather broader concept of probabilistic model validation. The broad aim of probabilistic model validation is the pursuit of quantified levels of confidence in prediction, as discussed in Chapter 4.

8.4.1 Features employed

As observed, both approaches have a training phase which requires the availability of experimental data. The exact type of data used will have a critical effect on the construction of the diagnostic. The most important point is that the data, expressed through multivariate data vectors called 'features', must be sensitive to damage.

For physics-based methods in general, the selection of the feature type is heavily dependent upon being able to identify the feature both from the model and the experimental data. In the model-updating literature, this has historically led to the broad adoption of modal properties (primarily frequencies and modeshapes) or FRFs. Comparison between experiment and the initial model leads to residuals that must be minimised when updating structural models. Often, the updating step will require analytical results, specific to given features e.g. sensitivity-based updating requires the formulation of derivatives of the error function with respect to the features; these derivatives have been computed and are available in the literature for the standard modal features [29]. This reliance upon modal properties has been

maintained despite the insensitivity of modal responses to damage being well documented. It should be noted that a third category of modal characteristic - damping - is rarely considered for forming residuals for model improvement or as a feature for damage detection, due to difficulties associated with characterising and measuring damping.

Despite the effort invested in reducing the residual between the responses of a physics-based model and the 'true' structure, some degree of offset will inevitably remain. This offset arises in part from the impossibility of achieving a true white-box level understanding of the physics of the structure; specific examples include joints and welds. Bias correction techniques may be used in an effort to circumvent this error.

The restriction to features that can be identified and correlated between model and structure is removed in the purely data-driven case. For the machine learning approach, 'restrictions' on the features are only to what can be measured from the structure: modal characteristics, time-domain data, spectral data, strain histories, images. Given this lack of restriction on the type of feature, emphasis is instead placed upon reducing the dimensionality of the problem as far as is possible (and thus reducing the requirement for infeasible amounts of training data) and selecting a feature set that is robustly indicative of damage.

The degree of 'restriction' on the feature set falls somewhere between these two levels for the hybrid approach. The decision on whether a particular response of the model should be considered a candidate feature will largely be dictated by whether it can be predicted with a satisfactory degree of confidence.

8.4.2 Feature selection

It is sensible to make a distinction between feature specification and feature selection. The broad class of features available will be determined by the choice of sensors and their number and distribution; as such this must be considered at the SHM operational evaluation stage as discussed in [93]. A great advantage of employing a physics-based approach is that the model may be used to guide the sensor placement and feature selection tasks (commonly applied sensor placement algorithms are provided in Appendix A), and it would in principle be possible to assess the sensitivity of each individual feature within the feature set to the

presence of damage: a preliminary example is given in the Case Study in Chapter 4. Examples are, however, relatively rare in the published literature related to model updating-based SHM, although in [16] parameter set reduction techniques have been applied. In the inverse model-based approach, as discussed above, the candidate feature set used is typically specified *a priori*, and is often the same for both the model-improvement stage and the structural monitoring stage. This is despite the objectives of these two stages being arguably somewhat different. The feature set used is defined largely by convention and convenience, rather than through analysis of suitability for damage detection.

In direct contrast, feature selection – the reduction of a candidate set to a maximally informative set - is very widely discussed in the machine learning context, where the features are essentially everything. A practical consideration here (returned to below) is that data are unlikely to be available for all damage states, placing a restriction on the domain for which the features are assessed. A solution making full use of the hybrid formulation may be to use the physics-based model to provide probabilistic representations of a candidate feature set, and to apply feature selection methods developed in the machine learning context to this set.

8.4.3 Treatment of uncertainty and variability

There are many potential sources of uncertainty in the SHM task. Several of these, including boundary condition variability and operational and environmental effects are common to both approaches. The use of physics-driven models introduces an additional set of model-form and parametric uncertainties. The two approaches handle uncertainties in markedly differing ways.

There are two elements to consider when assessing the importance of uncertain factors:

- Are the features that are being observed sensitive to the uncertain factors?
- What degree of uncertainty can be ascribed to the uncertain factors?

An advantage of the machine learning approach is that it can automatically accommodate any uncertainty in the specification of the system or structure. This is, however, contingent upon the learning algorithm being presented with training data that is representative of the variabilities and uncertainties that are expected in practice. A common, yet non-trivial, example is that of environmental effects. It has

been observed in numerous studies that temperature effects can have a significant effect upon the features employed for damage identification. In some cases, it has been found that the observed features display greater sensitivity to environmental effects than to the damage of interest [3]. The pitfalls of training a classifier using data gathered at a single temperature point in such a scenario are immediately apparent, and it would be expected that a classifier trained at one temperature point may perform poorly when extrapolating to changing environmental conditions. In this scenario, the possibility may exist to treat the variable factor (temperature) as observable, and to incorporate it into the feature vector presented to the machine learning algorithm. Where the variable factor is not readily observable, as may be the case for boundary conditions, a different approach may be taken. Effort may be directed towards recording training data that captures the range of variability expected in practice, for example through employing blocking and randomisation as appropriate in the test programme. This was the approach taken when developing test sequences in Chapter 5. An initial analysis of the sensitivity of candidate features to particular uncertain factors may aid this process.

In the inverse-problem approach to SHM, the rule has very largely been to build 'crisp' models which do not accommodate variation; departures from this rule have recently emerged via the use of stochastic FE [98] and fuzzy FE [99] etc. The quantification and propagation of uncertainty is in contrast a core activity in probabilistic model validation approaches. Sensitivity analysis is conducted using the initially-developed model to identify those factors to which model responses are sensitive. The experimental and modelling effort is directed at characterising the uncertainty in these factors, and the subsequent variability in the model outputs, The resulting probabilistic predictions may be used to supplement the available experimental data for machine learning in the hybrid approach, Even when uncertainty is explicitly handled in a law-based model; the form of the uncertainty must be guided or fixed on the basis of observational evidence; this may well change a white box model into a grey box in the same way that addition of nonlinear elements would.

8.4.4 Practical considerations

The discussion so far has been largely related to the advantages and disadvantages of models in terms of their explanatory power, their ability to encode uncertainty

and other matters which have been discussed in largely abstract terms. The intention with this section is to look at practical issues including those of resourcing.

In terms of cost (in terms of both time and money), physics-based models are undoubtedly time-consuming and challenging to build, develop and validate. They rely on the existence of expert and talented model builders with considerable training, both in the generalities of model building and the specifics of individual software packages. The software packages themselves may be costly in terms of initial licensing and maintenance. Large models will require intensive computing facilities with the associated operation and maintenance overheads. Robustness under uncertainty can be an issue for law-based models; these are typically calibrated at a single design point (the normal condition), and validated for responses not used for updating, rather than for other input states. Robustness and sensitivity are often left unassessed. Also, additional epistemic uncertainties (or model-form errors) will generally be introduced through modelling choices and simplifications. Law-based models are typically over-parameterised; often this makes updating a ill-conditioned task. In order to improve the stability of solution, regularisation may be employed [44]. Finally, during the monitoring phase, any assessment of integrity will require a full model update step. There are of course many points in favour of the law-based approach. First, law-based models can potentially be used to extrapolate; if the actual physical laws underlying the model extend beyond the initial context, then the constructed model will also extend. Secondly, the actual computational cost of the update step may not be excessive. A huge advantage of the law-based model is concerned with the possibility of observing multi-site damage. Consider the situation where the structure under investigation is a cantilever beam with a potential fatigue crack, as considered in [69]. Once the undamaged structural model has been validated, and the crack model has also been validated, there is little further problem with modelling the multiple crack scenario beyond inversion issues (inverse approach) or computational cost (forward approach). This is not at all the case for the data-based approaches.

In discussing the disadvantages of the data-based approach, one must begin with the ‘elephant in the room’. Sourcing data from the structure in its damaged state is unlikely to be feasible in scenarios like those involving large-scale structures. One clearly cannot conceive of physically damaging an aircraft in multiple ways in order to accumulate data for supervised learning. Even if multiple damage cases become

available, there is a potential requirement to account for statistical variation in the data. Finally, even if it is possible to acquire some data from damaged structures, It will certainly not be feasible to acquire data corresponding to multi-site damage, where the number of damage states for which data would be required grows exponentially with the number of discrete damage sites identified. Sourcing data is the main problem faced by data-based SHM. If only damage detection is required, things become much more positive as only measured data from an undamaged structure will be required. A related problem concerns the amount of training data required. If the machine learning model structure has many adjustable parameters, it may demand an unfeasible amount of training data. Data, even from undamaged structures, does not come without a cost. In terms of model development; this can be accomplished using machine learning software which is comparatively quick to master (compared with a nonlinear dynamic FE solver for example). Training the model may well be time-consuming, but then runs will be extremely fast during the monitoring phase.

8.5 Discussion

This chapter attempts to draw together the work presented in Chapters 2-7 on approaches to the SHM problem. As the chapter is essentially a discussion document, there is no real need for detailed conclusions. The chapter simply discusses the differences and commonalities between the law-based and data-based approaches to SHM. The main issues discussed here relate to: the relative explanatory power of types of models, the accommodation of uncertainty in models and finally, practical issues in implementing diagnostic strategies. While these issues surely do not exhaust issues for possible discussion, it is felt that their importance justifies their prominence here. It is observed that the two approaches discussed have their own individual strengths and weaknesses and that a hybrid approach may be possible that exploits the capabilities of both. The development of such a hybrid approach is the subject of ongoing research, with preliminary work having been presented in Chapter 4.

9 Conclusions and Future Work

9.1 Discussion and conclusions

In the course of the presented work, aspects of each of the three approaches to SHM that were identified in Chapter 1 have been treated in turn in the hope of building a broad picture of the options available for tackling the SHM problem. In carrying out this work advantages of each approach have been realised – in addition, several of the disadvantages and challenges reported in the relevant literature manifested themselves in the cases presented. The aim of this concluding chapter is express the author's opinion of the merits and ways-forward for each approach. Conclusions and recommendations for each approach are presented below. First, model-based methods are considered.

- Damage identification using model-updating methods was found to be a very challenging task. It demands strong skills in numerical modelling and inverse methods, plus an understanding of data acquisition techniques.

- Regularisation to allow the solution of the inverse problem was found to present a source of particular difficulty, and a factor whose importance is, arguably, often understated in published studies. For the case study presented, attempting to set a regularisation parameter using the L-curve method led to an inconclusive outcome- the distinct 'corner' of the L-curve that is sought for defining the optimal parameter value λ_{opt} was not apparent. It is recommended that reporting the details of the regularisation applied should be regarded as essential when presenting the results of model-updating-based SHM studies.

- Accurate mode-pairing based upon natural frequency and modeshape correlations proved non-trivial. From the forward modelling work presented in Chapter 4, it was found that erroneous mode-matching had occurred and been employed in Chapter 3, with the implication is that these errors will have adversely affected the damage location outcomes. This serves to highlight the importance of matching of the features predicted by the model (natural frequencies in this case) to those measured from the structure. Non-destructive structural modifications are advocated as one means of adding information during the mode-matching and validation exercises.

- Many options exist for the numerical modelling of damage, ranging from simple stiffness reduction to multi-parameter crack models. In this thesis stiffness reduction was the only method applied for prediction, with stiffness reduction having been shown to be adequate for damage detection and location [14]. However, for damage extent assessment the relationship between the damage extent and the parameter set used to characterise it must be established. In Chapter 4, this relationship was established through calibration of a model employing a stiffness-reduction damage model to one employing an element-removal damage scheme.

Data-based approaches allow SHM to be tackled as a problem of pattern recognition. A major advantage of taking a pattern recognition approach is that the need to develop and validate a numerical model is obviated. An overriding issue in applying such techniques, however, is the lack-of-data (or, more specifically lack-of-information) problem. This issue informs many of the conclusions drawn with regard to data-driven approaches.

- In an unsupervised learning mode, data is required from the structure in its undamaged state only. Acquiring such data should be feasible in real-world applications. However, even for this 'simple' case, caution should be taken. It appears intuitively correct that in order for a statistically discriminative model to generalise well to the classification of new data, it must be trained using data representative of the. full population from which new observations will be drawn. For example, for the study undertaken on an aircraft wing in Chapters 5-7, it was found that the features employed displayed sensitivity to relatively minor boundary condition changes. If a classifier were trained using

data gathered from a single representation of the boundary conditions, it would be unwise to expect it to perform well in cases where those boundary conditions change. It is recommended that time is spent at the *operational evaluation* [6] stage to identify non-damage factors that may affect the features used for damage identification. Having done so, a test sequence may be specified that allows a dataset that is representative of the conditions that will be observed in practice to be acquired. Design of experiments methods may be useful in constructing test sequences in a principled manner.

- In a supervised learning mode, data is required form the structure in both its undamaged and damaged states. This presents a rather more substantial lack-of-data issue than that experienced for the unsupervised learning mode. Data from the structure in its damaged state will rarely be available in practice. The approach developed in Chapters 5-7 is intended to reduce the amount of data required for the multiple-location damage problem, and is discussed below. However the problem of sourcing even single-site damage data remains. It is difficult to see how these problems may be overcome in a purely data-based paradigm; options that are currently under investigation include the use of pseudo-faults to artificially induce damage into the structure, aiding feature selection [39].

- An advantage of taking a data-based approach is that the pattern recognition scheme may be applied to a broad range of features. In the case presented in this thesis, transmissibility data were employed as raw data, and reduced to low-dimensional features using methods from outlier analysis. An initial aim during development of the wingbox model was to make predictions of transmissibility changes due to. This proved extremely challenging, in particular when seeking to compensate for inevitable test-model discrepancies. As a result, natural frequencies were ultimately employed as the primary feature for all the model-based work presented. The choice of feature type for model-based approaches is thus predicated to some extent by the ability to match features from the model to those from the experiment with confidence. This restriction is removed for the data-driven approach, allowing a broad variety of features to be evaluated based upon their information content.

A possible application of the hybrid approach is to couple a predictive law-based model with a statistical means of damage inference. The work presented in Chapter 4 focused on approaches to the predictive task. Damage identification on the basis of experimental data supplemented by model predictions remains an outstanding problem.

- Sensitivity analysis is a powerful tool for making the best use of law-based models. It may be applied both for guiding model development and for assessing the outputs of the 'finished' model. A preliminary investigation of the applications of sensitivity analysis for assessing feature 'robustness' was presented in the first case study of Chapter 4. Further investigation of how tools from sensitivity analysis may be applied to the model-based SHM problem form a major element of the future work recommended below.

- Uncertainty propagation is a further useful tool for the comparison of model outputs to experimental data and, potentially, for generating data to supplement experimental measurements. It should be noted that extending the use of predictive modelling to generate data suitable for training a statistical classifier would appear to be a difficult task, and would require a highly-trusted model to be developed. The computational expense of large numbers of model runs is a further consideration; however, this cost should be seen viewed in the context of the cost of the entire model development process. The cost associated with building and validating a given model may be considerable when compared with the cost of its execution.

9.2 Limitations of the presented work

Several limitations of the work presented (as opposed to the approaches investigated) are apparent. For the model updating study presented in Chapters 2 and 3, no significant 'new' contribution was added to that available in the literature. Important issues such as those related to larger parameter sets were not addressed, and issues of regularisation and iterative solution of the updating problem were largely circumvented.

The data-based approach presented in Chapters 5-7 dealt with a multiple-damage location problem, and the outcomes were deemed a success. However, the investigation was limited to a case which featured relatively large extents of damage,

and a small set of damage locations (five in total). No general rules regarding the validity of the approach were sought or established, and further work is required in this regard.

Only the first step in a possible 'hybrid' approach to SHM was covered in this chapter, and tools developed for the general non-deterministic verification and validation task appear to offer great promise for this stage. However, similar tools for implementing the statistical damage inference task are, as yet, elusive. Approaching this stage will likely require a re-evaluation of tools available from the fields of data-modelling and statistics.

Finally it should be stated that it would seem unwise to attempt to draw general conclusions regarding a very broad field of research on the basis of the small number of case studies presented here. However, it is felt that an interesting overview of aspects of the main approaches to SHM has been presented, and stimulus has been provided for future work in numerous directions.

9.3 Future work

The areas for future work highlighted here relate predominantly to the data-based and hybrid approaches- little future work relating to model updating-based methods is identified as a result of the study undertaken in thesis.

Further work for regarding the data-driven approach presented relates broadly to two questions: for what conditions is the approach valid? And were appropriate choices made with regard to the statistical classification method employed? As highlighted above, the case considered featured relatively large extents of damage and a small set of damage locations. The concern is that the degree of success observed will not hold for more challenging cases: for example, at lower levels of damage, higher numbers of locations, in the presence of environmental variability etc. Crucially, it should be investigated whether general rules can be learnt regarding the behaviour of the structure. It has been suggested that for the case investigated, the nature of the structure may have played a role in 'localising' the behaviour of some features to particular, bounded regions of the structure, and consequently making the multiple damage location problem tractable. The priority for future work related to this portion of the work undertaken should be to establish whether this

notion is valid. If so, an extension may be to investigate whether other sub-structures that exhibit comparable similar behaviour exist.

The second area for investigation is whether SVMs represented the most appropriate choice for classification, and if decisions related to the application of these classifiers were optimal. Several options for the architecture of multi-class SVMs exist, and further work with regard to the validation of hyperparameters may be of benefit. Further, comparison to other classification approaches was not pursued in this instance. While the particular SVM approach adopted worked well for the wing example, it is expected that the specifics of the classification approach adopted will grow in importance as the difficulty of the problem addressed increases.

The largest source of future work, however, is felt to be aspects of the hybrid approach. Realising a hybrid approach to the SHM task that follows the scheme discussed in this thesis will require significant further development of both methods for predictive modelling and for the statistical inference of damage.

The issue of handling uncertainty in complex numerical models is one that is pervasive across many fields of research. This has led to a renewed cross-disciplinary focus on approaches to numerical modelling and the handling of the uncertainties that arise, and the consequent development of tools to this end. For example, initial applications of Gaussian process (GP) emulation to problems in structural dynamics have been investigated and demonstrated to be effective in identifying factors to which model predictions are sensitive [54]. Emulation allows a fast-running surrogate of the full numerical model to be created using a limited (though carefully selected) number of model runs. This surrogate will typically run in several orders of magnitude less time than the original model, with the further defining characteristic that it is statistical (as opposed to law-based) in nature. Once created & validated, the emulator may be applied to tasks such as uncertainty propagation and global sensitivity analysis to an extent that would be precluded on time grounds using the original, slower-running model. Further, being informed by the numerical model yet statistical in nature, the emulator may be a suitable medium with which to tackle the inverse damage identification step. Finally, emulation may be a facilitating technology in approaches to the issue of test-model bias.

It is foreseen that by approaching the SHM problem through systematic application of the methods highlighted above, which have been developed and validated to a greater or lesser extent in other fields, a proper treatment of experimental and modelling uncertainty in SHM may be achieved.

References

[1] Fritzen CP, Kraemer P. Self-diagnosis of smart structures based on dynamical properties. Mechanical Systems and Signal Processing. 2009;23(6):1830-45.

[2] Rytter A. Vibration Based Inspection of Civil Engineering Structures [PhD Dissertation]: Aalborg University, Denmark; 1993.

[3] Sohn H. Effects of environmental and operational variability on structural health monitoring. Philosophical Transactions of the Royal Society a-Mathematical Physical and Engineering Sciences. 2007 Feb;365(1851):539-60.

[4] Doebling SW, Farrar CR, Prime MB, Shevitz DW. Damage identification and health monitoring of structural and mechanical systems from changes in their vibration characteristics: A literature review. Technical Report. Los Alamos, NM: Los Alamos National Lab; 1996.

[5] Sohn H, Farrar CR, Hemez FM, Shunk DD, Stinemates DW, Nadler BR. A Review of Structural Health Monitoring Literature: 1996-2001. Technical Report. Los Alamos, NM: Los Alamos National Laboratory; 2003. Report No.: LA-13976-MS.

[6] Farrar CR, Duffey TA, Doebling SW, Nix DA. A statistical pattern recognition paradigm for vibration-based structural health monitoring. Structural Health Monitoring 2000. 1999:764-73.

[7] Farrar CR, Worden K. Structural health monitoring - Preface. Philosophical Transactions of the Royal Society a-Mathematical Physical and Engineering Sciences. 2007 Feb;365(1851):299-301.

[8] Adams DE. Health monitoring of structural materials and components : methods with applications. Chichester: John Wiley & Sons Ltd. 2007.

[9] Balageas D, Fritzen C-P, Güemes A, eds. Structural Health Monitoring. 1 ed: ISTE Ltd 2006.

[10] Boller C, Chang F-K, Fujino Y, eds. Encyclopedia of Structural Health Monitoring: John Wiley and Sons 2009.

[11] Farrar CR, Lieven NAJ. Damage prognosis: the future of structural health monitoring. Philosophical Transactions of the Royal Society a-Mathematical Physical and Engineering Sciences. 2007 Feb;365(1851):623-32.

[12] Teughels A, De Roeck G. Structural damage identification of the highway bridge Z24 by FE model updating. Journal of Sound and Vibration. 2004;278(3):589-610.

[13] Moaveni B, Conte JP, Hemez FM. Uncertainty and Sensitivity Analysis of Damage Identification Results Obtained Using Finite Element Model Updating. Computer-Aided Civil and Infrastructure Engineering. 2009;24(5):320-34.

[14] Friswell MI. Damage identification using inverse methods In: Morassi A, Vestroni F, eds. *Dynamic Methods for Damage Detection (CISM Courses and Lectures vol 499)*. New York: Springer Wien 2008:13-66.

[15] Worden K, Friswell MI. Modal vibration-based damage identification. In: Boller C, Chang F-K, Fujino Y, eds. *Encyclopedia of Structural Health Monitoring*: John Wiley and Sons 2009.
[16] Fritzen CP, Jennewein D, Kiefer T. Damage detection based on model updating methods. Mechanical Systems and Signal Processing. 1998;12(1):163-85.
[17] Kaouk M, Zimmerman DC. Structural damage assessment using a generalized minimum rank perturbation theory. AIAA journal. 1994;32(4):836-42.
[18] Friswell MI, Mottershead JE, Ng GHT, Smart MG. Parameterisation of joints and constraints for finite element model updating. In: Wang KW, Yang B, Sun JQ, Seto K, Yoshida K, et al., editors. American Society of Mechanical Engineers, Design Engineering Division (Publication) DE; 1995; Boston, MA, USA; 1995.
[19] Friswell MI, Penny JET. Crack Modeling for Structural Health Monitoring. Structural Health Monitoring. 2002;1(2):139-48.
[20] Titurus B, Friswell MI, Starek L. Damage detection using generic elements: Part II. Damage detection. Computers & Structures. 2003;81(24-25):2287-99.
[21] Link M, Weiland M. Damage identification by multi-model updating in the modal and in the time domain. Mechanical Systems and Signal Processing. 2009;23(6):1734-46.
[22] Worden K, Manson G. The application of machine learning to structural health monitoring. Philosophical Transactions of the Royal Society a-Mathematical Physical and Engineering Sciences. 2007 Feb;365(1851):515-37.
[23] Worden K, Manson G, Fieller NRJ. Damage detection using outlier analysis. Journal of Sound and Vibration. 2000;229(3):647-67.
[24] Kullaa J. Damage detection of the Z24 bridge using control charts. Mechanical Systems and Signal Processing. 2003;17(1):163-70.
[25] Deraemaeker A, Reynders E, De Roeck G, Kullaa J. Vibration-based structural health monitoring using output-only measurements under changing environment. Mechanical Systems and Signal Processing. 2008;22(1):34-56.
[26] Worden K, Lane AJ. Damage identification using support vector machines. Smart Materials & Structures. 2001 Jun;10(3):540-7.
[27] Worden K, Manson G. Damage identification using multivariate statistics: Kernel discriminant analysis. Inverse Problems in Engineering. 2000;8(1):25-46.
[28] Görl E, Link M. Damage identification using changes of eigenfrequencies and mode shapes. Mechanical Systems and Signal Processing. 2003;17(1):103-10.
[29] Friswell MI, Mottershead JE. Finite Element Model Updating in Structural Dynamics. Dordrecht: Kluwer Academic Publishers 1995.
[30] Link M. Updating of Analytical Models - Basic Procedures and Extensions. In: Silva JMM, Maia NMM, eds. *Modal Analysis and Testing* Springer 1999.
[31] Friswell MI. Damage identification using inverse methods. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences. 2007;365(1851):393-410.
[32] Hansen PC, Oleary DP. The Use of the L-Curve in the Regularization of Discrete Ill-Posed Problems. Siam Journal on Scientific Computing. 1993;14(6):1487-503.
[33] Friswell MI, Mottershead JE, Ahmadian H. Finite-element model updating using experimental test data: parametrization and regularization. Philosophical Transactions of the Royal Society of London Series A-Mathematical Physical and Engineering Sciences. 2001;359(1778):169-86.
[34] Göge D, Link M. Assessment of computational model updating procedures with regard to model validation. Aerospace Science and Technology. 2003;7(1):47-61.

[35] Worden K, Manson G. Experimental validation of a structural health monitoring methodology: Part I. Novelty detection on a laboratory structure. Journal of Sound and Vibration. 2003 Jan;259(2):323-43.
[36] Imamovic N. Validation of large structural dynamics models using modal test data: PhD thesis: Imperial College London; 1998.
[37] Weiland M, Link M. Matfem 99 User's Guide, Version 26-May-2000: University of Kassel 2000.
[38] Barthorpe RJ, Manson G, Worden K. An Investigation into the Necessary Model Fidelity for SHM Feature Selection. *4th European Workshop on SHM*. Krakow 2008.
[39] Papatheou E, Manson G, Barthorpe RJ, Worden K. The use of pseudo-faults for novelty detection in SHM. Journal of Sound and Vibration. 2010;329(12):2349-66.
[40] Ahmadian H, Mottershead JE, Friswell MI. Physical Realization of Generic-Element Parameters in Model Updating. Journal of Vibration and Acoustics. 2002;124(4):628-33.
[41] Ibrahim RA, Pettit CL. Uncertainties and dynamic problems of bolted joints and other fasteners. Journal of Sound and Vibration. 2005;279(3-5):857-936.
[42] Mottershead JE, Friswell MI, Ng GHT, Brandon JA. Geometric parameters for finite element model updating of joints and constraints. Mechanical Systems and Signal Processing. 1996;10(2):171-82.
[43] Friswell MI, Penny JET, Garvey SD. Parameter subset selection in damage location. Taylor & Francis 1997:189 - 215.
[44] Ahmadian H, Mottershead JE, Friswell MI. Regularisation methods for finite element model updating. Mechanical Systems and Signal Processing. 1998;12(1):47-64.
[45] Hemez FM, Doebling SW. Validation of structural models at Los Alamos Natioanl Labaoratory. 41st AIAA/ASME/ASCE/AHS/ASC Structural dynamics and Materials Conference; 2000; Atlanta, Georgia; 2000.
[46] Hemez FM. Uncertainty Quantification and the Verification and Validation of Computational Models. In: Inman DJ, Farrar CR, Lopes Jr V, Steffen Jr V, eds. *Damage Prognosis for Aerospace, Civil and Mechanical Systems*. London, UK: John Wiley and Sons Ltd. 2004.
[47] Schwer LE. An overview of the PTC 60/V&V 10: Guide for verification and validation in computational solid mechanics: Transmitted by L. E. Schwer, Chair PTC 60V&V 10. Engineering with Computers. 2007;23(4):245-52.
[48] Rebba R, Mahadevan S. Validation of models with multivariate output. Reliability Engineering and System Safety. 2006;91(8):861-71.
[49] Rebba R, Mahadevan S, Huang S. Validation and error estimation of computational models. Reliability Engineering and System Safety. 2006;91(10-11):1390-7.
[50] Kardomateas GA. The initial post-buckling and growth behavior of internal delaminations in
composite plates. Journal of Applied Mechanics (ASME). 1993(60):903-10.
[51] Roache PJ. Verification and validation in computational science and engineering. Albuquerque, N.M.: Hermosa Publishers 1998.
[52] Press WH. Numerical recipes : the art of scientific computing. 3rd ed. ed. Cambridge: Cambridge University Press 2007.
[53] Saltelli A, Chan K, Scott EM. Sensitivity analysis. Chichester; New York: Wiley 2000.
[54] Becker W, Rowson J, Oakley J, Yoxall A, Manson G, Worden K. Bayesian Sensitivity Analysis of a Large Nonlinear Model. *ISMA2008 International Conference on Noise and Vibration Engineering*. Leuven, Belgium 2008.

[55] Zhang B, Waters TP, Mace BR. Identifying joints from measured reflection coefficients in beam-like structures with application to a pipe support. Mechanical Systems and Signal Processing.24(3):784-95.
[56] MacKay DJC. Information theory, inference and learning algorithms. Cambridge, UK: Cambridge University Press 2003.
[57] O'Hagan A. Bayesian analysis of computer code outputs: A tutorial. Reliability Engineering and System Safety. 2006;91(10-11):1290-300.
[58] Papatheou E, Barthorpe RJ, Worden K, Manson G. Fault induction using added masses for structural damage identification. *ISMA2008 International Conference on Noise and Vibration Engineering*. Leuven, Belgium 2008.
[59] Yi WJ, Zhou Y, Kunnath S, Xu B. Identification of localized frame parameters using higher natural modes. Engineering Structures. 2008;30(11):3082-94.
[60] Hemez FM, Wilson AC, Doebling SW. Design of computer experiments for improving an impact test simulation. *IMACXIX, International Modal Analysis Conference*. Kissimmee, FL 2001:977-85.
[61] Saltelli A, Ratto M. Global sensitivity analysis : the primer. Chichester: John Wiley & Sons, Ltd. 2008.
[62] Moens D, Vandepitte D. A survey of non-probabilistic uncertainty treatment in finite element analysis. Computer Methods in Applied Mechanics and Engineering. 2005;194(12-16):1527-55.
[63] Worden K, Manson G, Lord TM, Friswell MI. Some observations on uncertainty propagation through a simple nonlinear system. Journal of Sound and Vibration. 2005;288(3):601-21.
[64] Hedayat A, Sloane NJA, Stufken J. Orthogonal arrays: theory and applications. London: Springer 1999.
[65] Doebling SW, Hemez FM. Overview of Uncertainty Assessment for Structural Health Monitoring. 3rd International Workshop on Structural Health Monitoring; 2001; Stanford, California; 2001.
[66] Fang K, Li R, Sudjianto A. Design and modeling for computer experiments. Boca Raton, Fla. ; London: Chapman & Hall/CRC 2006.
[67] Helton JC, Davis FJ. Latin hypercube sampling and the propagation of uncertainty in analyses of complex systems. Albuquerque, NM: Sandia National Laboratories; 2002. Report No.: SAND2001-0417.
[68] Kennedy M. GEM-SA. 1.1 ed: Centre for Terrestrial Carbon Dynamics 2005.
[69] Ruotolo R, Surace C. Damage assessment of multiple cracked beams: Numerical results and experimental validation. Journal of Sound and Vibration. 1997;206(4):567-88.
[70] Contursi T, Messina A, Williams EJ. A multiple-damage location assurance criterion based on natural frequency changes. JVC/Journal of Vibration and Control. 1998;4(5):619-33.
[71] Guo HY, Li ZL. A two-stage method to identify structural damage sites and extents by using evidence theory and micro-search genetic algorithm. Mechanical Systems and Signal Processing. 2009;23(3):769-82.
[72] Lin R-J, Cheng F-P. Multiple crack identification of a free-free beam with uniform material property variation and varied noised frequency. Engineering Structures. 2008;30(4):909-29.
[73] Manson G, Worden K, Allman D. Experimental validation of a structural health monitoring methodology. Part II. Novelty detection on a Gnat aircraft. Journal of Sound and Vibration. 2003 Jan;259(2):345-63.
[74] Manson G, Worden K, Allman D. Experimental validation of a structural health monitoring methodology. Part III. Damage location on an aircraft wing. Journal of Sound and Vibration. 2003 Jan;259(2):365-85.

[75] Farrar CR, Doebling SW, Cornwell PJ, Straser EG. Variability of modal parameters measured on the Alamosa Canyon Bridge. Proceedings of the International Modal Analysis Conference - IMAC. 1997;1:257-63.
[76] Peeters B, De Roeck G. One-year monitoring of the Z24-bridge: Environmental effectsversus damage events. Earthquake Engineering and Structural Dynamics. 2001;30(2):149-71.
[77] Kullaa J, Heine T. Feature comparison in structural health monitoring of a vehicle crane. Shock and Vibration. 2008;15(3-4):207-15.
[78] Manson G, Papatheou E, Worden K. Genetic optimisation of a neural network damage diagnostic. Aeronautical Journal. 2008;112(1131):267-74.
[79] Worden K, Manson G, Allman D. Experimental validation of a structural health monitoring methodology: Part I. Novelty detection on a laboratory structure. Journal of Sound and Vibration. 2003 Jan;259(2):323-43.
[80] Maia NMM, Silva JMM, Ribeiro AMR. Transmissibility concept in multi-degree-of-freedom systems. Mechanical Systems and Signal Processing. 2001;15(1):129-37.
[81] Bishop CM. Neural Networks for Pattern Recognition 1995.
[82] Fawcett T. An introduction to ROC analysis. Pattern Recognition Letters. 2006;27(8):861-74.
[83] Hanley JA, McNeil BJ. The Meaning and Use of the Area under a Receiver Operating Characteristic (Roc) Curve. Radiology. 1982;143(1):29-36.
[84] Vapnik V. The nature of statistical learning theory. New York ; London: Springer 1995.
[85] Christianini N, Shawe-Taylor J. An Introduction to Support Vector Machines and other Kernel-based Learning Methods: Cambridge University Press 2000.
[86] Robinson J. Support Vector Machine Learning. Saarbrucken: VDM Verlag Dr. Muller Aktiengesellschaft & Co. KG 2008.
[87] Schölkopf B, Smola AJ. Learning with Kernels. Cambridge, Mass.: MIT Press Ltd 2002.
[88] Gunn SR. Support vector machines for classification and regression: School of Electronics and Computer Science, University of Southampton; 1998.
[89] Burges CJC. A tutorial on Support Vector Machines for pattern recognition. Data Min Knowl Discov. 1998 Jun;2(2):121-67.
[90] Hsu CW, Lin CJ. A comparison of methods for multiclass support vector machines. IEEE Transactions on Neural Networks. 2002;13(2):415-25.
[91] Gunn SR. Matlab Support Vector Machine Toolbox. 2.1 ed: School of Electronics and Computer Science, University of Southampton 2001
[92] Worden K, Dulieu-Barton JM. An overview of intelligent fault detection in systems and structures. Structural Health Monitoring. 2004 Mar;3(1):85-98.
[93] Farrar CR, Worden K. An introduction to structural health monitoring. Philosophical Transactions of the Royal Society a-Mathematical Physical and Engineering Sciences. 2007 Feb;365(1851):303-15.
[94] Worden K, Farrar CR, Dulieu-Barton JM. Principles of structural degradation monitoring. In: Boller C, Chang F-K, Fujino Y, eds. *Encyclopedia of Structural Health Monitoring*: John Wiley and Sons 2009.
[95] Saltelli A, Ratto M, Andres T, Campolongo F, Cariboni J, Gatelli D, et al. Global Sensitivity Analysis: The Primer. Chichester: John Wiley & Sons, Ltd. 2008.
[96] Schuster H-G. Deterministic Chaos: John Wiley and Sons Australia Ltd 1998.
[97] Colemanballs. [cited 20/08/2010]; Available from: http://en.wikipedia.org/wiki/Colemanballs
[98] Ghanem RG, Spanos PD. Stochastic Finite Elements: A Spectral Approach: Dover Publications Inc 2003.

[99] Lallemand B, Cherki A, Tison T, Level P. Fuzzy modal finite element analysis of structures with imprecise material properties. Journal of Sound and Vibration. 1999;220(2):353-64.
[100] Kammer DC. Sensor placement for on-orbit modal identification and correlation of large space structures. Journal of Guidance, Control, and Dynamics. 1991;14(2):251-9.
[101] Shah PC, Udwadia FE. A methodology for optimal sensor locations for identification of dynamic systems. Journal of Applied Mechanics, Transactions ASME. 1978;45(1):188-96.
[102] Papadimitriou C. Optimal sensor placement methodology for parametric identification of structural systems. Journal of Sound and Vibration. 2004;278(4-5):923-47.
[103] Shi ZY, Law SS, Zhang LM. Optimum sensor placement for structural damage detection. Journal of Engineering Mechanics. 2000;126(11):1173-9.

Appendix A: Sensor Placement Optimisation

The basic problem of fault detection is to deduce the existence of a defect in a structure from measurements taken at sensors distributed on the structure. The quality of these measurements and thus the quality of structural health monitoring achieved is to a large extent dependent upon where sensors are placed on the structure. Cost and practicality issues preclude the instrumentation of every point of interest on the structure and lead us to select a smaller set of measurement locations. A large variety of performance indices have been developed for the problem of sensor placement, but it is only comparatively recently that the problem as been considered from an SHM perspective. Some methods require a single calculation to be performed, some are iterative, and many others take the form of an objective function to which an optimisation technique must be applied. Here, the theoretical bases underlying several approaches to the problem of sensor placement optimisation for SHM are presented. The material presented is drawn from an extended piece of work on the topic of sensor placement (ref encyclopaedia). The descriptions have been kept intentionally mathematically light: full descriptions of the algorithms may be found in the references.

Effective independence (EI)

The method of Effective independence (EI or EFI) was introduced by [100] based on earlier work in [101]. It makes use of the Fisher information matrix (FIM), which offers a measure of the information a sampled random variable contains about an unknown parameter: formally, Fisher information is the variance of the score with respect to the unknown parameter. Where there are multiple unknown parameters it may be stated in matrix form with elements

$$\big(I(\theta)\big)_{ij} = E\left[\frac{\partial}{\partial\theta_i}\ln f(x;\theta)\,\frac{\partial}{\partial\theta_j}\ln f(x;\theta)\right] \tag{A.1}$$

Where $\theta$is the vector of unknown parameters; $\big(I(\theta)\big)_{ij}$is the Fisher information with respect to the unknown parameters θ_i and θ_j; x is the sampled random variable; $f(x;\theta)$=$L(\theta)$ is the likelihood function of θ; and $E[\dots]$ denotes the expectation.

For the SPO problem the target modeshapes may be regarded as the unknown, sought parameters, with the sampled data being that available from the given sensor distribution. Every DOF in the candidate set is ranked according to its contribution to the determinant of the FIM, and the lowest ranked DOF is eliminated. The new, reduced set is then re-ranked, and the process repeated in an iterative manner until the desired number of sensors remains. This is adopted as the optimal measurement set. Maintaining the determinant of the FIM leads to the selection of a set of sensor locations for which the mode shapes of interest are as linearly independent as possible, while retaining sufficient information about the target modal responses. The approach is based on the EI distribution vector E_D, defined as the diagonal of the matrix

$$E = \Psi_{\mathrm{fs}}\left\{\Psi_{\mathrm{fs}}^{\mathrm{T}}\Psi_{\mathrm{fs}}\right\}^{-1}\Psi_{\mathrm{fs}}^{\mathrm{T}} \tag{A.2}$$

where Ψ_{fs} is a matrix of FE target modes partitioned according to a given sensor distribution. Each diagonal element is the fractional contribution of each sensor location to the rank of E, which can only be full rank if the target mode partitions are linearly independent. The algorithm is iterative; at each step, terms in E_D are sorted to give the least important sensor which is then deleted. The corresponding elements in Ψ_{fs} are also deleted. The iteration concludes when the required number of sensors is obtained.

Average driving point residue (ADPR)

A drawback of the EI approach is that the algorithm can select sensor locations that display low signal strength, making the system vulnerable to noisy conditions. The average driving point residue (ADPR) offers a measure of the contribution of any point to the overall modal response. If *N* modes of interest are to be measured, the ADPR can be calculated from FE data as

$$ADPR_j = \sum_{i=1}^{N} \frac{\Psi_{ij}^2}{\omega_i} \tag{A.3}$$

Where Ψ_{ij} is the j^{th} element of the i^{th} mode shape and ω_i is the i^{th} modal frequency

Effective independence driving point residue (EI-DPR)

The values given by the EI algorithm are weighted by the ADPR values to give the EFI-DPR vector

$$E_{D_j} = E_D ADPR_j \tag{A.4}$$

This adaptation leads to a greater likelihood of sensors being placed in areas of high signal strength. In addition to improving signal to noise ratios, this tends to result in selection of relatively uniformly spaced sensor locations.

Kinetic energy method (KE)

The KE method assumes that the sensors will have maximum observability of the modes of interest if the sensors are placed at points of maximum kinetic energy for that mode, and accordingly ranks sensor locations based on their dynamic contribution to the target mode shapes. It follows a similar procedure to that used in the EI method, the key difference being that a kinetic energy measure, rather than the determinant of the FIM, is maximised. It is alternatively known as the Modal Kinetic Energy (MKE) method or the Kinetic Energy Method (KEM) in the literature.

KE indices are calculated for all candidate sensor locations as follows

$$KE_{pq} = \Psi_{pq} \sum_{s=1}^{n} M_{ps} \Psi_{sq} \tag{A.5}$$

where KE is the kinetic energy matrix; Ψ is the matrix of target mode shapes; KE_{pq} is the kinetic energy contribution of the DOF to the target mode shape and M_{ps} is the corresponding mass matrix element.

The kinetic energy of each candidate DOF is thus weighted by the corresponding component from the mass matrix M. The sensor locations that offer the highest KE indices are selected as the measurement locations. A further option is the Average Kinetic Energy (AKE) method, which places sensors according to KE averaged over

all modes of interest; this avoids bias against placing sensors at nodes of particular modes.

As the method selects those sensor locations with the largest available signal amplitudes signal-to-noise ratios tend to be high, making the method attractive for use in noisy conditions. However, in contrast to EI the KE method does not consider the linear independence of the target modes, an important consideration for both modal identification and test-analysis correlation.

Eigenvalue vector product (EVP)

The EVP method computes the product of the eigenvector components for candidate sensor location for the range of modes to be measure N. The maximum for this product is a candidate measurement point. Some modification may be required if a point is a node of one of the modes. The EVP is calculated as

$$EVP = \prod_{j=1}^{N} \left| \Psi_{ij} \right| \tag{A.6}$$

Mutual information

Mutual information gives a measure of how much information one sensor location 'learns' from another. If there are two sets of measurement locations, A and B, the amount of information learned by a_i about b_j is represented by the mutual information,

$$I(a_i, b_j) = \log_2 \left[\frac{P_{AB}(a_i, b_j)}{P_A(a_i) P_B(b_j)} \right] \tag{A.7}$$

where a_i and b_j are the measurements from locations A and B respectively; $P_A(a_i)$ and $P_B(b_j)$ are the individual probability densities for A and B and $P_{AB}(a_i, b_j)$ is the joint probability density for A and B.

If the measurement of a_i is completely independent of the measurement of b_j, $I(a_i, b_j)$ becomes zero. The average mutual information between A and B is calculated by averaging over all the sensor locations, and the optimal sensor location determined by minimising the mutual information between sensors.

Information entropy method

Optimal sensor placement is achieved by minimising the change in the information entropy $H(D)$, which is given by

$$H(D) = E_\theta[-\ln p\,(\theta|D)] = -\int p(\theta|D)\ln p\,(\theta|D)d\theta \tag{A.8}$$

where θ is the uncertain parameter set; D is the dynamic test data; and E_θ is the mathematical expectation with respect to θ. A rigorous mathematical description is given in [102]

Sensitivity based methods

In the SHM specific method proposed by [103], the matrix E used in the EI method is adapted to use the sensitivity matrix developed for damage location. The modified matrix is given by

$$E = F(K)\{F(K)^{\mathrm{T}}F(K)\}^{-1}F(K)^{\mathrm{T}} \tag{A.9}$$

where $F(K)$ is the vector of sensitivity coefficients of the mode shape changes with respect to a damage vector.

As for the EI approach, the diagonal terms of the matrix E provide the fractional contribution of the corresponding measurement location to the rank of E. The location that contributes least is removed, and the process is repeated iteratively until the required quantity of sensors remains.

Appendix B: Publications

Edited book chapters

R.J.Barthorpe, and K.Worden (2009) Sensor placement optimisation for structural health monitoring. *Encyclopedia of Structural Health Monitoring eds. C. Boller, F-K. Chang, Y. Fujino, John Wiley & Sons Ltd.*

Journal papers

E.Papatheou, K.Worden, G.Manson & **R.J.Barthorpe** (2010) The use of pseudo faults for SHM feature selection and pattern recognition. *Journal of Sound and Vibration, Vol. 329, Issue 12, p. 2349-2366.*

A.Deraemaeker, A.Preumont, E.Reynders, G.DeRoeck, J.Kullaa, V.Lämsä, K.Worden, G.Manson, **R.J.Barthorpe**, E.Papatheou, P.Kudela, P.Malinowski, W,Ostachowicz & T.Wandowski (2010) Vibration-based structural health monitoring using large sensor networks. *Smart Structures and Systems, Vol. 6, No. 3, p.335-347*

Conference papers

E.Papatheou, G.Manson, **R.J.Barthorpe** & K.Worden (2010) Damage location using added masses in a Piper Tomahawk aircraft wing. *Proceedings of ISMA2010 International Conference on Noise and Vibration Engineering, Leuven, Belgium*

R.J.Barthorpe, K.Worden, C.Surace, G.Demarie,(2009) A comparative study of approaches to damage detection. *Proceedings of XIX Congresso AIMETA, Ancona, Italy*

G.Manson, **R.J.Barthorpe**, (2009) Multiple damage location using independent features. *Proceedings of IWSHM2009, 7th International Workshop on Structural Engineering Dynamics, Stanford, California*

J.Hensman, **R.J.Barthorpe** (2009) Feature extraction from spectral data using the Bayesian evidence framework. *Proceedings of DAMAS2009, 8th International Conference on Damage Assessment of Structures, Beijing, China*

R.J.Barthorpe, G.Manson, K.Worden (2009) A forward approach to model-based structural health monitoring. *Proceedings of ICEDyn2009, International Conference on Structural Engineering Dynamics, Ericeira, Portugal*

G.Manson, **R.J.Barthorpe** (2009) Advanced feature selection for simplified pattern recognition within the damage identification framework. *Proceedings of ICEDyn2009, International Conference on Structural Engineering Dynamics, Ericeira, Portugal*

R.J.Barthorpe (2009) Identification of robust, damage sensitive features for model-based structural health monitoring: an effect screening approach. *Proceedings of USD2009 2nd International Conference on Uncertainty in Structural Dynamics, Sheffield, UK p.301-10*

R.J.Barthorpe, K.Worden, G.Manson (2009) Finite element model-based feature generation for structural health monitoring. *Proceedings of IMAC XXVII International Conference on Modal Analysis, Florida*

E.Papatheou, **R.J.Barthorpe**, K.Worden, G.Manson (2008) Fault induction using added masses for structural damage identification. *Proceedings of ISMA2008 International Conference on Noise and Vibration Engineering, Leuven, Belgium*

R.J.Barthorpe, K.Worden, G.Manson (2008) An investigation into the necessary model fidelity for SHM feature selection. *Proceedings of the Fourth European Workshop on SHM, Krakow, Poland p.980-9*

E.Papatheou, K.Worden, G.Manson, **R.J.Barthorpe** (2008) The use of pseudo faults for SHM feature selection and pattern recognition. *Proceedings of the Fourth European Workshop on SHM, Krakow, Poland. p.1104-12*