1

Introduction

Modern societies are heavily dependent upon structural and mechanical systems such as aircraft, bridges, power generation systems, rotating machinery, offshore oil platforms, buildings and defence systems. Many of these existing systems are currently nearing the end of their original design life. Because these systems cannot be economically replaced, techniques for damage detection are being developed and implemented so that these systems can continue to be safely used if or when their operation is extended beyond the design basis service life. Also, in terms of the design and introduction of new engineering systems, these often incorporate novel materials whose long-term degradation processes are not well understood. In the effort to develop more cost-effective designs, these new systems may be built with lower safety margins. These circumstances demand that the onset of damage in new systems can be detected at the earliest possible time in an effort to prevent failures that can have grave life-safety and economic consequences.

Damage detection is usually carried out in the context of one or more closely related disciplines that include: structural health monitoring (SHM), condition monitoring (CM), nondestructive evaluation (NDE) – also commonly called nondestructive testing, or (NDT), health and usage monitoring system (HUMS), statistical process control (SPC) and damage prognosis (DP).

The term *structural health monitoring* (SHM) usually refers to the process of implementing a damage detection strategy for aerospace, civil or mechanical engineering infrastructure. This process involves the observation of a structure or mechanical system over time using periodically spaced dynamic response measurements, the extraction of damage-sensitive features from these measurements and the statistical analysis of these features to determine the current state of system health. For long-term SHM, the output of this process is periodically updated information regarding the ability of the structure to continue to perform its intended function in light of the inevitable ageing and degradation resulting from the operational environments. Under an extreme event, such as an earthquake or unanticipated blast loading, SHM could be used for rapid condition screening, to provide, in near real time, reliable information about the performance of the system during the event and about the subsequent integrity of the system.

Condition monitoring is analogous to SHM, but specifically addresses damage detection in rotating and reciprocating machinery, such as that used in manufacturing and power generation (Worden and Dulieu-Barton, 2004).

Both SHM and CM have the potential to be applied on-line, that is during operation of the system or structure of interest. In contrast, *nondestructive evaluation* (NDE) is usually carried out off-line after the site of the potential damage has been located. There are exceptions to this rule, as NDE is also used as a monitoring tool for *in situ* structures such as pressure vessels and rails. NDE is therefore primarily

Structural Health Monitoring: A Machine Learning Perspective, First Edition. Charles R. Farrar and Keith Worden.

used for damage characterisation and as a severity check when there is a priori knowledge of the damage location (Shull, 2002).

*Health and usage monitoring system*s (HUMSs) are closely related to CM systems, but the term has largely been adopted for the specific application to damage detection in rotorcraft drive trains (Samual and Pines, 2005). In that context, the health monitoring portion of the process attempts to identify damage, while the *usage monitoring* records the number of load cycles that the system experiences for the purposes of calculating fatigue life consumption.

Statistical process control (SPC) is process-based rather than structure-based and uses a variety of sensors to monitor changes in a process, with one possible cause of a change being structural damage (Montgomery, 2009).

Once damage has been detected, the term *damage prognosis* (DP) describes the attempt to predict the remaining useful life of a system (Farrar *et al*., 2003).

Condition monitoring, NDE and SPC are without doubt the most mature damage detection disciplines as they have made the transition from a research topic to actual engineering practice for a wide variety of applications. However, it is a widely held belief that SHM is in the process of making the transition into the application domain. This book will focus primarily on SHM as the authors believe that the time has come for a comprehensive and fundamental exposition of the basic principles of this branch of damage detection.

1.1 How Engineers and Scientists Study Damage

Materials scientists and engineers are the primary classes of technologists that study damage; in this, they commonly approach the problem by asking one or more of the following questions (in no particular order):

1. What is the cause of damage?
2. What can be done to prevent damage?
3. Once present, how are the effects of damage mitigated?
4. Is damage present?
5. How fast will the damage grow and exceed some critical level?

The answers to these questions will depend on whether one takes a material science point of view or an engineering point of view. As an example, the materials scientist may address question 1 by studying the initial imperfections at the grain boundary scale as shown in Figure 1.1 and attempt to develop tools that predict how these imperfections coalesce and grow under various loading conditions. They might also study properties such as surface finish that result from the manufacturing process or develop an understanding of material ageing and degradation processes at the micro-scale. In contrast, the engineer may attempt to establish allowable strength, deformation or stability criteria associated with the onset of damage. A materials scientist might approach the second question by designing new materials that are less susceptible to a particular type of damage (e.g. use of stainless steel in corrosive environments) while the engineer might incorporate alternate design strategies for manufacturability and reliability or prescribe operational and environmental limits for system use. Damage mitigation strategies might be accomplished with the development of self-healing materials, which is currently a focus of materials science research (Zwaag, 2007). Alternatively, engineers will prescribe maintenance and repair or limit operations (e.g. slow the speed of a vehicle) as a damage mitigation strategy.

Questions 4 and 5 are the focus of SHM and here the difference between how the material scientists and engineers address the problem is related to the length scale on which they study the problem. Additionally, a distinction arises based on the ability to do the damage assessment with the system in operation or if the assessment needs to be performed with the system in or out of service. More

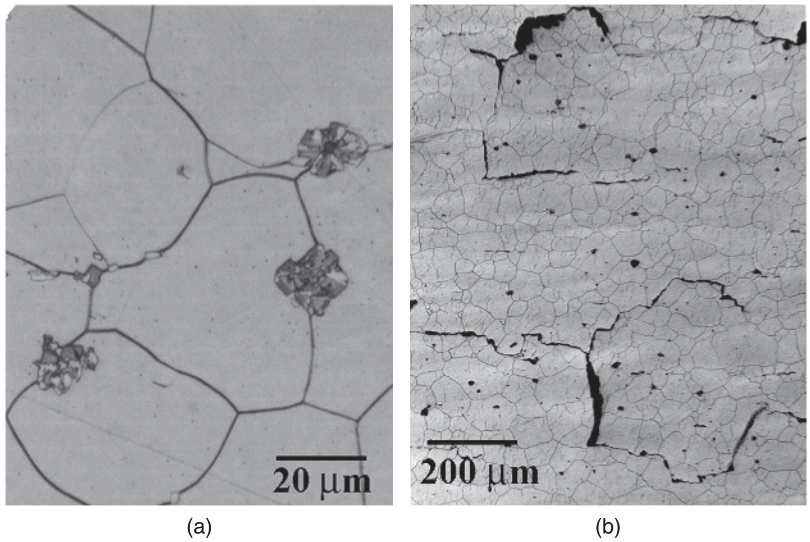

Figure 1.1 (a) Inclusions at the grain boundaries in U-6Nb. (b) A micrograph of a U-6Nb plate showing crack propagation along inclusion lines after shock loading (source: D. Thoma, Los Alamos National Laboratory).

drastically, the assessment may be carried out in a destructive manner. Materials scientists will often perform damage detection at the microscopic level using thin sectioning of the material to recreate a three-dimensional image of the microstructure. As previously mentioned, traditional NDE methods are applied to assess incipient macroscopic damage at the material and component level, typically with the system out of service. Wave propagation approaches to SHM, which can be used to assess damage with the system in operation, are also being used to assess incipient damage at the macroscopic material and component scale. Finally, other forms of SHM, like vibration-based approaches, can also be used to assess damage from the component to full system scale, as can CM, HUMS and SPC.

1.2 Motivation for Developing SHM Technology

Almost all private and government industries want to detect damage in their products as well as in their manufacturing infrastructure at the earliest possible time. Such detection requires these industries to perform some form of SHM and is motivated by the potential life-safety and economic impact of this technology. As an example, the semiconductor manufacturing industry is adopting this technology to help minimise the need for redundant machinery necessary to prevent inadvertent downtime in their fabrication plants. Such downtime can cost these companies on the order of millions of dollars per hour. Aerospace companies, along with government agencies in the United States, are investigating SHM technology for detection of damage to space shuttle control surfaces hidden by heat shields. Clearly, such damage detection has significant life-safety implications. Also, as an example from the civil engineering context, there are currently no quantifiable methods to determine if buildings are safe for reoccupation after a significant earthquake. SHM technology may one day provide a means of minimising the uncertainty associated with current visual post-earthquake damage assessments. The prompt reoccupation of buildings, particularly those associated with manufacturing, can significantly mitigate economic losses associated with major seismic events. Finally, many portions of our technical infrastructure are approaching or exceeding their initial design life. As a result of economic issues, these civil, mechanical and aerospace structures are being used in spite of ageing and the associated damage accumulation. Therefore, the ability to monitor the health of these structures is becoming increasingly important.

Maintenance philosophies have evolved to minimise the potential negative life-safety and economic impacts of unforeseen system failures. Initially, *run-to-failure* approaches to engineering system maintenance were used. With this approach the system is operated until some critical component fails and then that component is replaced. This procedure requires no investment in monitoring systems, but it can be extremely costly as failure can occur without warning. Clearly, this approach to maintenance is unacceptable when life-safety is a concern.

A more sophisticated maintenance approach that is used extensively today is referred to as *time-based maintenance*. This maintenance approach requires that critical components are serviced or replaced at predefined times or use intervals regardless of the condition of the component. A typical example is the recommendation that one changes the oil in their car after it has been driven a certain distance or at some prescribed time interval. This maintenance is done regardless of the condition of the oil. Another example is the requirement that a missile be retired after a certain number of captive-carry flight hours on the wing of an aircraft. Time-based maintenance is a more proactive approach than run-to-failure and it has made complex engineering systems such as commercial aircraft extremely safe. In some cases usage monitoring systems are deployed in conjunction with the time-based maintenance approach. Such a system might record the number of manoeuvres performed by a high-performance aircraft that exceed a certain threshold acceleration level. Maintenance would then be performed after the aircraft has accumulated some predefined number of these peak acceleration readings.

SHM is the technology that will allow the current time-based maintenance approaches to evolve into *condition-based maintenance* philosophies. The concept of condition-based maintenance is that a sensing system on the structure will monitor the system response and notify the operator that damage or degradation has been detected. Life-safety and economic benefits associated with such a philosophy will only be realised if the monitoring system provides sufficient warning such that corrective action can be taken before the damage or degradation evolves to some critical level. The trade-off associated with implementing such a philosophy is that it potentially requires more sophisticated monitoring hardware to be deployed on the system and more sophisticated data analysis procedures to interrogate the measured data.

Defence agencies are particularly motivated to develop SHM capabilities and to move to a condition-based maintenance philosophy in an effort to increase *combat asset readiness*. Military hardware is only effective if it is deployed for its combat mission. Minimising the maintenance intervals for the equipment maximises its availability for combat missions. Also, when such equipment is subjected to noncatastrophic damage from hostile fire, as shown in Figure 1.2, there is a need to rapidly assess the extent of this damage in an effort to make informed decisions about completing the current mission, about repair requirements and about subsequent use of the hardware.

Finally, many companies that produce high-capital-expenditure products such as airframes, jet engines and large construction equipment would like to move to a business model where they lease equipment as opposed to selling it. With these models the company that manufactures the equipment would take on the responsibilities for its maintenance. SHM has the potential to extend the maintenance cycles and, hence, keep the equipment out in the field where it can continue to generate revenues for the owner. Furthermore, the equipment owners would like to base their lease fees on the amount of system life used up during the lease time rather than on the current simple time-based lease fee arrangements. Such a business model will not be realised without the ability to monitor the damage initiation and evolution in the rental hardware.

1.3 Definition of Damage

In the most general terms, damage can be defined as changes introduced into a system, either intentionally or unintentionally, that adversely affect the current or future performance of that system. These

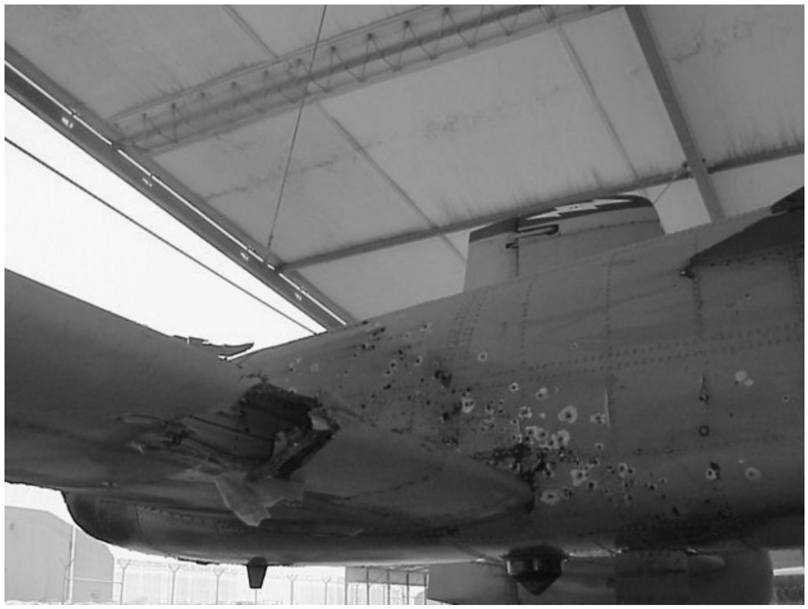

Figure 1.2 Damage sustained by an A-10 Thunderbolt during a 2003 Iraq War mission (source: US Air Force).

systems can be either natural or man-made. As an example, an anti-aircraft missile is typically fired to intentionally introduce damage that will immediately alter the flight characteristics of the target aircraft. Biological systems can be unintentionally subjected to the damaging effects of ionising radiation. However, depending on the levels of exposure, these systems may not show the adverse effects of this damaging event for many years or even future generations.

This book is focused on the study of damage identification in structural and mechanical systems. Therefore, *damage* will be defined as *intentional or unintentional changes to the material and/or geometric properties of these systems, including changes to the boundary conditions and system connectivity, which adversely affect the current or future performance of these systems*.

Thinking in terms of length scales, all damage begins at the material level as shown in Figure 1.1 and such material-level damage is present to some extent in all systems. Materials scientists and condensed-matter physicists commonly refer to such damage as *defects*; the term encompasses voids, inclusions and dislocations. Under appropriate loading scenarios, the material-level damage progresses to component- and system-level damage at various rates. *Failure* occurs when the damage progresses to a point where the system can no longer perform its intended function. Often failure is defined in terms of exceeding some strength, stability or deformation-related performance criterion.

Clearly, even though damage is present in all engineered systems at some level, modern design practices can account for this low-level damage and the systems perform as intended. The structure can often continue to perform its intended function when damage has progressed beyond the levels considered in design, but usually this performance is at some reduced level. As an example, the aircraft shown in Figure 1.2 was able to return to its base despite the severe damage it sustained. However, it is doubtful if it could perform at its original design levels during that return flight. Also, it may be some time before the structure experiences the appropriate loading conditions for the damage to cause a reduced level of performance. An extreme example of this situation occurred during the last flight of the space shuttle *Columbia*. Insulating foam impact during the launch caused damage to the shuttle. However, it was not

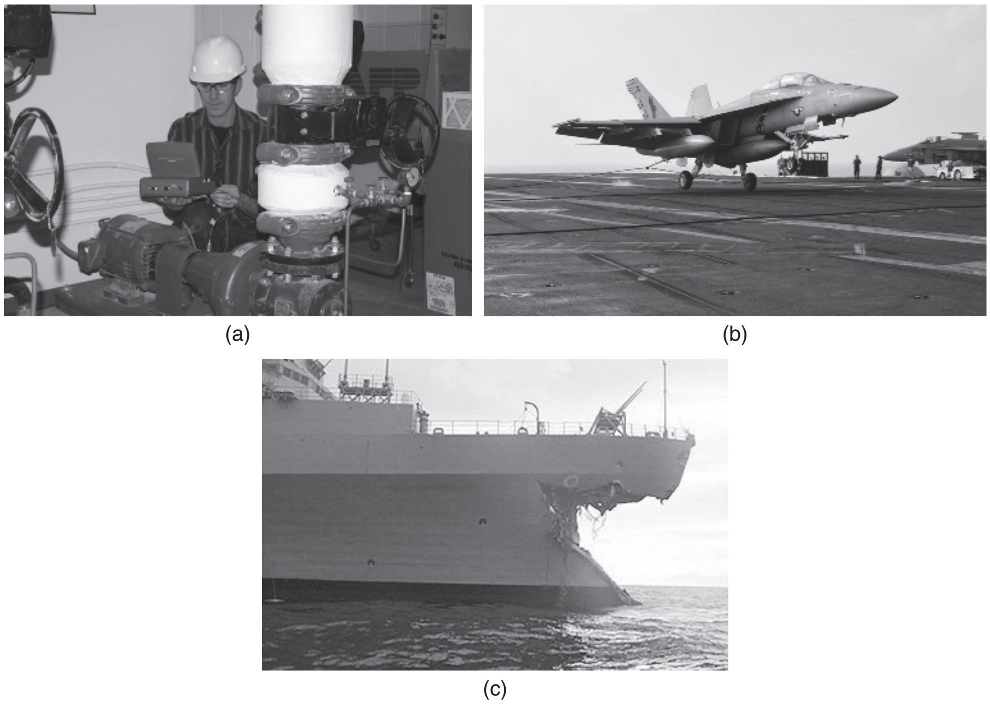

Figure 1.3 Illustration of three damage accumulation time scales. (a) Monitoring incremental damage accumulation in rotating machinery, (b) scheduled discrete damage accumulation resulting from a carrier landing (source: US Navy), (c) unscheduled discrete damage accumulation resulting from a ship-to-ship collision (source: US Navy).

until re-entry into the atmosphere when thermal environments were experienced that caused this damage to rapidly progress to catastrophic failure.

In terms of time scales, damage can accumulate incrementally over long periods of time, as in the case of damage associated with fatigue or corrosion. Damage can also progress very quickly, as in the case of critical fracture. Finally, scheduled discrete events such as aircraft landings and unscheduled discrete events such as birdstrike on an aircraft or transient natural phenomena such as earthquakes can lead to damage. Examples of damage developed over various time scales are shown in Figure 1.3.

A great deal of this book will be concerned with vibration-based approaches to SHM; this is amply justified by the fact that, in most practical scenarios, changes to a structural system caused by damage manifest themselves as changes to the mass, stiffness and energy dissipation characteristics of the system. Damage can also manifest itself as changes to the boundary conditions of a structure that reveal themselves as changes to the structure's dynamic response characteristics. As discussed earlier, the effects of the damage may become apparent on different time scales. The following examples illustrate situations where damage induces changes in one or more dynamical characteristics:

- A crack that forms in a mechanical part produces a change in geometry that alters the stiffness characteristics of that part while having almost no influence on the material characteristics or boundary conditions of the structure. Depending on the size and location of the crack, the adverse effects to the system can be either immediate or may take some time before they alter the system's performance.

Figure 1.4 Scour of bridge piers caused by increased flow rates that erode the supporting soil resulting in changes to the bridge's boundary conditions (source: US Geological Survey).

- Scour of a bridge pier is the process whereby increased flow rates around a pier erode the surrounding soil, as shown in Figure 1.4. This can be viewed as a change to the boundary conditions of the bridge that can compromise its structural integrity. However, this form of damage does not alter the local mass or stiffness properties of the structure itself.
- The loss of a lead balancing weight on a car wheel, the subsequent excessive wear of the tyre, loss of handling and loss of ride comfort is an example where the change of mass of the mechanical system can be viewed as the damaging event. In this case the stiffness and boundary conditions of the system are not altered by the damaging event.
- Finally, the loosening of a bolted connection in a structure is damage that alters the connectivity between elements of the structure while the stiffness and mass characteristics of the structural elements are not altered. Often this form of damage adds additional energy dissipation mechanisms to the structure, which would be reflected as an increase in measured vibrational damping properties.

1.4 A Statistical Pattern Recognition Paradigm for SHM

Implicit in the previous definition of damage is that the concept of damage is not meaningful without a comparison between two different states of the system, one of which is assumed to represent the initial, and often undamaged, state. This point is illustrated by Figure 1.5, which shows an apparently damaged highway bridge even though a close examination shows that pedestrians are still using this bridge to cross the river. Almost all readers, even if they have no background in damage assessment or bridge engineering, would affirm that Figure 1.5 shows a damaged bridge although there is no documentation indicating the initial state of this structure for comparison. This observation would appear to contradict the previous statement that a comparison with an undamaged state is needed to definitively conclude that the current observation represents a damaged condition. However, the readers' conclusion that this bridge is damaged is based on a mental comparison with the hundreds or thousands of examples of undamaged bridges that they have observed in their daily lives. Therefore, even for this very extreme case of damage some form of an initial condition comparison is needed before it can be stated that the pictured condition represents a damaged state for that structure. The other point to be made by this

Figure 1.5 A bridge located in Dagupan, Philippines, that was damaged by an earthquake in 1990.

example is that the reader employed pattern recognition in the process of mentally comparing the bridge image shown in Figure 1.5 to their internal database of previously observed healthy bridges. It will be argued in this book that pattern recognition provides a fundamental framework for carrying out SHM, although in most SHM applications this pattern recognition will need to be applied to mechanical or electrical sensor data, such as time-history readings as opposed to images. This book will also attempt to formalise this pattern recognition process by using the principles of machine learning.

Pattern recognition implemented through machine learning algorithms is a mature discipline. In abstract terms the theory provides mathematical means of associating measured data with given class labels. In the context of SHM, one wishes to associate the measured data with some damage state, the simplest – and arguably most important – problem being that of distinguishing between the states ‘healthy’ and ‘damaged’ for a structure. In mathematical terms there are a number of distinct approaches to pattern recognition, the main ones being the *statistical*, *neural* and *syntactic* approaches (Schalkoff, 1992). As all engineering problems are subject to various degrees of uncertainty, the statistical approach to pattern recognition appears to stand out as a natural approach for SHM purposes. As it will be seen in later chapters, neural network approaches can also be interpreted in statistical terms and also offer a robust means of dealing with SHM problems. In the example of the bridge earlier, the reader assigned the label *damaged* to the bridge by making a comparison of the structure with an internal database of healthy bridges representations. This database will have been accumulated, or *learned*, over an earlier period of time. The concept of learning representations from *training* data will be exploited throughout this book as the means of accomplishing pattern recognition. The mathematical framework needed for the problem is well established as the field of *machine learning* (Cherkassky and Mulier, 2007).

A general statistical pattern recognition (SPR) paradigm for an SHM system can be defined through the integration of four procedures (Farrar, Doebling and Nix, 2001):

1. Operational evaluation,
2. Data acquisition,
3. Feature selection and
4. Statistical modelling for feature discrimination.

Data normalisation, cleansing, compression and fusion are processes inherent in steps 2 to 4 of this paradigm. These processes can be implemented in either hardware or software and typically some combination of the two is used. The concept of machine learning enters into this paradigm primarily in steps 3 and 4.

Here, the idea of machine learning can be simply stated; it is to 'learn' the relationship between some features derived from the measured data (step 3) and the damaged state of the structure. If such a relationship between these two quantities exists, but is unknown, the learning problem is to estimate the function that describes this relationship using data acquired from the test structure – the *training* data. This estimation process is the focus of step 4. Learning problems naturally fall into two classes. If the training data comes from multiple classes and the labels for the data are known, the problem is one of *supervised learning*. If the training data do not have class labels, one can only attempt to learn intrinsic relationships within the data, and this is called *unsupervised learning*. Unsupervised learning can also be used to construct a model for a given single class that can then be used to test new data for consistency with that class; when used in such a manner, the process leads to novelty detection algorithms.

When one mentions the use of machine learning, there is often the misconception that this is an entirely data-driven process that makes no use of physics-based modelling. In fact, this need not be the case. In order to elaborate on this point it is useful now to discuss competing approaches to SHM. It is generally accepted that there are two main approaches, the 'inverse-problem' or 'model-based' approach and the 'data-based' approach.

The inverse-problem approach is usually implemented by building a physics-based or law-based model of the structure of interest; this is commonly a finite element (FE) model, although other modelling methods are used. Once the model is built, based on a detailed physical description of the system, it is usually updated on the basis of measured data from the real structure. This updating brings up an important point; it is very difficult to build an accurate model of a structure from first physical principles. Information or insight will be lacking in many areas, for example, and the exact nature of bonds, joints and so on can be difficult to specify. Another issue is that material properties may not be known with great accuracy; this is a common problem for civil engineers who will typically work with concrete. The updating step, then, adjusts the built model in such a way as to make it conform better with data from the real structure. The mathematical framework for this procedure is dominated by linear algebraic methods (Friswell and Mottershead, 1995). After updating, one has an accurate model of the structure of interest in its normal condition. When data from a subsequent monitoring phase become available, if any deviations from the normal condition are observed (e.g. the natural frequencies of the structure change), a further update of the model will indicate the location and extent of where structural changes have occurred, and this provides a damage diagnosis.

The data-based approach, as the name suggests, does not proceed from a law-based model. One establishes training data from all the possible healthy and damage states of interest for the structure and then uses pattern recognition to assign measured data from the monitoring phase to the relevant diagnostic class label. In order to carry out the pattern recognition, one needs to build a statistical model of the data, for example, to characterise their probability density function. This approach depends on the use of machine learning algorithms. In the data-based approach one can still make effective use of law-based models as a means of establishing good features for damage identification; this is discussed extensively in Chapters 7 and 8 later in this book.

There are pros and cons for both approaches; the reader can find a detailed discussion of these in Barthorpe (2011). In any case, the distinction between the two philosophies is not as clear-cut as one might wish. The model-based approach depends critically on the availability of training data for the initial update step; the data-based approach also establishes a model, but a statistical one. For various reasons, which will be elaborated later, the authors of this book believe firmly in the data-based approach.

The four steps of the SPR paradigm for SHM advocated in this book are briefly described below; the rest of the book is organised around this paradigm.

1.4.1 Operational Evaluation

The process of *operational evaluation* attempts to provide answers to four questions regarding the implementation of a damage identification investigation:

1. What is the life-safety and/or economic justification for performing the structural health monitoring?
2. How is damage defined for the system being investigated and, for multiple damage possibilities, which cases are of the most concern?
3. What are the conditions, both operational and environmental, under which the system to be monitored functions?
4. What are the limitations on acquiring data in the operational environment?

Operational evaluation begins to set limitations on what will be monitored and how the monitoring will be accomplished; it tries to tailor the damage identification process to features that are unique to the system being monitored and attempts to exploit unique features of the damage that is to be detected. Operational evaluation is discussed in more detail in Chapter 3.

1.4.2 Data Acquisition

The data acquisition portion of the SHM process involves selecting the excitation methods, the sensor types, number and locations, and the data acquisition/storage/transmittal hardware (see Chapter 4). This portion of the process will be application-specific. Economic considerations will play a major role in making decisions regarding the data acquisition hardware to be used for the SHM system. The interval at which data should be collected is another consideration that must be addressed. For earthquake applications it may be prudent to collect data immediately before and at periodic intervals after a large event. If fatigue crack growth is the failure mode of concern, it may be necessary to collect data almost continuously at relatively short time intervals once some critical crack has been identified.

1.4.3 Data Normalisation

The process of separating changes in the measured system response caused by benign operational and environmental variability from changes caused by damage is referred to as *data normalisation* (see Chapter 12). Examples of such variability are:

- An aircraft will change its mass during flight. A continuous change is caused by the burning of fuel; abrupt changes can be caused by the dropping of stores. Both of these effects are operational issues. If an in-flight SHM system were based on resonance frequencies, one would not wish to infer damage when a change occurred for benign reasons.

- The stiffness properties of a bridge can and do change with temperature. This variation can be quite complex; for example, the behaviour of the Z24 Bridge in Switzerland was observed to change when the ambient temperature dipped below the freezing point of the deck asphalt (Peeters, Maeck and Roeck, 2001). The variation described is a result of an environmental change; bridges are also susceptible to operational changes like variations in traffic loading.

Because system response data will often be measured under varying operational and environmental conditions, the ability to normalise the data becomes very important to the damage detection process; without this, changes in the measured response caused by changing operational and environmental conditions may be mistaken as an effect of damage. Additional measurements may be required to provide the information necessary to normalise the measured data and the need for this should be considered in the operational evaluation stage. When environmental or operational variability is an issue, the need can arise to normalise the data in some temporal fashion to facilitate the comparison of data measured at similar times of an environmental or operational cycle. Often the data normalisation issues will be key challenges to the field deployment of a robust SHM system.

1.4.4 Data Cleansing

Data cleansing is the process of selectively choosing data to pass on to or reject from the feature selection process. The data cleansing process is usually based on knowledge gained by individuals directly involved with the data acquisition. As an example, an inspection of the test setup may reveal that a sensor was loosely mounted and, hence, based on the judgement of the individuals performing the measurement; this set of data or the data from that particular sensor may be selectively deleted from the feature selection process. Signal processing techniques such as filtering and resampling can also be thought of as data cleansing procedures.

1.4.5 Data Compression

Data compression is the process of reducing the dimension of the measured data. The concept of data or feature dimensionality is discussed in more detail in Chapter 7. The operational implementation of the measurement technologies needed to perform SHM inherently produces large amounts of data. A condensation of the data is advantageous and necessary when comparisons of many feature sets obtained over the lifetime of the structure are envisioned. Also, because data will be acquired from a structure over an extended period of time and in an operational environment, robust data reduction techniques must be developed to retain feature sensitivity to the structural changes of interest in the presence of environmental and operational variability. To give further aid in the extraction and recording of the high-quality data needed to perform SHM, the statistical significance of the features will need to be characterised and used in the condensing process.

1.4.6 Data Fusion

Data fusion is the process of combining information from multiple sources in an effort to enhance the fidelity of the damage detection process. The fusion process may combine data from spatially distributed sensors of the same type such as an array of strain gauges mounted on a structure. Alternatively, heterogeneous data types including kinematic response measurements (e.g. acceleration) along with environmental parameter measurements (e.g. temperature) and measures of operational parameters (e.g. traffic volume on a bridge) can be combined to determine more easily if damage is present. Clearly, data fusion is closely related to the data normalisation, cleansing and compression processes.

1.4.7 Feature Extraction

The part of the SHM process that arguably receives the most attention in the technical literature is the identification of data features that allows one to distinguish between undamaged and damaged states of the structure of interest (Doebling *et al*., 1996; Sohn *et al*., 2004) (see Chapters 7 and 8). A damage-sensitive feature is some quantity extracted from the measured system response data that indicates the presence (or not) of damage in a structure. Features vary considerably in their complexity; the ideal is a low-dimensional feature set that is highly sensitive to the condition of the structure. Generally, a degree of signal processing is required in order to extract effective features. For example, if one wished to monitor the condition of a gearbox, one might start by attaching an accelerometer to the outer casing. This sensor would yield a stream of acceleration–time data. To reduce the dimension of the data without compromising the information content, one might use the time series to compute a spectrum. Once the spectrum is available, one can then extract only those spectral lines centred around the meshing harmonics, as these are known to carry information about the health of the gears. This specific feature extraction process is quite typical in that it involves both mathematical operations or transformations and the use of a priori engineering judgement. Another useful source of diagnostic features is to build (or learn) a physical or data-based parametric model of the system or structure; the parameters of these models or the predictive errors associated with these models then become the damage-sensitive features. Inherent in many feature selection processes is the fusing of data from multiple sensors (see Chapter 4) and subsequent condensation of these data. Also, various forms of data normalisation are employed in the feature extraction process in an effort to separate changes in the measured response caused by varying operational and environmental conditions from changes caused by damage (Sohn, Worden and Farrar, 2003).

1.4.8 Statistical Modelling for Feature Discrimination

The portion of the SHM process that has arguably received the least attention in the technical literature is the development of statistical models for discrimination between features from the undamaged and damaged structures. Statistical model development is concerned with the implementation of algorithms that operate on the extracted features to quantify the damage state of the structure; they are the basis of the SPR approach. The functional relationship between the selected features and the damage state of the structure is often difficult to define based on physics-based engineering analysis procedures. Therefore, the statistical models are derived using machine learning techniques. The machine learning algorithms used in statistical model development usually fall into two categories, as alluded to earlier. When training data are available from both the undamaged and damaged structure, *supervised learning* algorithms can be used; *group classification* and *regression analysis* are primary examples of such algorithms. In the context of SHM, *unsupervised learning* problems arise when only data from the undamaged structure are available for training. *Outlier* or *novelty detection* methods are the primary class of algorithms used in this situation. All of the algorithms use the statistical distributions of the measured or derived features to enhance the damage detection process. Chapters 9, 10 and 11 discuss in more detail the statistical modelling portions of the SHM process as implemented using machine learning principles.

The damage state of a system can in principle be arrived at via a five-step process organised along the lines of the hierarchy discussed in Rytter (1993). This process attempts to answer the following questions:

1. Is there damage in the system (existence)?
2. Where is the damage in the system (location)?
3. What kind of damage is present (type)?
4. How severe is the damage (extent)?
5. How much useful (safe) life remains (prognosis)?

Answers to these questions in the order presented represent increasing knowledge of the damage state. When applied in an unsupervised learning mode, statistical models can typically be used to answer questions regarding the existence (and sometimes, but not always, the location) of damage. When applied in a supervised learning mode and coupled with analytical models, the statistical procedures can, in theory, be used to determine the type of damage, the extent of damage and the remaining useful life of the structure. The statistical models are constructed in such as way as to minimise false diagnoses. False diagnoses fall into two categories: (1) *false-positive* damage indication (indication of damage when none is present) and (2) *false-negative* damage indication (no indication of damage when damage is present). If one wishes, one can design diagnostic systems that weight the costs of the two error types differently.

Statistical models are used to implement two types of SHM. *Protective monitoring* refers to the case when damage-sensitive features are used to identify impending failure and shut the system down or alter its use in some other manner before catastrophic failure results. In this case the statistical models are used to establish absolute values or thresholds on acceptable levels of feature change. *Predictive monitoring* refers to the case where one identifies trends in data features that are then used to predict when the damage will reach a critical level. This type of monitoring is necessary to develop cost-effective maintenance planning. In this case statistical modelling is used to quantify uncertainty in estimates of the feature's time rate of change.

1.5 Local versus Global Damage Detection

Interest in the ability to monitor a structure and detect damage at the earliest possible stage is pervasive throughout the civil, mechanical and aerospace engineering communities. Most current damage-detection methods are NDE-based using visual or localised experimental methods such as acoustic or ultrasonic methods, magnetic field methods, radiography, eddy-current methods and thermal field methods (Hellier, 2001; Shull, 2002). All of these experimental techniques require that the vicinity of the damage is known a priori and that the portion of the structure being inspected is readily accessible. Subject to these limitations, such experimental methods can detect damage on or near the surface of the structure. However, surface measurements performed by most standard NDE procedures cannot provide information about the health of the internal members without costly dismantling of the structure. As an example, micro-cracks were found in numerous welded connections of steel moment-resisting frame structures after the 1994 Northridge earthquake (Darwin, 2000). These connections are typically covered by fire-retardant and nonstructural architectural material. Costs associated with inspecting a single joint and then reinstalling the fire-retardant and architectural cladding can be on the order of thousands of dollars per joint. A typical twenty-storey building may have hundreds of such joints. Clearly, there is a tremendous economical advantage to be gained if the damage assessment can be made in a nonintrusive and more cost-effective manner.

In addition to the local inspection methods, there has been a perceived need for quantitative global damage detection methods that can be applied to complex structures. Among other things, this has led to the development of, and continued research into, methods that examine changes in the vibration characteristics of the structure. As discussed earlier, the basic premise of vibration-based damage detection is that damage will alter the stiffness, mass or energy dissipation properties of a system, which, in turn, alter the measured global dynamic response properties of the system. Although the basis for vibration-based damage detection appears intuitive, its actual application poses many significant technical challenges. The most fundamental challenge is the fact that damage is typically a local phenomenon and may not significantly influence the lower-frequency global response of a structure that is normally measured during vibration tests, particularly those where the response to ambient excitation is measured. Stated another way, this fundamental challenge is similar to that found in many engineering fields where there is a need to capture the system response *on widely varying length scales*, and such system modelling and measurement has proven difficult.

More recently researchers have been studying hybrid multiscale sensing approaches to SHM (Park *et al*., 2003). Such approaches rely on active sensing systems for local damage detection and the same sensor/actuators are used in a passive mode to measure the influence of damage on the global system response. Here the term *active* refers to systems where actuators are incorporated with the sensing system to provide a known input to the structure that is designed to enhance the damage detection process.

Fundamentally, there will always be a trade-off between the cost associated with deploying a local sensing system over a large area of the structure and the lack of fidelity associated with more global sensing systems. For most applications some hybrid system will most likely be employed that is based on a priori knowledge of specific areas on the structure that are most likely to experience damage.

1.6 Fundamental Axioms of Structural Health Monitoring

Because of the economic and safety implications associated with accurate damage identification, many new SHM studies and resulting technical advances have been seen in recent years. The advances that have been made in these more global approaches to damage detection, herein referred to simply as damage detection, are the result of coupling recent advances in various technologies such as computer hardware, sensors, computational mechanics, experimental structural dynamics, machine learning/SPR and signal analysis. Correspondingly, the number of papers dealing with this subject that have appeared in the technical literature has been rapidly growing for the past twenty years.

Based on information published in this extensive literature, the authors feel that the field has matured to the point where several fundamental axioms, or accepted general principles, have emerged (Worden *et al*., 2007; Farrar, Worden and Park, 2010). The axioms are first presented here, without further justification, for the reader to keep in mind while viewing the rest of this book. The authors believe that evidence will be presented to justify each axiom and some of the subtleties associated with each will be directly addressed in the subsequent chapters. These axioms will be revisited in Chapter 13 with more detailed justification as a means of summarising the material presented throughout this book. The axioms are:

Axiom I. All materials have inherent flaws or defects.

Axiom II. Damage assessment requires a comparison between two system states.

Axiom III. Identifying the existence and location of damage can be done in an unsupervised learning mode, but identifying the type of damage present and the damage severity can generally only be done in a supervised learning mode.

Axiom IVa. Sensors cannot measure damage. Feature extraction through signal processing and statistical classification are necessary to convert sensor data into damage information.

Axiom IVb. Without intelligent feature extraction, the more sensitive a measurement is to damage, the more sensitive it is to changing operational and environmental conditions.

Axiom V. The length and time scales associated with damage initiation and evolution dictate the required properties of the SHM sensing system.

Axiom VI. There is a trade-off between the sensitivity to damage of an algorithm and its noise rejection capability.

Axiom VII. The size of damage that can be detected from changes in system dynamics is inversely proportional to the frequency range of excitation.

Axiom VIII. Damage increases the complexity of a structure.

1.7 The Approach Taken in This Book

This book is designed to provide an overview of the current state of the art in SHM for those just entering the field and for those who want to begin applying this technology to real-world problems. More specifically, the authors believe this book will provide a systematic approach to addressing damage detection problems through the machine learning/SPR paradigm for SHM that forms the primary theme for this book. The second theme for this book is that successful implementation of a structural health monitoring process requires a synergistic, multidisciplinary approach. Finally, an important conclusion will be that there is no one damage detection method that is applicable to all structural and mechanical systems. These themes will be emphasised throughout this book and will be reinforced through numerous examples.

Depending on the application, SHM can be thought of as either a new emerging field, as in cases when applied to civil engineering and aerospace infrastructure, or as a fairly mature technology when it is viewed in the context of CM for rotating machinery. To further enhance the reader's understanding of this technology, descriptions of many methods that have been reported and applied beyond a laboratory setting will be presented along with summaries of methods currently under development at various research institutes. In addition, the current limitations of this technology and assumptions associated with a particular feature or statistical classification procedure will be continually emphasised throughout the book. Application perspectives from aerospace, civil and mechanical engineering communities will be used as examples. Also, practical issues related to implementation of damage identification technologies will be presented. In summary, the authors hope that this book will provide the reader with a general background in SHM technology, an understanding of the limitations of this technology, the areas of current research in damage detection and the appropriate reference material such that the reader can further study this subject in the context of their particular application.

References

Barthorpe, R. (2011) On model- and data-based approaches to structural health monitoring. PhD Thesis, Department of Mechanical Engineering, University of Sheffield.

Cherkassky, V. and Mulier, F.M. (2007) *Learning from Data: Concepts, Theory and Methods*, Wiley-Blackwell.

Darwin, D. (ed.) (2000) Steel moment frames after Northridge. Special Issue of *Journal of Structural Engineering*, **126**(1).

Doebling, S., Farrar, C., Prime, M. and Shevitz, D. (1996) Damage Identification and Health Monitoring of Structural and Mechanical Systems from Changes in Their Vibration Characteristics: A Literature Review. Los Alamos National Laboratory Report LA-13070-MS, available online from Los Alamos National Laboratory, http://library.lanl.gov/.

Farrar, C., Doebling, S. and Nix, D. (2001) Vibration-based structural damage identification. *Philosophical Transactions A*, **359**(1778), 131.

Farrar, C., Worden, K. and Park, G. (2010) Complexity: a new axiom for structural health monitoring? 5th European Workshop on Structural Health Monitoring, Sorrento, Italy.

Farrar, C., Sohn, H., Hemez, F.M. *et al*. (2003) Damage Prognosis: Current Status and Future Needs. Los Alamos National Laboratory Report LA-14051-MS, available online from Los Alamos National Laboratory, http://library.lanl.gov/.

Friswell, M.I. and Mottershead, J.E. (1995) *Finite Element Model Updating in Structural Dynamics*, Springer.

Hellier, C.J. (2001) *Handbook of Non-destructive Evaluation*, McGraw-Hill, New York.

Montgomery, D.C. (2009) *Introduction to Statistical Quality Control*, John Wiley & Sons, Inc., Hoboken, NJ.

Park, G., Sohn, H., Farrar, C.R. and Inman, D.J. (2003) Overview of piezoelectric impedance-based health monitoring and path forward. *Shock and Vibration Digest*, **35**(6), 451–463.

Peeters, B., Maeck, J. and Roeck, G.D. (2001) Vibration-based damage detection in civil engineering: excitation sources and temperature effects. *Smart Materials and Structures*, **10**, 518–527.

Rytter, A. (1993) Vibration Based Inspection of Civil Engineering Structures. Building Technology and Structural Engineering. Aalborg University, Aalborg, Denmark.

Samual, P.D. and Pines, D.J. (2005) A review of vibration-based techniques for helicopter transmission diagnostics. *Journal of Sound and Vibration*, **282**, 475–508.

Schalkoff, R.J. (1992) *Pattern Recognition: Statistical, Structural and Neural Approaches*, John Wiley & Sons, Ltd, Chichester, Sussex.

Shull, P. (ed.) (2002) *Nondestructive Evaluation: Theory, Techniques and Applications*, Marcel Dekker, New York.

Sohn, H., Worden, K. and Farrar, C.R. (2003) Statistical damage classification under changing environmental and operational conditions. *Journal of Intelligent Materials Systems and Structures*, **13**(9), 561–574.

Sohn, H., Farrar, C., Hemez, F. *et al*. (2004) A Review of Structural Health Monitoring Literature from 1996–2001. Los Alamos National Laboratory Report LA-13976-MS, available online from Los Alamos National Laboratory, http://library.lanl.gov/.

Worden, K. and Dulieu-Barton, J. (2004) An overview of intelligent fault detection in systems and structures. *Structural Health Monitoring*, **3**(1), 85.

Worden, K., Farrar, C., Manson, G. and Park, G. (2007) The fundamental axioms of structural health monitoring. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science*, **463**(2082), 1639.

Zwaag, S.v.d. (2007) *Self Healing Materials: An Alternative Approach to 20 Centuries of Materials Science*, Springer.