The Effects of Variability on Damage Identification with Inductive Learning

by

Thomas D. Snyder

Thesis submitted to the Faculty of the

Virginia Polytechnic Institute and State University

in partial fulfillment of the requirements for the degree of

Master of Science

in

Mechanical Engineering

APPROVED:

Prof. H. H. Robertshaw, Chairman

Prof. Z. Chaudhry

Prof. W. R. Saunders

December 15, 1994

Blacksburg, Virginia

The Effects of Variability on Damage Identification with Inductive Learning

by

Thomas D. Snyder

Committee Chairman: Harry H. Robertshaw

Mechanical Engineering

(ABSTRACT)

This work discusses the effects of inherent variabilities on the damage identification problem. The goal of damage identification is to detect structural damage before it reaches a level which will detrimentally affect the structure's performance. Inductive learning is one tool which has been proposed as an effective method to perform damage identification.

There are many variabilities which are inherent in damage identification and can cause problems when attempting to detect damage. Temperature fluctuation and manufacturing variability are specifically addressed. Temperature is shown to be a cause-effect variability which has a measurable effect on the damage identification problem. The inductive learning method is altered to accommodate temperature and shown experimentally to be effective in identifying added mass damage at several locations on an aluminum plate.

Manufacturing variability is shown to be a non-quantifiable variability. The inductive learning method is shown to be able to accommodate this variability through careful examination of statistical significances in dynamic response data. The method is experimentally shown to be effective in detecting hole damage in randomly selected aluminum plates from a manufactured batch.

Acknowledgments

I would like to take this opportunity to thank all those who were involved with making this project a reality. First, I would like to thank Harry Robertshaw for serving as my advisor and committee chair. It was your ability to lend guidance while still allowing me to take the research in its own direction which made this project more interesting and rewarding than I thought that it would be. Thanks for the help, the humor, and most of all the faith and encouragement along the way. I would also like to thank Zaffir Chaudhry and Will Saunders for their suggestions and for serving on my committee.

Special thanks goes to Peter Tappert for all the work he did along the way. Any time that the work began to get overwhelming and began to grind to a halt, you were there to sort out the confusion and get me back on track. And it was your inspiration that got the train moving in the first place. Thanks. I'd also like to thank everyone else from the lab (Darrell McAlister, Ben Fong, Dan Miller, and Dan Cole) for their assistance and friendship. You certainly helped make the hours in the lab enjoyable.

I would like to thank all my friends outside of classes (Keith, Alan, Romeo, Derek, Chris, and on and on) as well for helping make my decision to come back to school a good one. Without their assurance that I needed to participate with them in *many* study breaks, I'd have probably quit out of frustration long ago. Thanks guys.

And finally, I'd like to thank my family. It was your support and trust in my decisions through all the years that really made this possible. I couldn't have done it without you.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

"I have called this principle, by which each slight variation, if useful, is preserved, by the term natural selection."

- Charles Darwin

1.1 Motivation

WORK IN THE FIELD OF DAMAGE IDENTIFICATION IS MOSTLY GEARED TOWARDS NON-destructive techniques to identify structural problems before they become too costly to repair or cause catastrophic failure of the structure of interest. The need for the development of damage identification methods has arisen because of the problem that not all structures can be easily assessed for damage. Some structures are too costly and difficult to access for visual inspection. Others need to be checked on-line while they are operating. And often, the magnitude or type of damage which may be present simply cannot be detected by conventional means. As a result, a great deal of work has been performed in recent years to come up with a successful and practical damage identification method.

Almost all of the recent research in the field of damage identification is based on the premise that the response characteristics of a structure are functions of its physical properties and boundary conditions. In theory, any form of damage to a structure should present itself as a change in the local stiffness, mass or damping properties of the structure at the location of the damage. Using this assumption, damage should be evident from inspection of the response characteristics of the structure being analyzed. In this case, *response characteristics* refer to any type of representation of the dynamic response of a structure. Examples include, but are not limited to, impulse responses, frequency responses, or structural impedance responses. The task at hand, which is sometimes referred to as *signature analysis*, is to determine how to correlate the changes in response characteristics with their respective causes; and therefore identify damage.

This task, however, is not a straightforward one. This is because there are many factors other than damage which can cause changes in response characteristics. In fact, any factor which affects the physical properties or boundary conditions of a structure will change its dynamic response. Some of these factors include boundary condition variability, environmental changes, and manufacturing variability. For example, the dynamic response of an airplane wing will be different on a cold day versus a hot one. Additionally, two "identical" wings on different (but "identical") aircraft will have different response characteristics due to slight variations in their material compositions and manufacturing tolerances. Another source of variation comes from the method with which dynamic response data is collected. Sensor noise can be a factor, as can instrument precision. All of these variances are inherent in the damage identification problem. If research is to be composed on the premise that changes in dynamic response characteristics can be indicative of damage, then these inherent variances must be considered.

1.2 Literature Review

AS WAS STATED ABOVE, ALMOST ALL OF THE RECENT WORK IN DAMAGE IDENTIFICATION has centered around the idea of vibrational signature analysis. The basic concept is to come up with a set of response characteristics which represent the healthy or undamaged structure, and use them as a base of comparison for damage identification. The means by which researchers have come up with these response characteristics as well as the methods of looking for damage induced changes in the healthy response characteristics have varied widely however.

There are two basic schools of thought on how to attack the damage identification problem. One school depends on model-based data and simulation as a key element in the damage identification scheme. Those who subscribe to model-based methods claim the ability to simulate many expected (or unexpected) damage cases, in fact generating a large base of comparison response characteristics for testing, as one of its advantages. However, questions arise concerning the model output's correllation with experimental data and the ability to expand to more complex models. The other school relies on actual structure testing as its means of obtaining the necessary response characteristics. This school's proponents claim the ability to get real data and avoid model-induced errors brought on by model simplification and inaccuracy as some of it's advantages. Conversely, there is skepticism regarding the validity and ability of the method when there are no available damage cases to test.

Concentrating first on work performed on the theory of model-based methods, several approaches have been attempted. A great deal of work has been performed in an effort to use finite element models for damage detection. Kung et al. (1989) presented a method in which he fit experimentally determined mass and stiffness matrices to a finite element model. He then looked for maximum changes in these matrices as an indicator of damage. Hickman et al. (1990) described a method designed for aircraft structural health monitoring which involved using a finite element model to examine the sensitivities of its resonant frequencies to different types of structural damage. A modal analysis scheme was then implemented based on the model's results. Lindner et al. (1994) took the finite element methods a step further by computing estimations of damage coefficients for each element of the model based on experimental data. Then an element-by-element comparison could be carried out across all modes which would locate damage at a specific location on the structure. The method showed success in a numerical example of a cantilevered beam with a five percent reduction in stiffness at one element, but no experimental verification has been completed.

In 1992, Kudva et al. presented a paper describing a damage detection scheme for structural health monitoring of aircraft. His work was founded on using neural networks' pattern recognition ability to look for changes in the response characteristics of the healthy aircraft. The network he proposed was trained with response data obtained from finite element models of each expected damage state. Elkordy, et al. (1993) performed an experiment on a five story steel frame in which he successfully detected damage using experimentally obtained response characteristics in conjunction with a finite-element trained neural network. Manning et al. (1994) trained a neural network using transfer function poles and zeros for example damage cases of an active truss. The pole and zero information was obtained from an analytical model of the truss. He showed that this method was able to detect (in a numerical simulation) a 25% reduction in stiffness in selected members of a 25 bar transmission tower.

While there has been some success in the field of model-based damage identification research, there are some drawbacks to this method. The biggest problem which has been encountered is the inability to accurately model complex structures. The finite element modelling resolution has been sufficient for simple beam and truss models, but when the structure complexity increases, these models have not yielded accurate response characteristics. It also becomes increasingly difficult to create an analytical representation of more complex structures and systems. As a result, some researchers have attempted to develop damage identification methods which are model-independent.

Neural networks are a popular tool in the model-independent school of thought as well. In 1991, Roitman et al. showed how neural networks could be trained with experimentally obtained modal analysis data from a scale model of an offshore oil platform. His work was significant because he found that natural frequencies are not always the best place to look for damage-induced changes in response characteristics, contrary to some previous beliefs. Ganino successfully trained a neural network to detect delaminations of composite patches on aluminum beams (1993) and plates (1994). Through his work, he found that the networks were unable to detect damage properly when in the presence of temperature changes in the experimental environment. He determined the problem to be the sensitivity of piezoelectric ceramic (PZT) patches on his experimental testbed, but he was not able to train the network to accomodate for this sensitivity. Povich (1994) presented another neural network example in which he trained a network with frequency response functions experimentally obtained from accelerometers attached to a 20-bay planar truss. He was able to uniquely identify only one third of sixty possible damage cases (removal of one strut) with this approach, and convergence of the neural net was a problem due to the complexity and number of members in the structure.

Not all the model-independent work is based on neural networks. Hofer (1987) presented a method by which fiber optic strands were embedded crosswise in a composite. When the structure was damaged, it would sever some strands and damage detection was facilitated by looking for interruptions in the light patterns through the fiber optics. In 1993, Liu et al. investigated the use of real-time x-ray techniques to examine the damage fields around edge-cracked sheet specimens. Perhaps the most significant work to the research contained in this thesis was that performed by Tappert et al. (1993, 1994) on using inductive learning methods as a means of damage identification. In his work, comparisons of experimentally obtained impedance response data were made using statistical hypothesis testing. It is the reliance on statistical tools to perform the damage assessment which lends itself to a study of the effects of variance and cause-effect variability on the damage identification problem.

One problem with the research that has been performed so far is that not much has been done to specifically address the problem of inherent variances and their effect on the damage identification problem. This is a very important concept which must be studied if a damage identification method is to ever become practical. Ganino's (1993) work concerning the effects of temperature on PZT performance during damage detection would be one exception to the attitude to ignore these variances.

There has, however, been some work carried out in other fields which lends itself to the study of these inherent variances in one manner or another. One example would be the work presented by Deval et al. (1993). He describes a method of adding temperature compensation to a circuit design which acts as a PTAT (proportional to absolute temperature) temperature compensator without causing any structural change to the circuit. This is an example of temperature compensation (which is one of the requirements for an ideal damage identification method), but it is a far cry from how compensation might be accomplished in a damage identification method. Because most of these variability parallels are not in the field of damage identification, an extensive literature review of the topic of inherent variance has not been included.

1.3 Objective and Scope

AN IDEAL DAMAGE IDENTIFICATION METHOD SHOULD BE ABLE TO PERFORM SUCCESSFULLY in the presence of all possible variances inherent to the task. These include such things as environmental fluctuations, manufacturing tolerances and boundary condition variability.

An ideal method also should be adaptable to numerous different sensor and actuator types, manipulations of dynamic response data, levels of available computing power and complexities of the structure of interest. It is the goal of this thesis to recognize some of the variances inherent in the damage identification problem and incorporate them into a damage identification scheme.

The inductive learning method researched by Tappert et al. (1994) has been chosen as the method which will most easily accomodate this goal. Through modification of his algorithms, it can be shown that temperature compensation can be incorporated into the damage identification process. Also, boundary condition variability and manufacturing variability can be addressed. The elegance of inductive learning allows for simple modifications to be made which can be used to make the method robust to different sensor types and response characteristics. Many of the characteristics of an ideal damage identification method can be incorporated into a method based on inductive learning.

Chapter 2 will describe in some detail the inductive learning damage identification scheme outlined by Tappert. It will also specifically address several of the inherent variances of the damage identification problem. The role and importance of statistical tools will be discussed in the context of damage identification and factor variability. A crucial element of these statistical tools is the ability to set specific alpha levels of confidence and siginificance in the inductive learning algorithms. This will become an essential factor in eliminating some of the inherent, or cause-effect (CE), variabilities.

A specific experiment was run, as an extension of the work Tappert performed, which addresses the problem of temperature compensation. Chapter 3 presents a thorough discussion of this experiment and the results drawn from it. A temperature compensation algorithm is detailed and shown to be easily incorporated. Some projectionary discussion is also given concerning compensation of other environmental factors.

Chapter 4 is a study of the effects of manufacturing variability on the damage identification problem. An experiment is conducted to show how this variability can be included in the damage identification process. Finally, Chapter 5 discusses the conclusions and significance of the work.

Chapter 2

Inductive Learning

"In our description of nature the purpose is not to disclose the real essence of the problem but only to track down, so far as it is possible, relations between the manifold aspects of our experience."

- Niels Bohr

2.1 Approach

THE GOAL OF AN INDUCTIVE LEARNING BASED DAMAGE IDENTIFICATION METHOD IS TO discover empirically how the effects of damage manifest themselves in the response characteristics of a structure. Literally, *inductive learning* means learning from examples. Through data collection, an Inductive Learning Algorithm (ILA) can take many examples of response characteristics of a healthy structure and use them to "learn" about the structure. The more examples which are present, the more the algorithm will "know". By the same token, an ILA can take response characteristics of damaged structures and "learn" about them as well.

The domain in which inductive learning operates is that of statistics. In some applications, this foundation is rather loose. For example in the medical field, inductive learning has been used to aid in the prevention of heart disease (FirstMark Technologies

Ltd., 1990). Such data as diet, exercise regimen, family history, and sleeping habits are compiled for patients with heart disease. These *examples* are then studied to find (statistical) trends which are indicators that can be used to prescribe a lifestyle which will lead towards less risk of heart disease.

The need for statistics is much greater in the damage identification problem. The ILA which Tappert (1994) describes takes dynamic response data from a structure and examines it statistically to generate a set of *rules* that can be used to detect damage. Tappert's experiment involved the detection of an added mass (2% of the total plate mass) at various locations on an aluminum plate hung with free-free boundary conditions. The form of response characteristic which was used was a measure of the real and imaginary parts of the structural impedance response (measured in ohms) of the plate over a frequency range of 4200-5000 Hz with 1 Hz resolution.

The first step in the rule formation process is to gather sets (examples) of impedance response data from the healthy plate. It should be noted that any variabilities which are present during data collection will exhibit themselves in this response information. These examples are then averaged and a sample mean and sample variance is calculated for each frequency point. The same process is then followed for any expected damage cases which are available for testing. A set of sample means and sample variances for each damage case are calculated accordingly. Figure 2.1 shows an impedance response ensemble for Tappert's healthy plate. Each frequency point, or bin, is referred to as a *random variable*. Each data point, in turn, is a *random sample*.

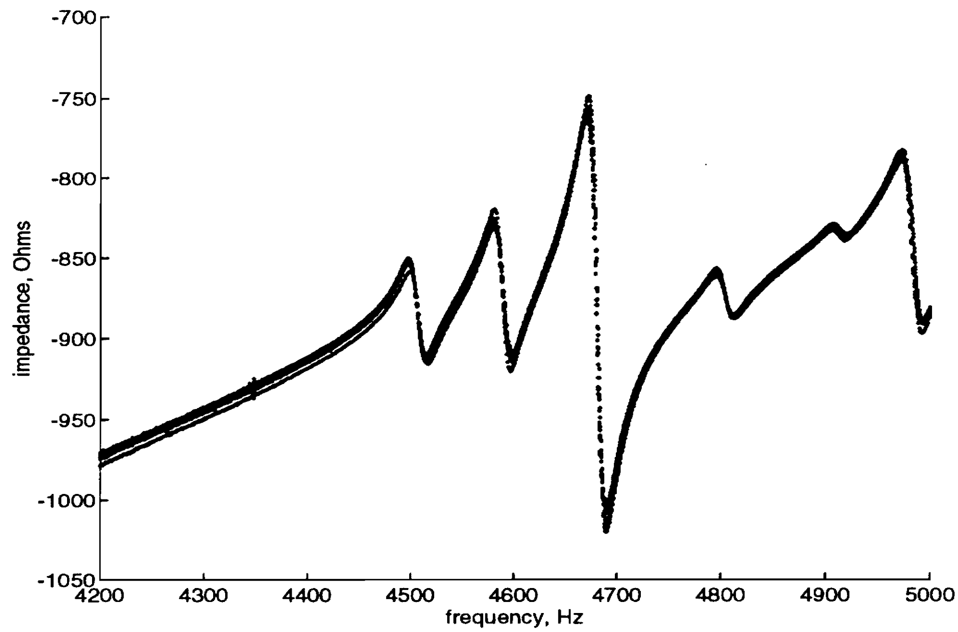

Figure 2.1: Impedance response ensemble (imaginary part)

Figure 2.2 shows the sample means and sample variances associated with portions of the data from Fig. 2.1. It should be noted that the variance is generally higher in the areas of resonance than in some of the other sections of the response. It will be shown later that this supports Roitman's (1991) claims that natural frequencies are not always the best regions of the dynamic response to detect damage.

So far, all of the data analysis has resulted in *sample* means and variances. Since the original premise of signature analysis was to look for changes in response data, it must be determined how well these *sample* means and variances estimate the *true* means and

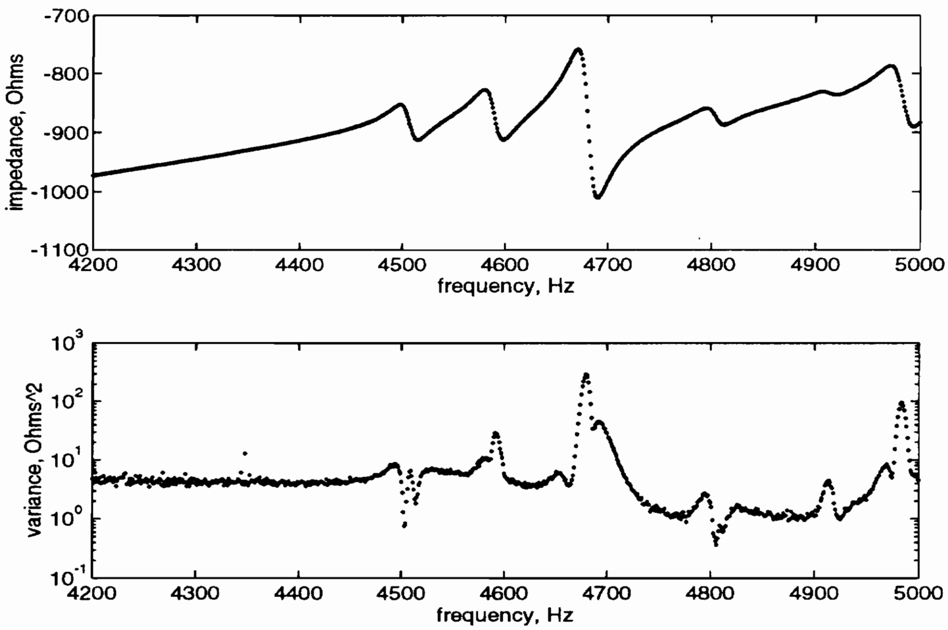

Figure 2.2: Sample means and variances of imaginary part of impedance response

variances. It is possible to set up confidence intervals on the sample means which will bound them to within a given α level. By doing this, it is possible to find a bounded region within which the true mean is expected to fall. The size of this region (and thus our confidence in it) is determined by the selection of α. The result is a confidence bound which has a calculated statistical confidence that the true mean falls within the bounds. Figure 2.3 shows the confidence bounds for an enlarged portion of the previous figure. It can be seen that as the sample variance decreases, the confidence bounds narrow (vertically). This fact will play an important role in monitoring the health of a structure.

It is now time to discuss how the ILA actually detects and identifies damage. Response data has been collected for the healthy structure and several expected damage

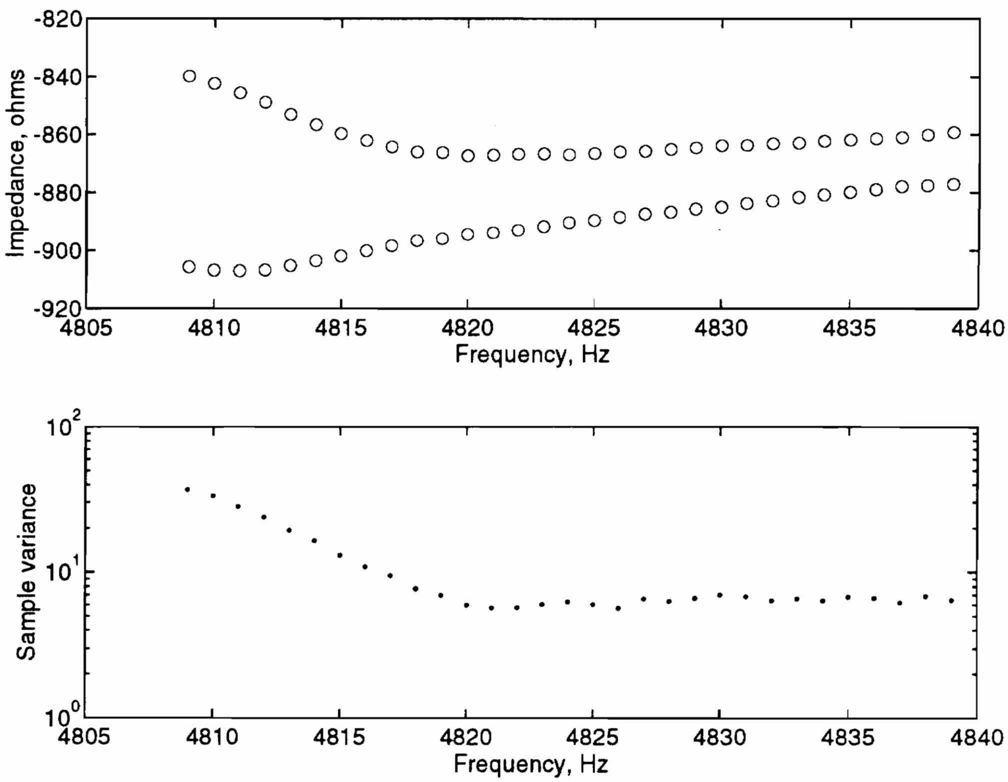

Figure 2.3: Effect of sample variance on confidence bounds

cases. The ILA now needs to compare the different health cases to look for the places (frequency bins) where the response has changed from one case to another. The method by which this comparison is made is a statistical *hypothesis test.* The null hypothesis which is chosen, is that of equal means.

For each random variable (frequency bin), a hypothesis test of the means of each binary combination of health states is performed. This hypothesis test is a function of the difference in sample means, the sample variance, and the number of samples (examples). The test returns a number (ranging from 0-1) which is referred to as the *significance level.*

The significance level is in a way a measure of how confidently it can be said that the two true means being compared are equal. A significance level of zero means that there is a 100% confidence that the true means are not the same. This corresponds to a total rejection of the null hypothesis. Conversely, a significance level of one indicates that the hypothesis should be accepted without question.

The ILA checks all the binary combinations of health states for each random variable and looks for instances where the significance level is very small. When the significance level is small, there is a high confidence that the means being compared are different, and therefore that a *statistically significant* difference is present in the response characteristics of the two different health states. As the significance level gets smaller, confidence in the ability to distinguish between damage states increases.

In Tappert's experiment, five different health states were examined (NO DAMAGE, DAMAGE CASES A-D). For each health state, ten sets of response data were experimentally collected. Each set consisted of 801 frequency bins with a real and imaginary data point; or a total of 1602 pieces of data. The sample mean and sample variance was then calculated for each of the 1602 random variables in each ensemble. Next, hypothesis tests were performed for each binary combination of health states at each frequency bin and for all combinations of data type and sensor. For example, a hypothesis test was performed for damage case A versus damage case B at frequency bin number one using the imaginary part of the impedance response collected with sensor number one. Then damage case A was tested versus damage case C and so on.

Once all the hypothesis tests were completed, the ILA selected the sixteen random variables which had the lowest significance levels for each of the possible health state comparisons. These variables are called the *predictor variables*. Because the significance levels were low, the predictor variables represent the areas where the confidence is highest that the response characteristics are different. It is from these selected predictor variables that the damage detection *rules* which were mentioned above were generated.

Each predictor variable was bounded with an α level of confidence as described above. Figure 2.4 shows an example of selected predictor variables and their bounds for

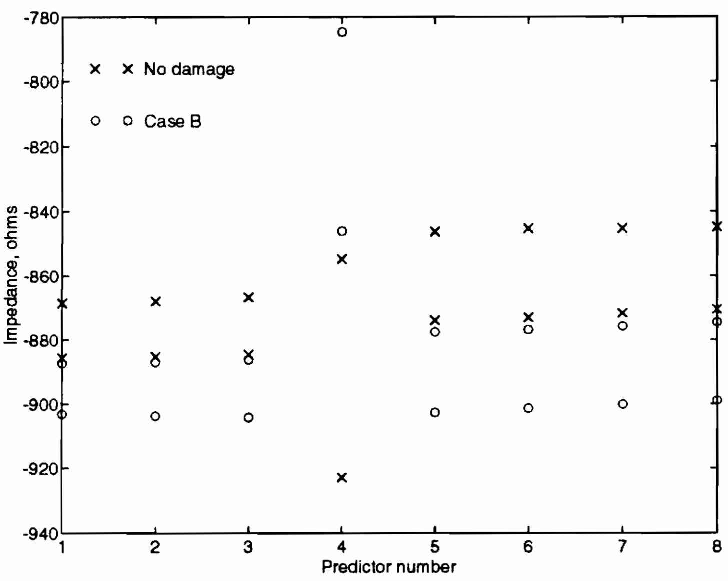

Figure 2.4: Selected rules for DAMAGE CASE B versus NO DAMAGE

the binary comparison of DAMAGE CASE B and NO DAMAGE of Tappert's plate. These bounds are what are called the damage identification *rules*. Thus, there is a set of rules which represent all the possible health state comparisons.

Everything which has been described about the method so far is a one-time operation. Once the rules have been initially generated there is no need to re-create them at a later time. However, it should be noted that if more examples are gathered, they can be added to existing data to refine the statistical accuracy of the rules. This is an extension of inductive learning in its most basic definition.

Now that the rules have been generated, it is time to apply them. A response characteristic from a structure must be input to the ILA. The algorithm then takes the data and looks at the random samples which correspond to the previously determined predictor variables. These points are then tested with the rules to determine if the experimental data points fall within the bounds of any of the health state data. If they match one of the health states for all (or some acceptable percentage of) the rule sets, then the structure's health status is reported correspondingly. For example, if a set of experimental data is sent to the ILA, it will be applied to the rules and the result of each test will be tabulated. If each test resulted in the data falling in the NO DAMAGE rule bounds, then the ILA will report the damage state as such. Figure 2.5 shows an example where the test data corresponded to damage case B.

If after applying the rules, the test is not decisive (i.e. the data did not fall in any of the rule bounds or it fell in different sets of rule bounds), then the ILA will report a health status of DAMAGED BUT OF AN UNKNOWN ORIGIN. This is an important result because it allows the method to detect damage which was not available for dynamic testing or was not expected.

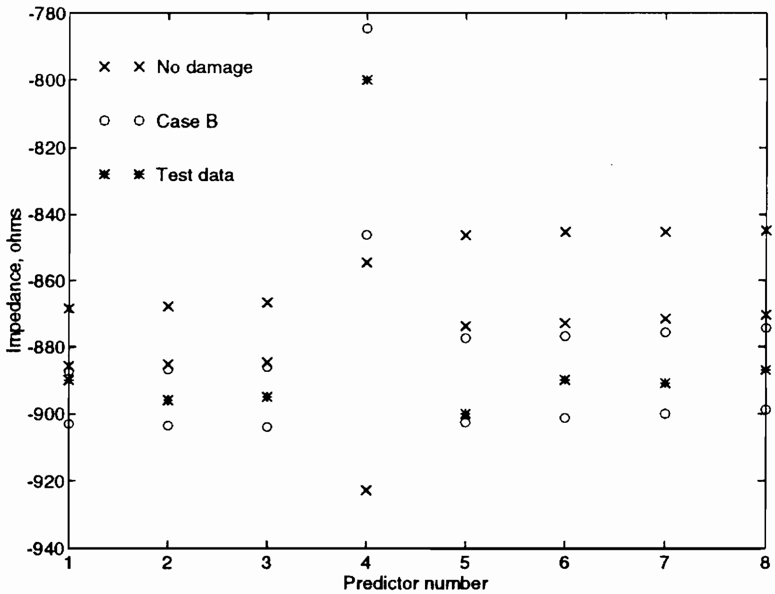

Figure 2.5: Rule application

2.2 Customizing the Method: Part I

ONE OF THE UNIQUE ASPECTS OF DAMAGE IDENTIFICATION USING INDUCTIVE LEARNING IS that the method can be easily customized to fit the needs of a particular problem. And, even with this customization, the final answer which the method outputs has the potential to be correct to a statistically calculable accuracy. For this to happen, all the variances inherent in the particular problem must be accounted for. Assuming that all the sources of variation have been controlled, there are five elements of the inductive learning algorithm's process which define this accuracy. They are the number of experimental data sets in each rule ensemble, the α level set for the confidence bounds, the maximum acceptable significance level chosen from the hypothesis test, the number of predictor variables selected, and the number of rule comparisons which are made in determining the structure's health status.

Each of these five factors are "knobs which can be turned" to customize the ILA to a particular damage identification problem. These, it will be shown later, are also knobs which can be used to assist in alleviating the detrimental effects of variabilities associated with the problem. Four of these knobs are algorithm adjustments which can be easily incorporated into the programming of the ILA. The last involves the testing of the structure itself.

As stated before, the more sets of experimental data which are available to the ILA, the more refined and accurate the damage detection rules will be. The additional sets of data could come from several sources. In the case of damage detection on a single structure, the additional sets could come from data acquisition repetition, the addition of more sensors and sensor locations, or the gathering of different types of response characteristics. The first option will increase the accuracy of the rules because the relative accuracy of the sample mean compared to the true mean is a function of the number of samples taken. The latter two options increase the odds that significant predictor variables can be found due to the addition of more random variables. Probably the limiting factor in choosing any of these methods to refine the damage identification method is how much time, equipment and money are available for the additional testing. Also, there is eventually a point of diminishing returns on the amount improved accuracy which is a result of larger ensembles.

The more important "knobs to turn" in the mind of this author are the algorithm adjustment knobs. The first two adjustments which can be made involve the significance level which is accepted from the hypothesis test and the α level which must be selected within the program. With regard to the significance level, the smaller the significance level is which is accepted, the more confidently it is possible to distinguish between health states. This level affects how reliable the variables which are selected as predictors will be. In some cases, the lowest N significance levels appearing will be accepted as the predictor variables. But in other cases it may be more desirable to select all the variables with a significance level below some threshold.

The α level is the one associated with the confidence bounds which are placed on the sample means. A larger α will result in wider bounds. This in turn means that there is a higher probability of the experimental data falling in one or more of the rule bounds which are created. One potential hazard of making the bounds too large is that mutual exclusivity between the bounds for one damage case versus another may be lost. Figure 2.6 shows an example where a smaller α may be desired. Of course a very small α will result in the damage state being returned as something other than DAMAGE OF UNKNOWN ORIGIN only when there is almost a near assurance that the answer is correct.

The final two knobs which can be used to customize the ILA are involved with

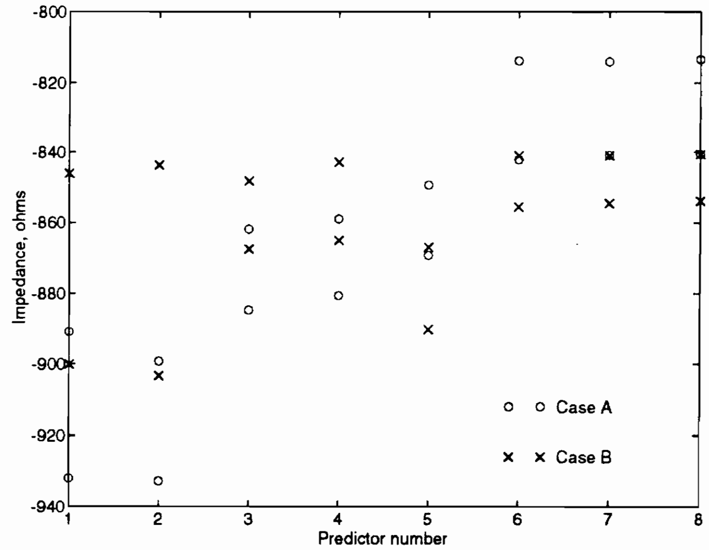

Figure 2.6: Lack of mutual exclusivity due to large α

rule generation and application process. The first is the selection of the number of predictor variables. As stated above, the predictors could be selected as the variables with the N number of lowest significance levels as Tappert did in his work. Or, they could be all the random variables which have a significance level below some given number. The number of predictors represents the number of possible rules which the algorithm can use to detect damage. A greater number of predictors is generally desired since this allows for a greater range of places to look for damage. It also helps to prevent from becoming dependent on a few particular points of a response characteristic to be correct for the task.

The number of predictors is also important in choosing how many of the rules will be applied when attempting to detect damage. It may be desirable to select a number of the rules at random and apply them. Another method might be to pick a health state and test all of the rules for it. If this state passes, then the algorithm is finished. If the health state fails, then the next one is tested and so on. Tappert used a method where all of the rules were applied and one particular damage case had to fall within a specific health state's rule bounds a certain percentage of the time in order for it to be outputted as the answer. The tradeoff involved in these last two customization methods concerns the algorithm's computation and return time.

One of the bonuses of this particular damage detection scheme is that it can be easily implemented on a personal computer. Very little in the manner of computing power is necessary for the algorithm to operate. Tappert's experiment and the work presented here was performed on a 486 MHz PC using MATLAB for Windows as the programming domain that the ILA operated from. The actual rule application time took from a few seconds to as many as a few minutes to return a health status depending on how many predictors were selected and how many rule comparisons were applied.

The amount of accuracy necessary in the damage identification process will depend primarily on the particular problem being acted upon. All of the methods described above can be used to customize the ILA to the problem. There is another usefulness in these "knobs" however. This usefulness becomes apparent when an attempt is made to control or accommodate many of the inherent variances of the damage identification problem which have been discussed before.

2.3 Variability Compensation

THERE ARE SEVERAL DIFFERENT METHODS WITH WHICH THE VARIANCES INHERENT IN THE damage identification problem might be controlled. The ultimate goal is to make the effects of these variabilities on the response characteristics of a structure insignificant with respect to the effects caused by damage. There are several possible solutions to this problem, and invariably they are accomplished through the ILA customization methods described above.

One possible solution would be to separate entirely the effects of the variabilities from the effects of the damage. Another solution might be to determine places where the magnitude of the changes caused by damage are much greater than the variability induced changes. In this way, rule bounds which are sufficiently exclusive from damage case to damage case will detect changes caused by damage (a large response change) but not those caused by unwanted variabilities (a small change). This variance compensation plan incorporates the manipulation of the alpha level in the ILA as well as the selection of significant predictor variables.

Another possible way to alleviate the variability problem might be to examine sets of data over a range of a particular variability and look for trends in the response characteristic changes it causes. Then perhaps the variability effect can be removed. This solution is really just another application of inductive learning. The difference now is that the algorithm is first learning how to identify the variability effect in the response data. Once this effect is found it looks for trends (statistics again) in the variability induced changes and then finds a mathematical way to remove this effect.

It is important to note at this point that in order for the algorithm to be used properly for damage identification in the presence of any variability, the variability must be present when the examples are collected which will be used for rule generation. It is therefore very important to design a data collection scheme which includes all of the variabilities which might be expected in a given problem. However, in order to properly compensate for all of these variabilities it may be necessary to first learn the effect of each individually. It is for this reason that each of the experiments described in the following chapters was designed carefully to isolate one variance at a time, if possible. The process of holding all the factors constant except the one being examined is a common one in statistics and is referred to as *blocking*.

With this said, another important point must be made. Though the data collection may have been carefully controlled, and it is hoped that all the variabilities present were known, one of the advantages in implementing a damage identification scheme on a digital computer is that the ILA is completely ignorant of the problem it is solving. From the point of view of the algorithm, it only sees collections of numbers and applies mathematical operations and statistical rules to them. It returns a completely unbiased answer to the problem since all its work is founded in mathematical principles.

The next chapter describes an experiment where variability compensation was implemented to remove the effect of temperature fluctuations on a damage identification process.

Chapter 3

Environmental Effect

"If you can't stand the heat, get out of the kitchen."

- Harry S. Truman

3.1 Experimental Setup

ONE OF THE INHERENT VARIABILITIES THAT POSES ITSELF IN ALL BUT THE MOST STRICTLY controlled of damage identification problems is that of environmental factors. Fluctuations in temperature, pressure and humidity, to name a few, are factors which need to be studied in the quest to create a practical damage identification method. An extension of Tappert's et al. (1994) experiment was conducted to specifically investigate the effects of temperature on this problem. All of the test apparatus and the testing structure which was used was the same as what Tappert had used for his experiment.

An experiment was performed on a 12" × 12" × 1/8" aluminum plate hung at four corners with fishing line to simulate free-free boundary conditions. The plate was hung horizontally from a wooden frame at a distance of four inches from the mounting points. Two piezoelectric ceramic (PZT) patches were bonded to the plate in diagonally symmetric locations and are referred to as the green sensor and red sensor corresponding to the color of the lead which was connected to each. Figure 3.1 shows the locations of the PZT's on the plate.

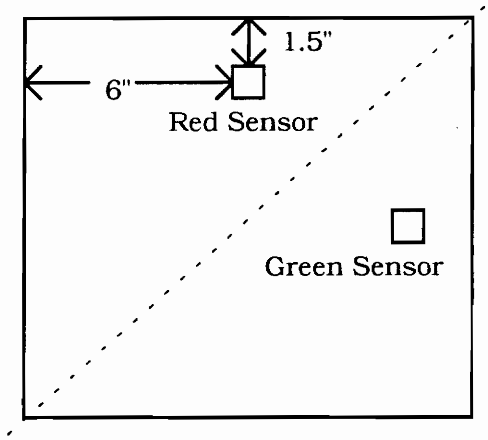

Figure 3.1: PZT locations

Data was collected using a Hewlett-Packard 4194A Impedance/Gain-Phase Analyzer. The analyzer used a single PZT patch as a sensoriactuator to send a one volt peak-to-peak sinusoidal signal, and record a frequency swept electromechanical impedance response. This response was measured in real and imaginary components and electromechanical impedance is measured in ohms. The analyzer's inputs were connected to the sensor lead and a ground wire which was grounded directly to the plate using copper conducting tape. A data transfer output from the analyzer was connected to a DOS based PC with a hardware interface board which facilitated the saving of the response data in .MAT files. The analyzer settings were set in exactly the same manner as was done in Tappert's experiment (medium integration times, eight averages, real and imaginary parts of impedance response). Figure 3.2 shows the entire experimental setup.

Once the data was collected, it was analyzed using MATLAB for Windows version 4.0. A number of .M files were written to perform the various tasks associated with the damage identification problem. Many of these .M files were written by Tappert and

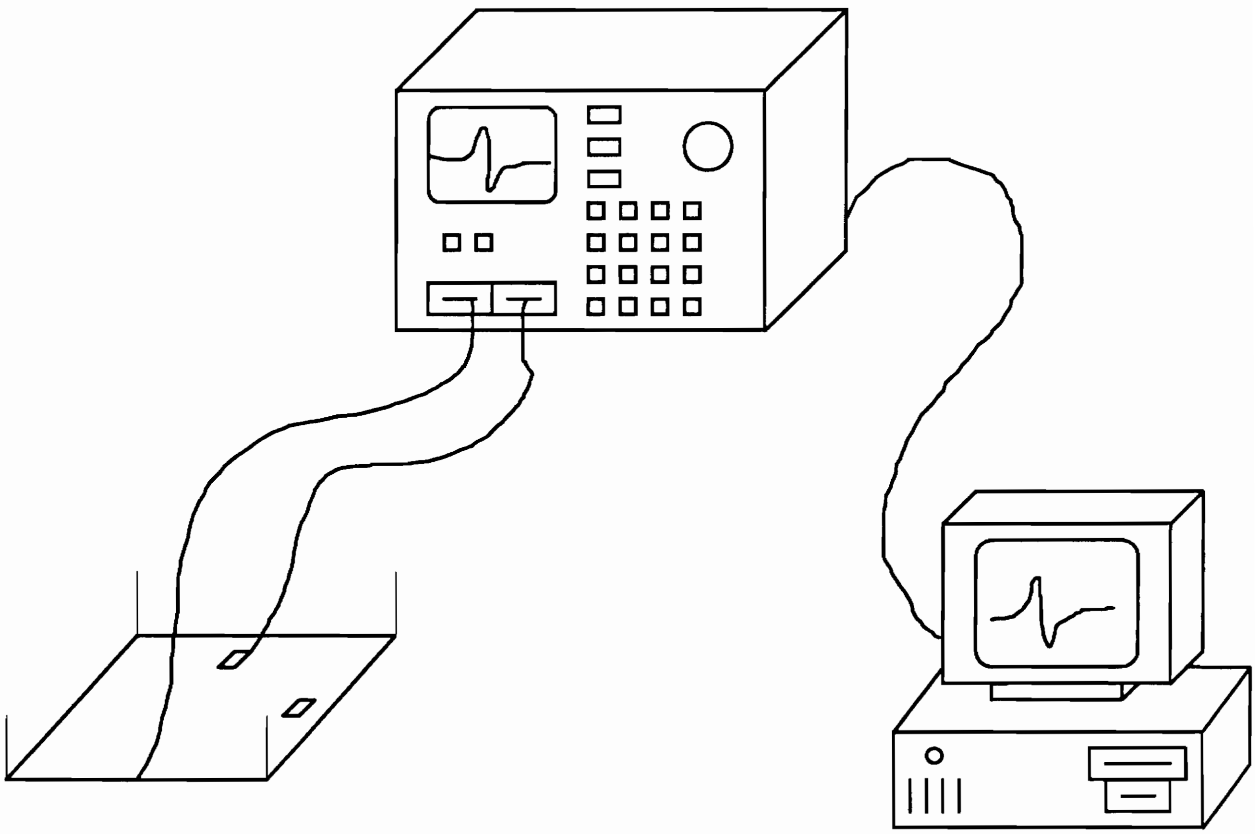

Figure 3.2: Experimental setup

were not changed for this experiment. Other files were changed and added to accommodate temperature compensation as will be explained in section 3.3.

3.2 Temperature Variability Study

THE GOAL OF THE TEMPERATURE VARIABILITY EXPERIMENT WAS TO FIND OUT WHAT EFfects temperature fluctuations had on the response characteristics of a structure and then find the simplest and most effective way to compensate for these effects. The first step in performing the experiment was to determine how to control as many of the inherent variances present in the problem as possible so that the temperature effect could be isolated.

A thermometer was mounted to the experimental testbed so that the temperature in close proximity to the plate could be monitored during experimentation. The plate was hung in its testing position and the green sensor was connected to the impedance analyzer. It was decided to collect response data from the plate at different ambient temperatures rather than attempting to heat or cool the plate to specific temperatures. The room in which the experiment was conducted had a typical ambient temperature fluctuation of about 20-26° Celsius during the course of the day and night. The temperature changes seemed to be a function of the time of day, outside weather, and the number of people working in the room at a given time. The thermostat settings in the room were not altered.

The temperature changes in the room were slow enough that it was assumed that the temperature was constant during data collection. Ten sets of data (the number Tappert used) were collected for each temperature that was examined. This data collection process took about 10-12 minutes (only the green sensor was used) and the temperature was monitored during this collection to assure that it remained constant.

It was desired to collect response data for each Celsius degree increment between 20 and 26 degrees. This process took several days, due to a need for the ambient temperature to be an untested one at the same time that the impedance analyzer was available for use. As a result, an unwanted variability was present in the problem. Ideally, the experiment could be set up at the beginning and not disturbed in any way (with the exception of temperature fluctuations) until all the data collection was completed. However, due to other work being performed using the impedance analyzer, it was necessary to disconnect the plate from the analyzer on several occasions. On these occasions, the plate had to be removed from the mounting structure and the structure and plate had to be moved. This introduced boundary condition variability into the problem.

In Tappert's experiment, the plate was hung and re-hung repeatedly to add boundary condition variability to his experiment. The rules that he generated had this variability incorporated into them, and the method was shown to work in its presence. As a result, it was decided to try to use the same rules which Tappert generated, if possible, for this experiment. By doing this, it was assumed that all the inherent variabilities could be reasonably controlled (by either exclusion or previous incorporation into the algorithm) as long as the plate was hung and re-hung in the same way that Tappert did.

Data collection was performed over a frequency range of 4000-5200 Hz. This range included the range from which the rules data had been collected. All of the data was taken for the case of the healthy plate. It was assumed that the effects of temperature changes on the response characteristics of the plate would be the same regardless of the health state which was present. This assumption will be verified by showing that the temperature compensation mechanism was effective when these damage states were tested later.

Figure 3.3 shows the effect a six degree Celsius temperature fluctuation had on the imaginary part of the impedance response over the frequency range which was tested. It

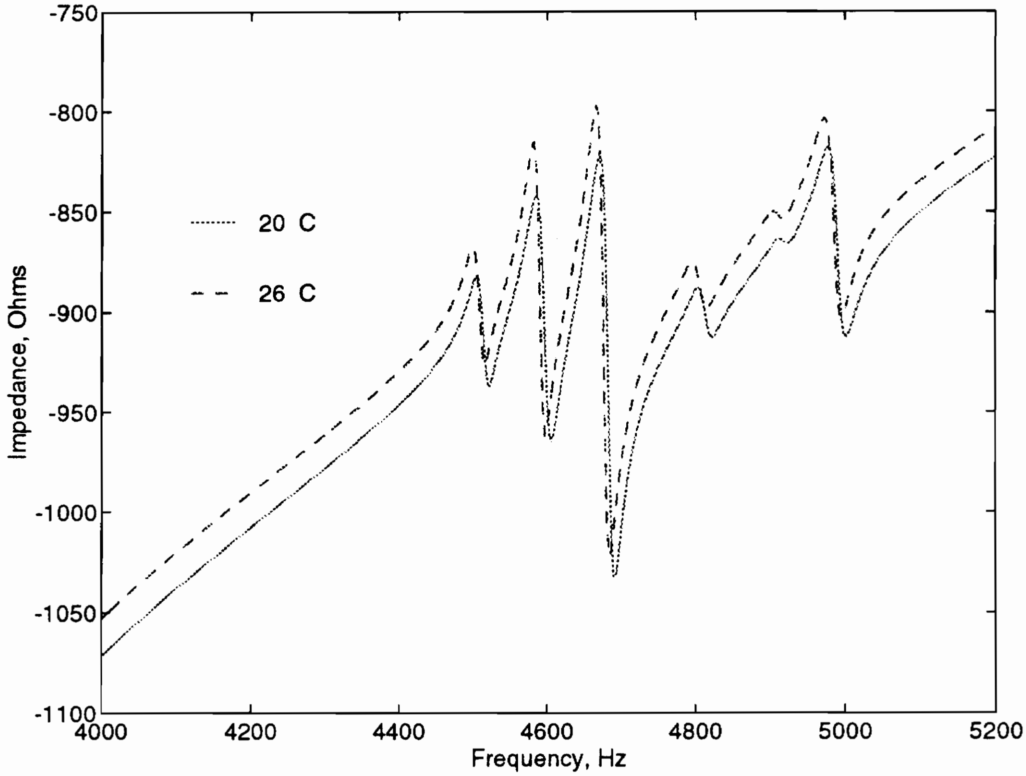

Figure 3.3: Temperature effect: imaginary impedance response

can be seen that the fluctuation caused a shift in the response. The shift appears to be fairly linear in regions away from the natural frequencies of the plate, but is more complex near these frequencies. Also, it has both magnitude dependent and frequency dependent components. Figure 3.4 shows the effect that the same fluctuation had on the real part of the impedance response. It can be seen that this shift appears be primarily horizontal, though a slight magnitude shift may be present as well.

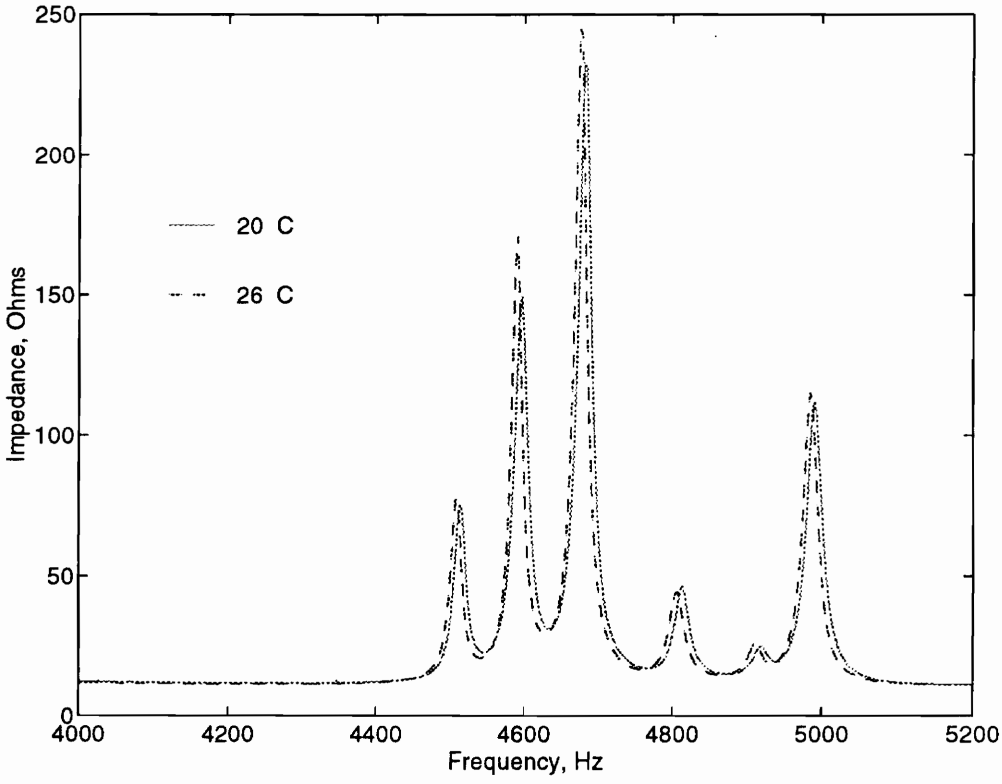

Figure 3.4: Temperature effect: real impedance response

3.3 Customizing the Method: Part II

SEVERAL APPROACHES WERE CONSIDERED AS A MEANS OF COMPENSATING FOR THIS temperature induced shifting. One method would be to add temperature into the definition of the random variables. This would incorporate the temperature effect directly into the damage identification rules the ILA creates. This method has some drawbacks though. Before, the random variables were a function of frequency, sensor, and data type (real or imaginary). With this method, they would now be a function of temperature as well. In order to have a complete set of damage identification rules, response data would have to be collected as it was in Tappert's experiment, except that the process would have to be repeated for every temperature which was expected to be encountered.

This method has one advantage. It would result in a method with a definable accuracy as long as the temperature being tested is one for which a set of damage identification rules had been made. However, the cost of this advantage is an increase in the amount of data collection and computing time which is a factor of how many different temperatures are included in the algorithm's rule generation. Also, the algorithm would not be able to identify damage with this statistical accuracy at any temperatures other than the ones which were specifically prepared for.

Another approach to attacking the temperature problem would be to apply some sort of calibration to the ILA's data in such a way that the temperature effect could be canceled out. This calibration could be either applied to the rules, to the test data, or to both. Since it was desired to use Tappert's rules in their original form, it was decided to work with the test data.

The temperature response information was studied to try to find a pattern in how the shifts were a function of the temperature change. This was not exactly a straightforward task. The shifting of the real and imaginary data was complex in nature. However, the complexities, for the most part, only manifested themselves near the natural frequencies of the plate. These are also the areas that typically had the highest sample variances. Having this prior knowledge of the response behavior from a mathematical point of view, it was known that the ILA typically did not choose points near the natural frequencies as its significant predictor variables.

Because of this knowledge it was decided to try and develop an equation for a linear shift which would approximate the actual shifting that occurred. If the shift was sufficient for areas away from the natural frequencies, perhaps it would be sufficient for all the areas that the significant predictors were being selected from. The data which was collected for Tappert's rules was gathered at a temperature of 25° C, so this was chosen as the reference temperature. For each degree below 25° C that test data was collected, the magnitude of the imaginary test data was shifted upwards by 2.5 units (Ω) and left one frequency bin (1 Hz). The reverse held for each degree above 25° C. No magnitude shift was applied to the real test data, but the same frequency shift applied.

This shift was applied to the temperature data which had been collected. The difference between the 25° C data and the calibrated data was plotted to show the effectiveness of this method. Figures 3.5 and 3.6 show this plot for the imaginary and real response data of the extreme case ($\Delta T = 5°$ C) tested. The calibration of test data was performed each time the algorithm was used. The file TEMPCOMP.M, which is included (along with all other pertinent .M files) in Appendix A, contains the code which applied the calibration.

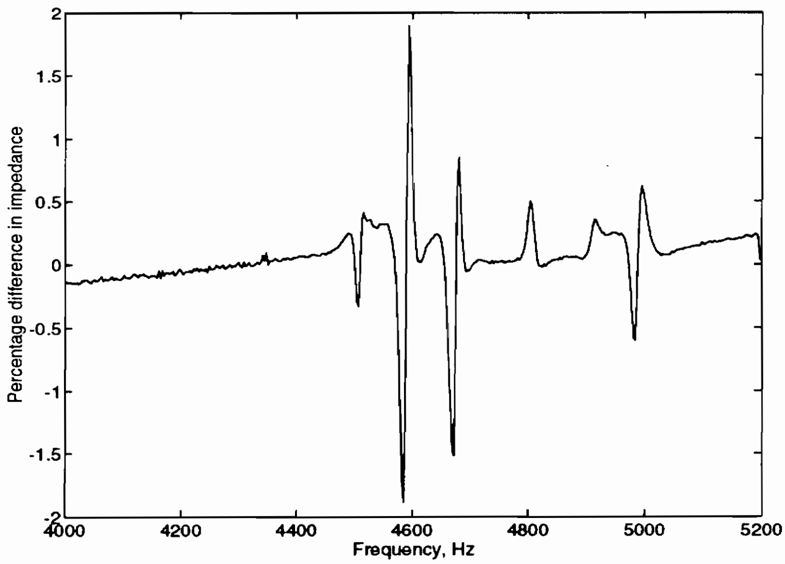

Figure 3.5: Difference between 25° C and calibrated 20° C imaginary impedance

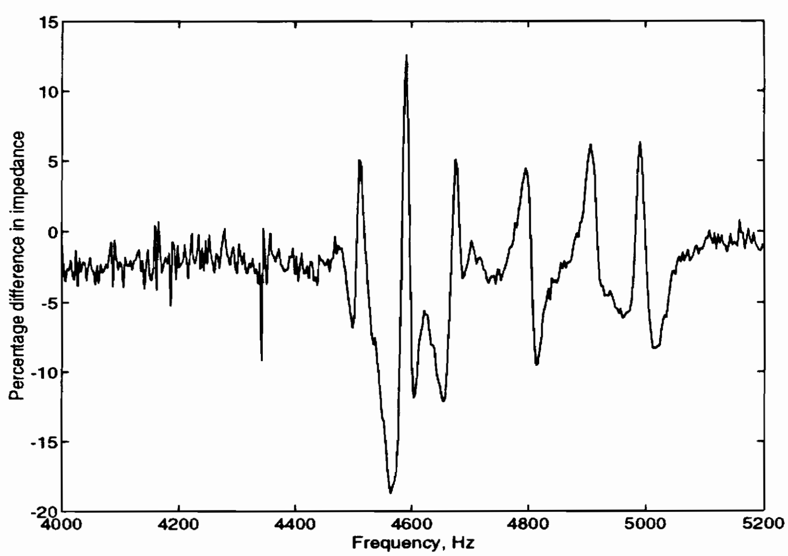

Figure 3.6: Difference between 25° C and calibrated 20° C real impedance

The ILA was modified to incorporate this new piece of the algorithm. Before temperature compensation, the ILA had two basic paths of operation. The first was the rule generation path. This was a one time operation and includes gathering of response data, application of statistical tools (hypothesis testing, confidence bounding), and selection of significant predictor variables. The second path was the test path which only had one function. The function was to take a set of test data, apply the rules and report the resultant health status. Figure 3.7 shows how the test path of the ILA was modified to include temperature compensation. Now, the test data is pre-processed before the rule application takes place.

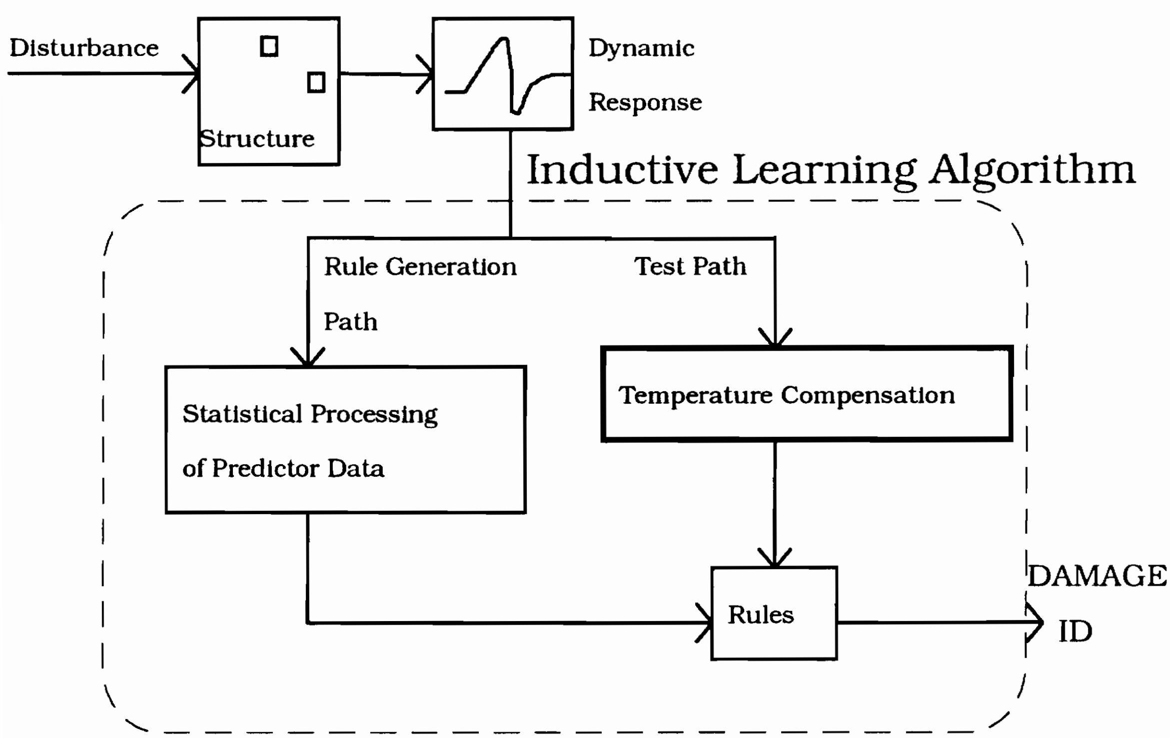

Figure 3.7: ILA path modification

Once the calibration had been developed and the code was written, the method was tested at several different temperatures and for all the damage cases. The temperatures included tests at 19° C, which was a temperature which had not been tested when developing the calibration equations. The tests included temperatures above and below the reference temperature. Also, the tests were performed for all of the possible health states of the plate. Figure 3.8 shows the locations for added mass which defined the damage cases.

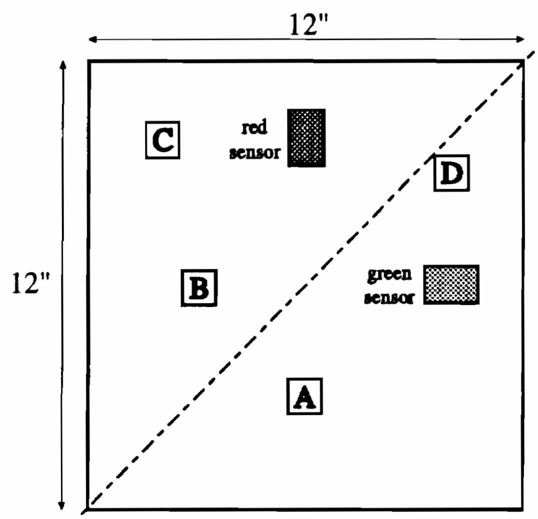

Figure 3.8: Damage case definition

The fact that the ILA with temperature compensation was able to successfully differentiate between these damage cases proves the assumption that temperature induced dynamic response changes would manifest themselves similarly regardless of the health state present on the plate. Table 3.1 shows a sampling of the tests that verified the temperature compensation algorithm.

Table 3.1: Sampling of experimental results with and without temperature compensation

Temperature at which test data was collected (Celsius)	Algorithm result before temperature compensation	Algorithm result after temperature compensation	Actual damage case present
19°	B	A	A
19°	Unknown	No damage	No damage
23°	Unknown	B	B
26°	D	C	C

Though no work has been done to study other environmental factors. It is the be lief of this author that in all likelihood, these factors will cause similar changes in the response characteristics of a structure, and thus they could be similarly compensated. However, it is vital to note that until actual research is conducted to study this, the above *assumption* is no more than that. One of the advantages of this method over some other damage identification techniques is that through the reliance entirely on statistics, no biases or "guesses" enter into the detection process. In order to keep this purity in the method, actual research needs to be conducted for every variance which it is desired to include in the damage identification problem.

It is with this in mind that the problem of manufacturing variability was approached...

Chapter 4

Manufacturing Variability Effect

"Where order in variety we see, and where, though all things differ, all agree."

- Alexander Pope

4.1 Experimental Setup

PERHAPS THE MOST IMPORTANT VARIABILITY TO CONSIDER WHEN ATTEMPTING TO CREATE a practical damage identification method is manufacturing variability. In many applications, it is desired to perform damage identification on a number of structures that are essentially identical. In other cases, a single structure may have old parts periodically replaced with new ones as they wear out. These are both cases where it is vital that the damage identification method being used is able to take into account the manufacturing variabilities which are unavoidably present.

Manufacturing variability can come from several sources. This variability includes differences in the size and shape of a structure due to manufacturing tolerances and methods. It also includes differences in material compositions. It can be brought on later in the life of a structure if the structure is disassembled and then reassembled, or if old parts are replaced with new ones. The reason why manufacturing variability is so important to implement into damage identification is that its presence is nearly unavoidable in practical applications. Because of this importance, an experiment was designed to study the effects of manufacturing variability on damage identification.

The experiment was performed with much of the same equipment as was used in the temperature study. Again, the Hewlett Packard 4194A Impedance/Gain Phase Analyzer was used to gather real and imaginary impedance data over a frequency range. However, this time several plates were used in an effort to introduce manufacturing variability information into the collected data.

Fifteen plates were manufactured out of a 1/16" sheet of aluminum. The plates were cut into 12" × 16" rectangles with a shearing machine. This is a fairly imprecise manufacturing method which was desired so that the manufacturing variability effect would be amplified in the experiment. To further "manufacture" the plates, 1/16" holes were drilled in the upper two corners of the plate at a distance of 1/4" from the top and sides. These holes would be used for mounting the plates. Finally, four PZT's were bonded to the plate in a parallelogram formation as shown in Fig. 4.1. Similarly to the temperature experiment, the PZT's are referred to by the color of the lead attached to them.

The impedance analyzer was used in exactly the same manner as in the temperature study with respect to its equipment settings (number of averages, integration time, etc.). This time, however, a different frequency range was selected. The range from 4000-4400

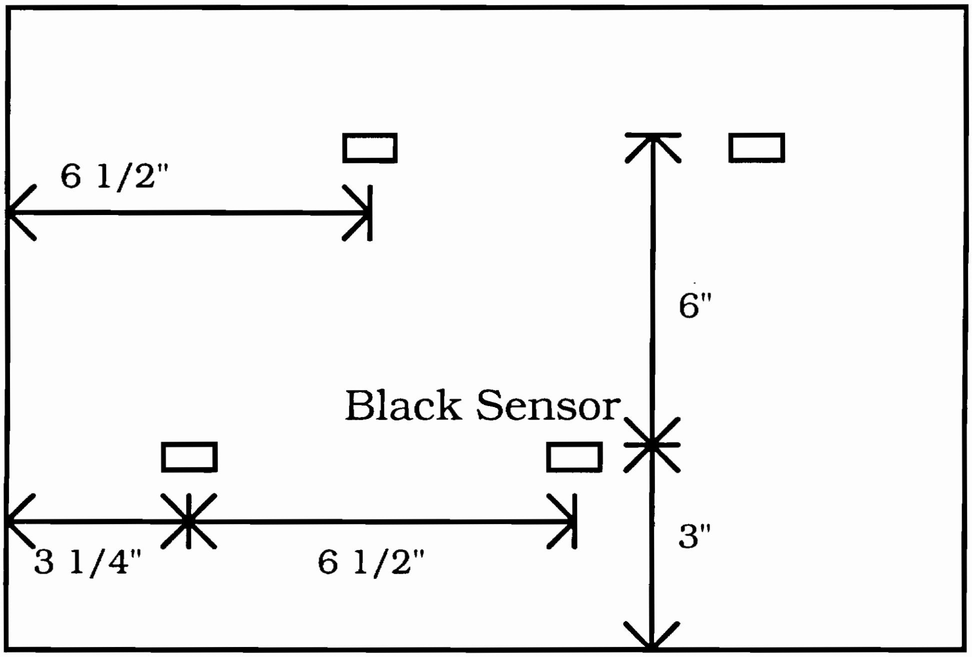

Figure 4.1: Plate sensor arrangement

Hz was used, and an explanation of how it was chosen will be given in the next section. Again, the data was collected with a 1 Hz resolution.

During data collection, the plate was hung vertically from fishing line to simulate free-free boundary conditions. The PZT lead was connected to one of the analyzer's inputs. The other input was connected to a ground wire which was attached to the top center of the plate with copper conducting tape. Each plate had its own ground wire, and this attachment contributed to the manufacturing variability as well. Figure 4.2 shows the experimental setup, including the DOS based PC which was interfaced with the impedance analyzer so that the collected data could be saved into .MAT files. Data analysis was performed using .M files coded to run on MATLAB for Windows version 4.0.

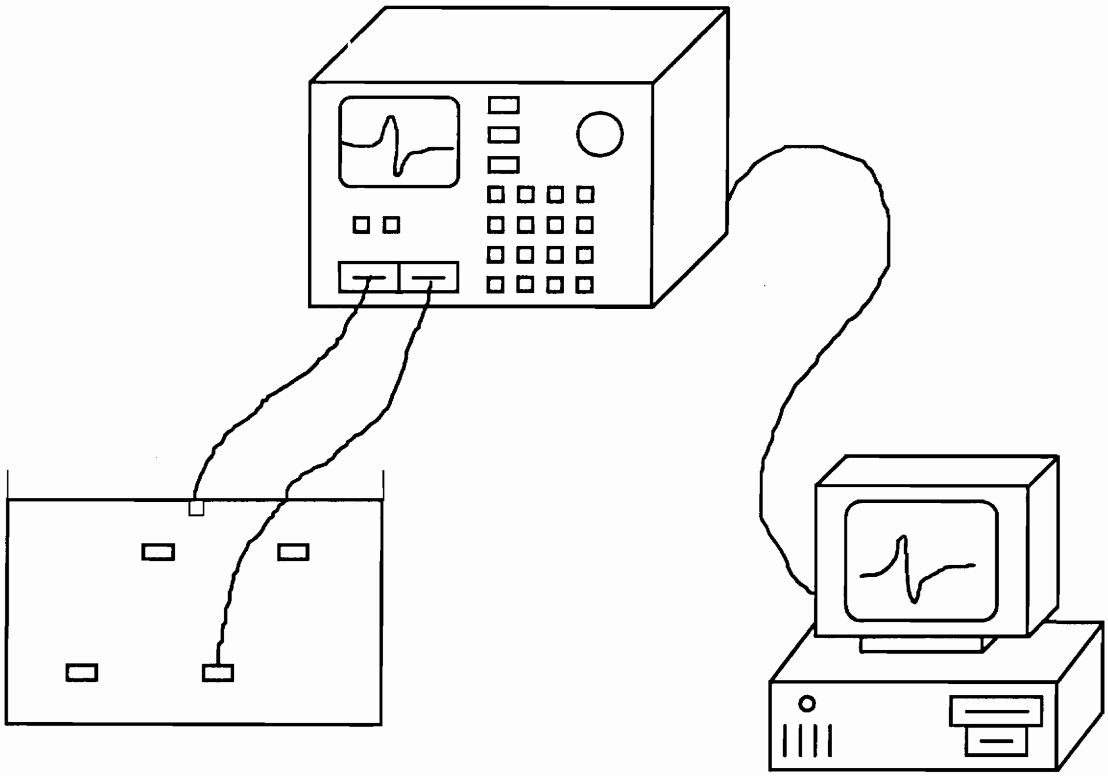

Figure 4.2: Experimental Setup

4.2 Manufacturing Variability Study

THE GOAL OF THIS PARTICULAR EXPERIMENT WAS TO FIND OUT WHAT EFFECT MANUFACTURING variability had on the response characteristics used to identify damage. The first step in determining this effect was to attempt to control all of the other inherent variabilities in the problem so that manufacturing variability would be isolated or blocked.

The first variabilities to control were the environmental ones. The temperature was carefully monitored during experimentation to assure that it remained a constant during data collection. Also, all of the data collection for all the plates used in the experiment was performed in one session in an effort to keep the environmental factors constant as well. In this way, it is assumed that the air pressure and humidity were constant since there were no drastic weather changes between tests. (The work was performed in an air-conditioned environment as well). The experimentation was performed at night when other activity in the building was at a minimum.

In previous work, the plate was hung and re-hung between data sets to add boundary condition variability into the response information. It was desired to remove this variability entirely. However, this was not possible since a number of different plates were being tested and they each had to be hung at least once. Therefore, it was decided to hang each plate only once and collect all of the data sets from them at one time. In this way, boundary condition variability was present, but hopefully minimized.

Before data could be collected, a frequency range had to be selected for testing. It was desired to choose a 400 Hz range due to the fact that the impedance analyzer could only take 401 points of data at a time. Since the 1 Hz bin spacing of past experiments had worked so effectively, it was hoped to continue with the same resolution in this experiment. The first step in selecting a range was to look at a large (~5000 Hz) frequency sweep. Then the range could be whittled down to 400 Hz by consistently stepping down the sweep size in "promising" areas. These areas were ones in which the response had a part which was particularly active (many resonances and anti-resonances) and a less active part (fairly flat response). By selecting an area with these characteristics, the ILA will have a good representation of the areas in which manufacturing induced (and damage induced) changes may exhibit themselves. Figure 4.3 shows an example response in the 4000-4400 Hz region which was finally selected.

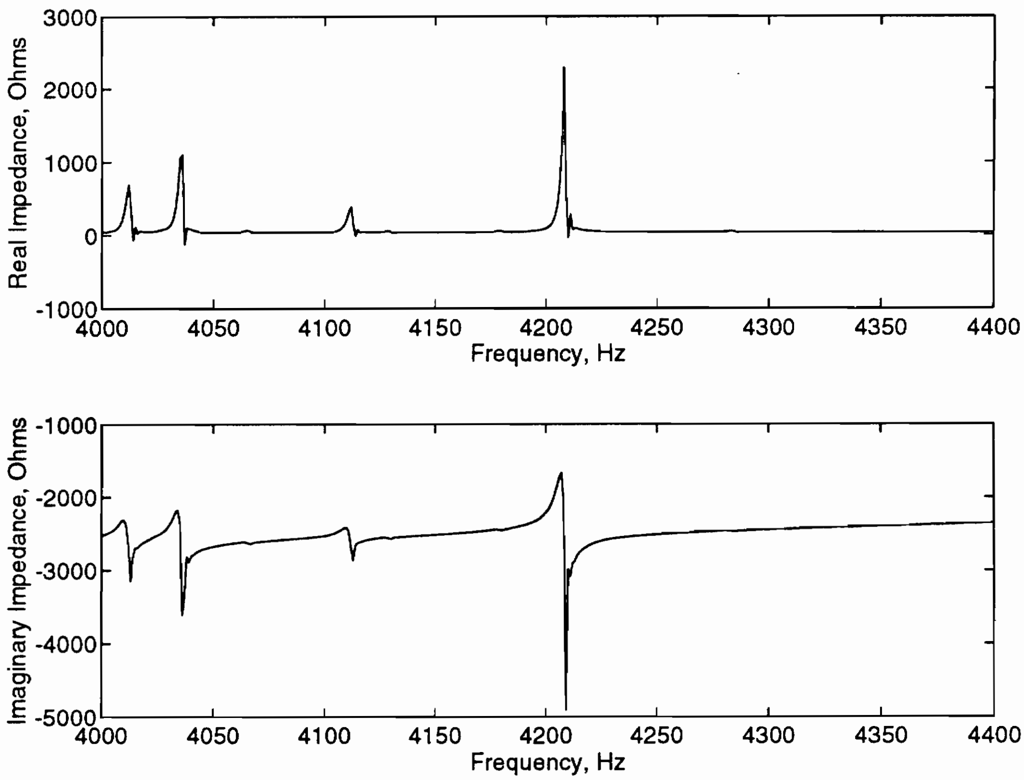

Figure 4.3: Response in 4000-4400 Hz frequency range

The next step was to decide the best way to introduce manufacturing variability into the response data which would be used to generate the damage identification rules. It was decided to select three plates at random to use for the rule generation. More plates would result in better rules due to the addition of more information for the ILA to sort through, but keeping our goal of practicality in mind, it was decided to use a smaller number. In many applications, it may be possible to test only a very small percentage of the structures of interest when generating rules.

Once the three plates had been randomly selected, each was hung and had ten sets of impedance response data collected from it. All the testing was performed when the ambient room temperature was 23.5°C. For the healthy (NO DAMAGE) case, each plate was hung only once and was not taken down until all the necessary data had been collected from it. The same mounting location, wires and hangers were used for each plate. Only the black sensor was used for data collection. This was done out of a desire for simplicity and because it was strongly believed (based on past experience) that only one sensor would be necessary to identify damage.

When the healthy data had been collected, each plate was damaged by drilling a 3/8" hole in it at a pre-determined location. The location was the same on each plate and is shown in Fig. 4.4. The same three plates were then re-hung and DAMAGE case data was

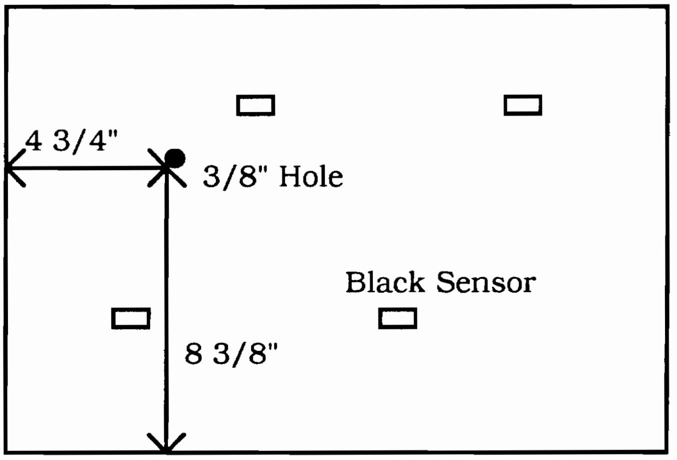

Figure 4.4: Hole location for DAMAGE case

collected in exactly the same manner as for the NO DAMAGE case. With all the necessary data collected, it was time to create the damage identification rules.

4.3 Customizing the Method: Part III

IN THE CASE OF TEMPERATURE'S EFFECT ON THE STRUCTURE'S RESPONSE CHARACTERIStics, there were several obvious choices of how to begin looking for compensation methods. This was due to the fact that temperature's effect was fairly simple. Temperature fluctuation is a variability that is easy to quantify and therefore, the compensation scheme can follow from this quantification. Basically, a correlation between ΔT and the change in response characteristic data was found and exploited. The effect of manufacturing variability is much more complex however.

Manufacturing variability is not a quantifiable variability like temperature is. There is no scale with which to measure it. Because of this, there is no reason to expect that the effect of manufacturing variability on a structure's response characteristics will be any easier to quantify. For this particular experiment this was the case.

Figures 4.5 and 4.6 show a few of the response sets from the different plates superimposed on one another to demonstrate that they are quite different even though they were taken from supposedly identical plates. Granted, there is a good bit of similarity in the curves, but to what specifically should the differences be attributed to? The PZT placement? Irregularities in the aluminum's composition? It is difficult to say.

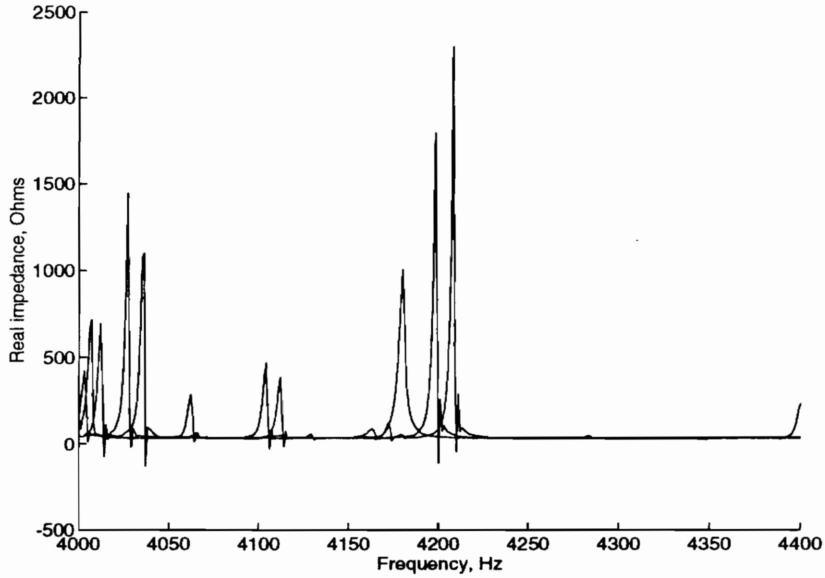

Figure 4.5: Superimposed real responses from three plates

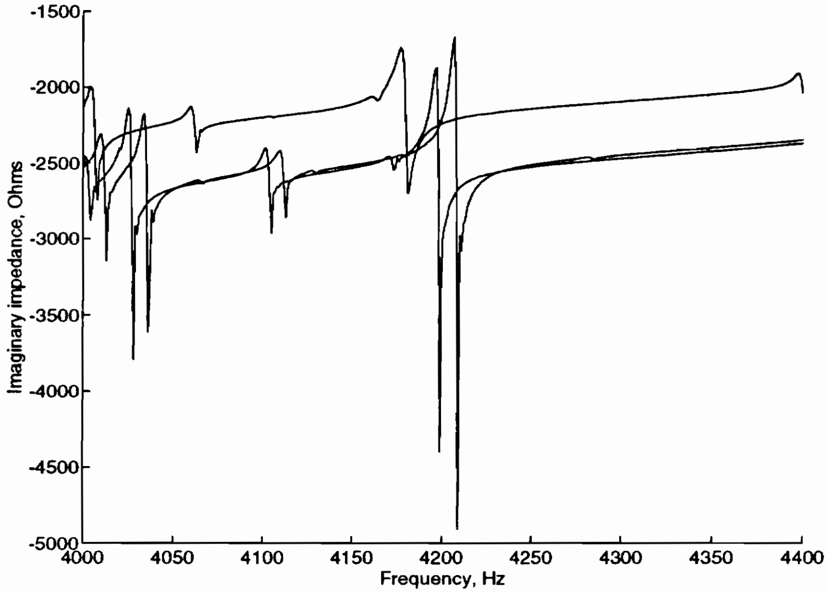

Figure 4.6: Superimposed imaginary responses from three plates

The task is much more difficult than before, but it is not impossible. It was decided to begin by creating a set of rules for each of the three plates individually and comparing them. The significance level was calculated for each frequency point for each plate. But it had to be decided how to select the predictors which would constitute the rules. In the past experiments, the N random variables with the lowest significance levels were selected as the predictor variables.

The five lowest significance levels for each of the real and the imaginary parts of the impedance response data were selected as the best places to discriminate between the DAMAGE case and the NO DAMAGE case. But when this was done for the three plates, each one had different predictors. Not a single predictor of the ones selected for one plate was even in close proximity to a predictor which was selected for another plate. This presented a problem. There was no way to tell that a good place to look for damage in one plate's response would yield the same success on another plate.

A new method of rule generation needed to be created. Plots of the significance levels showed that a large number of predictors were of a magnitude less than about 1×10^{-5}. Significance level is measured on a scale from 0-1 with zero representing a 100% assurance that the true means being compared in the hypothesis test are not equal. In terms of the experiment, a significance level of zero means that the true means of the data representing two different health states are not equal, and therefore can be discriminated from one another.

When the algorithm selected the lowest significance levels, some were on the order of 1×10^{-16}. This provided an extremely high assurance that different health states could be discriminated at that particular frequency bin. However, since the significant bins were not consistent from plate to plate, it was decided to relax our desire to use the absolute lowest significance levels in favor of trying to find frequency bins which were good for all of the plates.

A threshold of 1×10^{-5} was set for the significance level. While this is several orders of magnitude larger than some of the previously selected significance levels, it is still small enough that there is a high degree of assuredness that the damage cases could be discriminated between at points below this level. A routine called PVINTRSC.M was coded which would examine each frequency bin and check to see if the significance level for each plate was less than or equal to the threshold. If it was, then that bin was selected as a predictor location. If not, then it was passed. All of the code which was written for this experiment is included in Appendix A in the order in which it was used during the damage identification process.

Once the significant predictors were selected, it had to be decided how to use them. If the predictor variables were to be bounded for the data from each plate, then there would be as many complete sets of rules as there were plates which had been tested. But due to the disparity of the response characteristics (Figs. 4.5, 4.6), there was doubt that the rules would all be mutually exclusive for the NO DAMAGE versus DAMAGE comparison, and at the same time be nearly the same for the comparisons from plate to plate.

The bounds were plotted to verify this doubt. Figure 4.7 shows that the rules for the predictors selected do not make a lot of sense. In some cases, the healthy bounds for one plate fell in between the damage bounds for the other two.

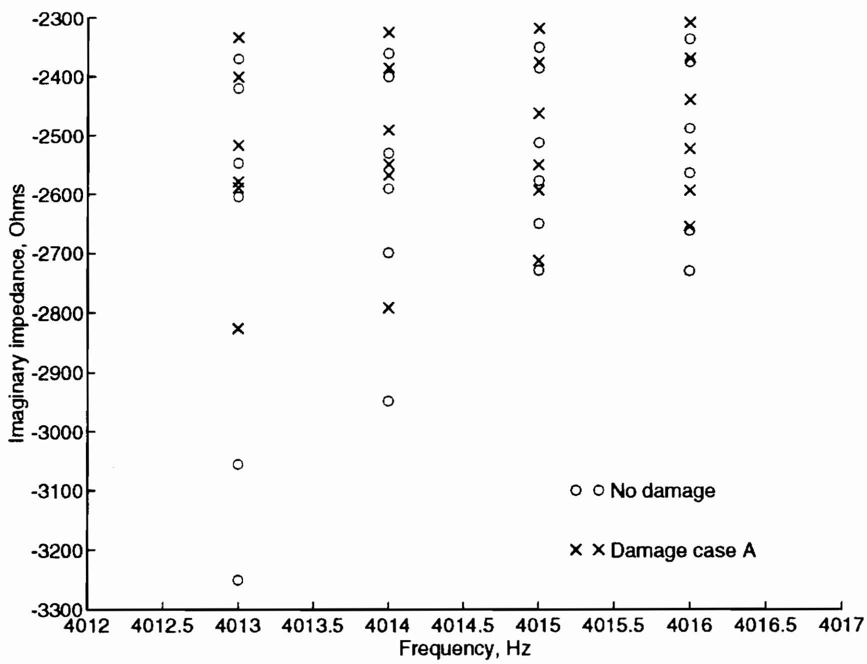

Figure 4.7: NO DAMAGE versus DAMAGE rules for imaginary data of three plates

To understand why the rules look this way, the meaning of the significance level must be fully understood. As stated before, the significance level is a measure of how well it is possible to discriminate between two means (or in our case sample means). The significance level is a function of the number of samples, the sample means, and the sample variances of the data being compared. When a low significance level is returned, it means that a certain combination of these factors has probably occurred. In general, a low significance level corresponds to an area where not only are the sample means fairly separated (an obvious indicator that the true means are different), but also where the sample variances are relatively small. In essence, the ILA picks areas of the response where the health states are *stably different.*

Getting back to Fig. 4.7, the rules seem fairly chaotic. The ILA has selected predictor locations where the health states are stably different on each plate. However, there is no assurance that the responses will also have the same magnitude at these points. Thus, while each set of rules might appear good, taken as a group they are not useful due to the significant variations caused by the manufacturing variability.

All is not lost however. The ILA has selected good areas to look for damage. For all of the plates, the differences between NO DAMAGE and DAMAGE are statistically sound. It is the rules themselves which need to be re-examined. A new approach to rule generation needs to be developed.

When a new structure is brought in to be tested, it is first tested to get a set of healthy response data. It is from this data that the rules will be formed. The data at the significant predictor locations is bounded and monitored. As new sets of response data are periodically taken, the rule locations can be checked to assure that the data has not moved outside the bounds. If this occurs, then the algorithm can report a DAMAGE status.

Three new plates were tested and then, after a time, damaged to test the method. There were nine rules which had been selected (four imaginary and five real). These were tested and if the test data fell in at least eight of the rule bounds, the plate's status was reported as NO DAMAGE. When the data began moving outside of the bounds, a DAMAGE flag went up. In each case, while the plate remained healthy, the ILA continued to respond with a NO DAMAGE answer. Once damage (a hole) was introduced, the method detected it each time. Table 4.1 shows a summary of the results.

Table 4.1: Manufacturing variability experimental results

Plate number	Damage state	Number of rule successes	Algorithm result
4	No damage	9	No damage
10	No damage	8	No damage
12	No damage	9	No damage
4	Damage	4	Damage
10	Damage	5	Damage
12	Damage	6	Damage

The modifications to the ILA were successful in adding manufacturing variability to the problem (the main goal of the experiment), but something was lost as well. Once again, practicality is an issue. Manufacturing variability compensation was a key goal in making the method more practical. But it came at the cost of being able to distinguish between various damage cases. In fact, the method described above is only valid for one damage case. In some applications this is sufficient, but in many others a more robust method will be required. It is important to keep this in mind as the conclusions are discussed. For while the experiments presented herein were successes, there is still much work to be done...

Chapter 5

Conclusion

“We should be careful to get out of an experience only the wisdom that is in it.”

- Mark Twain

5.1 Experimental Results

EXTENSIVE WORK HAS BEEN CONDUCTED IN RECENT YEARS TO FIND NEW DAMAGE IDENTIFICATION methods. In many applications, structures are too costly and difficult to access for visual inspection. In other applications it is desired to have a method which can constantly detect damage on-line while a system is operating. Inductive learning has been shown to be an effective method for performing damage identification. This method can effectively identify and discriminate between expected damage cases on a structure by performing a statistical signature analysis of the structure’s response characteristics. Inductive learning can also detect damage which is of an unexpected variety, and, of course, it can verify that a structure is healthy.

As research is conducted to create new damage identification methods, it is important to keep in mind the purpose and goals of the research. In the mind of this author, practicality is the most important goal in designing a damage identification method (aside from the obvious requirement that it works). For a method to be truly useful, and therefore for the research effort to have been worthwhile, it should be able to be implemented practically. It is with this in mind that a study of the effect of variabilities on the damage identification problem was begun.

A large part of making a damage identification method practical is making it robust to the many variabilities inherent in damage identification. These variabilities come from many sources including sensor and equipment variability, boundary condition variability, environmental variability, and manufacturing variability, to name a few. Through careful experimentation, these variabilities can be isolated and examined to find their effect on damage identification.

An experiment was conducted to examine the effects of temperature fluctuations on the problem. Data was collected in a manner such that response information at a range of temperatures could be studied. From this examination, a way was proposed to modify the Inductive Learning Algorithm (ILA) to account for this variability. A calibration equation was developed which shifted response data to a reference temperature before applying the damage identification rules. The method was proven to be effective in detecting several different damage cases at a range of temperatures. It was successful at temperatures which had been encountered during rule generation and at new temperatures as well.

A second experiment was conducted to study the effects of manufacturing variability on the damage identification problem. This is perhaps the most important variability to conquer in terms of making the method practical. In many applications, it would be desirable to have a method which can identify damage on a number of structures which are seemingly identical, but in reality have many differences. As in the temperature experiment, a method of blocking manufacturing variability was developed from which it was possible to collect response data with this variability's effect incorporated into it. A new rule selection method was developed which finds the best predictors over a range of structures rather than the best ones for each structure individually. A new rule application method was then developed where the rules were based on test data rather than coming entirely from algorithm training. This method was tested on several structures and worked effectively.

5.2 General Conclusions

SEVERAL SMALLER GOALS WERE ACHIEVED IN THE QUEST TO CREATE AN IDEAL DAMAGE identification method. In the second experiment, relatively small holes were drilled in the plate to simulate damage. In much of the past experimentation, either gross damage needed to be present for the methods to be effective or something like an added mass was used to simulate damage (as was done in the first experiment). Before too much further research was conducted on the use of inductive learning for damage identification, it was important to be sure that it could identify realistic damage (e.g. holes) of a small magnitude. This is the type of damage which is most commonly searched for in many applications.

Two other goals which were met were the addition and accommodation of two of the more common variabilities to the problem. Inductive learning proved to be a capable method for identifying damage in the presence of these variabilities. But more importantly, it is possible to begin to see a general method for adding any expected variability into the damage identification system.

It seems that there are really two different categories of variability which are inherent in the damage identification problem. The first category is the one which was referred to in the introduction as the cause-effect variabilities. Temperature is an example. These variabilities are measurable in some manner, and it seems that their effect on the response characteristics used in damage identification is, as a result, also measurable. Once it is determined how to measure the variability and its effect, it is a simple task to eliminate the detrimental effects.

The second group of variabilities is the non-quantifiable variabilities. These are the ones, such as manufacturing variability, which have an unobvious effect on the response characteristics. When dealing with these variabilities, it is necessary to rely much more heavily on the ILA's strength: statistics. The removal of the detrimental effect of this type of variability is facilitated by a careful selection of predictor variables. The significance level from the hypothesis test becomes even more important. If the right rules (the best rules for a given set of conditions) are selected and applied in a logical manner, then the probability of the method succeeding is very good.

Of course the most important factor for dealing with either type of variability is to be sure that it is properly represented in the data which the ILA analyzes. Because the ILA is a digital computer, it only sees a collection of (meaningless) numbers. It cannot make any assumptions about the numbers that researchers in the past have made (such as assumptions that resonance peak shifts were indicative of damage). Because the method is unbiased from the start, it is important to present it with data which is equally as unbiased. This is achieved by making sure that all the expected variabilities are present when data is collected.

If all the steps are properly implemented, then the true beauty of the use of inductive learning as a damage identification tool can shine through. The method has been proven to select areas, which are not necessarily obvious, to look for damage. And, it is possible to define how well the method works due to its reliance on statistics to generate an answer.

5.3 Customizing the Method: Part IV

NOW THAT INDUCTIVE LEARNING HAS BEEN SHOWN TO BE AN EFFECTIVE MEANS OF PERforming damage identification, there are several new directions in which related research may progress. In the same spirit as this work was performed, other inherent variabilities could be studied in an effort to make a more robust damage identification method. Another step towards adding robustness would be to show that the method works for different types of sensors and response data. Perhaps there is a new type of rule bound which would be effective when implemented with time domain response data. The algorithmic method itself is robust, and it is simple enough that it should be easy to modify for new data types.

Further work is needed on the manufacturing variability problem. Now that it has been shown to be a conquerable variability, work needs to be conducted in an effort to regain the algorithm's ability to distinguish between several damage cases in the presence of manufacturing variability.

Another interesting path of research might be the conversion of the method from a damage *identification* method to a damage *detection* method. Presently, the algorithm is only able to identify expected damage cases. Any other damage is simply reported as DAMAGE OF AN UNKNOWN ORIGIN. It would be interesting to investigate the possibility of teaching the algorithm to detect the presence of damage and then zero in on its location. Near-field signature analysis is one possible tool which could be studied towards this end.

One other bridge which must be eventually crossed in making the ILA a practical damage identification method is the dependence on binary comparisons in the hypothesis test. Because of how the hypothesis test operates, at least two health states are required in order to develop a significance level and ultimately the damage identification rules. But in some applications, it may not be feasible to test the structure in any form except its healthy state. Perhaps there is another way to generate rules in such a way that the statistical significance is not lost, but only the healthy structure is required.

Regardless of where the research heads, the damage identification problem will continue for a long time.

References

"The sole substitute for an experience which we have not ourselves lived through is art and literature."

- Alexander Solzhenitsyn

Bartlett, John, and Justin Kaplan. *Bartlett's Familiar Quotations*. Little, Brown and Company, 16th ed., 1992.

Bendat, Julius S., and Allan G. Piersol. *Random Data: Analysis and Measurement Procedures*. John Wiley & Sons, 2nd ed., 1986.

Deval, Y., S.G. Ducouret, and J.P. Dom. (1993). "Ratiometric Temperature Stable Current Reference." *Electronics Letters*. Vol. 29, no. 14, July 27, 1993, pp. 1284-1285.

Elkordy, M.F., K.C. Chang, and G.C. Lee. (1993). "Neural Networks Trained by Analytically Simulated Damage States." *Journal of Computing in Civil Engineering*, Vol. 7, no. 2, pp. 130-145.

FirstMark Technologies, Ltd. (1990). *Knowledge Seeker User's Guide*.

Ganino, Anthony J. (1993). "Damage Identification Using Artificial Neural Network Techniques." CIMSS. (unpublished)

Ganino, Anthony J. (1994). "Effect of Increased Damage and Temperature on the Structural Impedance of a Composite/Aluminum Plate Specimen." (unpublished)

Hickman, G.A., J.J. Gerardi, and Y. Feng. (1990). "Application of Smart Structures to Aircraft Health Monitoring." *International U.S./Japan Conference on Adaptive Structures*. Maui, Hawaii, November 1990, pp. 966-984.

Hinkelmann, Klaus, and Oscar Kempthorne. *Design and Analysis of Experiments: Volume I: Introduction to Experimental Design*. John Wiley & Sons, Inc., 1994.

Hofer, B. (1987). "Fibre Optic Damage Detection in Composite Structures." *Composites*. Vol. 18, no. 4, September 1987, pp. 309-316.

Hogg, Robert V., and Johannes Ledolter. *Applied Statistics for Engineers and Physical Scientists*. Macmillan Publishing Company, 2nd ed., 1987.

Kudva, J.N., N. Munir, and P.W. Tan. (1992). "Damage Detection in Smart Structures Using Neural Networks and Finite-Element Analyses." *Smart Mater. Struct.*, **1** (1992), pp. 108-112.

Kung, D.N. et al. (1989). "A Time Domain System Identification Technique Based on Impulsive Loading for Damage Detection." *Proceedings of the 8th International Conference on Offshore Mechanics and Arctic Engineering*. The Hague, Netherlands, March 1989, pp. 307-313.

Lindner, Douglas K., Gregory Twitty, and Richard Goff. (1993). "Damage Detection, Location, and Estimation for Large Truss Structures." AIAA-93-1484-CP, pp. 1539-1548.

Liu, C.T. and B. Tang, (1993). "Investigating Damage Fields in a Particulate Composite Material Using Real-Time X-ray Technique." *X-ray Detector Physics and Applications II*. SPIE Vol. 2009, pp. 240-245.

Manning, R.A. (1994). "Damage Detection in Adaptive Structures Using Neural Networks." AIAA-94-1752-CP. pp. 160-172.

Povich, Clinton R. and Tae W. Lim. (1994). "An Artificial Neural Network Approach to Structural Damage Detection Using Frequency Response Functions." AIAA-94-1751-CP, pp. 151-159.

Roitman, N. et al. (1992). "Identification of Offshore Platform Structural Damage Using Modal Analysis Techniques." *Mechanical Systems and Signal Processing*. 1992, **6**(3), pp. 287-295.

Tappert, Peter M., and H.H. Robertshaw. (1994). "Manipulations of Dynamic-Response Data for Use in Inductive Learning Methods to Identify Damage." *Proceedings of the 1994 North American Conference on Smart Structures and Materials*. Orlando, Florida, February 1994.

Tappert, Peter M. (1994). "Damage Identification Using Inductive Learning." Master's Thesis, Virginia Polytechnic Institute and State University.

Appendix A

TEMPCOMP.M

```
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% This routine performs the temperature calibration to 25°C (reference)
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

if temperature == 26
  tI = tI - 1.5;
elseif temperature == 25
  disp('Temperature is already at reference')
elseif temperature == 24
  tI = tI + 1.5;
elseif temperature == 23
  tI = tI + 3.5;
elseif temperature == 22
  tI = tI + 8;
  tI = [tI(4:401,:) tI(401,:) tI(401,:) tI(401,:)];
  tR = [tR(4:401,:) tR(401,:) tR(401,:) tR(401,:)];
elseif temperature == 21
  tI = tI + 13;
  tI = [tI(4:401,:) tI(401,:) tI(401,:) tI(401,:)];
  tR = [tR(4:401,:) tR(401,:) tR(401,:) tR(401,:)];
elseif temperature == 20
  tI = tI + 13;
  tI = [tI(4:401,:) tI(401,:) tI(401,:) tI(401,:)];
  tR = [tR(7:401,:) tR(401,:) tR(401,:) tR(401,:) tR(401,:) tR(401,:) tR(401,:)];
else
  disp('The temperature entered was not calibrated for compensation.')
  disp('No compensation has been performed')
end
```

LOADDATA.M

```
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% This routine loads the .dat files exported from the impedance analyzer and saves them
% in a MATLAB format as MVDATA.MAT
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

clear

for i = [11,14,15]
 for j = 1:10
  for k = 'na'
   eval(['load p',num2str(i),k,num2str(j),'.dat'])
  end
 end
end

freq = linspace(4000,4400,401);

clear i j k

save mvdata
```

NORMTEST.M

```
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% This routine imports the .MAT files generated by loaddata.m and tests the data sets for
% normality.
%
% The second part of the routine calculates the sample means and variances of the
% matrices and saves them as MEAN_VAR.MAT
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

clear

load mvdata

% Calculation of comparitive 'Q' value for the normality test. With this, a significance
% level is set as well as a number of degrees of freedom, which is based on a number of
% cells, K, defined below.

siglev = .99;
dof = 2;
Qcomp = chi2inv(siglev,dof);

% Load 'similar' data into (6) 10x401 matrices.

for i = [11,14,15]
 for j = 'na'
  for k = 1:10
   eval(['pR',j,num2str(i),'(',num2str(k),',:) = [p',num2str(i),j,num2str(k),'(:,2)]'';'])
   eval(['pI',j,num2str(i),'(',num2str(k),',:) = [p',num2str(i),j,num2str(k),'(:,3)]'';'])
  end
 end
end

% Test of normality

for i = [11,14,15]
 for j = 'na'
  for k = 'RI'
   for cc = 1:401

eval(['x = p',k,j,num2str(i),'(:,cc);'])
```

```
spread = max(x) - min(x);

% Sample mean and sample variance of random variable
m = mean(x);
s = sqrt(var(x));

% Number of cells (partitions) to break the pdf curve into for test
K = 5;

% The r vector indicates the lower bounds of cells
r(1) = -Inf;
for l = 2:K
  r(l) = min(x) + (l-1)/K*spread;
end

% Calculation of the probabilities for each cell
for l = 2:K
  p(l-1) = normcdf(r(l),m,s)-normcdf(r(l-1),m,s);
end
p(K) = 1-sum(p(1:4));

% Count the number of samples which fall in each cell
count = zeros(K-1,1);

for y = 2:K
  for z = 1:length(x)
    if r(y-1) < x(z) & x(z) < r(y);
      count(y-1) = count(y-1) + 1;
    end
  end
end

count(K) = length(x) - sum(count);

% Calculation of the number of Q residuals based on number of expected samples in each
%  cell and actual number of samples.
for l = 1:K
  Q(l) = (count(l)-length(x)*p(l))^2/(length(x)*p(l));
end

% Comparison of Q's based on Qcomp calculated above.  If the distribution is found not
% to be normal, the samples are redefined as infinite.
Qtot = sum(Q);
```

```
 if Qtot < Qcomp
  eval(['p',k,j,num2str(i),'(:,cc) = Inf*ones(10,1);'])
 end
     end
   end
  end
end

% End of normality test

% Calculation of sample means and sample variances for matrices

for i = [11,14,15]
 for j = 'na'
   for k = 'RI'
    eval(['p',k,j,num2str(i),'m',' = mean(p',k,j,num2str(i),');'])
    eval(['p',k,j,num2str(i),'s',' = var(p',k,j,num2str(i),');'])
    eval(['clear p',k,j,num2str(i)])
   end
 end
end

% Clear excess variables

for i = [11,14,15]
 for j = 'na'
   for k = 1:10
    eval(['clear p',num2str(i),j,num2str(k)])
   end
 end
end

clear i j k l y z m s x dof Qcomp siglev

freq = linspace(4000,4400,401);

save c:\tds\thesis\mvar\mean_var
```

PREDICT.M

```
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% This routine creates the predictor variables. It first verifies the assumption of equal
% variance necessary for the hypothesis test.
%
% Then, it calls HYPTEST.M to perform the hypothesis test with the null hypothesis:
%
%                      H0: μ1 = μ2
%
% The resultant significance levels are saved as PREDICTS.MAT
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

clear
load mean_var

% Define significance level for test of equal variance
alpha = .01;

% Define degrees of freedom based on (number of samples - 1)
dofx = 9;
dofy = 9;

% The ratio of the sample variances cannot be greater than 'test'
test = finv(1-alpha,dofy,dofx);

% This loop makes all of the comparisons needed, calling HYPTEST.M to calculate the
% significance level if it passes equal variance test

for i = [11,14,15]
 for j = 'RI'
  for k = 1:length(pIn11m)
   eval(['if p',j,'n',num2str(i),'s(k)/p',j,'a',num2str(i),'s(k)>test|
p',j,'a',num2str(i),'s(k)/p',j,'n',num2str(i),'s(k)>test, pvna',j,num2str(i),'(k) = Inf; else
pvna',j,num2str(i),'(k) =
hyptest(p',j,'n',num2str(i),'m(k),p',j,'n',num2str(i),'s(k),p',j,'a',num2str(i),'m(k),p',j,'a',num2st
r(i),'s(k),10,10);end;'])
  end
 end
end
```

```
% Clear excess variables

for i = [11,14,15]
  for j = 'RI'
    for k = 'ms'
      for l = 'na'
        eval(['clear p',j,l,num2str(i),k])
      end
    end
  end
end

clear i j k l Num alpha dof dofx dofy test

save c:\tds\thesis\mvar\predicts
```

HYPTEST.M

```
function [sig] = hyptest(xnew,sx,ynew,sy,Nx,Ny)
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% This function takes two random variables with sample means and sample standard
% deviations taken from N samples and performs a hypothesis test using the Student's t
% statistics. The function outputs two significance factors. These factors are based on
% two separate tests of the null hypothesis. That is that the means of the two samples are
% the same. The first method  assumes that xnew is the true mean of one random
% variable and performs a hypothesis test to determine to what degree of certainty we
% can say that the second random variable has the same mean. The lower the value of
% Sig, the more confidence we have in saying that the means of the two random variables
% are not equal. The second method, which outputs sig, uses the t-statistic to determine
% to what degree of certainty we can say that the means of the two random variables
% have the same mean. Once again the lower the value of sig, the more confidence we
% have in saying that the means of the two samples are not equal. The converse is also
% true, that the higher the value of sig, the more confidence we have in saying that the
% means of the two random variables are the same. A lot of this function was taken from
% TTEST.M and TTEST2.M in the STATS toolbox in MATLAB. They, however
% calculate from a given sample whereas this program calculates from a given sample
% mean and sample standard deviation.
%                                                     PMT 3/19/94
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

% This is the beginning of the two-way hypothesis test
% Degrees of freedom
dfx = Nx-1;
dfy = Ny-1;

% Effective degrees of freedom
dfe = dfx + dfy;

for i = 1:length(xnew)
  msx(i) = dfx*(sx(i));
  msy(i) = dfy*(sy(i));
  Sp(i) = (msx(i)+msy(i))/dfe;
  diff(i) = xnew(i) - ynew(i);
  AA = 1/(dfx+1) + 1/(dfy+1);

% Pooled variance
  pooleds(i) = sqrt(Sp(i)*AA);
  ratio(i) = diff(i)/pooleds(i);
```

```
  sig(i) = tcdf(ratio(i),dfe);
  sig(i) = 2*min(sig(i),1-sig(i));
end
```

PVINTRSC.M

```
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% This routine finds all of the random variables which have corresponding significance
% levels below a given threshold for all of the structures of interest.
%
% These intersections will become the damage identification rule locations and are saved
% as RULEPV.MAT
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

clear
load predicts

% Set maximum allowable significance level
cutoff = .00001;

% Check each frequency bin for acceptable significance levels
for i = 1:401
 for j = 'RI'
  eval(['if pvna',j,'11(:,',num2str(i),') < cutoff & pvna',j,'14(:,',num2str(i),') < cutoff &
pvna',j,'15(:,',num2str(i),') < cutoff; sig',j,' = [sig',j,'; freq(:,',num2str(i),')
pvna',j,'11(:,',num2str(i),') pvna',j,'14(:,',num2str(i),') pvna',j,'15(:,',num2str(i),')]; end;'])
 end
end

% Clear excess variables
for i = [11,14,15]
 for j = 'RI'
  eval(['clear pvna',j,num2str(i)])
 end
end

clear i j l cutoff

save c:\tds\thesis\mvar\rulepv
```

LOADDAM.M

```
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% This routine loads the test case information exported by the impedance analyzer and
% saves it as TEST.MAT
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

clear

for i = [4,10,12]
  eval(['load p',num2str(i),'.dat'])
end

clear i

save test
```

DAMRULES.M

```
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% This routine picks the test data which correspond to the significant predictor variables.
% The information is then saved as DAMTEST.MAT
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

clear
load rules
load meanvar

% Select data points for real data rules
for k = 1:length(sigR)
  eval(['freqpt = sigR(k,1) - 3999;'])
  for i = [4 10 12]
    eval(['pR',num2str(i),'test = [pR',num2str(i),'test; freqpt+3999
    pRn',num2str(i),'m(1,freqpt)];'])
  end
end

% Select data points for imaginary data rules
for k = 1:length(sigI)
  eval(['freqpt = sigI(k,1) - 3999;'])
  for i = [4 10 12]
    eval(['pI',num2str(i),'test = [pI',num2str(i),'test; freqpt+3999
    pIn',num2str(i),'m(1,freqpt)];'])
  end
end

clear freq sigI sigR i j k p1 p4 p8 num freqpt

save c:\tds\thesis\mvar\damtest
```

BOUNDPV.M

```
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% This routine applies confidence bounds to the significant predictors. These rules are
% then saved as RULES.MAT
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
clear
load rulepv
load mean_var
load damtest

% Number of samples
N = 10;

% Generate confidence bounds
for j = 'RI'
  eval(['num = length(sig',j,');'])
  for k = 1:num
    eval(['freqpt = sig',j,'(k,1) - 3999;'])
    for i = [4 10 12]
      for l = 'n'
        eval(['[lcl',l,',ucl',l,'] =
        toler(p',j,l,num2str(i),'m(:,freqpt),p',j,l,num2str(i),'s(:,freqpt),N);'])
      end
      eval(['p',j,num2str(i),'rule = [p',j,num2str(i),'rule; freqpt+3999 lcln ucln];'])
    end
  end
end

% Clear excess variables
for i = [4 10 12]
  for j = 'RI'
    for k = 'ms'
      for l = 'n'
        eval(['clear p',j,l,num2str(i),k])
      end
    end
  end
end
clear i j k l lcln lcla ucln ucla N del lcl ucl s2 xbar freqpt

save c:\tds\thesis\mvar\rules
```

TOLER.M

```
function [lcl,ucl] = toler(xbar,s2,N)
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% This function uses values of k from Table C.9 in Hogg and Ledolter with alpha = .01
% and p = .99. With inputs of sampled mean and sampled variance as well as number
% of samples, this function outputs confidence bounds for random samples of the random
% variables.
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

if N == 10
 k = 5.308;
end

for i = 1:length(xbar)
del = k*sqrt(s2(i));
lcl(i) = xbar(i) - del;
ucl(i) = xbar(i) + del;
end
```

HEALTHST.M

```
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% This routine checks test data versus the damage identification rules. A counter is
% incremented and a threshold of successes is set to determine the health status reported.
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
clear
load rules
load damtest
n = 0;
% Plate number being tested
testnum = 4;

% Check imaginary rule bounds
for i = [testnum]
 for k = 1:size(pI11rule,1)
  eval(['if pI',num2str(testnum),'test(',num2str(k),',2) >
pI',num2str(i),'rule(',num2str(k),',2) & pI',num2str(testnum),'test(',num2str(k),',2) <
pI',num2str(i),'rule(',num2str(k),',3); n=n+1; end;'])
 end
end

% Check real rule bounds
for i = [testnum]
 for k = 1:size(pR11rule,1)
  eval(['if pR',num2str(testnum),'test(',num2str(k),',2) >
pR',num2str(i),'rule(',num2str(k),',2) & pR',num2str(testnum),'test(',num2str(k),',2) <
pR',num2str(i),'rule(',num2str(k),',3); n=n+1; end;'])
 end
end

% Display number of successes
n

% Determine health status
if n >= 8;
 disp('No damage')
end
if n < 8;
 disp('Damage')
end
```

Vita

Tom Snyder was born in Elmira, New York on January 24, 1971. Most of his childhood was spent in Harrisburg, Pennsylvania after a short stint in Horseheads, New York. His parents moved to Manassas, Virginia (where they still reside) when he was 10 and he continued moving south when he decided to pursue a degree in Mechanical Engineering at Virginia Tech in 1989. After recieving the degree in May of 1993, he decided to continue on at Tech at the Masters level, and a year and a half later he successfully defended his thesis. Having avoided the real world for long enough, he has decided to continue his trek towards warmer climes and is now preparing to put his schooling to work designing cellular phones for GE-Ericsson in Research Triangle Park, North Carolina.