THE MODAL DISTRIBUTION METHOD FOR STATISTICAL ANALYSIS OF MEASURED VIBRATIONAL ACCELERATION DATA

Bert Sweetman PhD, P.E. [1] (Associate Member, ASCE) and Myoungkeun Choi [2]

ABSTRACT

The newly proposed modal distribution method statistically quantifies overall differences between measured time-histories. The method is general and may find a broad variety of applications, but seems particularly well suited for structural health monitoring because it can be used to infer changes in structural condition from measured response data with only limited knowledge of the excitation. In the new method, power spectra of measured structural response are interpreted as series of independent modal responses. Each modal response is isolated and rescaled to be a statistical distribution. A statistical comparison between data-sets (windows in a measured time-history) results in a quantitative significance level of differences between power spectra. An example is presented to validate the new method and to quantify how long a time-history is required for the new method to meet confidence level requirements. The modal distribution method is found to be very effective at detecting subtle changes of mean modal frequencies, which may be used to infer changes in structural condition.

CE Database subject headings: Acceleration, Modal analysis, Random vibration, Response spectra, spectral density function, Spectral analysis, Structural dynamics, Structural response, Data processing

INTRODUCTION

The use of sensors in structural applications has expanded dramatically in recent years. The cost and complexity of data collection has decreased to the point that sensor networks intended to detect structural failure are expected to eventually become a part of every major new engineered structure. Acceleration remains the preferred measurement because of low sensor cost and high reliability. Detecting variations in structure response, making decisions, and taking appropriate action all rely on interpreting measured data, which is inherently a statistical problem. Simply detecting that a measured quantity has changed is not sufficient for decision making: a method to quantify the statistical significance of observed changes in measured data is necessary.

The earliest work for damage detection through modal analysis was on offshore structures. Early investigators e.g. (Vandiver 1975), (Vandiver 1977) concluded that changes in modal frequencies due to member damage of offshore structures are detectable based on numerical simulation and a scale model, respectively. Three North Sea platforms were monitored for six to nine months and it was concluded that changes in the response spectrum can be used to monitor structural integrity (Loland and Dodds 1976). Since that time, numerous investigators have measured

[1] Assistant Professor, Maritime Systems Engineering Dept., Texas A&M Univ. Galveston, 200 Seawolf Parkway, Galveston, TX 77553-1675. E-mail: sweetman@tamu.edu

[2] Doctoral Candidate, Civil Structural Engineering, Dept. of Civil Engineering, Texas A&M Univ., College Station, TX 77843-3136. E-mail: mkchoi@tamu.edu

accelerations on various civil structures and used computed power spectra to make inferences about the condition of these structures. Conventional methods require detailed knowledge of both the excitation force and vibration response (Crema and Mastroddi 1998); (Harris and Crede 1995). In a different approach, (Smyth et al. 2000) apply vibration signature analysis and advanced statistical methods to structural health monitoring, but in this work detailed knowledge of the excitation is still required. White noise excitation is often assumed in practical applications in which the excitation force is unknown e.g. (Katafygiotis and Lam 1997); (Abe et al. 1999), but in some applications such an assumption may be unreasonable, especially if there is known to be a dominant frequency of excitation.

A major state of the art review for structural health monitoring (SHM) conducted as Sandia National Labs noted that little progress has been made in application of statistical methods to SHM (Sohn et al. 2003). The work presented here addresses the need for additional statistical methods for interpretation of power spectra computed from measured accelerations in SHM applications. This work makes use of a well-established practice from the field of acoustics: mapping power spectra into probability space (Hansler and Schmidt 2004). The unique contribution of this paper to the field of SHM is combining modal analysis with this mapping such that conventional statistical tools can be applied to assess the significance of observed changes in modal frequencies.

The proposed method compares modal response characteristics computed for separate segments of a measured time-history. The new method has important advantages over more conventional analyses: 1) The only measured time-history required as input is that of the response, 2) the acceleration data is separated into individual vibrational modes of response, enabling detection of vibrational changes in individual modes even if the overall response is relatively unaffected, 3) if there is a dominant excitation frequency which varies over time, the mode containing that frequency can be excluded from the analysis such that a shift in the excitation frequency is not misinterpreted as a change in structural behavior, and 4) the result is a numerical prediction of the confidence level that there has been some change in the response of the structure. Here, the T-statistic is used because the distribution provides a reasonable fit to the power spectra being compared in the example; other statistical tests including nonparametric methods could alternatively be used in place of the T-test with no fundamental change to the underlying methodology presented here. The T-statistic relies on an underlying Gaussian distribution, which is strongly violated in cases of very high noise or closely spaced peaks. Finally, the authors note that the method presented here is intended only to quantify the statistical significance of shifts in modal frequencies, and leaves the any inference regarding structural parameters such as mass, stiffness, or damage to others.

OVERVIEW OF THE NEW METHODOLOGY

This overview describes the methodology from a conceptual standpoint without computational detail; detail is presented later. In the newly proposed method, a measured acceleration time-history is divided into a series of segments, or windows, which are to be sequentially analyzed. Each of the segments is first converted into a power spectrum. It is implicitly assumed that the dynamic behavior of a structure in a given frequency range can be considered as a set of individual modes of vibration and that these vibrational modes can be considered to be independent. Accordingly, the power spectrum is divided into a series of response peaks using a penalty method to find the local minima. Each peak is then treated individually. Any response peak can be excluded from the analysis if the excitation has a dominant frequency near that peak.

Each remaining peak in each power spectrum is mapped into probability space, which in practice is a simple re-scaling that is commonly used in the field of acoustics e.g., (Hansler and Schmidt 2004). Mapping the vibrational modes into statistical distributions enables use of conventional statistical tools to compute the significance of changes in modal frequencies for each vibrational mode being considered. The mean and variance of each distribution are calculated directly from geometry.

The overall goal is to quantify the significance of subtle changes in the state of a vibrating structure by comparing modal frequency distributions between different segments of a measured response time-history. The comparison of two distributions is performed by statistical comparison of means of the distributions. The mean of each sequential modal distribution in the first window is statistically compared with that of the second window, i.e, the mean of the first mode of the first window is compared with the mean of the first mode of the second window, as are each of the subsequent modal means. Here, the T-statistic is applied: each mean difference is weighted by the magnitude of the energy associated with that mode and combined to compute an overall T-statistic. The statistical significance of any difference between individual modal means is reported, as well as a combined statistic quantifying differences between all modal means. The entire procedure could be repeated over consecutive segments (windows) of a time-history to monitor for changes in modal response.

COMPUTATIONAL DETAILS

Separation of Modal Distributions: The Penalty Method

The power spectrum is computed from the Fourier transform of a window of the measured time-history. Power spectra of measured data generally are not smooth, and identifying the correct minima to use as dividing points between modal distributions is non-trivial. The methodology presented here relies on an initial estimation of the frequency of each mode of vibration for the structure. This initial value would be expected to come from structural analysis, though it could also come from visual inspection of the power spectrum. This initial value need not be of high accuracy because it is only used to specify the endpoints of a penalty method. The penalty method searches for one local minimum between every two adjacent modal frequencies, plus one additional minimum above the highest and one below the lowest specified frequencies. The penalty method has been found to be robust in applications with high noise or closely spaced frequencies: it always finds the local minimum, regardless of the distance from zero on the spectral plot. The user is cautioned, however, that in these applications the Gaussian assumption is strongly violated and the T-statistic should not be applied. Basing the modal separation on local minima also provides additional robustness against noise in the signal because the squaring inherent to the calculation of the power from the Fourier spectrum amplifies maxima in the spectrum more than minima.

In searching for each local minimum, the method searches the region between two adjacent modal frequency estimates by progressively seeking the minimum of increasingly fine divisions of the region. First, the region is divided into three equal frequency intervals as shown in the second frame of Figure 1. The area under the power spectrum within each of the three intervals is then computed; every frequency in that interval having the lowest average energy is assigned a penalty value of 1.0. Next, the number of divisions between initial modal frequency guesses is increased to four. The same penalty approach is applied: again, every frequency in the region having the smallest average energy receives a penalty; on this second iteration the penalty is less than 1.0. The process is repeated with the number of divisions increasing by 1 and the penalty decreasing linearly for each subsequent penalty assessment until the average of each region includes not less than 25 of the original bars. Beyond four intervals the method is modified: a penalty is applied to every frequency in each interval having lower average energy (lower average bar height) than both of its contiguous neighbor intervals. After the final iteration, the total penalty is calculated for each of the frequencies between initial modal frequency estimates. The frequency having the largest total penalty is concluded to be the local minimum between energy peaks and is subsequently used as the dividing point. All energy between adjacent local minima is assumed to be associated with a single vibrational mode. Once these modes have been isolated and the peak associated with any dominant frequency in the excitation (if any) has been removed, distribution parameters are calculated for each mode.

Modal Distributions and Probability Density Functions

The modal frequencies resulting from the penalty method are then treated as random variables. As is common in acoustics, e.g. (Hansler and Schmidt 2004), the acceleration power spectrum is mapped into probability space. In this mapping, the power spectral density can be interpreted as the relative probability that the signal contains energy at any individual frequency. Thus, each bar in the power spectrum near a modal frequency provides an uncertain estimate of that modal frequency, with a relative probability that the estimate is correct. This is the classic definition of a probability density function. Here, the energy surrounding each modal frequency is treated as a probability distribution of candidates for the true modal frequency of the underlying structural vibration:

$$P_i(e_n) = \frac{a_n}{A_i} \tag{1}$$

where a_n is the area of the n'th frequency bar and A_i is the total energy associated with the i'th modal frequency (Figure 1), or the total area between relative minima:

$$A_i = \sum_{n=1}^{N_i} a_n = \sum_{n=1}^{N_i} S(n)df \tag{2}$$

in which $S(n)$ is the spectral offset at the n'th frequency, df is the frequency spacing and N_i is the total number of bars within the i'th modal distribution. The mean and variance are computed directly from the offsets of the power spectrum using conventional definitions of moments and central moments, e.g. (Bethea 1984), (Brandt 1970).

$$\mu_{i,w} = \frac{m_i}{A_i} = \frac{1}{A_i}\sum_{n=1}^{N_i} S(n)f(n)df \tag{3}$$

$$s^2_{i,w} = \frac{\theta_i}{A_i} = \frac{1}{A_i}\sum_{n=1}^{N_i} S(n)(f(n)-\mu_{i,w})^2 df \tag{4}$$

where $\mu_{i,w}$ and $s^2_{i,w}$ are the mean frequency and sample variance of the modal distribution, m_i is the first geometric moment, θ_i is the second central geometric moment and $f(n)$ is the n'th frequency.

Two important sources of variations in the mean and variance are considered. The first source is that the data-set being sampled is finite, so any computed mean and variance are sample statistics and not necessarily those of the underlying population; this source of variation is accounted for in the T-statistic. In this application to random vibrations there is a second important source of variation caused by non-ergodicity; it has been found here that the discontinuous ends of a time-history can have a meaningful impact on computed statistical moments that is not considered in the T-statistic. To minimize apparent differences between windows caused by this non-ergodic effect, power spectra are averaged over successive data-points. Specifically, for each window a power spectrum is computed 100 times, with both the start and end points shifted two points forward in time; the resulting 100 spectra are averaged and treated as a single representative spectrum for that window.

Comparison: Statistical Hypothesis Testing

Spectral peaks of field data generally appear as a cluster of energy around some frequency. Here, mean modal frequencies are treated as random variables and observed clusters of energy are treated as statistical distributions. The benefit is that rigorous statistical analysis can then be applied to assess the significance of apparent differences between distributions. Here, a T-test is applied.

The T-test is a conventional method to determine the statistical significance of a difference between two sample means from populations sampled from underlying Gaussian distributions. The test can be used when the number of samples is too small for the Central Limit Theorem to apply and where the true variances of the underlying processes are not known to be equal. Use of the T-test in this application is not strictly justified because the underlying distributions are not known to be Gaussian. In cases where peaks are broadly spaced and damping is sufficient to have "bell-shaped" modal peaks, use of the T-test may be reasonable. In some instances, especially cases with very high noise or closely spaced spectral peaks, the Gaussian assumption would be strongly violated and the T-test should not be used. Other options, such as non-parametric tests, may be applicable. If the T-test is deemed suitable, the T-statistic can be computed as e.g. (Navidi 2006):

$$T_i = \frac{\Delta\mu_i}{s_i} \tag{5}$$

where

$$\Delta\mu_i = \mu_{i,1} - \mu_{i,2} \tag{6}$$

$$s_i^2 = s_{i,1}^2/N_{i,1} + s_{i,2}^2/N_{i,2} \tag{7}$$

$$N_{i,w} = D_w\mu_{i,w} \tag{8}$$

in which $\mu_{i,w}$ is the mean of the i'th mode of the w'th segment, $s_{i,w}^2$ is the sample variance, and the number of samples $N_{i,w}$ is estimated as the number of cycles expected at each modal frequency. D_w is the duration of the w'th data window.

T_i is distributed approximately as a Student's T with the number of degrees of freedom for the i'th mode equal to e.g. (Navidi 2006):

$$DOF_i = \frac{(s_{i,1}^2/N_{i,1} + s_{i,2}^2/N_{i,2})^2}{(s_{i,1}^2/N_{i,1})^2/(N_{i,1}-1) + (s_{i,2}^2/N_{i,2})^2/(N_{i,2}-1)} \tag{9}$$

Overall Comparison of Response

Accurate assessment of the significance of changes in the observed vibrational response requires consideration of all modes. Here, the ensemble of observed differences between each pair of mean modal frequencies is treated as a set of repeated measurements with differing uncertainties, i.e., observed differences between the first pair of modal distributions are combined with observed differences between pairs of sequentially higher modes. An overall P-value is calculated by computing a combined T-statistic, which weights differences between the means by the fraction of energy represented for each individual mode. The fraction of energy associated with the i'th mode for each of the two windows is:

$$E_{i,w} = \frac{A_{i,w}}{\sum_{i=1}^{I} A_{i,w}} \tag{10}$$

where I is the total number of modes in the power spectrum and w is either 1 or 2, denoting either the first or second window of the time-series. The overall T-statistic is found using Equation 11, in which $\Delta\mu$ and s^2 are computed as averages weighted by the fraction of energy associated with each of the I modes:

$$T = \frac{\Delta\mu}{s} \tag{11}$$

where

$$\Delta\mu = \sum_{i=1}^{I} 0.5(E_{i,1} + E_{i,2})\Delta\mu_i \tag{12}$$

$$s^2 = \sum_{i=1}^{I} 0.5(E_{i,1} + E_{i,2})s_i^2 \tag{13}$$

The T-statistic resulting from Equation 11 represents the weighted average of differences between the means of each pair of modes. The number of degrees of freedom for use with a standard T-distribution is estimated as the total number of vibration cycles in all vibrational modes. Generally, this number will be larger than the observed number of peaks in the time-history because modal frequencies with small fractions of the total energy will not contribute to the number of directly measurable peaks. The total degrees of freedom to be used in conjunction with the T-statistic is estimated as the sum of individual degrees of freedom:

$$DOF = \sum_{i=1}^{I} DOF_i \tag{14}$$

TYPE I AND TYPE II ERRORS: MINIMUM DATA REQUIREMENTS

One significant issue that must be addressed prior to any practical implementation of any data analysis technique is assessing how much data is required for the method to be effective. There are exactly two ways a method can fail: 1) it could detect a difference between the two power spectra when there is in fact no underlying difference, or 2) it could fail to detect a difference that does in fact exist. The null hypothesis (H_0) implicitly underlying use of the T-test is that there is no difference between the two distributions, which makes category 1) a Type I error (reject H_0 when it is true) and category 2) a Type II error (fail to reject H_0 when it is false).

For many practical applications, it is desirable to calculate how long a data-set is required to detect a modal shift of a predetermined size with a pre-specified confidence level while maintaining an acceptably low probability of false alarms. To compute the required length of the data-set, the P-value is itself treated as a random variable. For any specific number of vibration cycles, a histogram of predicted P-values could be created and normalized to a PDF. As part of this work, numerous sets of P-values have been generated by application of the modal distribution method to pairs of time-histories, each independently realized from a pair of target spectra with known modal differences; the distribution of predicted P-values has been found to be well described by a theoretical gamma distribution. The PDF and CDF (g and G) can be expressed in terms of the gamma and incomplete gamma functions, respectively e.g. (Navidi 2006):

$$g(x;\alpha,\theta) = \frac{x^{\alpha-1}e^{-\frac{x}{\theta}}}{\Gamma(\alpha)\theta^{\alpha}} \tag{15}$$

$$G(x;\alpha,\theta) = \frac{\gamma(\alpha,\frac{x}{\theta})}{\Gamma(\alpha)} \tag{16}$$

in which $\Gamma(\alpha)$ denotes the gamma function: $\Gamma(\alpha) = \int_0^\infty t^{\alpha-1}e^{-t}dt$. The shape parameter ($\alpha$) and the scale parameter (θ) define the gamma distribution completely, and can be calculated from the mean (μ) and sample variance (s^2) of an ensemble of P-values as $\alpha = \mu^2/s^2$ and $\theta = s^2/\mu$.

The area under the gamma PDF below the pre-specified confidence level corresponds to the probability that any future prediction of a P-value will be at or below that confidence level, i.e., the probability that the method will detect changes in modal frequencies with the required confidence level. The probability of a Type I error equals this area if there is no actual difference between the underlying distributions; the probability of a Type II error can be computed as one minus this area if there is a difference between distributions.

It is relatively straightforward to determine the minimum required time-history for some new application using the gamma distribution. First, the mean and standard deviation of the predicted P-values must be known. These values can be determined by simulation from target spectra with spectral shapes comparable to those expected for the physical system to be studied and with modal frequency shifts equal to those sought to be detected. The mean of the P-values is the expected value of future P-value predictions; the standard deviation can be used in conjunction with the gamma distribution to assign confidence intervals to this mean. A worked example of each error type is provided.

EXAMPLE: MINIMUM DATA REQUIREMENTS

This example quantifies the probability of both Type I and Type II errors, then compares the results of this method with conventional tests of stationarity. The example is based on simulated data for which the actual underlying target spectra are precisely known so the effectiveness of the new method can be assessed. The underlying spectra are based on physical data of practical importance, vortex induced vibrations of marine drilling risers. Here, the riser is undamaged in all cases, and the dominant response frequencies are not related to environmental excitation frequencies: any change in modal response parameters is believed to be caused by hydrodynamic interactions, in particular to be caused by the mass of the water entrained with the vibration.

Marine risers are the pipes transmitting fluids between the sea-floor and a floating production platform. A riser, or any other slender structure subject to strong flow across its axis, may interact with the current to create vortex induced vibrations (VIV). Prediction of this complicated fluid-structure interaction has been a historically intractable problem, in part because two separate quantities are unknown: the excitation force due to fluid-structure interaction and the effective mass of the riser which is highly influenced by the mass of the water entrained with the riser motion, usually referred to as the "added mass."

Underlying Spectra and Data Simulation

In order to quantify the effectiveness of the new method, it is tested against time-histories simulated from ideal target spectra, which are precisely known. The initial target power spectrum is generated from measured field data by first computing the power spectrum from a measured time-history and then smoothing the first two modes into the ideal target spectrum S shown in Figure 2. Known changes in modal properties are then introduced to the smoothed spectrum. Time-histories needed for the example are then simulated from these spectra using the conventional method: $x(t) = \sum_{n=1}^{N} H_n cos[2\pi f(n)t + \theta_n]$, in which $x(t)$ is the time-history of the simulation, H_n is the amplitude of each frequency component, which is related to the energy at that frequency, $f(n)$ is the n'th circular frequency and θ_n is a random phase angle. As always, simulation from a power spectrum implies stationarity and use of the Fourier transform in simulation implies ergodicity. Here, the simulation includes as many frequency components as there are time-steps so the series does not repeat within a simulation. Finally, a Gaussian white-noise signal is added to each time-history. The level of noise added is 1.9 mg rms, which is typical of the noise level observed in the actual field data underlying this example; noise inherent to a high-quality ± 1 g accelerometer alone is typically around 0.5 mg rms.

Application of the Method

To assess the minimum amount of data necessary, the new method is applied to each pair of simulated time-histories for increasingly long data windows. The shorter time-histories are the early parts of the longer histories, i.e., the 30 minute data window test for target spectrum S_1 is the first part of the time-history used in the 60 minute test. Unless otherwise noted, all cases include the 1.9 mg Gaussian white noise and all plots are generated by performing 25 independent tests and averaging the results.

Type I Errors: No change in true underlying spectra

To quantify the probability of having a Type I error, 50 realizations of the original bimodal target spectrum S are simulated and the modal distribution method is applied to the resulting 25 independent pairs of time-histories. Figure 4 shows the average of the P-values from these tests, all of which are above 50%. These very high P-values correspond to the fact that observed differences result from random chance and not from differences in the underlying spectra. The error bars correspond to a double-sided 90% confidence interval; there is almost no chance of a false-alarm when at least 30 minutes of data is considered.

Type II Errors: Known changes in mean modal frequencies

In this part of the example, the mean of each mode is shifted to the left. Here, the change to the mean of each mode ($\Delta\mu_i$) is chosen such that each mean of each mode is shifted an amount consistent with a constant $\Phi(z)$ of a standard normal distribution:

$$\Delta\mu_i = zs_i; \qquad \Phi(z) = \int_{-\infty}^{z} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}u^2} du \tag{17}$$

For example, applying Equation 17 with $\Phi(z)$ = 5% of the area under a standard normal distribution yields $\Delta\mu$ = -0.1257 s_i. Transforming this value to each modal distribution of the original target spectrum, S, results in the new mean frequencies shown in Table 1. The resulting spectra S_1, S_2 and S_3 are target spectra with progressively greater shifts in mean-modal frequencies. In addition to this mean change, the variance of each mode of each target spectrum is increased from the original by 10% while holding the total energy of the spectra constant. Figure 3 shows the original target spectrum, S, with one of the modified spectra, S_1. The resulting magnitudes of each frequency shift is particulary convenient to illustrate calculation of the amount of data required for the new method to attain a pre-determined level of effectiveness. Specifically, to the extent the Gaussian assumption is not violated, the resulting confidence levels should be a function of only the number of cycles, without regard for vibrational mode with which these cycles are associated.

Minimum Required Time-History Duration

After the modified target spectrum has been developed, the new S_1 and original S are used to simulate a total of fifty sets of time-histories. The modal distribution method is then applied to each resulting twenty-five pairs of time-histories to calculate a P-value for each. These results are averaged and presented as Figure 5. The monotonic decrease shown in Figure 5 shows greater confidence levels (smaller P-values) are associated with longer time-histories. Confidence levels of 95% (5% P-values) are often considered to indicate statistical significance. The upper two lines on the figure indicate confidence levels based on each of the first two modes. The lowest line on the figure indicates the method's capability to detect the change in the underlying spectrum using all available data, which in this case means both modes. The very subtle shift shown in Figure 3 can be detected with about 95% confidence (5% P-value) using only 30 minutes of data.

Figure 6 shows the probability of failing to detect the known shift in modal frequencies from Figure 3. These probabilities are calculated by treating the P-value as a random variable (Equation 16). Using only 30 minutes of data, there is about an 18% chance the method will not detect the change in the underlying spectrum with a P-value of 5% or less. However, given 90 minutes of data, the method is almost certain to detect even the very subtle change indicated in Figure 3.

Minimum Required Number of Cycles

The main factor for determining the necessary data set length is in fact the number of vibrational cycles rather than the temporal duration. Systems with higher frequencies or with more vibrational modes have more cycles in a fixed duration. The data of Figure 5 is again presented as Figure 7, but with the x-axis presented as cycles. The number of cycles is estimated as the simulation duration times the sum of mean modal frequencies. The fact that the P-value predictions based on the first mode, second mode and combination nearly collapse to a single line when plotted against cycles is a property of the relative shift of each of the modal frequencies corresponding to a fixed percent of an ideal Gaussian distribution (Equation 17) and not of the method itself.

Figure 7 shows that approximately 1,500 cycles are needed to obtain a statistically significant detection of the very subtle shift between the S and S_1 power spectra. The figure shows results both for an ideal noise-free signal and for a signal containing a realistic 1.9 mg of white noise. The added noise has only minimal effect on the P-values predicted for each mode and the overall comparison. Results for each mode and combined results all tend to collapse toward a single line as would be expected for a Gaussian distribution. Inclusion of the noise creates only a slight change in computed non-Gaussianity for each mode: average skewness for both modes is reduced from 0.082 to 0.067 and average kurtosis is increased from 3.10 to 3.15, compared with the ideal Gaussian values of 0 and 3. In Figure 8, however, results do not collapse to a line, which qualitatively shows the effect of strong violations of the Gaussian assumption. The level of noise included here is 19 mg, resulting in average skewness of -0.52 and kurtosis of 4.51. The T-statistic is clearly not valid for this strong violation of the Gaussian assumption. Returning to the low-noise (1.9 mg) case, Figure 9 shows that at least 3,000 cycles are needed to obtain a 5% probability of failing to detect such a shift, based on the gamma distribution. Similar results are expected for predictions using three or more modes.

Thus far the example has sought to detect extremely subtle shifts in the underlying power spectra and detecting these shifts has required considerable time durations. Detecting larger changes between underlying spectra requires less data. Figure 10 shows the amount of data required to detect increasingly large shifts in modal frequencies. All shifts are made to both modes in accordance with Equation 17; the details of the underlying target spectra are shown in Table 1. The curve corresponding to shifts in mean modal frequencies corresponding to 5% of the area under the standard normal is fit through the same data as presented in Figure 7. The very subtle shift corresponding to 5% of the area under a normal distribution requires around 1,400 cycles to detect at a 5% P-value; if the size of the modal shifts are approximately doubled, the required number of cycles drops to around 400 cycles.

Comparison with Conventional Tests

The newly proposed method is compared with two conventional tests of statistical stationarity: the reverse arrangements test and the runs test e.g. (Newland 1993) and (Bendat and Piersol 1991). Numerous pairs of 150 minute time-series were used to test these conventional methods. These tests were made for both Type I and Type II errors using the same data as previously used to verify the new method. Neither of these conventional methods was found to be able to meaningfully detect the change in the target spectrum

The reverse arrangements test never detected a statistical change in the time-series, regardless of whether or not there is a change between underlying spectra. Results from the runs test varied with the number of segments into which the time-series was divided as part of the test, but in no case were the results useful. When 40 segments were used, the probability of a Type I error (false alarm) was found to be 24% and the probability of a Type II error 76%. Using 100 segments, the probability of a Type I error was found to be 52% and probability of a Type II error 44%. Recall that the new modal distribution method applied to these same time-series yielded a probability of a Type I error and of a Type II error near zero (Figures 4 and 6).

CONCLUSIONS

A newly proposed modal distribution method of quantifying the significance of subtle changes in modal vibrations based on power spectra has been presented. The new method is very effective at detecting changes in mean frequencies of individual modal vibrations (Figure 5–10), even if the data is simulated from underlying target spectra with only very subtle differences (Figure 3). The method is also very robust against false-alarms. The amount of data required decreases considerably if larger changes in underlying spectra are to be detected (Figure 10). The new method is considerably more effective than the conventional runs and reverse arrangements tests.

The modal distribution method is based on transforming individual modes of a power spectrum into probability space, which enables application of conventional statistical tools. Response modes containing frequencies of dominant excitation can be explicitly excluded from the analysis. After transformation into probability space, each modal response is treated as an independent statistical distribution for which the spectral moments are directly calculated from geometry. The mean and variance are calculated independently for each individual mode of each of two segments of the time-history and then used in a statistical comparison. Combining differences between the means of individual modal frequency pairs results in an overall quantitative significance level of the difference between power spectra. (Minimum data requirements for the new method are discussed in detail, and a method is outlined to determine minimum data requirements for new applications.

The newly proposed modal distribution method is general and is applicable to any number of modes of vibration. The development of the method presented here makes use of the T-statistic, so this presentation is directly applicable to only well-separated modes with a generally Gaussian shape and relatively low levels of noise in the signal. In concept, any distribution could be used in place of the Student-T. The new method may find a broad variety of applications, though it seems particularly well suited for structural health monitoring because detailed knowledge of the excitation is not required as input.

ACKNOWLEDGMENTS

This work was supported by Texas A&M University and by the National Science Foundation, Division of Civil and Mechanical Systems under Agreement Numbers CMS-0428585 and CMS-0448730. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the view of the National Science Foundation. Riser acceleration data was generously made available for this work by British Petroleum.

REFERENCES

Abe, M., Fujino, Y., Kajimura, M., Yanagihara, M., and Sato, M. (1999). "Monitoring of a long span suspension bridge by ambient vibration measurement." *Structural Health Monitoring 2000*, Stanford University, Palo Alto, California. 400–407.

Bendat, J. S. and Piersol, A. G. (1991). *Random Data: Analysis and Measurement Procesures*. John Wiley & Sons, New York.

Bethea, R. M. (1984). *Statistical Methods for Engineers and Scientists*. M. Decker.

Brandt, S. (1970). *Statistical and Computational Methods in Data Analysis*. North-Holland Pub. Co.

Crema, L. B. and Mastroddi, F. (1998). "A direct approach for updating and damage detection by using frf data." *Proceedings of ISMA23, Noise and Vibration Engineering*, Leuven, Belgium.

Hansler, E. and Schmidt, G. (2004). *Acoustic Echo and Noise Control: A Practical Approach*. Wiley-IEEE Press.

Harris, C. M. and Crede, C. E. (1995). *Shock and Vibration Handbook*. McGraw-Hill, New York.

Katafygiotis, L. and Lam, H. (1997). "A probabilistic approach to structural health monitoring using dynamic data." *Structural Health Monitoring, Current Status and Perspectives*, Stanford University, Palo Alto, California. 152–163.

Loland, O. and Dodds, J. (1976). "Experience in developing and operating integrity monitoring system in north sea." *Proceeding of Eighth Annual Offshore Technology Conference*. 313–319.

Navidi, W. (2006). *Statistics for Engineers and Scientists*. McGraw-Hill Higher Education, New York.

Newland, D. E. (1993). *An Introduction to Random Vibrations, Spectral & Wavelet Analysis*. Longman, Scientific & Technical and John Wiley & Sons, third edition.

Smyth, A., Masri, S., Caughey, T., and Hunter, N. (2000). "Surveillance of mechanical systems on the basis of vibration signature analysis." *ASME Journal fo Applied Mechanics*, 67, 540–551.

Sohn, H., Farrar, C. R., Hemez, F. M., Shunk, D. D., Stinemates, D. W., and Nadler, B. R. (2003). "A review of strucutural health monitoring literature." *Report No. LA–13976–MS*, Los Alamos National Laboratory.

Vandiver, J. K. (1975). "Detection of structural failure on fixed platforms by measurement of dynamic response." *Proc. of the 7th Annual Offshore Technology Conference*. 243–252.

Vandiver, J. K. (1977). "Detection of structural failure on fixed platforms by measurement of dynamic response." *Journal of Petroleum Technology*, 305–310.

Table 1: Target Spectra Parameters with Percent Change from Original Idealized Spectrum

Target	1st Mode				2nd Mode			
Spectrum	Mean	% Change	Variance	% Change	Mean	% Change	Variance	% Change
S	0.12500	N/A	0.00060	N/A	0.28500	N/A	0.00120	N/A
S_1	0.12192	-2.46%	0.00066	10%	0.28064	-1.53%	0.00132	10%
S_2	0.12037	-3.70%	0.00066	10%	0.27845	-2.30%	0.00132	10%
S_3	0.11879	-4.96%	0.00066	10%	0.27622	-3.05%	0.00132	10%

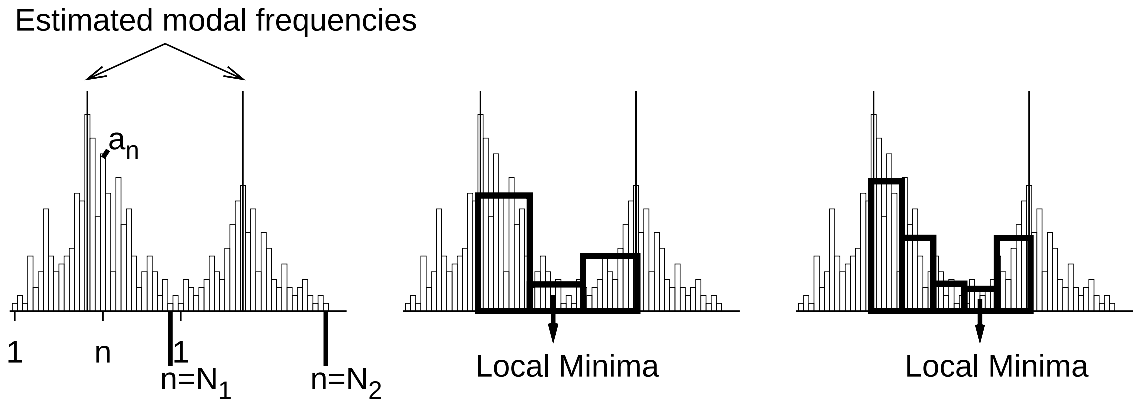

Fig. 1: Separation of modal distributions: The penalty method

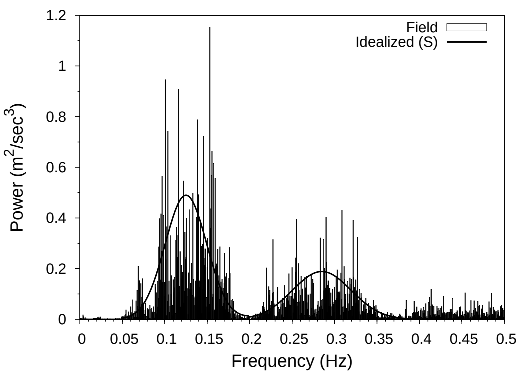

Fig. 2. Field Data Spectrum and Idealized Target Spectrum S

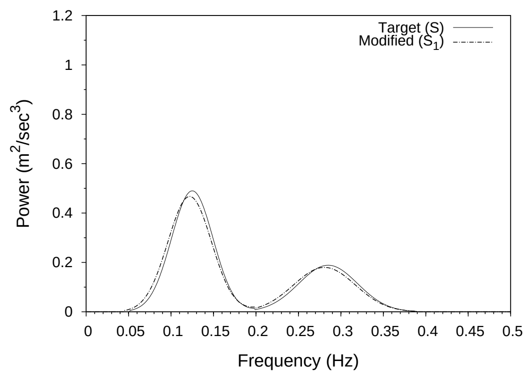

Fig. 3. Original Spectrum (S) and Modified Spectrum (S_1)

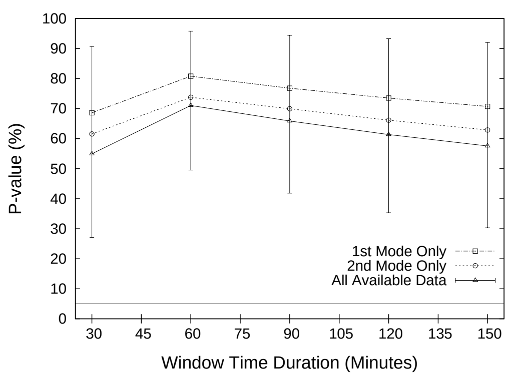

Fig. 4. P-value that the observed differences between the time series are due to differences between underlying spectra when in fact no such differences exist

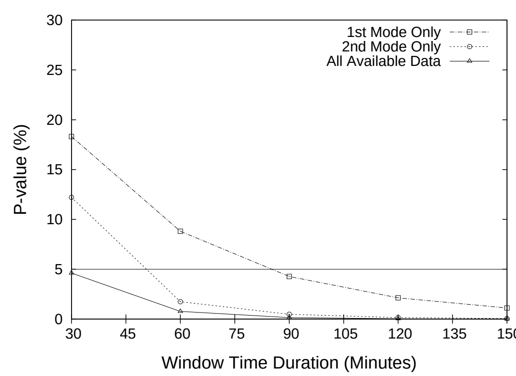

Fig. 5. Ability to detect underlying changes in the target spectrum as a function of time-history duration (smaller P-values indicates greater ability)

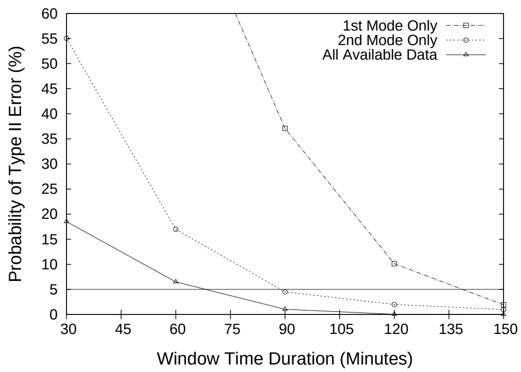

Fig. 6. Probability of failing to detect a known difference between underlying spectra as a function of time-history duration

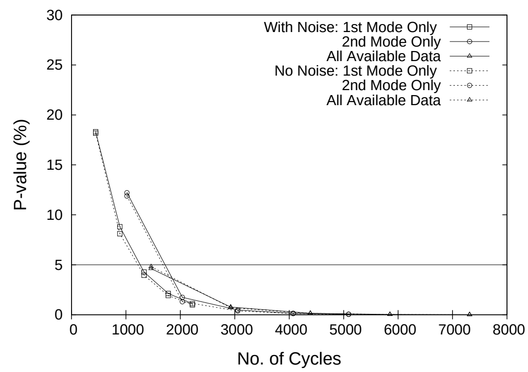

Fig. 7. Ability to detect underlying changes in the target spectrum as a function of number of cycles with and without noise

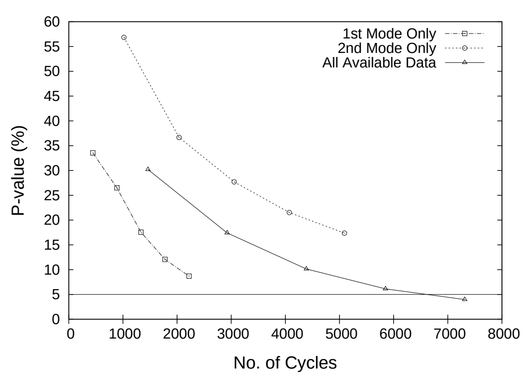

Fig. 8. Very high noise case (19 mg): Detection of underlying changes in the target spectrum with Gaussian assumption strongly violated

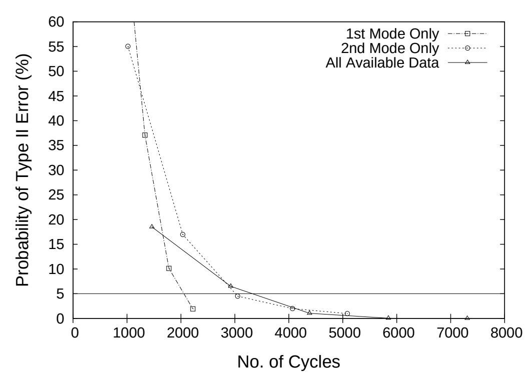

Fig. 9. Probability of failing to detect a known difference between underlying spectra as a function of number of cycles

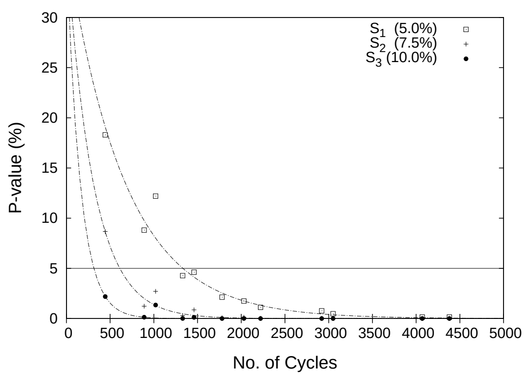

Fig. 10. Number of cycles required to detect increasingly large shifts in mean modal frequencies. Percents refer to fractions of the area under a standard normal distribution.