



ELSEVIER

Contents lists available at ScienceDirect

Probabilistic Engineering Mechanics

journal homepage: www.elsevier.com/locate/probengmech

The modal distribution method for statistical analysis of measured structural response

B. Sweetman^{a,*}, M. Choi^{b,1}^a Texas A&M University, Maritime Systems Engineering Department, 200 Seawolf Parkway, Galveston, TX 77553-1675, United States^b Texas A&M University, Department of Civil Engineering, College Station, TX 77843-3136, United States

ARTICLE INFO

Article history:

Received 14 May 2008

Received in revised form

22 October 2009

Accepted 2 November 2009

Available online xxxx

Keywords:

Acceleration

Modal analysis

Random vibration

Response spectra

Spectral density function

Spectral analysis

Structural dynamics

Structural response

Data processing

ABSTRACT

The newly proposed modal distribution method statistically quantifies overall differences between measured time-histories. In this method, power spectra of measured structural response are interpreted as collections of independent modal responses. Each modal response is isolated, re-scaled, and interpreted as a statistical distribution. Data-sets (windows in a measured time-history) are then compared using standard statistical methods, resulting in a quantitative significance level of the 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. 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.

© 2009 Published by Elsevier Ltd

1. 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 in which modal analysis was applied to damage detection was for offshore structures. Early investigators e.g. [1,2] 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 [3]. Since that time, numerous investigators have measured 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 [4,5]. In a different approach, vibration signature analysis and advanced statistical methods are applied to structural health monitoring [6], 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. [7,8], 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 [9]. The work presented here addresses the need for additional statistical methods for interpretation of power spectra computed from measured accelerations in SHM applications.

The method proposed here compares modal response characteristics computed for separate segments of a measured time-

* Corresponding author. Tel.: +1 409 740 4834; fax: +1 409 741 7153.

E-mail addresses: sweetman@tamu.edu (B. Sweetman), myoungchoi@gmail.com (M. Choi).¹ Present address: Texas A&M University, Maritime Systems Engineering Department, 200 Seawolf Parkway, Galveston, TX 77553-1675, United States.

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. 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.

2. Theory: The modal distribution method

2.1. Overview

This section gives an overview from a conceptual standpoint; computational details are presented in Section 3. 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 segment is first converted into a power spectrum through use of the FFT. As always, use of the power spectrum implies stationarity within any window of the time-history, and use of the Fourier transform implies ergodicity. 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 frequency ranges, each of which includes a distribution of energy around a single peak frequency. Hereafter, this distribution of energy around a single peak frequency is referred to as a modal distribution. The process of dividing the modal distributions is based on using a penalty method to find local minima between energy peaks (Section 3.1). After this initial division, each modal distribution is treated individually. Any frequency ranges excepting those including a dominant frequency of excitation are presumed to be modal responses; those ranges including dominant frequencies of excitation can be explicitly excluded from the analysis.

Each remaining mode is mapped into a statistical distribution in which the random variable is the modal frequency. The theory underlying this new methodology is further discussed in Section 2.2. The overall goal is to quantify the statistical significance of subtle changes in the state of a vibrating structure by comparing modal frequency distributions representing two different segments of a measured response time-history. 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. In the implementation presented here, the T -statistic is applied to quantify the significance of observed differences. The significance of each mean difference is computed independently, and then results are weighted by the magnitude of the energy associated with that mode to compute an overall combined T -statistic. Applicability of the T -statistic and computation of degree of freedoms are further discussed in Section 2.3. 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.

2.2. Interpretation of a power spectrum as a statistical distribution

The heart of the new method lies in interpretation of the vibrational response spectrum as a statistical distribution of modal frequencies. The individual frequency bars making up the discrete power spectrum represent the energy occurring at that frequency, so the discrete power spectrum can be considered as a discrete distribution in which the abscissa is the frequency and the ordinate is the amount of energy at that frequency. Areas under the power spectrum that are not clustered around distinct peaks can be explicitly excluded from the analysis; exclusions might include energy associated with very closely spaced modal peaks and vibrations not associated with any mode shape. The region of the power spectrum associated with any single mode is then normalized to unit area, such that each bar represents the fraction of total energy of the individual mode. The resulting distribution has the same properties as a conventional relative histogram, in which any bar commonly represents the fraction of total occurrences. This interpretation of the power spectrum enables use of conventional statistical tools to compute the significance of changes in mean modal frequencies.

Interpretation of power spectra as statistical distributions for use in statistical inference is quite novel, but the inverse of this interpretation is commonly used in simulation. An infinite time-history of a stationary process would result in a theoretically perfect (statistical) distribution of the energy of the response, typical of a target power spectrum used in simulation. Mapping of a power spectrum into statistical space is precisely the inverse of simulating a time-series using a power spectrum with random phase. In first-order simulation, a time-history is commonly realized from a target power spectrum as:

$$x(t) = \sum_{n=1}^N H_n \cos[2\pi f(n)t + \theta_n] \quad (1)$$

where $x(t)$ is the time-history of the simulation. H_n is the amplitude of each frequency component, the square of which is directly proportional to energy density at that frequency, i.e., the height of an individual bar of the power spectrum. $f(n)$ is the n th frequency (Hz), and θ_n is a random phase angle.

The power spectrum is appropriate for interpretation as a statistical distribution, rather than the Fourier amplitude spectrum, because the total area under the power spectrum has physical interpretation as the variance of the process. Separate areas under the power spectrum can be added directly. The combined area of two frequency bars of the power spectrum represents the total power of the process at those two frequencies, analogous to combining frequencies of occurrence in a traditional frequency histogram; the combined area of two frequency bars of a Fourier spectrum has no particular physical interpretation.

2.3. Comparison of statistical distributions

An infinite time-history would result in an ideal statistical distribution of energy, but in practice a finite time-history is applied to approximate the true underlying distribution. Various statistical methodologies have been developed to compare the means of two data-sets including corrections for relatively small sample sizes. In the response spectra considered here, the Gaussian distribution provides a reasonable fit to the power spectra being compared, and so the T -statistic provides a convenient method of comparison between the two distributions. The authors note that there is no theoretical reason to expect modal peaks to conform to the Gaussian distribution. In some cases the Gaussian assumption is strongly violated, including those with very high noise or closely spaced peaks; in such cases the T -statistic should not be used.

Other statistical tests including nonparametric methods based on sampling from an observed distribution could alternatively be used in place of the T -test with no fundamental change to this underlying methodology; investigation of extending this method beyond Gaussian shaped modal peaks is the subject of ongoing work by the authors.

In a more typical application, the T -statistic is used with a histogram by first selecting bin widths for the raw data, then binning the data into a frequency histogram to form a statistical distribution. The standard deviation of the raw data can be closely estimated from the geometry of the resulting histogram. The purpose of the T -statistic is to compensate for the effect of small sample size on the standard deviation as compared with an ideal underlying Gaussian distribution. Here, the frequency intervals implicit to the FFT are equivalent to the bin widths. The energy associated with any frequency bar relates to the energy of each sinusoidal cycle at that frequency. As in a more typical application, the standard deviation of the distribution is closely estimated from its geometry; note, however, that this standard deviation is not equivalent to the square root of the variance of the process, but instead represents the distribution of energy about the mean frequency of an individual modal response. The number of degrees of freedom must also be estimated for use with the T -distribution. Here, the number of degrees of freedom is estimated as the number of cycles over which the energy is averaged, which is approximated by the mean frequency in Hertz times the duration of the time-history in seconds. This procedure is equivalent to dividing the total duration by the peak period since the mean and peak frequencies coincide under the Gaussian assumption.

3. Computational details

3.1. 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 pre-specified 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 and/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.

To find 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 Fig. 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.

3.2. Modal distributions and statistical distributions

The first step in calculating distribution parameters is normalizing each modal distribution to have unit energy such that it can be treated as a frequency histogram, as outlined in Section 2.2. Within each modal distribution, the fraction of energy associated with each frequency bar is computed as:

$$P_i(e_n) = \frac{a_n}{A_i} \quad (2)$$

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 (Fig. 1), or the total area between relative minima:

$$A_i = \sum_{n=1}^{N_i} a_n = \sum_{n=1}^{N_i} S(n)df \quad (3)$$

in which $S(n)$ is the offset at the n th frequency of the power spectrum, 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. [10,11].

$$\mu_{i,w} = \frac{m_i}{A_i} = \frac{1}{A_i} \sum_{n=1}^{N_i} S(n)f(n)df \quad (4)$$

$$s_{i,w}^2 = \frac{\theta_i}{A_i} = \frac{1}{A_i} \sum_{n=1}^{N_i} S(n)(f(n) - \mu_{i,w})^2df \quad (5)$$

where $\mu_{i,w}$ and $s_{i,w}^2$ 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 implicitly 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.

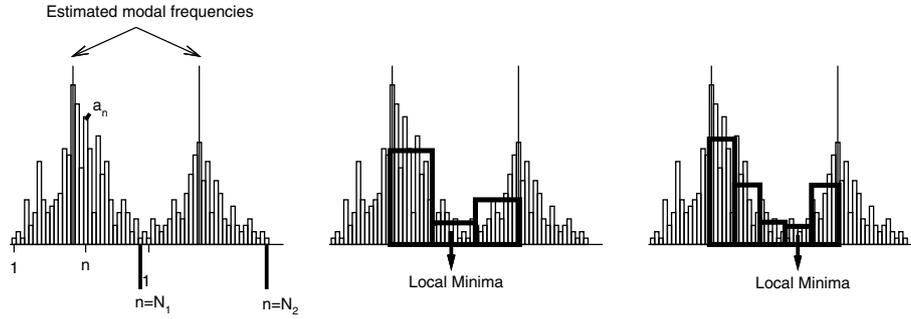


Fig. 1. Separation of modal distributions: The penalty method.

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.

3.3. Comparison: Statistical hypothesis testing

Spectral peaks of field data generally appear as a cluster of energy around some frequency. Here, mean frequencies of individual modes 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.

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, though in the example presented here the energy peaks are broadly spaced and damping is sufficient to have “bell-shaped” modal peaks, so use of the *T*-test may be reasonable. A power spectrum must be computed to assess the degree to which individual peaks conform to the Gaussian distribution because the shape of these peaks is independent of the relative Gaussianity of the underlying process. As noted in Section 2.3, for some cases the Gaussian assumptions are strongly violated and use of the *T*-test is not appropriate. If the *T*-test is deemed suitable, the *T*-statistic can be computed as e.g. [12]:

$$T_i = \frac{\Delta\mu_i}{s_i} \tag{6}$$

where

$$\Delta\mu_i = \mu_{i,1} - \mu_{i,2} \tag{7}$$

$$s_i^2 = s_{i,1}^2/N_{i,1} + s_{i,2}^2/N_{i,2} \tag{8}$$

$$N_{i,w} = D_w\mu_{i,w} \tag{9}$$

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. [12]:

$$DOF_i = \frac{(s_{i,1}^2/N_{i,1} + s_{i,2}^2/N_{i,2})^2}{(s_{i,1}^2/N_{i,1})^2/(N_{i,1} - 1) + (s_{i,2}^2/N_{i,2})^2/(N_{i,2} - 1)} \tag{10}$$

3.4. 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:

$$E_{i,w} = \frac{A_{i,w}}{\sum_{i=1}^I A_{i,w}} \tag{11}$$

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 Eq. (12), in which $\Delta\mu$ and s^2 are computed as averages weighted by the fraction of energy associated with each of the *I* modes:

$$T = \frac{\Delta\mu}{s} \tag{12}$$

where

$$\Delta\mu = \sum_{i=1}^I 0.5(E_{i,1} + E_{i,2})\Delta\mu_i \tag{13}$$

$$s^2 = \sum_{i=1}^I 0.5(E_{i,1} + E_{i,2})s_i^2 \tag{14}$$

The *T*-statistic resulting from Eq. (12) 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. This number is expected to be larger than the observed number of zero-upcrossings in the time-history because modal frequencies with small fractions of the total energy will not directly contribute to the number of upcrossings. The total degrees of freedom to be used in conjunction with the *T*-statistic is estimated as the sum of individual degrees of freedom:

$$DOF = \sum_{i=1}^I DOF_i \tag{15}$$

4. 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 pre-determined 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 function e.g. [12]:

$$g(x; \alpha, \theta) = \frac{x^{\alpha-1} e^{-\frac{x}{\theta}}}{\Gamma(\alpha, \infty) \theta^\alpha} \quad (16)$$

$$G(x; \alpha, \theta) = \frac{\Gamma(\alpha, \frac{x}{\theta})}{\Gamma(\alpha, \infty)} \quad (17)$$

in which $\Gamma(\alpha, y)$ denotes the gamma function: $\Gamma(\alpha, y) = \int_0^y t^{\alpha-1} e^{-t} dt$. The shape parameter (α) 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.

5. Example 1: Simulated data

In this section and Section 6, two examples are presented, both relevant to vortex induced vibrations of marine drilling risers, and both based on physical data of practical importance. 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”.

The first family of examples uses data simulated from target power spectra originally estimated from measured data for a case with poorly defined spectral peaks; the second example applies the new method directly to measured time-histories for a case with well-defined peaks. The use of simulated data is helpful because the underlying target spectra are known precisely, and so can be used to assess the effectiveness of the method under various conditions. The example quantifies the probability of both Type I and Type II errors, then compares the results of this method with conventional tests of stationarity. In the second example, the method is applied to measured data to demonstrate applicability to measured riser data. Results are broadly as expected, though effectiveness cannot be well-quantified because the statistical characteristics of the parent processes are not precisely known.

5.1. Data simulation from known underlying target spectra

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 (Fig. 2). The field data shown has observable energy above the first two modes that is not clearly separable into distinct peaks; this energy may be indicative of higher modes of vibration at very closely spaced frequencies. In this example, energy higher than the first two modes is not believed to be useful for detecting changes in the underlying process, and so is excluded from the analysis in both data simulation and method application. After the measured data has been idealized into a target power spectrum, known changes in modal properties are introduced to the smoothed spectrum. Time-histories needed for the example are then simulated from these spectra using the conventional method as in Eq. (1). In this example, the simulation includes as many frequency components as there are time-steps to ensure randomness. 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.

5.2. 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 min data window test for target spectrum S_1 is the first part of the time-history used in the 60 min 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.

5.2.1. 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. Fig. 4 shows the average of the P -values from these tests, all of which are above

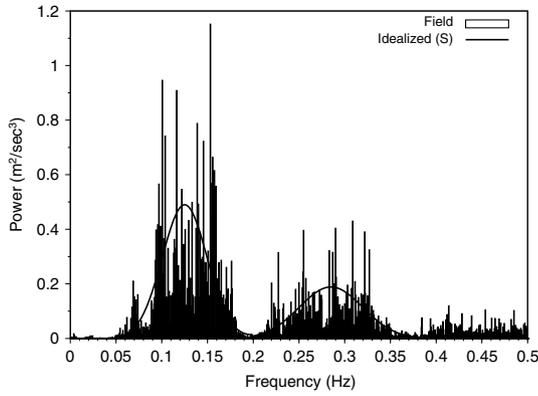


Fig. 2. Field data spectrum and idealized target spectrum S.

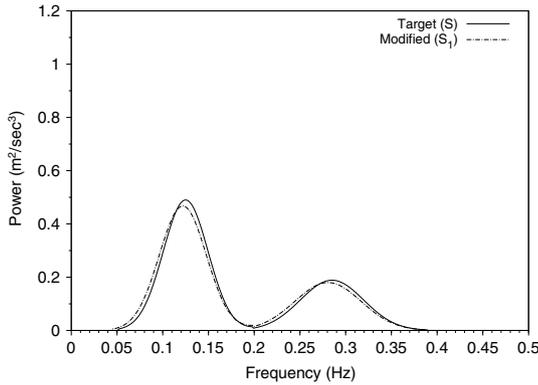


Fig. 3. Original spectrum (S) and modified spectrum (S₁).

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 min of data is considered.

5.2.2. 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:

$$\Delta\mu_i = z s_i; \quad \Phi(z) = \int_{-\infty}^z \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}u^2} du. \quad (18)$$

For example, applying Eq. (18) with $\Phi(z) = 5\%$ of the area under a standard normal distribution yields $\Delta\mu = -0.1257s_i$. 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. Fig. 3 shows the original target spectrum, *S*, with one of the modified spectra, *S*₁. The resulting magnitudes of each frequency shift is particularly 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.

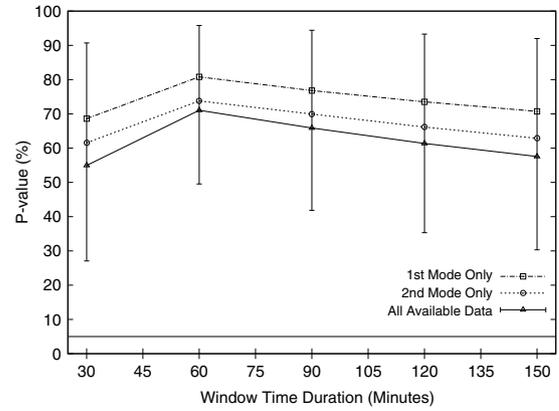


Fig. 4. *P*-value that the observed differences between the time-series are due to differences between underlying spectra when in fact no such differences exist.

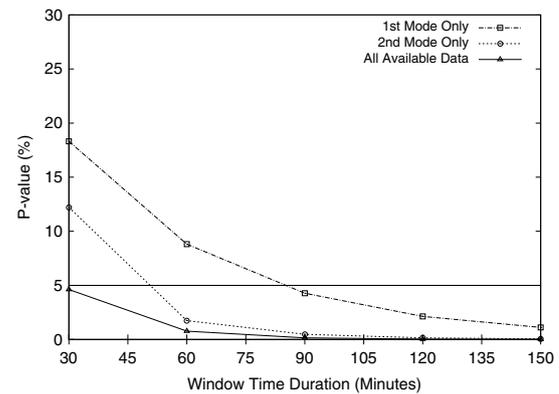


Fig. 5. Ability to detect underlying changes in the target spectrum as a function of time-history duration (smaller *P*-value indicates greater ability).

5.2.3. 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 Fig. 5. The monotonic decrease shown in Fig. 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 Fig. 3 can be detected with about 95% confidence (5% *P*-value) using only 30 min of data.

Fig. 6 shows the probability of failing to detect the known shift in modal frequencies from Fig. 3. These probabilities are calculated by treating the *P*-value as a random variable (Eq. (17)). Using only 30 min 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 min of data, the method is almost certain to detect even the very subtle change indicated in Fig. 3.

5.2.4. 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 Fig. 5 is again presented as Fig. 7, but with the x-axis presented

Table 1
Target spectra parameters with percent change from original idealized spectrum.

Target spectrum	1st mode				2nd mode			
	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

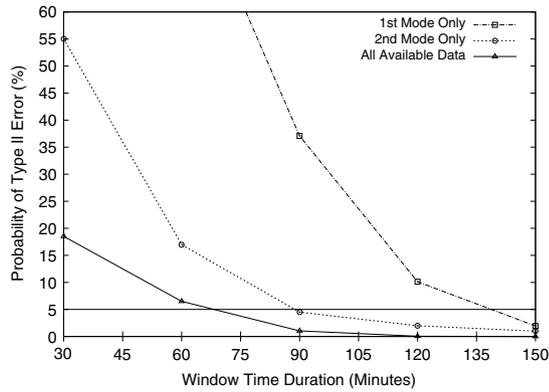


Fig. 6. Probability of failing to detect a known difference between underlying spectra as a function of time-history duration.

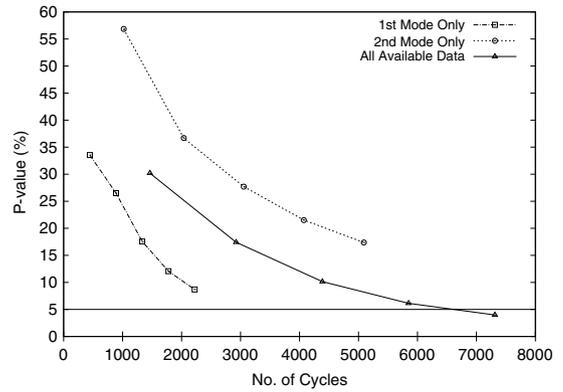


Fig. 8. Very high noise case (19 mg): Detection of underlying changes in the target spectrum with Gaussian assumption strongly violated.

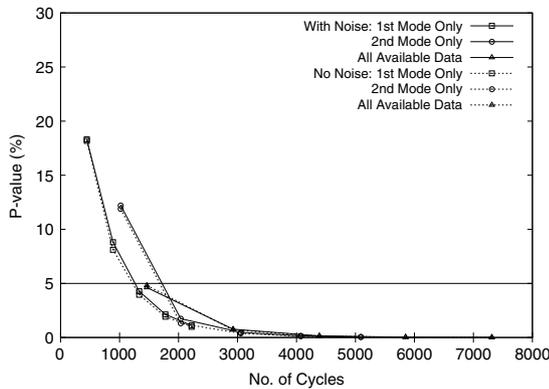


Fig. 7. Ability to detect underlying changes in the target spectrum as a function of number of cycles with and without noise.

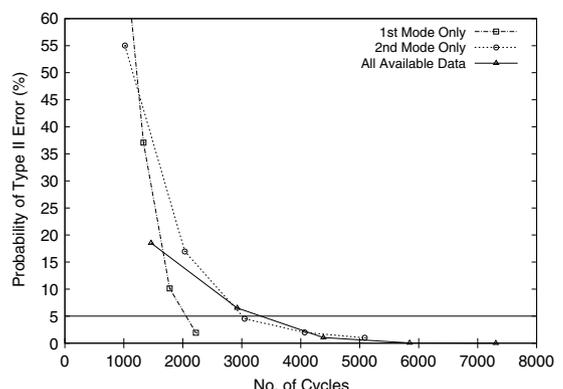


Fig. 9. Probability of failing to detect a known difference between underlying spectra as a function of number of cycles.

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 (Eq. (18)) and not of the method itself.

Fig. 7 shows that approximately 1500 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 Fig. 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, Fig. 9 shows that at least 3000 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. Fig. 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 Eq. (18); 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 Fig. 7. The very subtle shift corresponding to 5% of the area under a normal distribution requires around 1400 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.

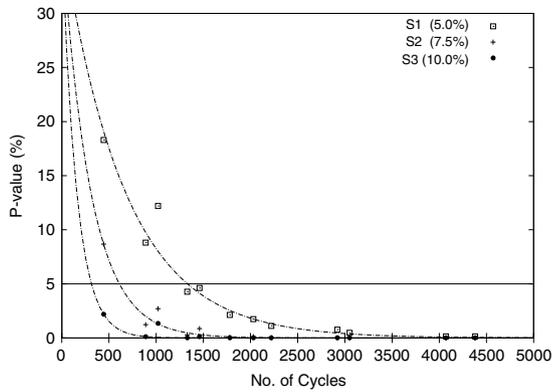


Fig. 10. Number of cycles required to detect increasingly large shifts in mean modal frequencies. Percents refer to fractions of the area under a standard normal distribution.

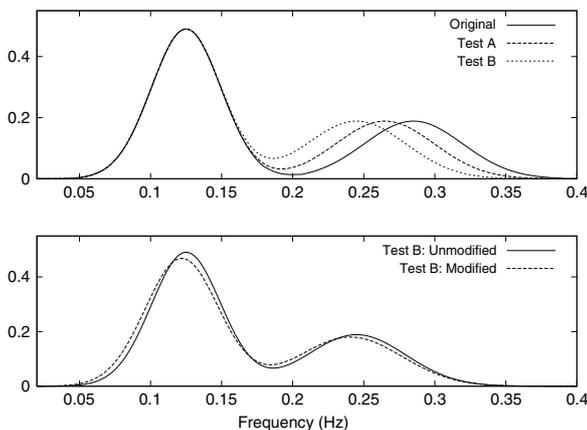


Fig. 11. Target spectra for use in assessing sensitivity to closely spaced peaks.

5.3. Closely spaced peaks

The newly proposed method relies on several underlying assumptions, one of which is that the peaks are well spaced. Closely spaced peaks violate the Gaussian shape assumption because the method effectively cuts off the tail of each distribution at the bottom of the trough between peaks. Here, the importance of widely spaced peaks is investigated by application of the method to spectra in which the peaks are shifted together. Fig. 11 shows the original target spectrum used in earlier parts of the example, plus two modified spectra in which the peaks have been moved together. In the earlier parts of the example, the means of the first and second modes (Fig. 11, upper) were 0.125 Hz and 0.285 Hz, respectively. Here, the mean of the first mode is left at 0.125 Hz, and in Test A the second mode is shifted to 0.265 Hz, and in Test B the second mode is shifted to 0.245 Hz. For each of these cases, a slight shift in the power spectrum is then introduced. Case B is shown in Fig. 11, lower. A constant magnitude of this slight shift is maintained for each case at 5% of the area under a standard normal curve to demonstrate the test sensitivity to decreasing peak spacing. Fig. 12 shows the test results in higher p -values (less sensitivity) with decreasing peak spacing, but the trend is not dramatic.

5.4. Comparison with conventional tests

The newly proposed method is compared with two conventional tests of statistical stationarity: the reverse arrangements test

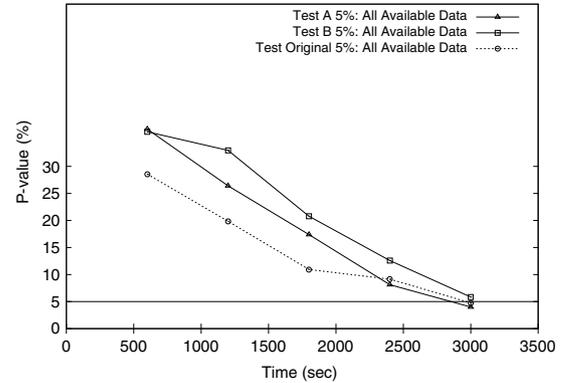


Fig. 12. Ability to detect underlying changes in the target spectrum for progressively more closely spaced peaks.

and the runs test e.g. [13,14]. Numerous pairs of 150 min 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 (Figs. 4 and 6).

6. Example 2: Measured data

In this section, the modal distribution method is applied to measured acceleration time-histories of marine risers. Assessment of the effectiveness of any new statistical method applied to field-measured data is challenging because the true answer, whether or not the data-sets represent the same underlying process, generally is not known. Here, four individual measured time-series are selected for which the actual process is relatively simple and well understood.

Figs. 13 and 14 show power spectra of the measured acceleration time-histories, each of which was computed as an SRSS of accelerations along the three principal axes of a single accelerometer for a single 30 min data-set. All four of these spectra exhibit exceptionally well-defined peaks, which gives further insight into the process, but unfortunately do not conform well to the Gaussian assumption. Applying the MDM to time-histories measured at the same time but at different locations along the riser length gives good assurance that the vibrations represent the same parent process. The motion associated with Fig. 13 is a vibration having transverse motion along the length characteristic of a string vibrating in its second and fourth physical modes, as deduced from a series of accelerometers along the riser length. The motion associated with Fig. 14 was measured later and the associated motion is characteristic of the third and fourth physical modes of vibration.

Figs. 15 and 16 show the results of application of the MDM to the measured time-histories. In this test, the frequency range was manually limited to 0.035 Hz to 0.075 Hz, which effectively

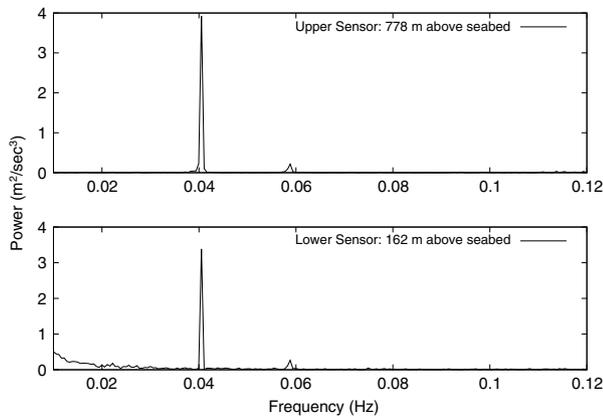


Fig. 13. Observed power spectra for riser vibration characteristic of modes 2 and 4.

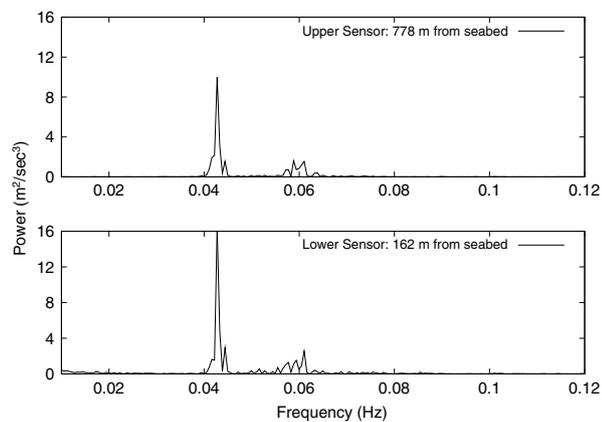


Fig. 14. Power spectra observed later for the lower riser location (162 m above seabed). Vibration characteristic of modes 2 and 4 compared with modes 3 and 4.

eliminates considerable low-frequency energy from the deeper accelerometer location. The method was found to be far less effective without this manual intervention. Fig. 15 shows the results of the MDM being exercised on the pair of acceleration time-histories that both represent the same parent process (Fig. 13). As expected, the MDM predicts there is only a very small overall chance that these time-histories do not represent the same process. The method did find a statistically significant difference between the two higher peaks of the spectra, most probably because of the substantially higher noise level at the more deep-water sensor location. The lower mode, however, showed virtually no statistical difference, and its substantially larger area caused it to dominate the overall MDM result.

In contrast, Fig. 16 shows results for two time-histories, both measured at more deep-water sensor, but at different times (the lower plots on Figs. 13 and 14). Significantly different frequencies in the first peak are expected because different physical modes of vibration were observed at the two times. As expected, the MDM found a large difference between the lower peaks of the two spectra, and the relatively larger energy of that vibration caused it to dominate the overall MDM prediction.

The cases presented in this example were selected from a large set of available data because the physical vibration modes are clear and the peaks in the power spectrum are well defined. In addition to these cases, the investigators also exercised the MDM on more complex cases including: high degrees of non-stationarity, modal peaks clearly corrupted by wave motion of the vessel, and spectra with poorly defined peaks. In cases including very complicated

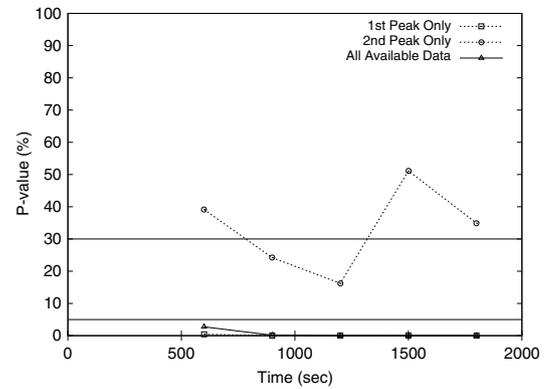


Fig. 15. MDM results from comparison of time-histories measured at upper and lower locations at the same time. Behavior characteristic of modes 2 and 4.

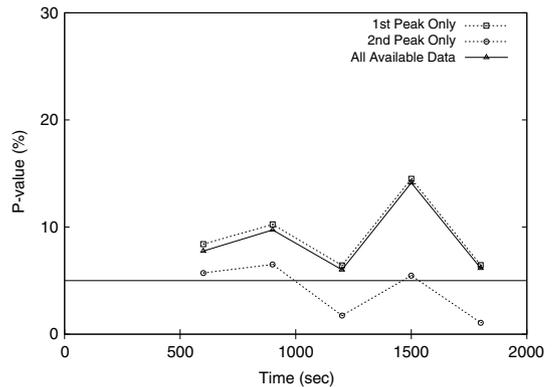


Fig. 16. MDM results comparing time-histories measured at the same location but at separate times. Behavior characteristic of modes 2 and 4 versus 3 and 4.

riser behavior, it was at times difficult to ascertain exactly what should be concluded from the numerical results produced by the MDM. These complex cases were not believed to be helpful as published examples, but the reader is cautioned that interpretation of MDM results for very complex structural behaviors may not be straightforward.

7. 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 (Figs. 5–10), even if the data is simulated from underlying target spectra with only very subtle differences (Fig. 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 (Fig. 10). The new method is considerably more effective than the conventional runs and reverse arrangements tests. The method has been demonstrated to be effective on measured data exhibiting relatively simple vibratory behavior (Figs. 15 and 16), but application on very complex behaviors may not be straightforward.

The modal distribution method is based on transforming individual modes of a power spectrum into a statistical distribution, 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. An example is presented in which minimum data requirements for the new method are investigated 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.

Acknowledgements

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 made available for this work by British Petroleum.

References

- [1] Vandiver JK. Detection of structural failure on fixed platforms by measurement of dynamic response. In: Proc. of the 7th annual offshore technology conference. 1975. p. 243–52.
- [2] Vandiver JK. Detection of structural failure on fixed platforms by measurement of dynamic response. *Journal of Petroleum Technology* 1977;305–10.
- [3] Loland O, Dodds J. Experience in developing and operating integrity monitoring system in north sea. In: Proceeding of eighth annual offshore technology conference. 1976. p. 313–19.
- [4] Crema LB, Mastroddi F. A direct approach for updating and damage detection by using frf data. In: Proceedings of ISMA23, Noise and vibration engineering. 1998.
- [5] Harris CM, Crede CE. *Shock and vibration handbook*. New York: McGraw-Hill; 1995.
- [6] Smyth A, Masri S, Caughey T, Hunter N. Surveillance of mechanical systems on the basis of vibration signature analysis. *ASME Journal of Applied Mechanics* 2000;67:540–51.
- [7] Katfygiotis L, Lam H. A probabilistic approach to structural health monitoring using dynamic data. In: *Structural health monitoring, current status and perspectives*. Palo Alto (California): Stanford University. 1997. p. 152–63.
- [8] Abe M, Fujino Y, Kajimura M, Yanagihara M, Sato M. Monitoring of a long span suspension bridge by ambient vibration measurement. In: *Structural health monitoring 2000*. Palo Alto (California): Stanford University. 1999. p. 400–7.
- [9] Sohn H, Farrar CR, Hemez FM, Shunk DD, Stinemates DW, Nadler BR. A review of structural health monitoring literature, Tech. rep. LA-13976-MS, Los Alamos National Laboratory. 2003.
- [10] Bethea RM. *Statistical methods for engineers and scientists*. M. Decker; 1984.
- [11] Brandt S. *Statistical and computational methods in data analysis*. North-Holland Pub. Co.; 1970.
- [12] Navidi W. *Statistics for engineers and scientists*. New York: McGraw-Hill Higher Education; 2006.
- [13] Newland DE. *An introduction to random vibrations, spectral & wavelet analysis*. 3rd ed. Longman, Scientific & Technical and John Wiley & Sons; 1993.
- [14] Bendat JS, Piersol AG. *Random data: Analysis and measurement procedures*. New York: John Wiley & Sons; 1991.