Available online at www.sciencedirect.com
SCIENCE DIRECT®
Journal of Sound and Vibration 283 (2005) 295–310
JOURNAL OF SOUND AND VIBRATION
www.elsevier.com/locate/jsvi

A multivariate, attractor-based approach to structural health monitoring

L. Moniz[a,*], J.M. Nichols[b], C.J. Nichols[c], M. Seaver[b], S.T. Trickey[b], M.D. Todd[d], L.M. Pecora[a], L.N. Virgin[c]

[a]*U.S. Naval Research Laboratory, Code 6362, 4555 Overlook Avenue SW, Washington, DC 20375, USA*
[b]*U.S. Naval Research Laboratory, Code 5673, 4555 Overlook Avenue SW, Washington, DC 20375, USA*
[c]*Department of Mechanical Engineering, Duke University, Durham, NC 27708, USA*
[d]*Department of Structural Engineering, University of California San Diego, La Jolla, CA 92093-0085, USA*

Received 12 September 2003; accepted 14 April 2004

Abstract

In this work, recent advances in the use of nonlinear time-series analysis for structural health monitoring are extended to incorporate multivariate data. Structural response data recorded at multiple locations are combined using a multivariate time delay embedding in order to reconstruct the structure's dynamical attractor. Using this approach, a global phase-space representation of the dynamics may be realized for spatially extended systems. A new attractor-based metric, chaotic amplification of attractor distortion (CAAD), is then introduced as a damage sensitive feature. The approach is implemented using data acquired from a composite beam, bolted at either end to steel plates. Degradation to the system is introduced as a loosening of the bolts at one end of the structure. Results based on multivariate attractor reconstruction show a clear ability to detect both the presence and magnitude of damage to the connection. Comparisons are then drawn between this approach and one where the same feature is extracted from attractors reconstructed using data acquired from the individual sensor locations. These features are combined "post-extraction" using a linear discriminant coordinant analysis. Performing the analysis separately at the individual sensor locations results in a significant reduction in discriminating power.

*Corresponding author. Tel.: +1-202-767-6259; fax: +1-202-767-1697.
E-mail address: moniz@anvil.nrl.navy.mil (L. Moniz).

0022-460X/$ - see front matter © 2004 Elsevier Ltd. All rights reserved.
doi:10.1016/j.jsv.2004.04.016

1. Introduction

Vibration-based Structural Health Monitoring (SHM) has arisen as one possible approach to the problem of structural damage identification and localization. The structure of interest is excited through either ambient or applied loading and the resulting time series analyzed for damage-induced changes to the dynamics. Much of the research in this field has therefore focused on the problem of feature extraction, which attempts to answer the question: *what metric computed from that time series accurately quantifies changes to the dynamics*? Ideally the chosen feature will be able to discern the magnitude, type, and location of the degradation. This problem belongs to the general class of inverse problems where the practitioner attempts to infer something about a structure through the system's vibrational response.

The feature extraction problem has largely tended to focus on a modal analysis of the response. The structure is excited with a broadband input and the response analyzed for changes in frequencies, mode shapes, damping, flexibility, etc. [1–5]. Other non-modal-based features have included cross-correlation [6], auto-regressive approaches [7,8], neural networks [9,10], time-series dimensionality [11,12], wavelets [13–15] and genetic algorithms [16]. Many of the studies conducted to date have focused on extracting the relevant information from a single sensor. Even in the case where multiple sensors are present, the features extracted from the corresponding time series are often analyzed separately. Recent efforts have been directed toward treating the problem of damage detection in a multivariate framework, that is, looking for *combinations of* features (gleaned from different locations) capable of assessing the health of a structure [7,17].

This work represents an effort to combine information from different spatial locations prior to the feature extraction. The approach builds on earlier efforts aimed at transforming the feature extraction problem into one that lends itself to attractor-based analysis. Recent studies have shown that exciting a structure with a chaotic input can force the system (structure + forcing) dynamics to occur on a low-dimensional attractor. Various attractor-based features have been highly successful in determining both the presence and magnitude of several forms of structural degradation (see Refs. [18–20]). Here the attractor reconstruction is generalized to incorporate multiple time series. Time series collected at each of the sensor locations are analyzed using a recently developed multi-dimensional false nearest-neighbors approach to choose the number of delay vectors required from each of the sensors. The time series are then combined using the prescribed multivariate time-delay embedding in order to produce a reconstruction of the system's attractor. The resulting representation of the system dynamics is then analyzed for damage-induced changes using a new feature, chaotic amplification of attractor distortion (CAAD). The approach is demonstrated effective in detecting both the presence and magnitude of degradation in a composite/steel joint. We compare the technique to results obtained by performing the same analysis on attractors reconstructed from the individual sensors. The same feature (CAAD) is computed at each location and combined a posteriori using a linear multivariate discriminant analysis. Results clearly indicate that in this instance there is a large advantage to combining time series information *before* proceeding with the feature extraction.

2. Attractor reconstruction from multivariate data

Assume a dynamical system evolving according to

$$\dot{\mathbf{x}} = \mathbf{F}(\mathbf{x}), \quad x \in \mathbb{R}^d. \tag{1}$$

The state of the system at any point in time is specified by the vector $\mathbf{x}$ in d-dimensional *phase space*. An initial condition $\mathbf{x}(0)$ will, under the action of $\mathbf{F}$, asymptotically approach the system's dynamical attractor, a subset of phase space which may be thought of as a geometric object to which all trajectories belong. The central goal of an attractor reconstruction is to qualitatively preserve the dynamics of the "true" underlying attractor by reconstructing the solution of Eq. (1) i.e. re-create the state vectors $\mathbf{x}$. In experiment one will typically not have access to each of the system's state variables and must instead rely on a set of powerful mathematical embedding theorems to reconstruct the missing information qualitatively. Based on the early work of Whitney [21] and later by Takens [22] and Sauer [23] it can be shown that delayed copies of the original time series may be used to form an embedding

$$X \equiv \mathbf{x}(n) = (x(n)), x(n+T), \ldots, x(n+(m-1)T)), \tag{2}$$

which faithfully represents the geometry of the system which exhibited the dynamics in the time series.

By taking advantage of the coupling that occurs naturally in the dynamics, information pertaining to the unobserved variables may be extracted from a single $x(n)$. This remarkable result has formed the basis for an entire class of attractor-based system identification tools in the broad field of nonlinear time-series analysis. Theoretically the choice of delay is unimportant, however in practice the choice may have important consequences. As the delay approaches $T = 0$ the resulting pseudo-state vectors will be identical while for large delay the resulting vectors may become completely unrelated and destroy any connection to the underlying dynamics. In order to minimize redundancy the delay is often selected in accordance with a pre-defined drop in the autocorrelation function or as the first minimum of the average mutual information function [24]. Here, the delay is taken as the time at which the autocorrelation loses $\frac{2}{3}$ of its value. Multiple time series taken from different sensors often exhibit a broad range of autocorrelation delay times. Because redundant information is less harmful than the omission of information in a time-series embedding, we compute autocorrelation separately for each time series and use the minimum time (among all time series) in which autocorrelation decays to $\frac{1}{3}$ of its value. Although the first minimum of the average mutual information function is often used to find embedding delay, it has been shown that mutual information can occasionally give incorrect information, especially in the context of multivariate embedding [25].

Improper selection of the embedding dimension m can also produce a poor reconstruction. The goal in selecting m is to minimize the number of crossings of reconstructed trajectories due to embedding in too small a space. If information in higher dimensions is projected onto too few coordinates, the trajectories in the reconstructed space will not reflect the dynamics correctly. For example, consider an attempt to embed a helix into a two-, rather than a three-dimensional space. In two dimensions the helical trajectory is reduced to a cycle and information concerning the expansion of the trajectory in one direction is lost.

If one knows the fractal dimension d_f of the system, recourse is made to the theoretical criteria of Sauer [23] that $m > 2d_f$. In most cases such information is not available and an empirical approach is taken. The false nearest-neighbor technique of Kennel et al. [26] works by slowly "unfolding" the attractor, that is, increasing the embedding dimension starting from $m = 1$ and searching for the value at which the number of false projections reaches zero. A false projection is defined as a relative separation in distance greater then some pre-defined threshold as a trajectory is extended from m to $m + 1$ coordinates. Again, think of the example of the helix; points in separate cycles are on top of each other and thus near neighbors in two dimensions. As the embedding dimension increases to three they are no longer neighbors. Those would be counted as false near neighbors as the dimension increased from 2 to 3. If the dimension is increased from 3 to 4, no points that are neighbors in 3 dimensions fail to be neighbors in 4 dimensions; 3 dimensions is therefore sufficient. The false nearest-neighbors approach has become the standard approach for selecting an embedding dimension for a single time-series response and has also been used in various forms to select embedding dimension in multivariate time-series analysis [27–29].

In the instance of multiple observations, as it often occurs with spatially extended systems, the process remains largely the same. However, in the multivariate case the false nearest-neighbors approach must incorporate delayed copies of all time series data.

We denote the multivariate time series $x_i(n)$ as the ith time series sampled at discrete time n. In order to find an appropriate embedding dimension, we performed a multivariate false nearest neighbors [29,30] examination. Because the data from the sensors were strongly coupled, as opposed to the weakly coupled data in Ref. [29] we performed a suite of false nearest-neighbors examinations (described in Table 1) in order to investigate more completely the possible dimension of the attractor. We begin with one coordinate from a time series ($x_i(n)$) (a base dimension of 1; see Table 1, case ID A–E) and then add delay coordinates $x_i(n)$ from each time series, one at a time. We consult Fig. 1 for cases A–E to see the total fraction of false nearest

Table 1
Multivariate false nearest neighbors base dimension scenarios, including dimension of each time series

Case ID	Total base dimension	Dimension of time series					Case ID	Total base dimension	Dimension of time series				
		1	2	3	4	5			1	2	3	4	5
A	1	1	0	0	0	0	N	6	2	1	1	1	1
B	1	0	1	0	0	0	O	6	1	2	1	1	1
C	1	0	0	1	0	0	P	6	1	1	2	1	1
D	1	0	0	0	1	0	Q	6	1	1	1	2	1
E	1	0	0	0	0	1	R	6	1	1	1	1	2
F	2	1	1	0	0	0	S	7	2	2	1	1	1
G	3	1	1	1	0	0	T	8	2	2	2	1	1
H	4	1	1	1	1	0	U	9	2	2	2	2	1
I	5	1	1	1	1	1	V	10	2	2	2	2	2
J	2	1	0	0	0	1	W	7	2	1	1	1	2
K	2	0	0	0	1	1	X	7	1	1	1	2	2
L	3	0	0	1	1	1	Y	8	1	1	2	2	2
M	4	0	1	1	1	1	Z	9	1	2	2	2	2

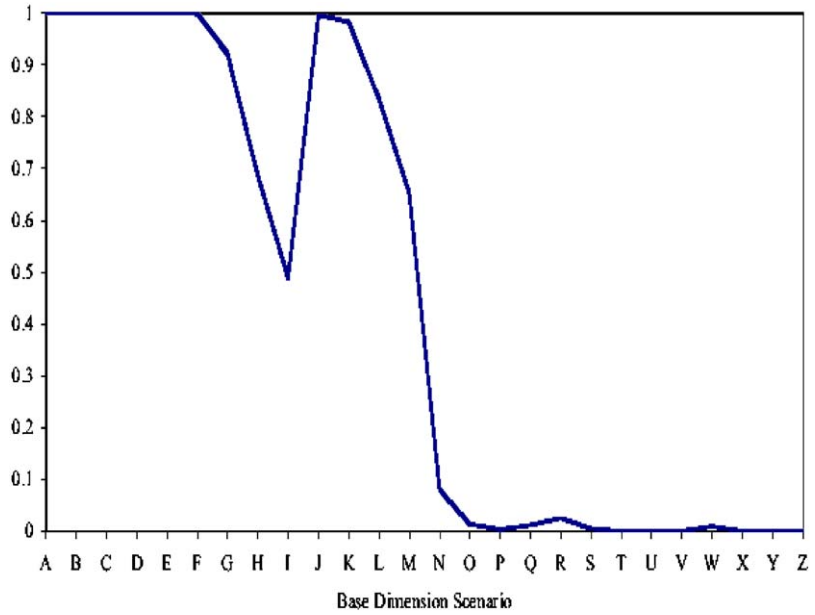

Fig. 1. Total fraction of false nearest neighbors with addition of delays.

neighbors after addition of one delay *per time series*. Embedding one time series only (cases A–E) is clearly inadequate; the fraction of false nearest neighbors is nearly 1.0. This is echoed in cases F–N. We see that at least one coordinate per time series is necessary for the fraction of false nearest neighbors to be less than 0.5. For case I, with delay coordinate in dimension 5 $((x_1(n), x_2(n), \ldots, x_5(n)))$ we see some improvement to a fraction of 0.5 false nearest neighbors. Thus, we concentrate on cases N–Z. For these, we begin with $(x_1(n), x_2(n), \ldots, x_5(n))$ and add additional delays $(x_i(n+T))$, $i = 1, \ldots, 5$ for the base dimension.

For example, case S's base vectors look like $(x_1(n),\ x_1(n+T),\ x_2(n),\ x_2(n+T),\ x_3(n),\ x_4(n),\ x_5(n))$. During the false nearest-neighbors examination, additional delay coordinates $(x_i(n+jT),\ j = 1, 2, 3)$ were added, one time series at a time. For example, to add a delay to Case S from time series 1 we add $x_1(n+2T)$, to add a delay to case S from time series 5 we add $x_5(n+T)$. We again consider Fig. 1, cases N–Z, to see the total fraction of false near neighbors for adding a additional delay in all 5 time series, starting with each base dimension. We investigated scenarios N–Z whether or not the previously added delays or scenarios showed an increase in false nearest neighbors (this is in contrast to Ref. [29]). Thus, we were able to see the changes in the number of false nearest neighbors over a comprehensive range of possible delay embeddings. Using this data, we determined that a sufficient embedding dimension was 10 (in Table 1, case V), using 2 delays per time series in our multivariate embedding (see, e.g. Refs. [28,29]). A vector at time n using this scheme is constructed as follows:

$$\mathbf{x}(n) = (x_1(n), x_1(n+T), x_2(n), x_2(n+T), \ldots, x_5(n), x_5(n+T)). \tag{3}$$

The base dimension scenarios are listed in Table 1. Note that the false near neighbors are computed as dimension is *increased* from the beginning, or base dimension. As in the univariate

false nearest neighbors, if the fraction of false near neighbors going from the lower dimension to the next higher dimension is small, the lower embedding dimension is sufficient. Since it is possible for one sensor's time series to include information that another time series does not, we fully explore all output from the scheme in order to determine a sufficient embedding dimension.

The next section assumes a reconstructed attractor based on either single or multivariate observations.

3. Chaotic amplification of attractor distortion

One of the hallmarks of a chaotic system is that the trajectories of two nearby points will eventually diverge; this is the sensitive dependence on initial conditions. In a stable filtered chaotic system, the expansion of the drive signal governs the evolution of trajectories; the characteristic exponents of the filter (real parts of the logarithms of the eigenvalues) must be negative (contracting). However, there is at least one positively expanding direction in the chaotic driver. Thus if we take two different filters of the same chaotic drive signal, the expansion in both, governed by the expansion of the drive, is the same.

In our system we excite two different structures with the same chaotic signal. These structures act as stable linear filters of the signal. Although it is not possible to measure the direct dynamical effect of different filters on a chaotic signal because of the stability of the filter mentioned above, the geometry of two attractors will differ if the filter changes even with an identical excitation. However, if the change in the filter is small, geometric changes may not be statistically detectable. We exploit the sensitive dependence in order to amplify the small geometric differences that we may see in the two different attractors.

We also exploit the fact that although there is some noise inherent in any physical measurement, the characteristics of the linear filter provided by the structure do not change unless the structure itself changes. Since we are studying the steady state of the system rather than transients, we can expect that attractors reconstructed from identically filtered signals should be essentially geometrically identical. If we consider a representative point on one attractor, and find its nearest neighbor on a geometrically identical attractor, the points will be close. These close points will have trajectories that will remain close for at least a small amount of time; they will *shadow* each other. If two attractors are not geometrically identical, the nearest neighbor to a point on one attractor will be farther (though perhaps not statistically detectably so) from the point. However, these two trajectories will diverge sooner than the trajectories of the points that were close. This divergence is caused by the drive signal's amplification of small distortions in attractor geometry after several time steps.

This situation is entirely analogous to the so-called "butterfly effect." Imagine the ability to measure one scenario with a butterfly and an otherwise identical scenario without the butterfly. Differences in the scenarios at the time of a movement of the butterfly's wings are nearly undetectable. However, after several hours, some changes could be evident in the steady state of the system. It is this *evolved* difference that we are using to identify changes in the geometry of the state space.

We call the number of time steps used to evolve the trajectories the *shadowing time*. The shadowing time is determined by examining trajectories on attractors that are known to exhibit

identical dynamics, such as those from several identically excited undamaged or pristine structures. The shadowing time is calculated as the number of time steps within which all trajectories stay within a given multiple of the distance between trajectories at time 0, with a maximum of 10 time steps. Theoretically, the sensitivity of this method could be adjusted up or down by varying the number of time steps in the shadowing time, but in our experiments the shadowing time was kept constant.

In order to perform this test based on the differences in evolved trajectories, we choose a fiducial point from one of the attractors. We then find the nearest neighbor to this point on the other attractor. We note with incremental differences in the geometry, the distance of the nearest neighbor will increase with increased change in the filter. We follow the trajectory of the initial points on both attractors, and measure the distance of one trajectory to the other as the trajectories evolve. We then record the distance from one trajectory to the other after an appropriate number (the shadowing time) of time steps. This is the *chaotic amplification of attractor distortion* (CAAD) for this representative point. We repeat this for a set of representative points on the attractor, and record the distribution of CAAD values. We compare distributions of the CAAD metric between attractors that are known to exhibit identical dynamics to CAAD distributions between attractors which we conjecture to exhibit different dynamics due to damage.

Typically, we compute many CAAD metric distributions between attractors from pristine or undamaged structures' responses. To assess damage, we then compute (several instances of) the CAAD metric between the attractors from an undamaged structure's response and attractor from the response of a structure that we suspect is damaged. In this method, *the attractors from the responses of the undamaged structures provide the model for the system*. Distortion of the model, which we measure with the CAAD metric, indicates damage or change to the structure.

4. Statistical considerations

The problem of assessing damage based on a collection of feature values is inherently one of statistical pattern recognition [31] and can be approached on several levels of sophistication. For this assessment, we collect a set of features coming from an undamaged structural response. We then may ask if subsequently collected features are part of an "undamaged" set or part of a statistically different, presumably damaged, set? The next level of sophistication requires that the features be classified as coming from one or more different damaged conditions, i.e. determine the magnitude of the damage. Further classification schemes may possibly indicate type and/or location of damage but here we are concerned with detecting the presence and the magnitude of the damage.

A collection of N_m CAAD values extracted from time series coming from one of N_d damage levels is denoted Φ_{ij}, $i = 1, \ldots, N_m$, $j = 0, \ldots, N_d$. The $j = 0$ set of features corresponds to the undamaged or baseline feature set. Each of the j collections is coming from some unknown distribution. These distributions may be transformed to Gaussian by using a resampling or bootstrapping procedure and invoking the central limit theorem [32]. Some number of elements from distribution j, call this Φ_j, are chosen at random and their mean computed. Repeating the process produces a collection of N_m means which, according to the central limit theorem, will approach a normal distribution as the sample size (number of elements randomly selected

from Φ_j) becomes large. Confidence intervals for the mean CAAD value may then be assigned as

$$UCL_j = \mu_j + Z_{\alpha/2}\sigma_j, \quad LCL_j = \mu_j - Z_{\alpha/2}\sigma_j, \tag{4}$$

where μ_j, σ_j are the means and standard deviations for the resampled data at damage level j and α represents the prescribed level of confidence. Here we take $\alpha = 0.05$ corresponding to a value of $Z_{0.025} = 1.96$. The CAAD values obtained by comparing the undamaged data to another undamaged data set produces the baseline set of feature values. CAAD values from data sets that lie outside UCL_0, LCL_0 indicate damage at the 95% confidence level.

As described in the previous section, computation of the CAAD metric involves combining all of the sensor data into a single attractor before the algorithm is implemented. However, the same algorithm may be applied to data coming from each of the sensors *individually*. Denoting the number of sensors by N_s, the features values are now given by the N_s-vector $\mathbf{\Phi}_{ij}$ where boldface denotes a vector. Confidence intervals for each element of these vectors may be formed individually using the process just described. A better solution, however, is to use some combination of the feature values to optimize the discriminating power. One frequently used approach is to search for the optimal linear combination of feature values for discriminating between several different classes (damage levels) using multiple observations. The goal is to find a vector of coefficients $\mathbf{a} = (a_1, a_2, \ldots, a_{N_s})^{\mathrm{T}}$ which maps a set of N_s sample means onto N_d independent intervals via the transformation

$$\tilde{\Phi}_{ij} = \mathbf{a}\mathbf{\Phi}_{ij}, \tag{5}$$

where each of the $\tilde{\Phi}_{ij}$ is a scalar reflecting the diagnosis, i.e. which damage scenario has occurred. Choice of the coefficients is made by minimizing the resulting within-group variance (keeping the distributions of the $\tilde{\Phi}_{ij}$ narrow) while maximizing the between-group variance (maximizing the distance between the distributions of the $\tilde{\Phi}_{ij}$). Let $\bar{\Phi}_j$ be the mean feature value extracted for the jth damage level (within group mean). Let $\bar{\Phi}$ be the global mean, averaged over all damage levels. Following the derivation given by Seber [33] we define

$$\mathbf{W} = \frac{1}{N_m N_d} \sum_{j=1}^{N_d} \sum_{i=1}^{N_m} (\Phi_{ij} - \bar{\Phi}_j)(\Phi_{ij} - \bar{\Phi}_j)^{\mathrm{T}} \tag{6}$$

as the in-group portion of the covariance matrix and

$$\mathbf{B} = \sum_{j=1}^{N_d} \frac{1}{N_d} (\bar{\Phi}_j - \bar{\Phi})(\bar{\Phi}_j - \bar{\Phi}) \tag{7}$$

as the between-group portion. Note that both $\mathbf{W}$ and $\mathbf{B}$ are $N_s \times N_s$ matrices. Choosing the coefficients to meet the previously stated objectives (minimize in-group variance, maximize between-group variance) can be accomplished by maximizing the ratio

$$R(\mathbf{a}) = \frac{a' B a}{a' W a}. \tag{8}$$

This is recognized as the familiar eigenvalue problem

$$\mathbf{W}^{-1}\mathbf{B}\mathbf{a}_r = \lambda_r \mathbf{a}_r \tag{9}$$

producing $r = 1, \ldots, N_s$ different sets of discriminant coordinates (eigenvectors) and the corresponding eigenvalues λ_r which relate the effectiveness, or discriminating power of the $\mathbf{a}_r$. The collections of feature values at the individual sensors may then be transformed according to Eq. (5) using any of the r sets of coordinates. Confidence intervals for the transformed feature values are then assigned in the same fashion as for the multivariate attractor-based approach.

5. Experimental procedures

The approach was tested on a composite beam structure, bolted at either end to two steel plates as shown in Fig. 2. The beam dimensions were 1.219 m in length $\times$ 17.15×10^{-2} m in width and 19.05×10^{-3} m thick. Each of the end plates were made from stainless steel and measured 20×10^{-2} m $\times$ 20×10^{-2} m $\times$ 20×10^{-3} m. Two 4×1.905 cm thick bolts measuring 8.89 cm in length were used to fasten the beam to the steel frame. Each of the bolts were instrumented thus allowing for direct measurement of the axial force in the connection. The beam was manufactured using a quasi-isotropic layup consisting of (0/90) and ($\pm$45) 0.608 kg knit EGlass fabric. The specific layup is $[(+-45),\ (0/90)]6S$ meaning there are six sets of $(+-45),\ (0/90)$ plies stacked on top of each other in the first half of the laminate. The "S" is used to denote a symmetric laminate meaning that the other half of the laminate is six sets of $(90/0),\ (-+45)$ plies stacked on top of each other. Excitation was provided by means of a B&K electrodynamic shaker located at the mid-span of the beam and oriented such that the stinger was pressing down on the beam (see Fig. 2). In addition, an OMEGA LCFD-25 load cell was placed in-line with the stinger allowing for the input force to the beam to be recorded.

The system chosen for the excitation was the chaotic Lorenz oscillator, given by the three first order equations

$$\eta \dot{z}_1 = 16(z_2 - z_1), \quad \eta \dot{z}_2 = 40z_1 - z_2 - z_1 z_3, \quad \eta \dot{z}_3 = -4z_3 + z_1 z_2. \tag{10}$$

Each of the state equations is multiplied by a constant, η, allowing for both the frequency content and the Lyapunov exponents to be easily adjusted. The Lyapunov exponents, or LEs, describe the average rate of growth or decay of a perturbation to the dynamical system in each phase space direction. This Lorenz oscillator exhibits a large disparity between the magnitudes of the positive (expanding) and negative (contracting) LEs, a desirable property in controlling the

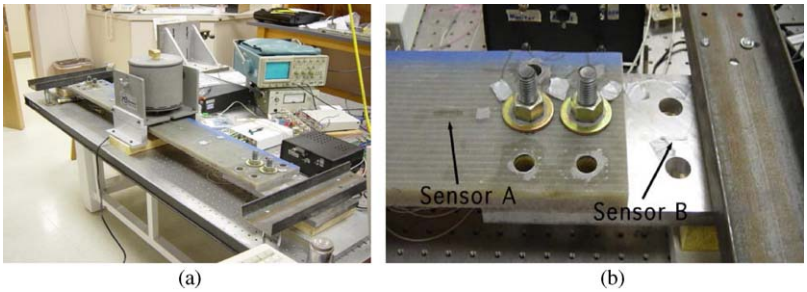

Fig. 2. (a) Photo of composite beam; (b) close-up of joint.

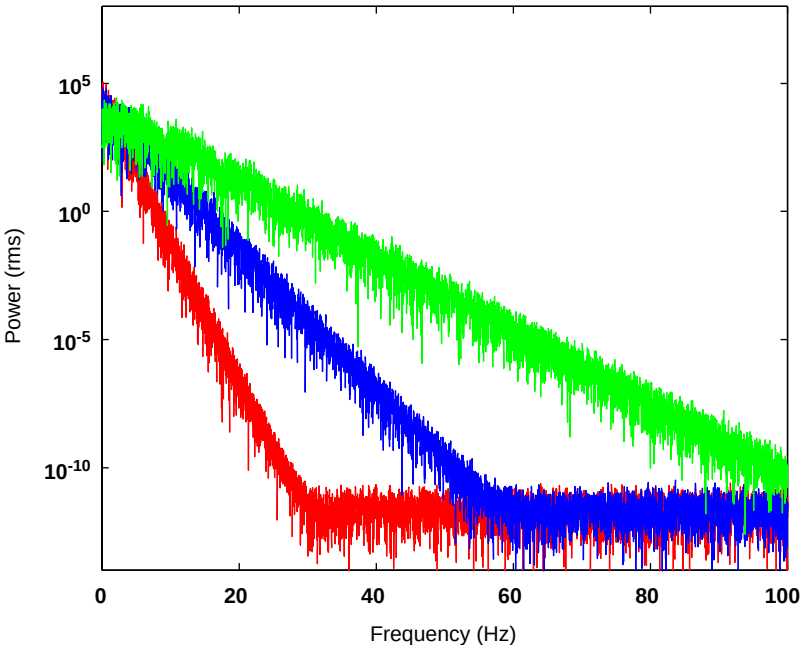

Fig. 3. Power spectrum of Lorenz system with varying coefficient η. From leftmost to rightmost, $\eta = 1.0$, $\eta = 0.5$, $\eta = 0.25$.

dimension of the structure's response (see, e.g. Ref. [34]). We note that the expansion and contraction rates of the drive/response system include both the characteristic exponents of the linear filter provided by the structure and the LEs of the drive [35]. Thus, sensitivity of the system can easily be tuned.

While other chaotic oscillators may be used to control the dimension of the structural response, this particular oscillator possesses characteristic time scales commensurate with many structures. Fig. 3 illustrates the power spectral density of the first state variable, z_1, for several values of η. Each plot shows the power spectral density of the first state variable z_1.

Structural response data were collected at five different locations, four along the center of the beam and one located on the steel plate at the degraded joint using fiber Bragg grating (FBG) strain sensors. These are labeled Sensor A through Sensor E respectively and are shown schematically in Fig. 4. Details of the strain sensing system may be found in Ref. [36]. Damage to the system is introduced as a degradation to the connection strength at one end of the beam (see right panel of Fig. 2). The two bolts (A,B) are slowly loosened from an initial clamping strength of 44.48 to 0.8896 kN. Experiments are performed at each of ten different values of axial load shown in Fig. 5. An experiment consists of five runs where a run is defined as exciting the beam with the Lorenz output and collecting a total of 50,000 observations at a rate of 850 Hz. Multiple runs were collected at each damage level in an effort to assess repeatability and generate a larger database of features for subsequent statistical analysis. Fig. 6 shows a sample of the driving signal, the reconstructed driving attractor, and a sample attractor reconstruction based on the response at sensor B.

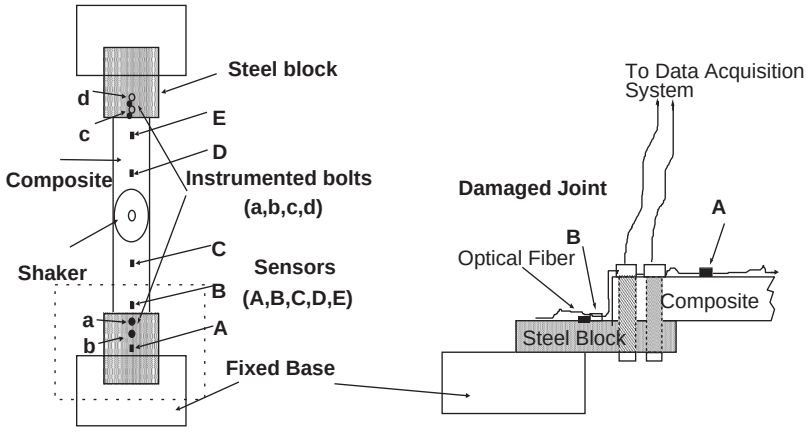

Fig. 4. Schematic of experimental set-up.

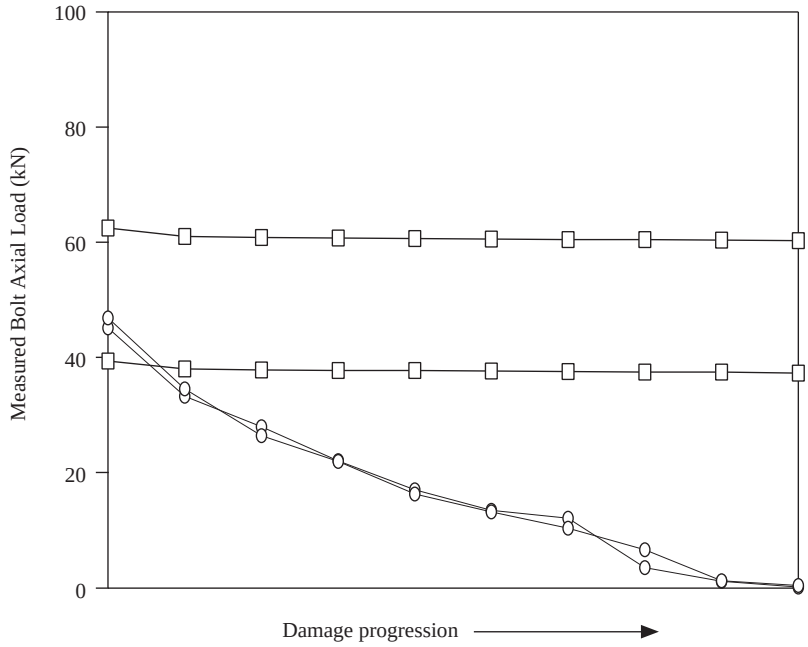

Fig. 5. Progression of axial load (damage) in each of the four instrumented bolts: undamaged end bolts (squares), damaged end bolts (circles).

6. Results

The autocorrelation function was computed for each of the five time series, measured from the structure in the undamaged condition. Delays for each of these series, as indicated by a $\frac{2}{3}$ drop in autocorrelation were found to be 15 time steps. Based on these delays, the multivariate false

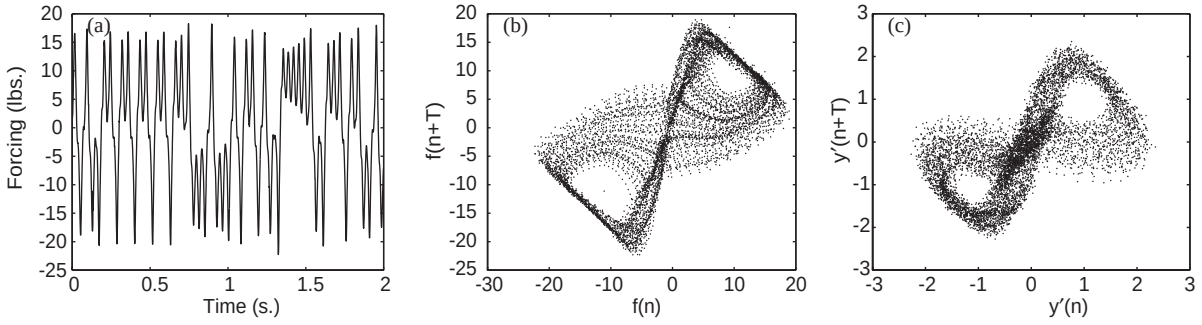

Fig. 6. (a) Driving signal, (b) driving attractor, (c) sample response attractor.

nearest-neighbors approach described in Section 2 was employed. This results suggests that the attractor reconstruction incorporate 2 delay copies of all the signals resulting in a 10-dimensional attractor. As mentioned previously the resulting object is most likely over-embedded as there are clearly some redundancies in the various time series. However, the cost associated with over-embedding is purely computational. Under-embedding can skew results, as mentioned in Section 2.

We reconstructed the attractors for the fully torqued scenario and determined that an appropriate shadowing time was 10 time steps. We found that most of the trajectories in one trial of this scenario stayed within 10 times the original separation between initial conditions; recall from the definition of CAAD in Section 3 that this is the criterion we defined for choosing shadowing time. We used 100 representative points on each attractor. We performed 10 trials for the fully torqued scenarios in order to obtain a probability distribution function for shadowing distances in the baseline attractor. We then performed the CAAD test between the attractors reconstructed from the fully torqued scenarios' data and those reconstructed from the other scenarios' data, again using 10 separate runs per experiment. We then compiled the probability distributions using the resampling procedure described in Section 4 and compared them. Fig. 7 shows the progression of mean CAAD value with damage. Confidence limits based on $\alpha = 0.05$ for the undamaged feature values are highlighted in gray. The CAAD metric is able to clearly resolve changes to the structures dynamics at the 26.69 kN axial load level. This represents a 50% improvement over previous efforts to detect the presence of damage in this same structure. The approach detailed in Ref. [37] using nonlinear prediction error was able to resolve change at the 17.79 kN level while an AR-based modeling approach featured in that same work could only resolve damage at 13.34 kN level. Furthermore, this particular metric is able to give information pertaining to the magnitude of the damage. Based on the confidence intervals damage may be classified in one of four statistically independent regions: $\sim 20.016 - 26.69$ kN, $\sim 8.006 - 17.79$ kN, $\sim 0.8896 - 4.448$ kN, and < 0.8896 kN. Not only can the presence of axial load loss be detected, but the magnitude can be discerned as well. The ease with which such a mapping is obtained is in part due to the monotonic increase in feature value. Trends in features that are non-monotonic can produce obvious difficulty in attempting to separate out the damage class to which the corresponding features belong.

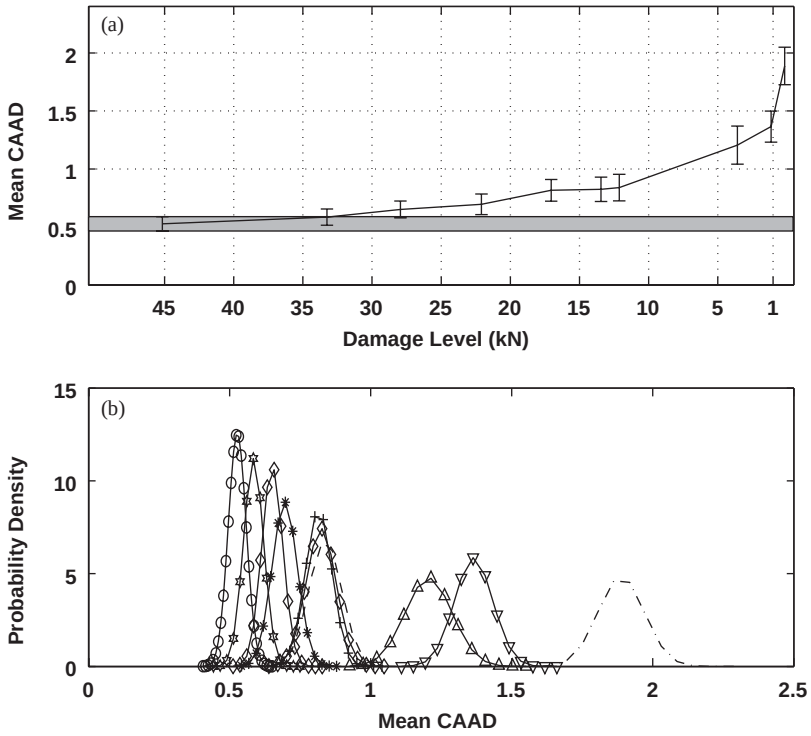

Fig. 7. (a) Mean CAAD trend with loss in axial load; (b) probability density.

As a comparison, we also computed the CAAD feature on each of the individual attractors, reconstructed from the time series collected at each of the five sensor locations. The five sets of feature values were then used to perform the discriminant coordinate analysis described in Section 4. Of the five resulting sets of discriminant coordinates, only two resulted in an ability to classify any of the damage scenarios (certain discriminant coordinates corresponding to low λ_r values are typically produced and are not able to separate many, if any, of the classes). The original progression of feature vectors, along with the two best sets of transformed feature values obtained from the discriminant analysis are shown in Fig. 8.

The top row of plots show the performance of the individual sensors while the bottom two show the combined discriminating power. Discriminant coordinates can at best distinguish the last four damage scenarios in any statistically meaningful sense. The failure of this type of analysis is not surprising considering the inability of the individual sensors to resolve the damage. For this particular experiment there is a clear advantage in combining time series data "up front" during the attractor reconstruction process. Such an approach incorporates all of the sensor information and therefore represents a more complete model of the system dynamics. While there is a computational cost associated with handling higher dimensional data, it can be minimized through implementation of a fast neighbor searching routine such as the kd-tree [38] and prioritized spatial search [39] algorithms used here.

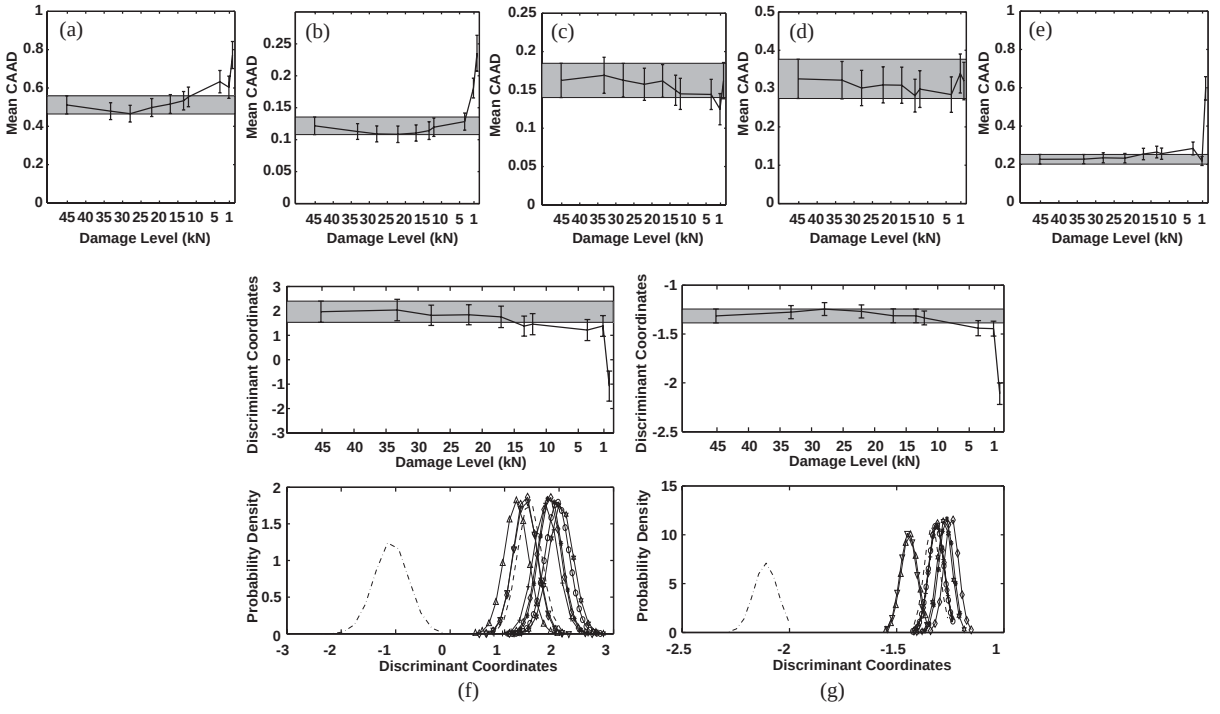

Fig. 8. (a–e) CAAD values for individual attractors and (f,g) two optimal discriminant coordinates as a function of damage.

7. Conclusions

We have described and implemented a new approach to multivariate structural health monitoring. The approach works by forming the phase space of the structure using delayed components from each of the measured structural response time series. Using this method we were able to detect statistically significant changes in the structural vibration at 26.698 kN axial load representing 60% of the fully clamped condition. Furthermore, it was shown that there exists a clear advantage to combining structural response data before the feature extraction. Linear combinations of the feature values collected from individual sensor's attractors were unable to resolve many of the damage cases.

Acknowledgements

The authors wish to acknowledge the American Society for Engineering Education (ASEE) and the National Research Council (NRC), for funding the first and second author's, respectively, post-doctoral fellowships and the National Science Foundation (NSF) for funding the third author's graduate research fellowship. We also thank Roshdy Barsoum of the Office of Naval Research (ONR) for funding this work and Keith Berube of the University of Maine for

providing the experimental test materials. This work was performed under a Naval Research Laboratory (NRL) Advanced Research Initiative 6.1.

References

[1] C.R. Farrar, T.A. Duffey, S.W. Doebling, D.A. Nix, A statistical pattern recognition paradigm for vibration-based structural health monitoring, in: *Structural Health Monitoring 2000*, 1999, pp. 764–773.
[2] D. Göge, M. Link, Results obtained by minimizing natural frequency and mode shape errors of a beam model, *Mechanical Systems and Signal Processing* 17 (1) (2003) 21–27.
[3] K. Bohle, C.-P. Fritzen, Results obtained by minimizing natural frequency and mac-value errors of a plate model, *Mechanical Systems and Signal Processing* 17 (1) (2003) 55–64.
[4] P. Verboven, E. Parloo, P. Guillaume, M. Van Overmeire, Autonomous structural health monitoring—part I: modal parameter estimation and tracking, *Mechanical Systems and Signal Processing* 16 (4) (2002) 637–657.
[5] S.W. Doebling, C.R. Farrar, M.B. Prime, A summary review of vibration-based identification methods, *Shock and Vibration Digest* 205 (5) (1998) 631–645.
[6] A.C. Okafor, K. Chandrashekhara, Y.P. Jiang, Location of impact in composite plates using waveform-based acoustic emission and gaussian cross-correlation techniques, in: *Proceedings of the SPIE*, Vol. 2718, 1996, pp. 291–302.
[7] H. Sohn, C.R. Farrar, N.F. Hunter, K. Worden, Structural health monitoring using statistical pattern recognition techniques, *Journal of Dynamic Systems, Measurement, and Control* 123 (2001) 706–711.
[8] J.B. Bodeux, J.C. Golinval, Modal identification and damage detection using the data-driven stochastic subspace and armav methods, *Mechanical Systems and Signal Processing* 17 (1) (2003) 83–89.
[9] C.M. Bishop, Novelty detection and neural network validation, in: *IEE Proceedings*: *Vision*, *Image*, *and Spectral Processing* 141 (4) (1994) 217–222.
[10] K. Worden, G. Manson, D. Allman, Experimental validation of a structural health monitoring methodology—part I: novelty detection on a laboratory structure, *Journal of Sound and Vibration* 259 (2) (2003) 323–343.
[11] D. Logan, J. Mathew, Using the correlation dimension for vibration fault diagnosis of rolling element bearings—i. Basic concepts, *Mechanical Systems and Signal Processing* 10 (3) (1996) 241–250.
[12] W.J. Wang, Z.T. Wu, J. Chen, Fault identification in rotating machinery using the correlation dimension and bispectra, *Nonlinear Dynamics* 25 (2001) 383–393.
[13] A.C. Okafor, A. Dutta, Structural damage detection in beams by wavelet transforms, *Smart Materials and Structures* 9 (6) (2000) 906–917.
[14] Z. Hou, M. Noori, R.St. Amand, Wavelet-based approach for structural damage detection, *Journal of Engineering Mechanics* 126 (7) (2000) 677–683.
[15] J.S. Owen, B.J. Eccles, B.S. Choo, M.A. Woodings, The application of auto-regressive time series modelling for the time–frequency analysis of civil engineering structures, *Engineering Structures* 23 (2001) 521–536.
[16] A. Nag, D.R. Mahapatra, S. Gopalakrishnan, Identification of delamination in composite beams using spectral estimation and a genetic algorithm, *Smart Materials and Structures* 11 (2003) 899–908.
[17] C.R. Farrar, S.W. Doebling, D.A. Nix, Damage identification with linear discriminant operators, in: *Proceedings of the 17th International Modal Analysis Conference*, 1999.
[18] J.M. Nichols, M.D. Todd, M.E. Seaver, L.N. Virgin, Use of chaotic excitation and attractor property analysis in structural health monitoring, *Physical Review E* 67 (2003) 016209.
[19] L. Moniz, L.M. Pecora, J.M. Nichols, M.D. Todd, J.R. Wait, Dynamical assessment of structural damage using the continuity statistic, *International Journal of Structural Health Monitoring* 3 (3) (2004) 199–212.
[20] J.M. Nichols, M.D. Todd, J.R. Wait, Using state space predictive modeling with chaotic interrogation in detecting preload loss in a frame structure experiment, *Smart Materials and Structures* 12 (4) (2003) 580–601.
[21] H. Whitney, Differentiable manifolds, *Annals of Mathematics* 37 (1936) 645.
[22] F. Takens, Detecting strange attractors in turbulence, in: D.A. Rand, L.-S. Young (Eds.), *Dynamical Systems and Turbulence*, Lecture Notes in Mathematics, Vol. 898, Springer, New York, 1981, pp. 366–381.

[23] T. Sauer, Time series prediction: forecasting the future and understanding the past, in: A.S. Weigend, N.A. Gershenfeld (Eds.), *SFI Studies in the Sciences of Complexity*, Vol. 15, Addison-Wesley, Reading, MA, 1993.
[24] A.M. Fraser, H.L. Swinney, Independent coordinates for strange attractors from mutual information, *Physical Review A* 33 (1986) 1134–1140.
[25] L. Pecora, S. Boccaletti, D.L. Valladares, L. Moniz, H.P. Geffert, T. Carroll, Multiple time series and attractor reconstructions, in: S. Boccaletti, T. Carroll, B. Gluckman, V. In, L. Kocarev, J. Kurths (Eds.), *Proceedings of the Seventh Experimental Chaos Conference*, American Institute of Physics, 2002.
[26] M.B. Kennel, R. Brown, H.D.I. Abarbanel, Determining embedding dimension for phase-space reconstruction using a geometrical construction, *Physical Review A* 45 (6) (1992) 3403–3411.
[27] L.J. Moniz, T.L. Carroll, L.M. Pecora, M.D. Todd, Assessment of damage in an 8-oscillator circuit using dynamical forcing, *Physical Review E* 68 (2003) 036215.
[28] L.J. Moniz, L.M. Pecora, T.L. Carroll, M.D. Todd, S. Trickey, J. Nichols, Vibration-based damage assessment using novel function statistics with multiple time series, in: S. Boccaletti, T. Carroll, B. Gluckman, L. Kocarev, V. In, J. Kurths (Eds.), *Proceedings of the Seventh Experimental Chaos Conference*, American Institute of Physics, 2002.
[29] S. Boccaletti, D.L. Valladares, L.M. Pecora, H.P. Geffert, T.L. Carroll, Reconstructing embedding spaces of coupled dynamical systems from multivariate data, *Physical Review E* 65 (2002).
[30] M. Kennel, R. Brown, H. Abarbanel, Determining embedding dimension for phase-space reconstruction using a geometrical construction, in: E. Ott, T. Sauer, J.A. Yorke (Eds.), *Coping with Chaos*, Wiley, New York, 1994.
[31] C.R. Farrar, S.W. Doebling, D.A. Nix, Vibration-based structural damage identification, *Philosophical Transactions of the Royal Society of London A* 359 (2001) 131–149.
[32] B.F. Manly, *Randomization and Monte Carlo Methods in Biology*, Chapman & Hall, London, 1991.
[33] G.A.F. Seber, *Multivariate Observations*, Wiley, New York, 1984.
[34] J.M. Nichols, M.D. Todd, M. Seaver, S.T. Trickey, L.M. Pecora, L. Moniz, Controlling system dimension: a class of real systems that obey the Kaplan–Yorke conjecture, *Proceedings of the National Academy of Sciences* 100 (26) (2003) 15299–15303.
[35] M.E. Davies, K.M. Campbell, Linear recursive filters and nonlinear dynamics, *Nonlinearity* 9 (1996) 487–499.
[36] M.D. Todd, G.A. Johnson, B.L. Althouse, A novel bragg grating sensor interrogation system utilizing a scanning filter, a mach-zehnder interferometer, and a 3×3 coupler, *Measurement Science and Technology* 12 (2001) 771–777.
[37] J.M. Nichols, L. Moniz, M.D. Todd, S.T. Trickey, M. Seaver, C.J. Nichols, L.N. Virgin, Detection of fastener preload loss in a hybrid composite-to-metal bolted joint, in: *Proceedings of the Fourth International Conference on Structural Health Monitoring*, 2003.
[38] J.L. Bentley, Multidimensional binary search trees in database applications, *IEEE Transactions on Software Engineering* SE-5 (4) (1979).
[39] G.R. Hjaltason, H. Samet, Ranking in spatial databases, in: M.J. Egenhofer, J.R. Herring (Eds.), *Proceedings of the Fourth Symposium on Spatial Databases*, Portland, ME, August 1995, Lecture Notes in Computer Science, Vol. 951, Springer, Berlin, 1995, pp. 83–95.