Use of chaotic excitation and attractor property analysis in structural health monitoring

J. M. Nichols,* M. D. Todd, and M. Seaver
U.S. Naval Research Laboratory, Code 5673, Washington, D.C. 20375

L. N. Virgin
Pratt School of Engineering, Duke University, North Carolina 27708-0300

(Received 14 May 2002; published 23 January 2003)

This work explores the utility of attractor-based approaches in the field of vibration-based structural health monitoring. The technique utilizes the unique properties of chaotic signals by driving the structure directly with the output of a chaotic oscillator. Using the Kaplan-Yorke conjecture, the Lyapunov exponents of the driving signal may be tuned to the dominant eigenvalues of the structure, thus controlling the dimension of the structural response. Data are collected at various stages of structural degradation and a simple nonlinear model, constructed from the undamaged data, is used to make predictions for the damaged response data. Prediction error is then introduced as a ''feature'' for classifying the magnitude of the damage. Results are presented for an experimental cantilevered beam instrumented with fiber-optic strain sensors.

DOI: 10.1103/PhysRevE.67.016209 PACS number(s): 05.45.Gg

I. INTRODUCTION

The field of structural health monitoring (SHM) is concerned with accurately monitoring the integrity of structures, both civil and military, in an effort to reduce ownership costs, improve operational lifetime, and most importantly, protect human life. Because the dominant mode of failure is often catastrophic, the associated economic and health costs are typically very high. One of the more common approaches to the problem is vibration-based SHM, a procedure whereby the structure in question is excited and the dynamical response observed for changes in ''features'' that are indicative of damage. The assumption here is that damage will manifest itself as a change in the dynamical properties of the structural response so that by quantifying these changes, the practitioner may detect, locate, or even diagnose the type of the damage. A vast majority of the literature considers features derived from a modal analysis of the structure, e.g., resonant frequencies, mode shapes, mode curvature, modal damping, flexibility, etc. Good summaries of previously proposed features, their application, and their effectiveness may be found in Refs. [1] and [2].

Recently, progress has been made by using chaotic excitation signals and attractor analysis to detect damage in structures [3,4]. This approach is fundamentally new in that the steady-state dynamics of the structure are considered, rather than the transient dynamic properties associated with a modal analysis. We will show in this work that a chaotic wave form may be tailored to yield a structural response, which is low dimensional, regardless of the number of degrees of freedom in the structure. The effect of damage is to alter the dynamics of the response, giving rise to different response attractors. Any algorithm that can accurately quantify these differences in the presence of uncertainty (such as due to nonstationary environmental fluctuations) becomes a valid method for detecting structural damage. A variety of attractor-based metrics exist for the purpose of making quantitative statements about the underlying dynamical process, e.g., correlation dimension [5], Lyapunov spectrum [6], false-nearest neighbors [7], and certain prediction schemes [8] are widely used.

The approach used here is a variation of the method used in Ref. [9] for examining nonstationarity in time-series data. The idea is to use data taken from the undamaged ''pristine'' structure to reconstruct reference, or base line, attractors. These data are to be compared to data collected from the structure as damage is incurred. By using a simple prediction scheme, points on the ''damaged'' attractors are forecast using the base line data as a model. Higher levels of damage will alter the structure's dynamic response, causing these base line models to lose their ability to make predictions. Prediction error then becomes a good feature for quantifying both the presence and the magnitude of structural degradation. The effectiveness of the approach is demonstrated on experimental data taken from a cantilevered beam where the level of damage is known, quantified by varying the clamping strength at the fixed end.

II. METHODOLOGY

A. Chaotic interrogation

Traditional vibration analysis involves exciting the structure with broadband random signals. The broadband nature of noise ensures a full modal response, ideal for frequency domain approaches to system identification or feature extraction. Chaotic signals also tend to possess broadband frequency spectra; however, unlike noise, chaos is deterministic. In fact many chaotic systems can be as low as three dimensional when described as a continuous time process. This determinism, coupled with the fractal nature of chaotic attractors, gives rise to a low-dimensional ''spread'' of trajectories that allow for attractor-based classification. Periodic orbits, for example, adhere to well-defined paths and will produce one-dimensional attractors, regardless of damage level. In addition, a chaotic system is defined by a positive

*Email address: pele@ccs.nrl.navy.mil

Lyapunov exponent implying extreme sensitivity to small changes in system parameters. The subtlety of damage induced changes to a structure further motivates this choice as the mechanism of excitation.

The difficulty in applying attractor-based methods to experimental systems is usually tied to the dimensionality of the process. High-dimensional ($D>4$) systems tend to produce poor results when an attractor-based analysis is used due to an inability to sufficiently populate the attractor. The proposed method takes advantage of the Kaplan-Yorke conjecture in order to produce low-dimensional response thus allowing for the standard tools of nonlinear time-series analysis to be successfully employed. The Kaplan-Yorke conjecture [10] relates the Lyapunov exponents (LEs) of a system to the dimension of that system via

$$D_L = K + \frac{\sum_{m=1}^{K} \lambda_m}{-\lambda_{K+1}}, \tag{1}$$

where K is the number of exponents that may be added before the sum becomes negative, the λ_m are the LEs, and D_L is the Lyapunov dimension. This relationship is important for two specific reasons. First, it implies that by adjusting the LEs the dimensionality of the system's response may be effectively controlled. Second, it states that *only the most weakly contracting* state space directions play a role in the dimensionality of the process. All negative exponents greater in magnitude than λ_{K+1} are effectively filtered out of the dynamics. The state-space directions associated with these exponents are contracting too quickly to be seen.

Consider a simple linear N-degree-of-freedom structure forced with the output of a separate dynamical process governed by the function **F**.

$$\dot{\mathbf{z}} = \mathbf{F}(\mathbf{z})$$

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{z}(t). \tag{2}$$

The coordinates **x** are N vectors that can be any dynamic measurable (strain, displacement, etc.), depending on how the constant coefficient matrix **A** is formed. This matrix contains all of the mass, stiffness, and damping properties present in the structure. Forcing for the system is provided by the input $\mathbf{Bz}(t)$ where the matrix **B** selects both the component of $\mathbf{z}(t)$ to be used for the forcing and the structural location(s) (at which coordinates **x**) where the forcing is to be applied. A one-degree-of-freedom version of this same system was studied in Ref. [11] by Pecora and Carroll where the forcing was taken to be the output of a chaotic oscillator. The goal of that work was to illustrate how controlling the LEs of a linear time-invariant filter could lead to a response with varying dimension via Kaplan-Yorke conjecture. Here we have extended the concept to a multiple-degree-of-freedom filter, represented here by the structure **A**.

Returning to Eq. (2), both the structure and the forcing will each contain their own sets of LEs so that the complete spectrum for the filtered chaotic signal $\mathbf{x}(t)$ is given by

$$\left.\begin{array}{ll} \lambda_u^C & : \quad u = 1, \ldots, d_1 \\ \lambda_v^L & : \quad v = 1, \ldots, d_2 \end{array}\right\} \Rightarrow \lambda_m^S,$$

$$\lambda_1^S > \lambda_2^S > \cdots > \lambda_{(d_1+d_2)}^S, \tag{3}$$

where the λ_u^C are the exponents associated with the d_1-dimensional forcing and the λ_v^L are the exponents of the d_2-dimensional filter. Regardless of the number of degrees of freedom associated with the filter, the fractal dimension of the entire system is controlled by Eq. (1). It should be noted that in order for these statements to be true the coupling between the forcing and the structure *must* be monodirectional. Any feedback from the structure which affects the dynamics of the forcing, could result in fundamentally different dynamics through bifurcation.

B. Tuned chaos

Damage to the system will result in changes to the eigenstructure of **A**. This in turn will alter the structure's Lyapunov spectrum, which, for a linear system, consists of the real parts of the eigenvalues of **A**. Changing the λ_v^L may thus alter the dimension of the filtered signal. Therefore, detecting changes in dimensionality becomes a valid candidate for detecting any type of damage that serves to alter the eigenvalues of a structure. Two criteria must be met, however, in order for this approach to work. First, the Lyapunov spectrum of the oscillator must overlap that of the structure. This ensures that changes to the LEs of the structure, i.e., by damage, will alter the dimension of the filtered signal. The *degree* of overlap, d_o, determines the extent to which the structure's dynamics are excited or, alternatively, the number of dimensions the structure is adding to the phase space. Second, the dominant exponent associated with the oscillator must be minimized for a given degree of overlap in order to maintain the lowest possible dimensionality. Considering Eq. (3), these criteria become

$$|\lambda_{d_1}^C| > |\lambda_{d_o}^L|$$

$$\sum_{m=1}^{d_o} \lambda_m^L > |\lambda_1^C| > \sum_{m=1}^{d_o-1} \lambda_m^L. \tag{4}$$

The interaction between chaotic and structural LEs for $d_o = 1$ is depicted in Fig. 1. If the above criteria are met, and the forcing is restricted to the output of a continuous time process ($\lambda_2^C = 0.0$), the dimension of the filtered chaotic signal will be

$$D_L = 1 + d_o + \frac{\lambda_1^C + \sum_{m=1}^{d_o-1} \lambda_m^L}{|\lambda_{d_o}^L|}, \quad d_o > 0. \tag{5}$$

The result of an appropriately tuned chaos and structure system is therefore a response that is low dimensional, and one whose dynamics will be fundamentally changed by the occurrence of damage to the structure. The focus is therefore

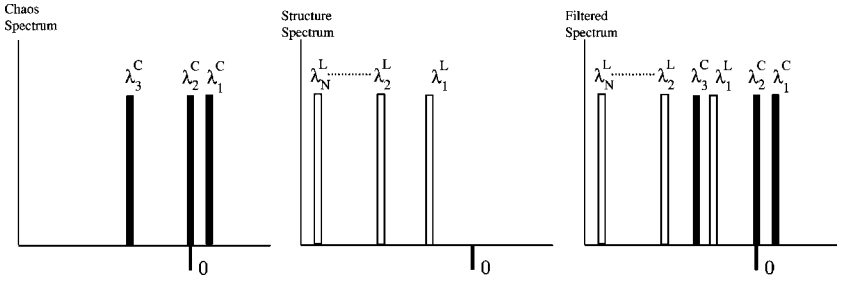

FIG. 1. Interaction of LEs for filtered dynamics: chaotic input Lyapunov spectrum (left); structural Lyapunov spectrum (middle); combined spectrum (right).

on finding a feature that can accurately quantify differences in the dynamical attractors produced by the system at various stages of structural degradation.

III. PREDICTION ERROR AS A FEATURE

In this paper we explore the use of prediction error as a feature in quantifying the level of structural degradation. Suppose that N_r independent time series (N_r independent experiments, or runs) are recorded. Recording N_r such responses allows for the inclusion of ambient variation in the set of data. Such variation is known to occur in practice and must be accounted for if damage-induced changes are to be distinguished from those due to environmental factors. We will represent this base line set of N_r time series as $X=\{x_1(n),x_2(n),\ldots,x_{N_r}(n)\}$, where each $x_i(n), i=1,2,\ldots,N_r$, is a vector of $n=1,\ldots,N$ discretely sampled values of undamaged structural response. Similarly, at some later time, when the structure is presumed to be damaged relative to the condition when the set X was collected, a new set of data $Y=\{y_1(n),y_2(n),\ldots,y_{N_r}(n)\}$ are collected in the same way (i.e., N_r independent runs with N discrete time samples in each time series). The object of the feature-based damage assessment is to extract some metric capable of discriminating between the base line (undamaged) set X and any future (damaged) dataset Y. Here, the base line data set X is used to empirically generate an attractor-based model of the pristine structure's dynamics. These models are then used to make predictions for any subsequent damaged dataset Y. The resulting *nonlinear cross-prediction error* becomes an excellent candidate for a feature. As the structural dynamics is altered, the expected value of the cross-prediction error will increase. The strength of the approach is that it directly addresses the question of interest: *"Has the dynamical response of the structure been altered due to damage?"*

The algorithm used here is adopted from Schreiber [9], where it was used to detect nonstationarity in time-series data. A simple attractor-based prediction scheme is used on the undamaged data to forecast the values of the damaged data some number of time steps s into the future. Using the familiar delay coordinate approach [12], each of the N_r base line time series in X may be used to construct a corresponding base line *attractor*

$$\mathbf{x}_i(n)=[x_i(n),x_i(n+T),\ldots,x_i(n+[M-1]T)],$$

$$i=1,\ldots,N_r \tag{6}$$

with embedding dimension M and time delay T, where we have adopted the boldface now to indicate the ith *attractor* rather than the ith time series in the dataset X. Similarly, any subsequent (damaged) dataset Y will yield N_r attractors, with each one denoted $\mathbf{y}_j(n)$. The method proceeds by selecting the ith base line attractor $\mathbf{x}_i(n)$ from X and the jth "test" attractor $\mathbf{y}_j(n)$ from Y. Given a randomly selected fiducial trajectory (point) at time index f on $\mathbf{y}_j(n)$, or $\mathbf{y}_j(f)$, the algorithm selects a corresponding set or "neighborhood" of all the points on the base line attractor $\mathbf{x}_i(n)$ that are within some radius ϵ of the randomly selected fiducial point. This set may be explicitly defined as

$$U_\epsilon^{\mathbf{x}_i}(\mathbf{y}_j(f))=\mathbf{x}_i(p):||\mathbf{x}_i(p)-\mathbf{y}_j(f)||<\epsilon. \tag{7}$$

The time index p associated with $\mathbf{x}_i(n)$ is completely arbitrary, and it does not necessarily have any correlation to the time index f, since this neighborhood is constructed purely from geometric considerations. The idea is to describe the evolution of $U_\epsilon^{\mathbf{x}_i}(\mathbf{y}_j(f))$ and to use this description as a predictor for how the data should evolve in time. The predicted value for $\mathbf{y}_j(f)$ at s time steps into the future, denoted $\hat{\mathbf{y}}_j(f+s)$, then becomes

$$\hat{\mathbf{y}}_j(f+s)=\frac{1}{|U_\epsilon^{\mathbf{x}_i}(\mathbf{y}_j(f))|}\sum_{\mathbf{x}_i(p)\in U_\epsilon^{\mathbf{x}_i}(\mathbf{y}_j(f))}\mathbf{x}_i(p+s), \tag{8}$$

where the quantity $|U_\epsilon^{\mathbf{x}_i}(\mathbf{y}_j(f))|$ simply denotes the number of points in the neighborhood. In this formulation, then, the predicted value is simply the average of predicted values for the neighborhood. In this sense the base line attractors are used as "look-up" tables that contain the various patterns present in the data. The working hypothesis is that these tables will lose their ability to serve as an accurate database as the dynamics is altered by damage. The prediction horizon s will depend on the rate at which the data are acquired and the specific application. For health monitoring purposes, assuming reasonably sampled data, $s=1$ will suffice. While more complicated prediction schemes exist, this is among the simplest models one can use to quantify the evolution of the dynamics. Since we only seek to distinguish one attractor from another, the quality of predictions is of diminished importance and the simplest, most computationally efficient scheme is considered optimal. Variations of this algorithm have been used for prediction and data cleansing [8,13].

Once the predictions have been made, the average prediction error between the ith base line attractor and the jth test attractor at one time step in the future may be defined by

$$\gamma_{i,j}=\frac{1}{N}\sum_{f=1}^{N}[\hat{\mathbf{y}}_j(f+1)-\mathbf{y}_j(f+1)]^2, \tag{9}$$

where the average is taken over all N points on the attractor. In practice, such a large summation is often time consuming, and a good estimate of the average may be obtained by evaluating Eq. (9) over some randomly selected subset of N. This error quantity is then normalized by the variance (denoted σ_i^2) of the ith base line attractor to yield the *normalized average cross-prediction error*

$$\hat{\gamma}_{i,j}=\frac{\gamma_{ij}}{\sigma_i^2}. \tag{10}$$

This computation is performed for each of the base line ($i=1\ldots N_r$) and damaged ($j=1,\ldots,N_r$) attractors for all $i<j$ such that the total feature set has $N_m=N_r(N_r-1)/2$ members. Excluding duplicate pairings, e.g., $\hat{\gamma}_{1,2}\hat{\gamma}_{2,1}$, was done so that each of the resulting N_m members of $\hat{\gamma}_{i,j}$ may be treated as independent random variables. Similar indices, e.g., $\hat{\gamma}_{1,1},\hat{\gamma}_{2,2}$ are also excluded from consideration. The variation the method seeks to capture is occurring from run to run so that adding similar pairs serves no purpose other than to potentially bias the data. Finally, the *auto prediction error* is computed for the base line attractor by replacing $\mathbf{y}_j(f)$ with $\mathbf{x}_i(f)$ in Eq. (7). The resulting set $\hat{\gamma}_{i,j}^A$ (where A denotes ''auto'') gives some idea of the prediction error one would expect to find in the instance when the dynamics are not changing. This entire process is then repeated for any subsequent data set Y obtained from the structure over time. If the practitioner collects N_d datasets over time, then each $\hat{\gamma}_{i,j}$ dataset, now denoted $\hat{\gamma}_{i,j}^k, k=1,\ldots,N_d$, describes the error generated by using the base line data to predict how subsequent datasets should evolve. It is expected that as damage-induced dynamics of the structure change, this error will grow to reflect the change.

IV. STATISTICAL CONSIDERATIONS

Although the prediction error features presented above are computed in a deterministic manner, there exist several sources of variability which add elements of stochastic behavior to any real data. Measurement noise, unaccounted variables, environmental fluctuation (nonstationarity), and other sources cause the data to populate *distributions* rather than reflect exact repeatability. Thus, the damage diagnosis problem is transformed into distinguishing among various distributions of data regarding their proper classification as coming from a damaged structure or not. The practitioner seeks to make a confidence-based, quantitative engineering judgment in this capacity so that proper corrective action may be taken.

A useful procedure for comparing datasets is the analysis of variance (ANOVA). In an ANOVA procedure, the parameters of interest are the *treatments*, which refers to the quantity (or quantities) assumed to be changing between datasets, and the *responses*, which are the measured or computed data to be analyzed. In this work, the driving influence of the difference among datasets is assumed to be due to damage progress in the structure, and thus only a single-factor treatment is considered (1D ANOVA). Different amounts of damage are known as treatment *levels*. If the data population mean of the kth treatment is denoted μ_k, the following null hypothesis may be established:

$$H_0:\mu_k=\mu_{k+1}, \tag{11}$$

where $k=1,\ldots,N_d$, and N_d is the number of different damage levels (treatments) being considered. The implied complementary alternative hypothesis is thus that in at least one pairwise comparison, the means differ. Acceptance of H_0 indicates that no distinguishable difference among data sets may be inferred, and thus no damage to the structure has occurred. Specific details regarding how a 1D ANOVA procedure is completed may be found in any basic statistical methods text, e.g., Ref. [14].

Traditional ANOVA considers single inferences for each μ_j by establishing confidence intervals. If the parent data populations are assumed to be normally distributed, it may be shown that the $100(1-\alpha)\%$ confidence interval about the difference between two means μ_j and μ_k may be given by

$$\mu_j-\mu_k=(\overline{X_j}-\overline{X_k})\pm t_{\alpha/2}\sqrt{E_{\mathrm{MS}}\left(\frac{1}{n_j}+\frac{1}{n_k}\right)}, \tag{12}$$

where $\overline{X_j}$ and $\overline{X_k}$ are the means of the two measured datasets, E_{MS} is the global mean-squared error as obtained from the ANOVA procedure, n_j and n_k are the number of data observations in each set, and $t_{\alpha/2}$ is the student-t test statistic at confidence level $\alpha/2$.

The interval constructed in Eq. (12) is deficient in one key way: each pairwise inference applies individually, and the method cannot be used to draw a family of inferences among several datasets. The deficiency may be alleviated by either constructing wider intervals, meaning less precise estimates, or reducing the confidence level. One popular method for properly estimating pairwise differences is the Bonferroni method. The method requires that the number of pairwise comparisons to be made be specified *a priori*. The Bonferroni interval estimates are still given by Eq. (12), but the student-t test statistic is modified from $t_{\alpha/2}\rightarrow t_{\alpha/2q}$, where q is the number of pairwise estimates to be considered. For N_d treatments (damage levels), there are $C_2^{N_d}=N_d!/2!/(N_d-2)!$ total pairs that may be considered. One advantage of the Bonferroni method, however, is that not all these pairs need to be considered. The final decision whether to accept or reject H_0 depends on whether the intervals constructed by the modified Eq. (12) contain zero or not; intervals that contain zero indicate an acceptance of H_0, and intervals that do not contain zero indicate rejection of H_0, i.e., the two means are significantly different.

Two important assumptions are inherent in using the ANOVA procedure as detailed above. First, it must be known

in some sense how many levels of data (i.e., damage levels) are contained within all the datasets. This may be thought of as *supervised learning*, where a controlled damage study is being performed such that known discrete levels of damage (even if they are not quantifiable as such) are observed. Second, the confidence intervals constructed by the modified Eq. (12) rely upon the assumption that the parent data distributions from which the measured data are sampled are normal (Gaussian). This second assumption is usually acceptable, provided that the data are not significantly skewed or multimodal. Even if the parent distributions are skewed or multimodal, recourse may be made to the central limit theorem by applying some sort of subsampling, resampling, or bootstrapping technique [15]. The general procedure in these techniques is to take large numbers of randomly sampled subsets of the collected data, take the subset mean, and use the collection of resampled means as the ''new'' dataset. This procedure was implemented on the data taken in this work, and it will be shown later that strongly Gaussian distributions of data resulted.

The first assumption poses a more difficult challenge. In a practical situation, true supervised learning is quite rare, as observations of data from many ''known'' damage levels are usually unavailable. Without this information, a proper quantitative ANOVA is impossible. One useful procedure in this instance is the use of *statistical process control*, sometimes known as *novelty* or *outlier* detection [16–18]. Here, data are collected from some base line state of the system, confidence limits are determined by standard mean-based hypothesis testing, and any future data are compared to these confidence limits. Data that fall within the confidence limits are assumed to come from the same population as the base line data, while data outside the limits, called ''outliers,'' are used as an indicator of ''other'' data. As damage occurs in the structure, it is reasonable to postulate, then, that the number of outliers would increase. This procedure is illustrated in Fig. 2. Base line data are obtained, and a certain confidence is assigned to the resulting distribution, as in panel (a). The raw data are shown in panel (b) with the upper and lower confidence limits drawn in. Some of the data fall outside the limits, since it is never possible to obtain all the real parent population data. With proper determination of the confidence limits from usual hypothesis testing, the outlier fraction should equal the balance of the confidence limit from 100%; in other words, for a confidence limit of 95%, the base line data should have about 5% of the data appearing as an outlier. Then, as new data are obtained, they are compared to the base line confidence limits, and outliers are counted. The dataset in Fig. 2(c), for example, shows a significant increase in the number of outliers observed. The outlier count could be trended over time to obtain a semiquantitative measure of degradation.

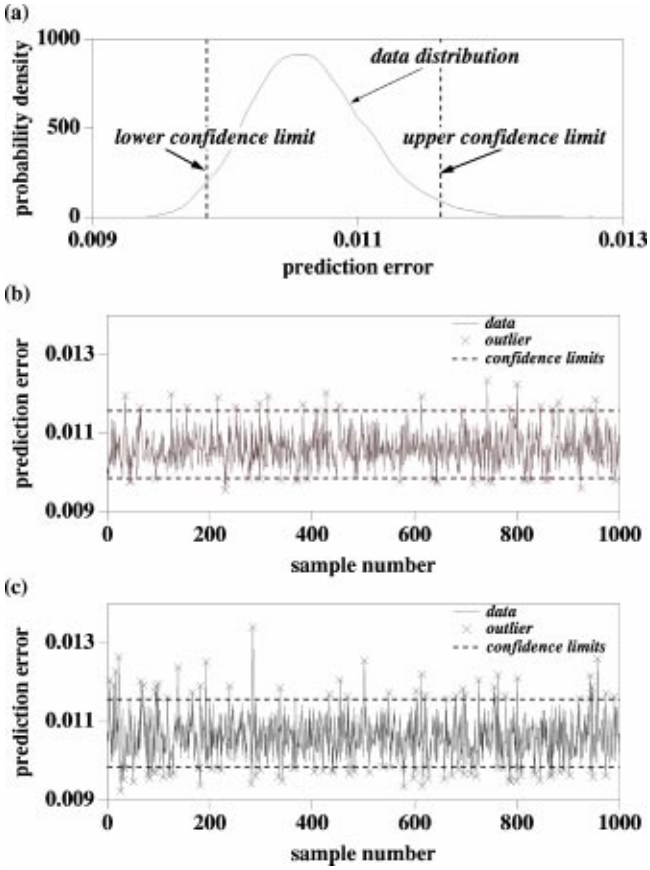

FIG. 2. Statistical process control: (a) probability density estimation and typical confidence limits for base line data; (b) data reflecting the distribution in (a); (c) new data showing an increased outlier count possibly indicative of damage.

V. EXPERIMENT SETUP

The system under study is a cantilevered aluminum beam, clamped at the fixed end by four clamping bolts threaded into an aluminum clamping stage. The beam itself had a length of 5.000×10^{-1} m, a width of 5.000×10^{-2} m, and a thickness of $t=3.175\times10^{-3}$ m. Springs were placed between the clamping stage and the clamping bolts so that the clamping strength could be varied. This provided a controlled mechanism by which the structure may be ''damaged.'' Seven fiber-optic strain sensors (based on fiber Bragg gratings) were evenly spaced along the central axis of the beam [19] and an accelerometer was attached to the forcing mechanism in order to record the excitation. The entire setup is shown in Fig. 3. The shaker is a M.B. Dynamics ''Modal 50'' attached to the base by a shaft fixed in place by set screws. It was determined experimentally that the first two LEs of the beam were

$$\lambda_1 \approx -7,$$

$$\lambda_2 \approx -11. \tag{13}$$

These estimates were obtained by exciting the structure with white noise, measuring the response, and implementing a version of the eigensystem realization algorithm (ERA) [20] in order to estimate the structure's state matrix **A**. The eigenvalues of this matrix were extracted and the real part was retained as the structure's Lyapunov exponents. Due to the limits of the excitation mechanism and the uncertainty associated with the ERA algorithm, the quantities (13) should be

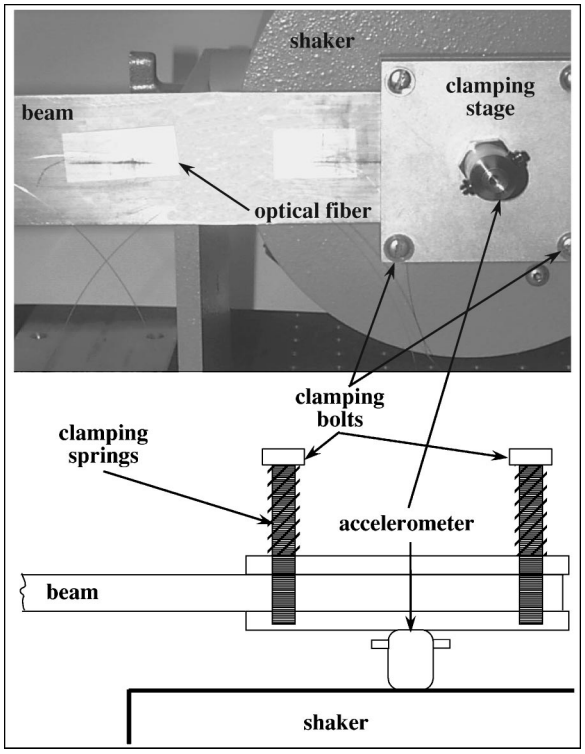

FIG. 3. Experimental setup.

regarded as rough approximations. Different forms of excitation and different algorithms may produce varying results. This is a consequence of some well-known difficulties in obtaining reliable estimates of negative LEs [21]. However, since the method only requires approximate values in order to be effective, this poses no limitation.

Excitation was chosen as the first state variable, z_1 of the chaotic Lorenz oscillator,

$$\epsilon \dot{z}_1 = 16(z_2 - z_1),$$

$$\epsilon \dot{z}_2 = 40z_1 - z_2 - z_1 z_3 .$$

$$\epsilon \dot{z}_3 = -4z_3 + z_1 z_2 , \tag{14}$$

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{z}, \tag{15}$$

where the tuning parameter ϵ is used to speed up or slow down the oscillator, depending on the LEs of the structure. The matrix **B** of Eq. (2) incorporates $\dot{z}_1, z_1$ due to the fact that the excitation is applied as base motion (as opposed to a point force). One obvious difference between the model (2) and the experiment is the number of coordinates required to describe the structure. Numerically such problems are discretized to some finite, usually small, number of degrees of freedom, while the actual beam contains an infinite number of degrees of freedom. In a practical sense this does not pose a problem as the time scales (or alternatively, LEs) become very small (very negative) for only a few state-space coordinates of the structure. The Kaplan-Yorke conjecture is essentially providing a means by which the higher-order structural coordinates (dimensions) are filtered out of the dynamics.

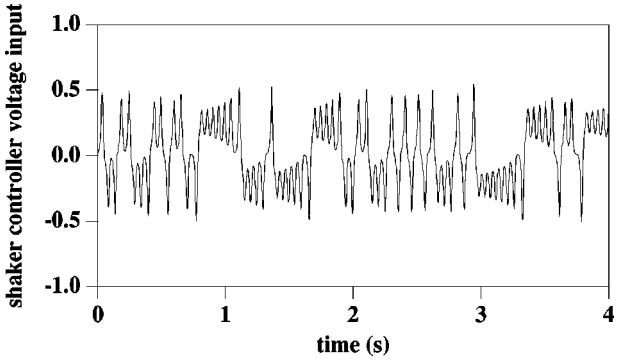

FIG. 4. Driving Lorenz wave form.

For this structure, two different forcing scenarios were considered. For the first, the ϵ were each set to 0.8, giving a complete spectrum of $\lambda_m^1 = (1.10, 0.00, -17.90)$ so that only one LE of the structure was overlapped by those of the driving signal ($d_o = 1$). The dimension of the response in this case is theorized to be $D_L \approx 2.2$. In the second case the degree of overlap stayed the same, but the oscillator was sped up by requiring that each of the $\epsilon = 1.8$ yielding a spectrum of $\lambda_m^2 = (2.47, 0.00, -40.27)$ and a dimension of $D_L \approx 2.4$. For each of these cases, the driving signals were generated using a fourth-order Runge-Kutta algorithm. LabVIEW data acquisition software was then used to convert the data files to the voltage that was output to the shaker controller. A sample of the driving wave form is shown in Fig. 4. Damage was simulated on the beam by loosening the compression on the springs holding the clamp. Under a fully clamped (undamaged) condition, all four springs were compressed to 1.18 cm, and three damage levels were produced by relaxing the springs to 1.40 cm (damage level 1), 1.95 cm (damage level 2), and 2.50 cm (damage level 3). Displacements were precision controlled with a micrometer. As the springs were lengthened, the elastic force imparted on the clamp decreased, simulating the relaxation of a bolted connection due to fastener degradation.

VI. RESULTS

This experiment yielded $N_r = 12$ runs for each $N_d = 4$ (one undamaged and three damaged) scenarios with each run consisting of $N = 45\,000$ points, sampled at 2 kHz. Final values for the average prediction error were obtained by evaluating Eq. (9) over 5 000 randomly selected trajectories *f*. The sets of prediction error therefore consist of $N_m = 66$ values which were then resampled with the means of 20-element random subsets to generate a 10 000-sample data set. Probability density functions of the resampled data using a kernel density estimation technique [22] with a Gaussian kernel are shown in Fig. 5 for both the slow (a) and fast (b) forcing cases. It is clear from the figure that increased levels of damage lead to a different set of dynamics and hence a higher prediction error. The distributions at all damage levels in both excitation time scales appear strongly Gaussian. To the naked eye, all

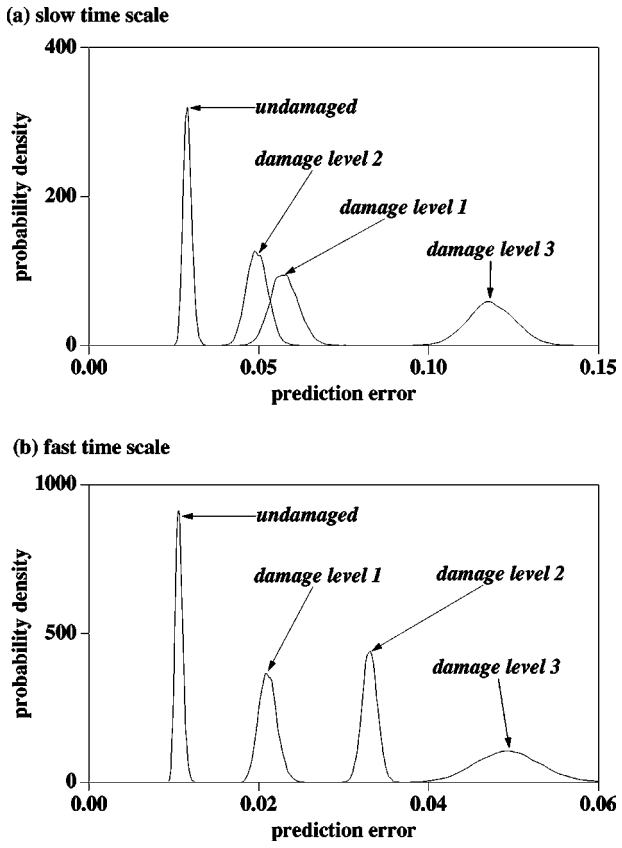

FIG. 5. Probability density estimations for data at all damage levels for (a) the slow time scale and (b) the fast time scale.

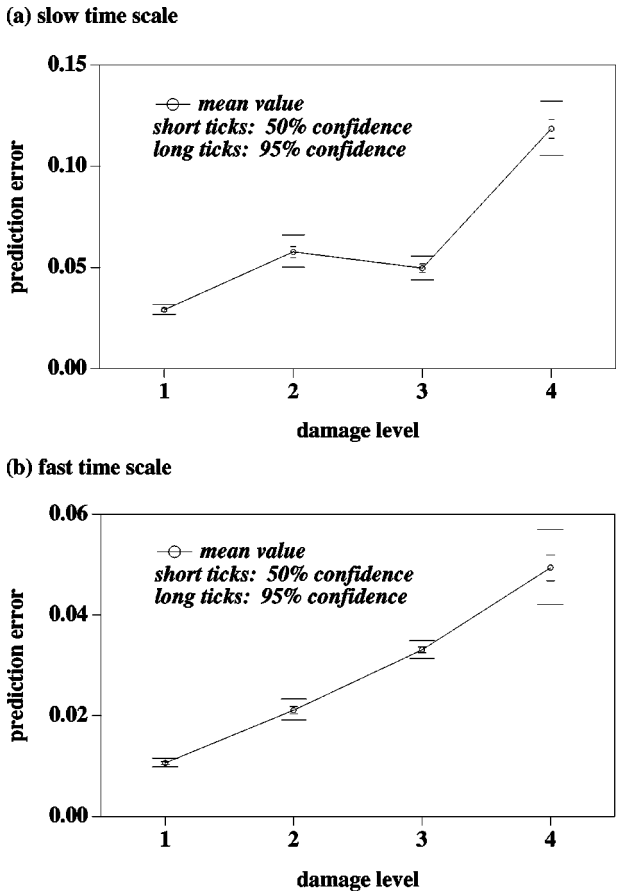

FIG. 6. Mean normalized cross-prediction error values at each damage level for (a) the slow time scale and (b) the fast time scale. Both 50% and 95% confidence limits are shown as indicated.

distributions appear distinguishable, with the possible exception of damage levels 1 and 2 at the slow excitation time scale. It is interesting to note that the prediction error is larger for this case. The reason for this concerns the time scales of the two different forcing processes. Because the first case involves the structure being driven at slower velocities, for a given number of samples, fewer oscillations take place leaving a more sparsely or less populated attractor. For the faster process, the N points will cover more oscillations and give a more complete geometric portrait of the attractor, resulting in a lower prediction error.

In order to quantify the differences between the distributions, 1D ANOVA was performed considering all six possible pairwise differences ($q=C_2^6=6$). In all six cases, the intervals constructed by the modified Eq. (12) did not contain zero, meaning that all damage levels are distinct from each other. This result was true for both excitation time scales, although the interval at the slow time scale when comparing damage levels 1 and 2 very nearly contained zero. Outlier analysis, while not as quantitative, revealed 100% outliers between all ''damaged'' datasets and the corresponding undamaged, base line dataset. This implies that even a simple online monitoring of outliers would likely detect changes to the base line very quickly. Figure 6 shows the progression with damage level of the distribution mean prediction error values for both time scales. Confidence limits are placed on the data at both 50% and 95% levels. These confidence limits are based on the appropriate quantiles for each distribution; this implies, in the 95% confidence case, for example, that 95% of the data in that distribution fall between the confidence limits. Regardless of time scale, it appears that at a 50% confidence level, all damage levels are distinguishable. At 95% confidence, the second and third damage levels have overlapping confidence bands for the slower excitation. Moreover, the trend is not monotonic, rendering damage magnitude classification more difficult.

VII. CONCLUSIONS

A framework for vibration-based SHM using attractor-based methods has been presented and demonstrated effectively in an experimental context. By using chaos to interrogate the structure, the dimension of the response is kept low, allowing for attractor-based classification. The random excitation commonly favored for vibration-based SHM, on the other hand, will always produce infinite (at least in a practical sense) dimension attractors, eliminating this particular approach from consideration. Nonlinear cross-prediction error is an effective feature for quantifying the level of degradation to a given structure. The method essentially quantifies the likelihood that a particular empirical model, taken from the base line data, describes the damaged data. A poor model is an indicator that the data have begun to deviate from those

produced by ''healthy'' dynamics, hence damage is identified. An effort has been made to incorporate the effects of ambient variation into the procedure by including multiple datasets in the analysis. Any method which is to be used in practice must take into account the distributions of features rather than analyzing one single realization. The nonlinear cross-prediction error was shown to be robust under variation, as both 1D ANOVA and outlier analysis showed rapid distinction between the damaged and undamaged datasets. In fact, damage magnitude classification also was possible, especially at the faster time scale. It is expected that the slower time scale could be equally as effective if the attractor had been populated more thoroughly; the lower dimension at this time scale would provide computational advantages and robustness as well. Finally, because this method relies on steady-state dynamics at relatively slow time scales (when compared to structural resonances), energy input to the system is minimal, in contrast to many of the sustained resonant driving methods used in common structural health monitoring practice.

ACKNOWLEDGMENTS

The authors acknowledge the National Research Council for financial support to J.M.N. and acknowledge the U.S. Naval Research Laboratory for funding this work under a 6.1 Advanced Research Initiative.

[1] S.W. Doebling, C.R. Farrar, and M.B. Prime, Shock Vib. Dig. **25**, 631 (1998).
[2] *Structural Health Monitoring: Demands and Challenges*, edited by F. K. Chang (CRC Press, Boca Raton, 2001).
[3] M.D. Todd, J.M. Nichols, and L.N. Virgin, Smart Mater. Struct. **10**, 1000 (2001).
[4] J.M. Nichols, M.D. Todd, S.T. Trickey, and L.N. Virgin, Meccanica (to be published).
[5] C. Craig, R.D. Neilson, and J. Penman, J. Sound Vib. **231**, 1 (2000).
[6] R. Brown, P. Bryant, and H.D.I. Abarbanel, Phys. Rev. A **43**, 2787 (1991).
[7] H.D.I. Abarbanel and M. Kennel, Phys. Rev. E **47**, 3057 (1993).
[8] H. Kantz and T. Schreiber, *Nonlinear Time Series Analysis* (Cambridge University Press, Cambridge, 1999).
[9] T. Schreiber, Phys. Rev. Lett. **78**, 843 (1997).
[10] J.L. Kaplan and J.A. Yorke, in *Functional Difference Equations and Approximations of Fixed Points*, edited by H.-O. Peitgen and H.-O. Walther, of Lecture Notes in Mathematics Vol. 730 (Springer-Verlag, Berlin, 1979).
[11] L.M. Pecora and T.L. Carroll, Chaos **6**, 432 (1996).
[12] F. Takens, in *Dynamical Systems and Turbulence*, edited by D. Rand and L.-S. Young, Vol. 898 of Lecture Notes in Mathematics, (Springer-Verlag, New York, 1981), pp. 366–81.
[13] J.M. Nichols and J.D. Nichols, Math. Biosci. **171**, 21 (2001).
[14] W.L. Quirin, *Probability and Statistics* (Harper & Row, New York, 1978).
[15] B. Efron and R.J. Tibshirani, *An Introduction to the Bootstrap* (Chapman and Hall, New York, 1993).
[16] K. Worden, J. Sound Vib. **201**, 85 (1997).
[17] M.L. Fugate, H. Sohn, and C.R. Farrar, Mech. Syst. Signal Process. **15**, 707 (2001).
[18] M.D. Todd, S.T. Trickey, J.M. Nichols, and M. Seaver, in *Proceeding of the 3rd World Conference on Structural Health Monitoring*, edited by F.K. Chang (Technomic Inc., Lancaster, PA, 2000).
[19] M.D. Todd, G.A. Johnson, and B.L. Althouse, Meas. Sci. Technol. **12**, 771 (2001).
[20] J.-N. Juang, *Applied System Identification* (Prentice-Hall, Englewood Cliffs, NJ, 1994).
[21] P. Bryant, *Proceedings of the 1st Experimental Chaos Conference* (World Scientific, Singapore, 1991), pp. 11–23.
[22] B.W. Silverman, *Density Estimation for Statistics and Data Analysis* (Chapman and Hall, London, 1986).