Sensitivity vector fields in time-delay coordinate embeddings: Theory and experiment

A. R. Sloboda* and B. I. Epureanu†
Department of Mechanical Engineering, University of Michigan, Ann Arbor, Michigan, USA
(Received 8 July 2012; revised manuscript received 7 November 2012; published 8 February 2013)

Identifying changes in the parameters of a dynamical system can be vital in many diagnostic and sensing applications. Sensitivity vector fields (SVFs) are one way of identifying such parametric variations by quantifying their effects on the morphology of a dynamical system's attractor. In many cases, SVFs are a more effective means of identification than commonly employed modal methods. Previously, it has only been possible to construct SVFs for a given dynamical system when a full set of state variables is available. This severely restricts SVF applicability because it may be cost prohibitive, or even impossible, to measure the entire state in high-dimensional systems. Thus, the focus of this paper is constructing SVFs with only partial knowledge of the state by using time-delay coordinate embeddings. Local models are employed in which the embedded states of a neighborhood are weighted in a way referred to as embedded point cloud averaging. Application of the presented methodology to both simulated and experimental time series demonstrates its utility and reliability.

DOI: 10.1103/PhysRevE.87.022903 PACS number(s): 05.45.Gg, 05.45.Tp

I. INTRODUCTION

Change in the dynamic response of a system indicates that its parameters have changed in some way. Thus, by monitoring a system's dynamic response it is often possible to identify and characterize any parametric variations the system undergoes. Knowing about parametric variations can be valuable in diagnostic or sensing applications, particularly when identifying incipient changes is paramount, as is in cases of damage detection.

Vibration-based modal methods [1] have been the traditional means of identifying parametric variations via the dynamic response. These methods rely on relating measured changes in modal parameters, such as resonant frequencies or mode shapes, to underlying variations in the parameters. Frequently employed in mechanical systems, modal methods are the basis of both dynamic atomic force microscopy (dAFM) and microcantilever based sensing. For these systems, accurate measurements of frequency shifts of vibrating microcantilevers are the basis of successful devices [2–4]. Continued efforts to improve these and related technologies attest to the general effectiveness of modal methods.

Under certain circumstances, however, modal methods can be less effective. For example, when a system is significantly nonlinear, when it has a low quality factor, or when multiple simultaneous parametric variations need to be distinguished from one another, other methods may be preferable. A specific alternative championed here is based on sensitivity vector fields (SVFs) [5–7], which quantify how dynamical system attractors deform under parametric variations. Because a SVF consists of a field of vectors distributed in state space, it can be successful even when a frequency-shift method would fail. To date, SVFs have been used in conjunction with a variety of dynamical systems [8] and have proven to be an effective method of identifying parametric variations. Carefully designed nonlinear feedback can further adjust a system's sensitivity to the parametric variations of interest.

Although SVFs have been successfully employed in conjunction with data from both simulations and experiments [9], previous SVF research has been limited to systems where a full set of state variables is accessible. In all but the simplest real systems, this kind of complete access is impossible. For some systems, measuring certain dynamical variables may be difficult. For others, it may be cost prohibitive to measure the full state because the system has a large dimension. For these kinds of systems, it is typical to have records of only a few state variables in the form of one or more time series; a time series being a sequence of scalar data points measured at successive times. So long as the system has a low-dimensional attractor, further analysis is feasible. Previous research concerned with methods of reconstructing dynamical system attractors when only time series data are available has shown reconstruction is possible using an embedding. The most common form is a time-delay coordinate embedding, in which the embedded vectors are composed of individual time series observations separated by a fixed delay time. The mathematical basis of time-delay coordinate embedding has been rigorously established [10–13] and several texts [14,15] illustrate methods for performing such reconstructions.

The focus of this work is creating SVFs within the reconstructed state space established by a time-delay coordinate embedding of time series data. Specifically, for an initial condition in the reconstructed space, it is necessary to determine two future embedded states: one for the system retaining the nominal set of parameters, and a second for the system having some parametric variation(s). The difference between these two future states is defined as an embedded sensitivity vector (eSV). A collection of eSVs across the entire attractor is an embedded sensitivity vector field (eSVF). Thus, the fundamental problem involved in constructing eSVFs is one of prediction. Making predictions using embedded nonlinear time series is an established research area, and we draw on the previous body of work [16–19]. Typically, we choose initial conditions that are on the nominal attractor so that only the future state of the varied system needs to be

*asloboda@umich.edu
†epureanu@umich.edu

predicted. Making a good prediction then requires gathering neighborhoods of states for the varied system surrounding a given initial condition belonging to the nominal system and constructing local models to fit these neighborhoods. Once an eSV is estimated, it must be validated to ensure its accuracy. This is accomplished by requiring accurate predictions of near-neighbor surrogates, checking local modeling coefficients, and ensuring the correct linearity and proportionality of eSVs that are generated by known parametric variations of different magnitudes. A method of quantifying the error in eSV predictions based on local modeling is outlined for cases where the equations of motion of the system are known. Applications of the methodology to various simulated time series, including series for a Duffing oscillator and for the Lorenz attractor demonstrate the effectiveness of the technique. These systems are also used to explore how additive noise influences the results. Further application to an experimental time series generated by a Chua's oscillator demonstrates how the methodology can be applied to a real, physical system.

II. EMBEDDED SENSITIVITY VECTOR FIELDS

A. Sensitivity vector field definition

A SVF is a collection of vectors that capture how a dynamical system attractor deforms as a result of a given parametric variation. The individual SVs making up a SVF are constructed by sampling trajectories of nominal and varied systems as they diverge over time. Two trajectories that are initially coincident in state space but differ by some parametric variation will evolve differently. By sampling these trajectories a time ΔT after their coincidence, a SV is generated that connects the sampled point on the nominal attractor to the sampled point on the varied attractor. This vector quantifies the divergence of the two trajectories. It depends on the underlying dynamical system, the nature of the parametric variation, and the evolution time ΔT. For short ΔT, the SV will be proportional to the parametric variation. Thus, by obtaining SVFs corresponding to several known parametric variations, a set of basis vectors can be constructed against which further unknown parametric variations can be compared, allowing them to be identified.

More mathematically, consider a dynamical system described by the flow $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, p, t)$, where $\mathbf{x}$ is the state vector and p is a system parameter that can vary. Using a Taylor series to develop a variational equation about the nominal trajectory $\mathbf{x}(t) = \mathbf{x}_0(t)$ and nominal parameter value $p = p_0$ and retaining only the linear terms results in

$$\delta\dot{\mathbf{x}}(t) = \mathbf{A}(t)\delta\mathbf{x}(t) + \mathbf{b}(t)\delta p, \tag{1}$$

where

$$\mathbf{A}(t) = \left.\frac{\partial \mathbf{f}}{\partial \mathbf{x}}(t)\right|_{\substack{\mathbf{x}=\mathbf{x}_0 \\ p=p_0}}, \tag{2}$$

and

$$\mathbf{b}(t) = \left.\frac{\partial \mathbf{f}}{\partial p}(t)\right|_{\substack{\mathbf{x}=\mathbf{x}_0 \\ p=p_0}}. \tag{3}$$

Here, the state variation (from the nominal trajectory) is represented by $\delta\mathbf{x}$ and the parameter variation (from the nominal parameter value) is represented by δp. If Eq. (1) is integrated over the evolution time ΔT, the result is an equivalent map

$$\delta\mathbf{x}(t + \Delta T) = \Phi(t + \Delta T, t)\delta\mathbf{x}(t) + \mathbf{q}(t + \Delta T)\delta p, \tag{4}$$

where the state transition matrix for the dynamical system Φ depends only on $\mathbf{A}(t)$. Trajectory divergence, expressed in the state variation $\delta\mathbf{x}(t + \Delta T)$ that develops over the evolution time ΔT, can be interpreted as the sensitivity vector $\mathbf{q}(t + \Delta T)\delta p$ when $\delta\mathbf{x}(t) = \mathbf{0}$ at the initial time (that is, if the trajectories are truly initially coincident at time t). In general, $\mathbf{q}$ will depend on both $\mathbf{A}(t)$ and $\mathbf{b}(t)$. The linearity of Eq. (4) highlights the fact that a specific parametric variation will elicit a proportional change in the generated SVF. This proportionality is what enables quantification of parametric variations through the comparison of newly generated SVFs to known, reference SVFs.

B. Embedded sensitivity vector field definition

Consider again the flow $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, p, t)$, where the state is given by $\mathbf{x} = [x_1\, x_2 \cdots x_n]^T$, with n being the dimension of the system. In its simplest form, performing a time-delay coordinate embedding of the flow involves sampling a single state component and then constructing time-delay coordinate vectors having the form

$$\mathbf{s}(t) = [x_*(t)\, x_*(t - \tau) \cdots x_*(t - m\tau)]^T, \tag{5}$$

where $*$ indicates a single state of the system (from the total of n states), τ is the delay time, and m is the number of time-delay coordinates. It has been proven [13] that $m \geqslant 2d_A$ is a sufficient condition to reconstruct an attractor of dimension d_A, although the necessary dimension may be less. In this time-delay coordinate embedding space, the definition of an eSVF is similar to that of a SVF in ordinary state space, but what is measured is the time-delay coordinate variation $\delta\mathbf{s}$ rather than the state variation $\delta\mathbf{x}$. Time-delay coordinate variation is variation solely in the component of the state x_* selected for constructing the time-delay coordinate vectors. Figure 1 shows the flow and its variation for a two dimensional embedding in x_*, illustrating graphically how the time-delay coordinate

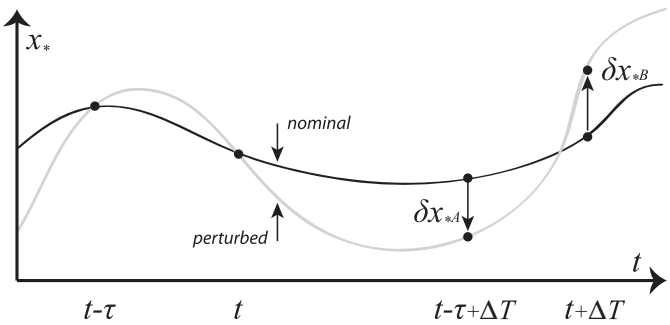

FIG. 1. Consider a flow with two state variables x_1 and x_2 embedded two dimensionally in x_*. This figure illustrates the required initial coincidence of the nominal and varied trajectories and the variations δx_* at later times which comprise an eSV.

variation has the general form

$$\delta\mathbf{s}(t+\Delta T) = [\delta x_*(t+\Delta T) \\ \delta x_*(t-\tau+\Delta T)\cdots\delta x_*(t-m\tau+\Delta T)]^T. \tag{6}$$

In the embedding space, time-delay coordinate vectors are guaranteed to have a map of the form $\mathbf{s}(t+\Delta T) = \mathbf{H}(\mathbf{s}(t))$, where $\mathbf{H}$ depends on ΔT (which we assume is some multiple of τ) and p. This means we can define an eSV $\mathbf{q}_*\delta p$ as

$$\delta\mathbf{s}(t+\Delta T) = \frac{\partial\mathbf{H}}{\partial\mathbf{s}}(t+\Delta T,t)\bigg|_{\substack{\mathbf{s}=\mathbf{s}_0\\ p=p_0}} \delta s(t) + \mathbf{q}_*(t+\Delta T)\delta p, \tag{7}$$

where

$$\mathbf{q}_*(t+\Delta T) = \frac{\partial\mathbf{H}}{\partial p}(t+\Delta T,t)\bigg|_{\substack{\mathbf{s}=\mathbf{s}_0\\ p=p_0}}. \tag{8}$$

Note that one can numerically compute the eSVs using Eq. (4) when the state space and the time-delay embedding space have the same dimension *and* the equations for the original flow $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x},p,t)$ are available. This idea is further developed in Sec. III E in the context of eSV validation.

In practical cases, one is unlikely to know the function $\mathbf{H}$ and must rely on some form of modeling to generate eSVs. The natural choice is to use local modeling, where neighborhoods of states are collected that serve as analogues for the initial conditions of interest. The future images of these states, used in conjunction with the local modeling, allow unknown trajectories to be predicted.

III. METHODOLOGY

There are four steps to generate an eSV for a given initial condition drawn from the nominal system's data set; namely:

(1) Gather a neighborhood of nearby states from the varied system's data set.

(2) Adjust the size of the neighborhood to optimize its accuracy in making predictions.

(3) Use local modeling in conjunction with the resized neighborhood to make a prediction of the varied system's trajectory beginning at the given initial condition.

(4) Take the difference between the varied trajectory (predicted by local modeling) and the nominal trajectory (known) in order to construct an eSV.

This process of eSV construction is illustrated in Fig. 2.

Once eSVs have been constructed, they are validated. There are three steps in this validation; namely:

(1) Rank eSV neighborhoods based on the distance between the initial condition and the nearest state of the varied system along with the error in predicting that state's future trajectory.

(2) Examine the coefficients resulting from the local modeling used to generate a given eSV and discard those eSVs whose coefficients are poor. For example, discard all neighborhoods when any coefficients are larger than 1.

(3) If the eSVs are being used to build up basis vectors for future testing (and hence have known parametric variations), check their colinearity and proportionality in conjunction with the eSVs of other calibration sets.

This process of eSV validation is also illustrated in Fig. 2.

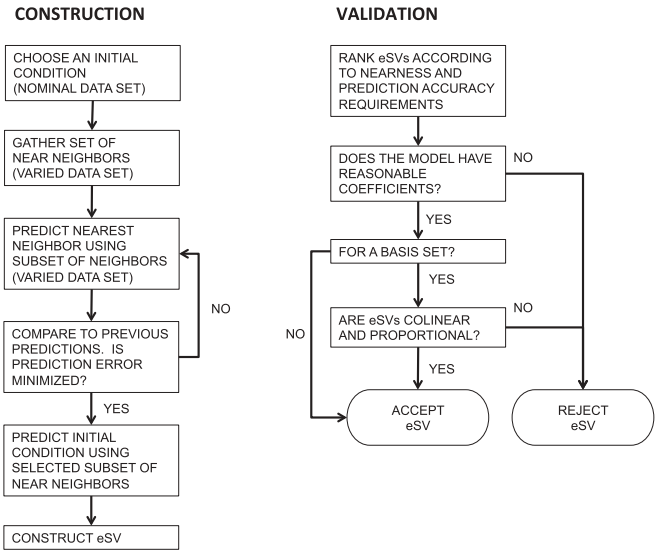

FIG. 2. Process of generating and validating eSVs that are suitable for analyzing parametric variations.

Below some important details regarding both the eSV construction process and the eSV validation process are provided.

A. Local modeling using embedded point cloud averaging

The question of how to effectively predict nonlinear time series using local models has been studied by several authors [16,18,19]. Global approaches for nonlinear time series prediction have also been proposed, including methods based on radial basis functions [17]. However, global methods are more difficult to implement and less likely to give accurate results. Thus, our focus is constructing eSVs using local models exclusively, specifically using embedded point cloud averaging (ePCA) (similar to the original PCA [7]).

For an initial condition on the nominal attractor, ePCA involves solving an underdetermined problem involving the neighboring embedded states from the varied data set given by

$$\begin{bmatrix} 1 \\ \mathbf{x}^t_{k,\text{center}} \end{bmatrix}_{m+1\times 1} = \mathbf{X}^t_k\alpha, \tag{9}$$

with

$$\alpha = [\alpha_1\cdots\alpha_n]^T, \tag{10}$$

and

$$\mathbf{X}^t_k = \begin{bmatrix} 1 & 1 & \cdots & 1 \\ \mathbf{x}^t_1 & \mathbf{x}^t_2 & \cdots & \mathbf{x}^t_n \end{bmatrix}_{m+1\times n}. \tag{11}$$

Here $\mathbf{x}^t_{k,\text{center}}$ is an initial condition of interest on the nominal attractor (i.e., the center of neighborhood k at time t), and α is a vector of unknowns. $\mathbf{X}^t_k$ holds the n embedded states from the varied data set in neighborhood k, and m is the embedding dimension. This is the original PCA formulation [7]. For ePCA, we enforce the additional condition that

$$\sum_{i=1}^{n}\alpha_i x^t_{i,r}x^t_{i,s} = x^t_{\text{center},r}x^t_{\text{center},s}\ \forall\, r,s = 1,\ldots,m, \tag{12}$$

where $x_{i,s}^{t}$ refers to the sth component of the measured embedded state $\mathbf{x}_i^t$. Solving Eqs. (9) and (12) together in the least-norm sense determines the weights each measured embedded state should be given so that their weighted average matches the initial condition as closely as possible while also minimizing quadratic cross terms.

Minimizing the terms in Eq. (12) minimizes the second-order terms in the Taylor series expansion (7). This is important and can improve eSV accuracy. A (future) prediction at time $t + \Delta T$ can then be made using the previously determined α by calculating

$$\mathbf{x}_{k,\text{center}}^{t+\Delta T} = \mathbf{X}_k^{t+\Delta T}\,\alpha, \tag{13}$$

with

$$\mathbf{X}_k^{t+\Delta T} = \left[\mathbf{x}_1^{t+\Delta T}\ \mathbf{x}_2^{t+\Delta T}\cdots\mathbf{x}_n^{t+\Delta T}\right]. \tag{14}$$

Here $\mathbf{X}_k^{t+\Delta T}$ holds (future) states that correspond to the embedded states in the initial neighborhood k at time $t + \Delta T$. The eSV $\mathbf{q}$ is given by the difference between future points on the nominal and varied trajectories, i.e., the known nominal system trajectory known and the varied system trajectory predicted by the model. That is,

$$\mathbf{q} = \delta\mathbf{x} = \mathbf{y}_{k,\text{center}}^{t+\Delta T} - \mathbf{x}_{k,\text{center}}^{t+\Delta T}, \tag{15}$$

where $\mathbf{y}_{k,\text{center}}^{t+\Delta T}$ is the predicted embedded future state of the varied system and $\mathbf{x}_{k,\text{center}}^{t+\Delta T}$ is the known embedded future state of the nominal system.

More typical linear (first-order) approximations found in the literature [19] have the form

$$\left[\mathbf{X}_k^{t+\Delta T}\right]^T = \left[\mathbf{X}_k^{t}\right]^T\beta + \epsilon, \tag{16}$$

with

$$\beta = \begin{bmatrix} \beta_{1,1} & \cdots & \beta_{1,m} \\ \vdots & \vdots & \vdots \\ \beta_{m,1} & \cdots & \beta_{m,n} \end{bmatrix}. \tag{17}$$

Here, β is, in general, a collection of vectors of unknowns, and ϵ represents error terms. Solving this overdetermined system for β in the least-squares sense, and then using the β vector in conjunction with an initial state allows computation of a predicted (future) state at $t + \Delta T$. Using this general framework, more accurate predictions may be possible by using higher-order models [20], but these quickly require a large number of neighborhood points, as the required coefficients increase as $m(s+m)!/(s!m!)$, where m is the embedding dimension and s is the order of the model. Related methods with additional complexities such as filtering [21] or methods which avoid a least-squares solution [22] have also been proposed.

The primary advantage of ePCA over these other methods is that the α vector provides a metric that can be used to judge the relative contribution of each neighborhood point in constructing eSVs. Different magnitudes in the α entries can indicate a poor neighborhood that will generate an inaccurate eSV.

B. Neighborhood sizing

Once a neighborhood of nearby states from the varied system's data set has been gathered, the closest of these states is sequestered as a surrogate for the initial condition. Predictions are made for this surrogate using local modeling in conjunction with neighborhoods of various sizes to determine what neighborhood size to use when making a prediction from the actual initial condition. Because we know the future (embedded) state of the surrogate, it is easy to determine what neighborhood size results in the most accurate prediction of its future trajectory.

This optimal neighborhood sizing capability is incorporated into the algorithm in two different ways. The first is suitable for cases where there are a large number of measured embedded states. Using the points closest to the initial condition, one can set up a series of increasing radii around the initial condition and consider the accuracy of predictions based on the neighborhood of embedded states contained within each radius. The radius with the lowest prediction error is selected for making predictions from the actual initial condition. The second way, appropriate when data are sparser, eschews the need for radii and simply drops the embedded state furthest from the initial condition in making consecutive predictions. One again chooses the neighborhood that results in the minimum prediction error.

C. Validation

To avoid regions of the state space giving poor eSVs, one can calculate the distance between the surrogate embedded states and the initial condition of interest. Then, one can rank eSV generating neighborhoods according to the product of this distance and the prediction error. Thus, close surrogate embedded states with good predictions are highly ranked whereas neighborhoods with surrogates that are distant or poorly predicted are not. This helps to avoid neighborhoods with especially bad predictions or disparate state space structure.

When the state space is highly structured, as is the case for many chaotic attractors, an initial condition that lies on the attractor of the nominal data set likely will not lie on the attractor of the varied data set. This occurs, for instance, when the parametric variation is large enough to significantly deform the varied attractor with respect to the nominal attractor. Simply using the neighborhood of embedded states in the varied set that is closest to the initial condition could lead to an incorrect eSV in this case because the selected neighborhood does not provide reasonable analogues for the initial condition.

In general, the ranking methodology places the worst performing neighborhoods near the bottom of the ranking. However, there is still considerable variability in eSVF accuracy for highly ranked neighborhoods so it is not sufficient in and of itself. The method's value lies in the fact that it quickly provides a list of reasonably good neighborhoods, and thus can be used to save time when checking for proportionality (discussed next) by providing a basis for limiting the number of neighborhoods that need to be checked in this subsequent stage.

The second validation of constructed eSVs is ensuring reasonable neighborhood coefficients. The α coefficients for

successful eSVs must all be less than 1 for the eSV to meet the requirements of this test. This ensures that all of the states in the neighborhood are contributing relatively equally in predicting the future trajectory.

The third validation of constructed eSVs is a proportionality check (in conjunction with a colinearity check). It only applies to eSVs for which we know the corresponding parametric variations, such as those generated to serve as part of a basis set. The proportionality check requires that eSVs meet the linearity requirement of Eq. (7). That is, any change in a given parameter perturbation δp should elicit a proportional change in the magnitude of the eSV. Checking proportionality requires having at least two different parameter perturbations for each parameter of interest available when generating SVFs belonging to the basis sets. The proportionality check can be used to reduce the error in the acceptable eSVs to very low levels for noiseless data.

D. Combination

Once a collection of validated eSVs is obtained, one can combine them into an eSVF in the form of a column vector of individual eSVs. If one is concerned with identifying only a single parameter, simply examining a test vector's proportionality to this eSVF can determine the unknown variation. However, to identify linearly independent changes in the attractor that occur under multiple simultaneous parameter variations, proper orthogonal decomposition (POD) is required [5–9]. When different eSVFs corresponding to several different parameter variations are collected as column vectors, they can be used to form a matrix $\mathbf{Q}$. The correlation matrix $\mathbf{C} = \mathbf{Q}\mathbf{Q}^T$ can then be constructed. The dominant eigenvalues of the correlation matrix indicate the number of linearly independent parameter changes that can be identified and the corresponding eigenvectors provide the basis for doing so.

E. Measuring prediction error

To determine the accuracy of eSVs, it is necessary to have some means of determining eSV error. When predictions are made for a single time series, it is typical to report the root mean square (rms) prediction error for out-of-sample data. This means that while most of the time series is used as a training set to determine the parameters to be used in the local modeling, the remainder of the series is sequestered to serve as a test set. This provides an independent repository of states for which predictions can be made using local models, but for which one also knows future states exactly. To determine the modeling error, one simply compares model predictions to known future states.

However, in constructing eSVs, we are making predictions for initial conditions on the nominal attractor using data from the varied attractor. This means that we have no test set to sequester. Thus, for cases when time series data alone is available, it seems that estimates of an eSVF's efficacy must be based solely on its ability to identify parameter variations accurately. However, when generating time series data from flow equations, there is an advantage in that the variables of the dynamical state that are not being embedded (and that would normally be unknown) can be retained. Having knowledge of these normally hidden state variables allows eSVs to be checked analytically.

An example serves to illustrate the methodology. Consider a two-dimensional dynamical system with a state $\mathbf{x} = [x_1\ x_2]^T$ that is also embedded in two dimensions using the x_1 variable. One can rewrite Eq. (1) as

$$\begin{bmatrix} \delta\dot{x}_1 \\ \delta\dot{x}_2 \end{bmatrix} = \begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \frac{\partial f_1}{\partial x_2} \\ \frac{\partial f_2}{\partial x_1} & \frac{\partial f_2}{\partial x_2} \end{bmatrix} \begin{bmatrix} \delta x_1 \\ \delta x_2 \end{bmatrix} + \begin{bmatrix} \frac{\partial f_1}{\partial p} \\ \frac{\partial f_2}{\partial p} \end{bmatrix} \delta p. \tag{18}$$

The equivalent of Eq. (4), the map involving the state transition matrix, is then given by

$$\begin{bmatrix} \delta x_1(t+\tau) \\ \delta x_2(t+\tau) \end{bmatrix} = \begin{bmatrix} \phi_{x_1x_1} & \phi_{x_1x_2} \\ \phi_{x_2x_1} & \phi_{x_2x_2} \end{bmatrix} \begin{bmatrix} \delta x_1(t) \\ \delta x_2(t) \end{bmatrix} + \begin{bmatrix} b_{x_1} \\ b_{x_2} \end{bmatrix} \delta p. \tag{19}$$

Here, the time t corresponds to the initial condition in the embedded state, and the time $t+\tau$ corresponds to a future state, which occurs at a time lag τ later. Since both the nominal and the varied trajectories are required to have the same initial embedded state, one has a boundary value problem, as $\delta x_1(t) = \delta x_1(t+\tau) = 0$, in Eq. (19). One obtains two equations for variations in the variable x_2 as

$$\delta x_2(t) = \frac{-b_{x_1}\delta p}{\phi_{x_1x_2}} \tag{20}$$

and

$$\delta x_2(t+\tau) = \phi_{x_2x_2}\delta x_2(t) + b_{x_2}\delta p. \tag{21}$$

Integrating Eq. (18) between t and $t+\tau$ in conjunction with the original flow $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, p, t)$, allows one to obtain the unknown entries of Eqs. (20) and (21). One integrates once with the initial conditions $[\delta x_1\ \delta x_2]^T = [0\ 1]^T$ and $\delta p = 0$ to obtain $\delta x_1(\tau) = \phi_{x_1x_2}$ and $\delta x_2(\tau) = \phi_{x_2x_2}$, and again with initial conditions $[\delta x_1\ \delta x_2]^T = [0\ 0]^T$ and $\delta p = 1$ to obtain $\delta x_1(\tau) = b_{x_1}$ and $\delta x_2(\tau) = b_{x_2}$. Using these to solve Eqs. (20) and (21), one finds the δx_2 values for $\delta x_1 = 0$ at t and $t+\tau$. Integrating from either of these states with the known δx_2 values and a specified δp allows the calculation of δx_{1A} and δx_{1B} at $t+\Delta T$ and $t+\tau+\Delta T$. These quantities, δx_1 and δx_2, are the two components of the eSV. Figure 1 illustrates this when $x_* = x_1$. Note that this technique for validating eSV predictions only works when the dimensions of the original state space and the embedded state space are equal and is susceptible to singularities if the parametric variation is large or the evolution time is long.

IV. TIME SERIES SOURCES

A. Simulated time series sources

Two familiar dynamic systems are used to generate simulated time series for analysis. The first system is the well-known Lorenz attractor whose equations are given by

$$\dot{x} = \sigma(y - x), \quad \dot{y} = -xz + rx - y, \quad \dot{z} = xy - bz, \tag{22}$$

where x, y, and z are state variables, and σ, r, and b are system parameters. For these parameters, the standard nominal values $\sigma = 10$, $r = 28$, and $b = 8/3$ were chosen. The system was sampled every $\Delta t = 0.05$. Collected data consisted of 10^5 samples for each data set with 10^4 samples on the nominal attractor serving as initial conditions.

The second system for demonstrating the approach is a Duffing oscillator whose representation is

$$\begin{aligned} \dot{x} &= y, \\ \dot{y} &= -x^3 + x - by + A\sin z, \\ \dot{z} &= \omega, \end{aligned} \tag{23}$$

where x, y, and z are state variables, and A, b, and ω are system parameters. The nominal parameter values were $A = 0.4$, $b = 0.25$, and $\omega = 1$. The system was sampled 100 times each driving period, and 10^5 samples for each data set were collected with 10^4 samples on the nominal attractor serving as initial conditions.

The data sets generated computationally were also corrupted by having various levels of Gaussian white noise added to them. This additive noise has the form

$$x_n = s_n + \eta_n, \tag{24}$$

where x_n is the noisy signal, s_n is the clean signal, and η_n is the additive white noise included at each time instant indicated by subscript n. The relative noise level can then be expressed as

$$\frac{\langle \eta^2 \rangle}{\langle (s - \langle s \rangle)^2 \rangle}, \tag{25}$$

with $\langle\,\rangle$ representing the mean.

B. Experimental time series sources

The source of experimental time series is a variant on the well-known Chua's circuit where the inductor is replaced by an operational-amplifier realization of a gyrator [23]. Figure 3 shows the layout of the circuit. The relevant component values are given in Table I. These are identical to those in Ref. [23] except that the capacitors are adjusted to 10 and 100 nF, respectively. Within the gyrator is a variable resistance R_L. One of the resistors within the realization of Chua's diode is also replaced with a variable resistor R_N. Thus, the circuit has three parameters (R, R_L, and R_N) which can be varied (to construct eSVFs) and also identified.

TABLE I. Chua's circuit properties.

Component	Value
R1	220 Ω
R2	220 Ω
R3	2.2 kΩ
R4	22 kΩ
R5	22 kΩ
R7	100 Ω
R8	1 kΩ
R9	1 kΩ
R	2 kΩ pot
R_L	2 kΩ pot
R_N	5 kΩ pot
C	100 nF
C1	10 nF
C1	100 nF

During each test, the system was sampled for approximately 10^6 data points using a digital to analog converter having a 16-bit resolution over a ± 10 V range. A set of 10^4 samples was drawn from the nominal system to serve as the set of initial conditions.

C. Initial data processing

Before any eSVFs can be constructed in an embedded state space, appropriate dimension [24,25] and time delay [26,27] with which to construct time-delay coordinate vectors need to be determined. These two quantities depend on both the dynamical system being studied and the way in which the system is sampled. Both factors can be determined using available software, such as the TISEAN package [28] or other specialized programs, such as those for determining dimension based on false strands [25]. The TISEAN package is also useful

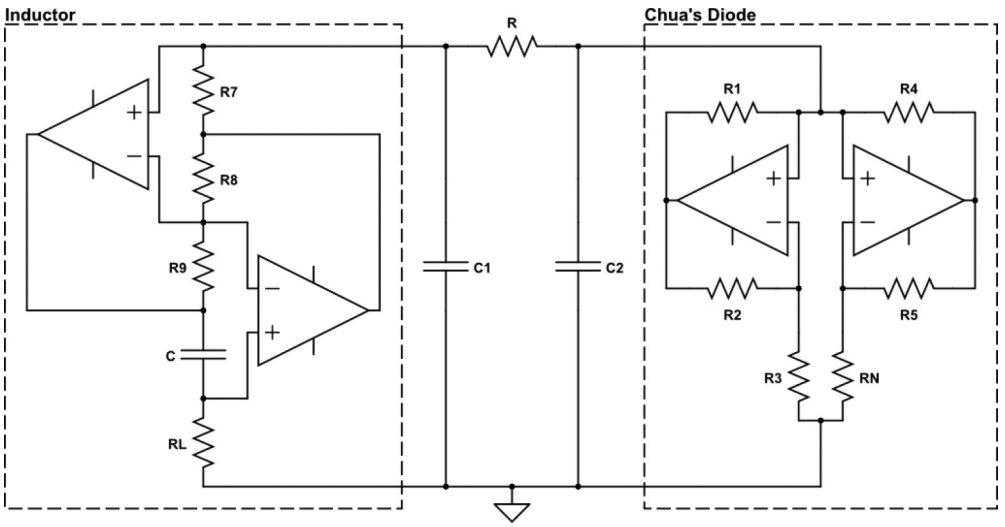

FIG. 3. Layout of the Chua circuit used to acquire the experimental data.

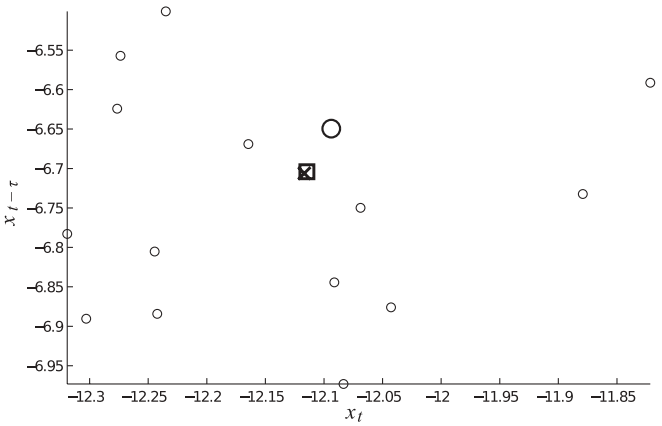

FIG. 4. Example of a good SV prediction for Lorenz attractor. The large ○ represents the nominal trajectory; small circles indicate the varied data in the neighborhood nearby. The × represents a prediction of the varied trajectory via local modeling and the □ via semi-analytical methods. Coincidence of the × and □ indicate a good prediction.

for determining whether or not given dynamics are chaotic via a calculation of the largest Lyapunov exponent (a system is chaotic when this exponent is positive).

V. RESULTS AND DISCUSSION

A. Simulated time series

For the simulated time series, using the data processing techniques in Sec. IV C, we found that an embedding dimension $m = 3$ appears sufficient for the Lorenz system, and the best time lag based on mutual information is approximately $\tau = 3$ ($\Delta t = 0.15$). For the Duffing oscillator, we choose either to embed x and y in an $m = 2$ embedding, treating the phase information (z) as known, or to separately fully embed in a $m = 3$ dimension embedding. For the Duffing system, the appropriate time delay is approximately $\tau = 10$.

The choice of time delay is a critical factor in constructing eSVs. If the time delay selected in the case of the Lorenz attractor is not optimal according to the first minimum of the mutual information, the error in eSV magnitudes increases significantly; in some cases doubling versus the error at the optimal time delay.

With embedding and time delay determined, one can apply the proposed methods to construct embedded eSVs. An example of such a construction is shown in Fig. 4. The states indicated by small circles represent varied trajectory points while the nominal trajectory is indicated by the large ○. The semi-analytical prediction of the varied trajectory based on the state transition matrix formulation is indicated by the □, while the prediction of the local model is indicated by ×. A coincident □ and × indicate an accurate prediction. A 1% variation in σ is being captured by this eSV for the Lorenz attractor.

The Lorenz system affords one the opportunity to compare the locations of initial conditions which are appropriate for constructing eSVs (as determined by the process outlined in Sec. III) with the locations of initial conditions generating low-error eSVs (as determined via the analytical methodology outlined in Sec. III E). Figure 5 illustrates one way to make this comparison. On the left, validated eSV initial conditions for 1% and 2% variations in σ are indicated by the darker points; the remainder of the initial conditions are shown by the lighter points. We consider the eSV validated if its proportionality is within 5% of the expectation. On the right, the same sets of initial conditions are shown, but in this case they are darker if the error in constructing a vector corresponding to 1% variation in σ is less than 10% and lighter otherwise. Here, we are making predictions for $\Delta T = 9$. The similarity between the two plots confirms that the methodology we have outlined for validating eSVs does select the eSVs with the lowest error. Moreover, these plots show that the distribution of low-error or successfully validated eSVs is nonuniform across the attractor, often being concentrated in bands or clusters based on the deformation. Careful tightening of the proportionality requirement in the validation step will reduce the number of acceptable initial conditions but those that remain will generally produce eSVs with less error. Similar results hold for the Duffing oscillator, and when one only embeds two of the state vectors and retains the phase as

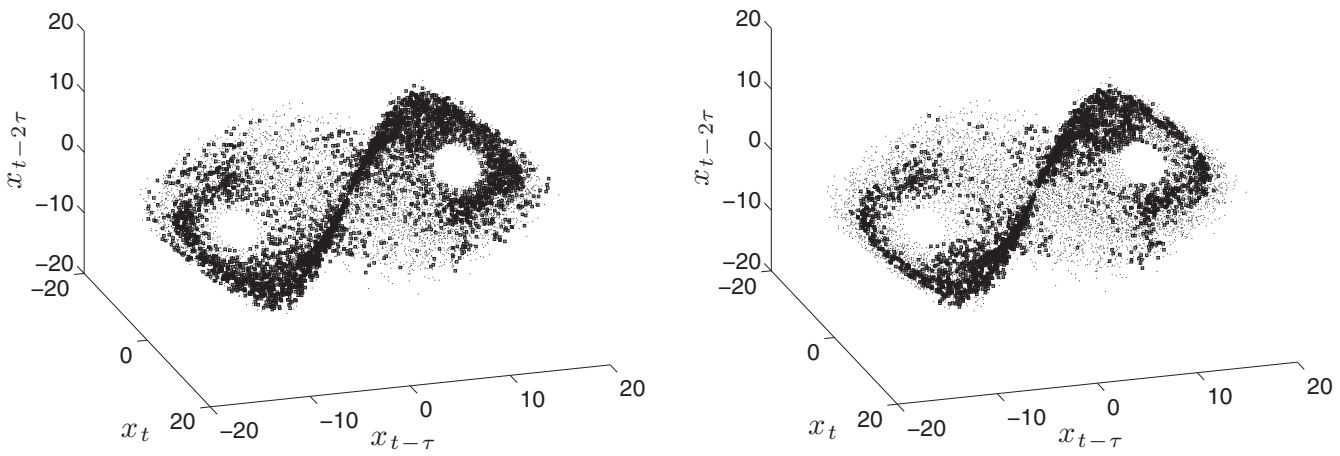

FIG. 5. Initial conditions used in generating eSVs for the Lorenz data; darker points indicate valid eSVs (left). The same set of initial conditions; darker points indicate eSVs having less than 10% error (right).

TABLE II. Simulated time series parameter reconstruction.

Case	Parameter	Variation (%)	eSVF prediction (%)
1	σ	1.50	1.47
	r	0.00	0.01
2	σ	0.00	0.04
	r	1.50	1.65
3	σ	1.50	1.46
	r	1.50	1.62
4	σ	3.00	2.97
	r	3.00	3.18

TABLE III. Simulated time series parameter reconstruction with noise.

Noise level	Parameter	Variation	eSVF prediction
0	σ	−0.15	−0.149
	b	−0.20	−0.191
	r	0.00	0.001
0.0001	σ	−0.15	−0.144
	b	−0.20	−0.194
	r	0.00	−0.001
0.001	σ	−0.15	−0.132
	b	−0.20	−0.127
	r	0.00	−0.008

additional information, the average eSV magnitude prediction error is generally lower.

With both surrogate ranking and proportionality validation one can consistently generate eSVFs with low error in low-noise environments. These accurate eSVFs allow for excellent parameter reconstructions. For example, in the case of the Lorenz attractor after having collected initial eSVFs for changes in σ, r, and b of 1% and 2% one can perform the reconstructions shown in Table II. These required the calibration eSVs to have a proportionality within 5% of the expectation based on the specified parameter variations. The eSVs were recorded for an evolution time of one time step ($\Delta T = 1$).

These results demonstrate the ability to make good predictions of parameter variations in this system, even when they occur simultaneously. The case where both parameters are changed by 3% also shows that, even in cases where one extrapolates beyond the eSVFs collected initially (here validations were performed only up to 2%), one can still generate good results.

The accuracy of these types of parameter variation predictions is affected by several factors. The first of these is the level of proportionality required to validate calibration vectors. Figure 6 illustrates how tightening the proportionality requirement influences the predictions over a range of different evolution times. Here, the actual values of the variations are -0.15 in σ, -0.2 in b, and 0 in r. It is obvious that, as the proportionality requirement is tightened, the predictions generally improve, markedly from 5% to 2% and only slightly from 2% to 1%. The number of initial conditions that are considered valid also declines.

Figure 6 also demonstrates that the length of the evolution time ΔT can have a large effect on eSV accuracy. If this time is too short, the eSVs will have small magnitude and will be prone to large error for any small errors in the individual time series predictions. For large times, a dynamic system that is chaotic will lose predictability and the error will again be large. Our predictions of b with a 5% proportionality requirement become inaccurate for $\Delta T > 10$ because a few of the initial conditions that pass the proportionality test generate eSVs that are in error. This suggests that it is prudent to make variation predictions over several different ΔT values and look for a consistent result before drawing any conclusions about how parameters may have changed.

A third factor affecting the accuracy of parameter predictions is quantity and density of the time series themselves. Having more data from which to construct a local model about a given initial condition can lead to an improvement in the predictions. The results are more subtle, but Fig. 7 shows how increasing the number of points in the time series from 10^5 to 10^6 can slightly reduce some of the variability in the variation predictions. The proportionality enforced for this plot was 2%.

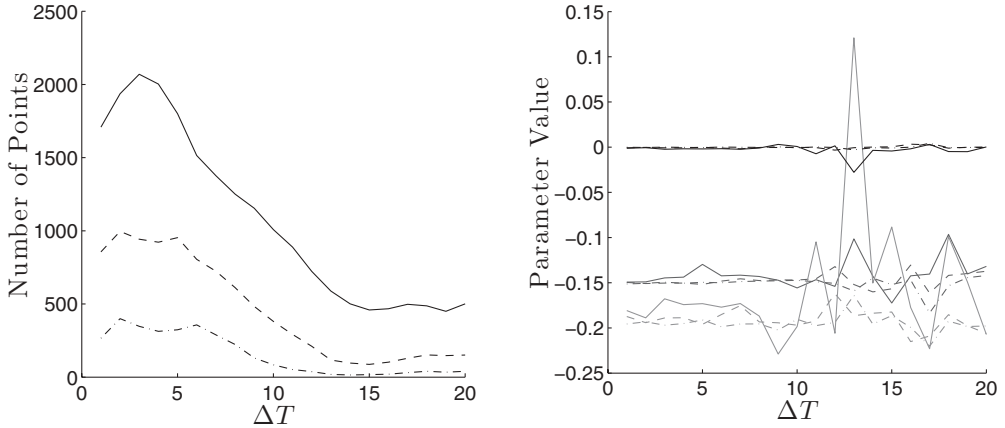

FIG. 6. Number of initial conditions validated for eSV construction (out of 10^3) (left). Predictions of the parameter variations (right). Solid line corresponds to 5% proportionality, dashed line corresponds to 2% proportionality, dashed-dotted line corresponds to 1% proportionality.

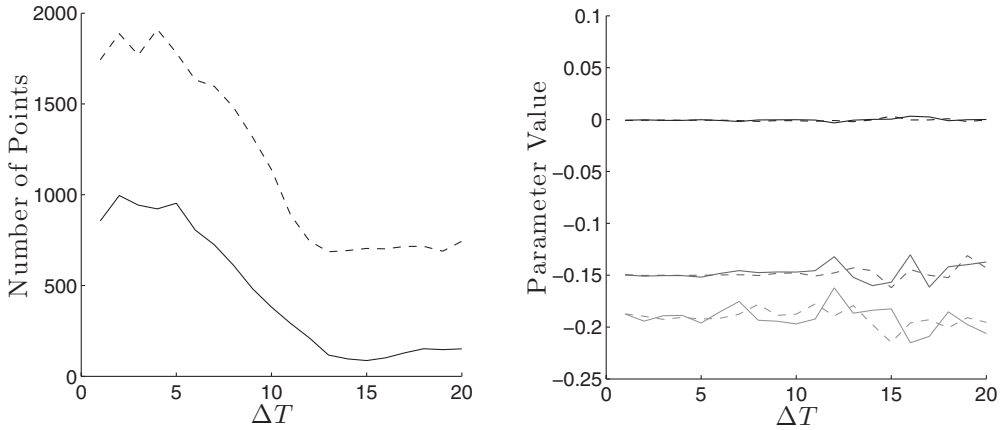

FIG. 7. Number of initial conditions validated for eSV construction (out of 10^3) (left). Predictions of the parameter variations (right). Solid line corresponds to 10^5 points in time series, dashed line corresponds to 10^6 points.

The error in embedded eSVs is also significantly influenced by noise. To investigate this, noise was added to the simulated time series as described in Eq. (24). The levels added ranged from 0.0001 to 0.1. For levels of noise 0.0001 and below, errors in eSV magnitude were not significantly influenced. However, for noise around 0.001 the variation prediction began to be degraded. For higher levels of noise it was often impossible to find any initial conditions that generated eSVs meeting the proportionality requirements. This suggests that minimizing noise in any experimental system is very important in order to be able to construct eSVs. It is possible that by predicting both the future points on the nominal trajectory as well as on the varied trajectory (rather than just accepting the nominal trajectory data at face value) the impact of noise could be further reduced through the effects of neighborhood averaging, but this has not been explored here. Table III shows the impact of noise on eSV parametric variation reconstruction for $\Delta T = 3$.

B. Experimental time series

The investigation of eSVs using time series data generated by simulated systems shows that the methodology laid out for eSV construction is reliable. However, it is also important to be able to achieve similar results with experimental data to confirm that embedded SVF construction is applicable to real systems. For this reason, we examine the effects of parametric variation by adjusting R and R_L in Chua's circuit.

TABLE IV. Chua's circuit time series parameter reconstruction.

Case	Parameter	Variation ($\delta\Omega$)	eSVF prediction ($\delta\Omega$)
1	R	32	34.3
	R_L	0	0.3
2	R	0	3.9
	R_L	-47	-44.2
3	R	-33	-27.2
	R_L	48	49.8

Applying the techniques for determining state dimension and time delay, we found that the experimental data should be embedded in a state space having $m = 3$ and a time delay of $\tau = 9$. The reconstructed attractor using these parameter values is shown in Fig. 8.

Once initial calibration eSVFs are taken about the nominal parameter values one can perform the reconstructions shown in Table IV. The fact that the data has low noise (the noise to signal ratio could be as low as 10^{-10} based on the resolution of the analog to digital converter) no doubt contributes to the predictions being accurate.

VI. CONCLUSIONS

This paper presents a methodology for the construction of SVFs in time-delay embedded coordinates based on local linear modeling of time series. A specific method for weighting neighborhood states in local modeling called ePCA was introduced. This method minimizes quadratic cross terms in an effort to improve eSV accuracy. Local modeling proved to

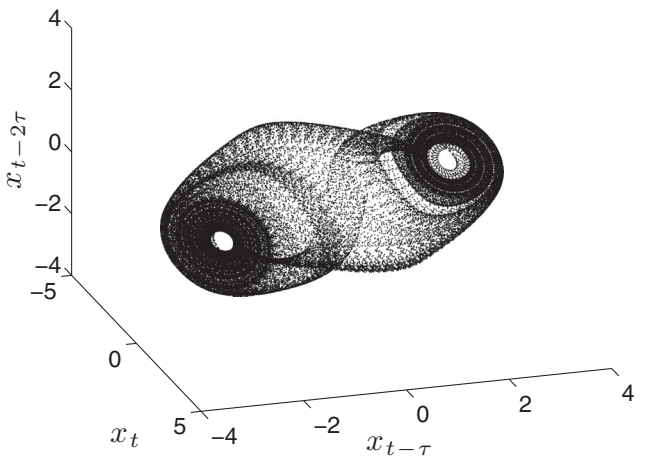

FIG. 8. A three-dimensional projection of the Chua oscillator attractor.

be a reliable means of constructing eSVFs and identifying parametric variations for both simulated and experimental time series. A technique for checking the accuracy of these SVF predictions was also presented for cases where the system equations are known and the dimension of the nonembedded state and the embedded state are equal. Being able to construct SVFs in embedded coordinates will make the application of SVFs to a wide variety of physical systems feasible, even when it is not possible to measure the complete state.

ACKNOWLEDGMENTS

The authors wish to recognize the NSF CMMI Dynamical Systems Program (Grant No. 0625011) for support of this work. The additional support of the NSERC Postgraduate Scholarship Program (PGS-D) is also acknowledged.

[1] S. W. Doebling, C. R. Farrar, and M. B. Prime, Shock and Vibration Digest **30**, 91 (1998).
[2] R. Garcia and R. Perez, Surf. Sci. Rep. **47**, 197 (2002).
[3] B. Ilic, H. Craighead, S. Krylov, W. Senaratine, C. Ober, and P. Neuzil, J. Appl. Phys. **95**, 3694 (2004).
[4] T. Burg, A. Mirza, N. Milovic, C. Tsau, G. Popescu, J. Foster, and S. Manalis, J. Microelectromech. Syst. **15**, 1466 (2006).
[5] S. Yin and B. Epureanu, Philos. Trans. R. Soc., A **364**, 2515 (2006).
[6] B. Epureanu and A. Hashmi, J. Vib. Acoust. **128**, 732 (2006).
[7] A. Hashmi and B. Epureanu, Nonlinear Dyn. **45**, 319 (2006).
[8] J. Lim and B. Epureanu, Nonlinear Dyn. **59**, 113 (2010).
[9] S. Yin and B. Epureanu, J. Vib. Acoust. **129**, 763 (2007).
[10] N. H. Packard, J. P. Crutchfield, J. D. Farmer, and R. S. Shaw, Phys. Rev. Lett. **45**, 712 (1980).
[11] F. Takens, in *Dyanamical Systems and Turbulence*, edited by D. Rand and L. Young (Springer, Berlin, 1981), pp. 366–381.
[12] D. Broomhead and G. King, Phys. D **20**, 217 (1986).
[13] T. Sauer, J. A. Yorke, and M. Casdagli, J. Stat. Phys. **65**, 579 (1991).
[14] H. Abarbanel, *Analysis of Observed Chaotic Data* (Springer-Verlag, New York, 1996).
[15] H. Kantz and T. Schreiber, *Nonlinear Time Series Analysis* (Cambridge University Press, Cambridge, 1997).
[16] J. D. Farmer and J. J. Sidorowich, Phys. Rev. Lett. **59**, 845 (1987).
[17] M. Casdagli, Phys. D **35**, 335 (1989).
[18] L. Smith and R. Bhansali, Philos. Trans. R. Soc., A **348**, 371 (1994).
[19] E. Bolt, Int. J. Bifurcation Chaos Appl. Sci. Eng. **10**, 1407 (2000).
[20] R. Brown, P. Bryant, and H. D. I. Abarbanel, Phys. Rev. A **43**, 2787 (1991).
[21] T. Sauer, in *Time Series Prediction: Forecasting the Future and Understanding the Past*, edited by A. Weigend and N. Gershenfeld (Santa Fe Institute, Santa Fe, 1994), pp. 175–193.
[22] D. Kugiumtzis, Phys. D **112**, 344 (1998).
[23] L. Torres and L. Aguirre, Electron. Lett. **36**, 1915 (2000).
[24] M. Kennel, R. Brown, and H. Abarbanel, Phys. Rev. A **45**, 3403 (1992).
[25] M. Kennel and H. Abarbanel, Phys. Rev. E **66**, 026209 (2002).
[26] A. M Fraser and H. L. Swinney, Phys. Rev. A **33**, 1134 (1986).
[27] T. Buzug and G. Pfister, Phys. Rev. A **45**, 7073 (1992).
[28] R. Hegger, H. Kantz, and T. Schreiber, Chaos **9**, 413 (1999).