Local false nearest neighbors and dynamical dimensions from observed chaotic data

Henry D. I. Abarbanel
Institute for Nonlinear Science, Department of Physics, and Marine Physical Laboratory, Scripps Institution of Oceanography, University of California, San Diego, Mail Code 0402, La Jolla, California 92093-0402

Matthew B. Kennel
Institute for Nonlinear Science and Department of Physics, University of California, San Diego, Mail Code 0402, La Jolla, California 92093-0402

(Received 22 October 1992)

The time delay reconstruction of the state space of a system from observed scalar data requires a time lag and an integer embedding dimension. The minimum necessary *global* embedding dimension d_E may still be larger than the actual dimension of the underlying dynamics d_L. The embedding theorem only guarantees that the attractor of the system is fully unfolded using d_E greater than $2d_A$, with d_A the fractal attractor dimension. Using the idea of *local* false nearest neighbors, we discuss methods for determining the integer-valued d_L.

PACS number(s): 05.45.+b

I. INTRODUCTION

The method of false nearest neighbors [1] determines the minimum embedding dimension necessary to reconstruct the state space of a dynamical system from observed data with time delay embedding [2,3]. This method seeks that dimension d_E in which vectors in $\mathbb{R}^{d_E}$

$$\mathbf{y}(k)=(x(k),x(k+T),\ldots,x(k+(d_E-1)T)) , \quad k=1,\ldots,N \qquad (1)$$

composed from scalar observations $x(k)=x(t_0+k\tau_s)$ with sampling time τ_s describe points whose state-space neighbors are a result of the dynamics rather than from being projected near one another as an artifact of using too low an embedding dimension. The embedding theorem [2,3] tells us that $d_E>2d_A$, where d_A is the "box-counting" fractal dimension, is sufficient to unfold the geometry of the attractor. The false neighbor method [1] tells, from the data itself, the necessary dimension d_E for globally unfolding the attractor. Choosing an insufficient d_E will cause parts of the attractor which are widely separated in the original, but unknown, state space to overlap spuriously in the reconstructed space. To detect this, the false nearest neighbors test constructs vectors from the data in dimension $d=1$, then $d=2$, and so forth, asking at each stage what fraction of nearest neighbors in the data set, as seen in dimension d, fail to remain close in dimension $d+1$. When all nearest neighbors are true, that is, do not significantly move apart when we go the next dimension, we have found the minimum embedding dimension $d=d_E$. In this paper we discuss variations of the false nearest-neighbor idea that estimate how many dimensions d_L are *locally* required to describe the dynamics generating the data, without knowing the equations of motion. Globally, one may require a dimension d_E larger than d_L, the dimension of the *active dynamics*, to successfully reconstruct a dynamical system.

If we are successful in establishing a value for d_L, then we can proceed to making models for prediction or control of the observed system in this dimension. Also if we are interested in the predictability of the system, then by locally estimating $d_L\times d_L$ dimensional Jacobian matrices, we may directly compute only the d_L true Lyapunov exponents [4–6]. As are the Lyapunov exponents, the local dimension d_L is *invariant* under smooth, invertible diffeomorphisms of state-space coordinates, and should thus be useful in classifying the dynamics. The minimum global embedding dimension d_E is not a similar invariant, however, as one may find, for example, different d_E for scalar time series data taken from different variables of a dynamical system, even though the underlying dynamics are the "same."

We want to examine the true neighbors of a data point $\mathbf{y}(k)$, as identified in d_E, locally in state space to see if we can describe the local geometry and dynamics with fewer than d_E coordinates. In the ideal case d_L is the dimensionality of the manifold, embedded in $\mathbb{R}^{d_E}$ that contains the data, and so is often called the topological dimension. If we find a $d_L<d_E$, it means that, depending on the global topology, we may be able to apply a smooth state-space transformation to the data to further unfold the attractor from a set in $\mathbb{R}^{d_E}$ to one in $\mathbb{R}^{d_L}$ that preserves all the important dynamical features. At this time we have no constructive method to find the required transformation. In other circumstances, the global topology might be that of a torus of dimension d_L or some combination of a torus and flat Euclidean space, but always with d_L degrees of freedom.

By using the term "active dynamics" we have in mind the notion that even in a very-high-dimensional system, perhaps even infinite dimensional, as in the continuum description of a fluid, points in the full high-dimensional state space will be drawn to an attractor which will occu-

py zero volume in the full space. The dynamical dimension is that finite integer which quantifies the number of degrees of freedom that captures motion on the observed attractor. An infinite-dimensional system may relax onto an attractor with finite degrees of freedom, but the active number is unknown *a priori*. In a simpler situation, we expect that a dynamical system in d variables, whether continuous $[\dot{\mathbf{x}}(t)=\mathbf{F}(\mathbf{x}(t))]$ or discrete $[\mathbf{x}(n+1)=\mathbf{F}(\mathbf{x}(n))]$, always to produce behavior that we can capture in $d_L \leq d$ dimensions. For example, in a system of three autonomous first-order ordinary differential equations running in a chaotic regime, we expect to find a local dimension of 3, though sometimes one would need embedding dimensions up to 6 in order to reconstruct the state space using data from one of the three scalar variables. As another simple example, note that a signal consisting of N independent sine waves has $d_L=N$, but requires at least $d_E=N+1$ to embed.

Note that we do not intend our nomenclature to imply that there is a different local dimension for differing regions of state space, but rather, that there is a single number quantifying the degrees of freedom necessary, for *all* neighborhoods, to successfully capture the deterministic dynamical evolution in local regions of the state space. If one were to define somehow integer valued "localized dimensions," one for each individual region of state space, our intuition is that our local dimension ought to be the maximum of these. However, we cannot immediately conceive of a reliable and useful definition for such truly localized quantities.

II. LOCAL FALSE NEIGHBORS

We assume that, as in the Introduction, we are presented with scalar data sampled at evenly spaced time intervals. For all our investigations, we choose the time delay $T\tau_s$ used in constructing vectors to be the first minimum, when there is one, in the average mutual information [7] over the attractor. If there is no minimum, as is the case for many maps, we choose $T=1$.

The essential idea of our methods is to find the nearest neighbor, as viewed in dimension d, of a reference point $\mathbf{y}$, and to see if it moves a "large" distance from $\mathbf{y}$ by examining the $(d+1)$st coordinate of $\mathbf{y}$ and its nearest neighbor in dimension $d+1$. If one selects the nearest neighbor from among the total data set, one recovers the global embedding dimension test. Now, we restrict the selection of the nearest neighbor to lie inside a medium-sized neighborhood surrounding our reference point. We first choose a fixed number N_B of neighbors of $\mathbf{y}$, as determined in the global embedding dimension d_E. Then we project the points in this neighborhood down to our trial dimension d, compute the *nearest* neighbor in dimension d and apply a test to see if it is a "local false neighbor." We look for errors that occur when we locally project a neighborhood down to a space of insufficient dimension. One may also define neighborhoods to encompass a fixed radius instead of a fixed number of points without substantially changing the results, but then one encounters some neighborhoods that are too sparse to extract good statistics, and they somehow have to be identified and discounted. So, for computational ease, we employ a fixed number of neighbors.

A. The simple approach

The simplest approach is a direct application of the tests in Ref. [1]. We embed our data in a space dimension at least as large as the minimum required dimension d_E. In this dimension, which we call the working dimension $d_W \geq d_E$, we identify N_B neighbors of each point $\mathbf{y}(n)$, $n=1,2,\ldots,N$, in our time series. Our focus is now on these N_B neighbors. We label the N_B+1 vectors in this neighborhood as $\mathbf{y}^{(r)}(n)$, $r=0,1,2,\ldots N_B$. $\mathbf{y}^{(0)}(n)=\mathbf{y}(n)$. From the time delay construction we know the components and the time index for each neighbor.

We proceed then in the same fashion as with the global false neighbor construction. First, we ask if the nearest neighbor to $\mathbf{y}(n)$ as determined in trial local dimension $d_L=1$, then $d_L=2$, and so forth until $d_L=d_W-1$ remains close when we go to local dimension d_L+1.

We identify the nearest neighbor in dimension d_L among the N_B total neighbors,

$$\mathbf{y}_{d_L}^{(\mathrm{NN})}(n)=(x(m),x(m+T),\ldots,x(m+(d_L-1)T)) . \tag{2}$$

The notation states that we are talking about the nearest neighbor (NN), as seen in dimension d_L with Euclidean metric along the first d_L components of the embedded vector, and locally chosen from the ball surrounding the point with time index n. The nearest neighbor has time index m in the original data stream. We have projected the neighborhood down from $\mathbb{R}^{d_E}$ to $\mathbb{R}^{d_L}$ simply by taking the first d_L components and finding the point closest to the center: $\mathbf{y}^{(0)}(n)$.

If $|x(m+d_LT)-x(n+d_LT)|/|\mathbf{y}_{d_L}^{(\mathrm{NN})}-\mathbf{y}_{d_L}^{(0)}|>R_T$, for some large cutoff R_T, then we designate $\mathbf{y}_{d_L}^{(\mathrm{NN})}(n)$ a false nearest neighbor. This means that points which are originally nearest neighbors have moved apart a large distance in comparison to their initial separation.

In the second test, we compare the "extra distance" $|x(n+d_LT)-x(m+d_LT)|$ to a characteristic size of the entire neighborhood:

$$R_{N_B}(n)=\frac{1}{N_B}\sum_{r=1}^{N_B}|y_1^{(r)}(n)-y_{\mathrm{av}}(n)| , \tag{3}$$

with

$$y_{\mathrm{av}}(n)=\frac{1}{N_B}\sum_{r=1}^{N_B}y_1^{(r)}(n) . \tag{4}$$

If $|x(n+d_LT)-x(m+d_LT)|/R_{N_B}(n)>R_T'$, then we also designate this point $\mathbf{y}(n)$ as having a false nearest neighbor in local dimension d_L.

We repeat this for neighborhoods centered around all $\mathbf{y}(n)$, and we record, as a function of d_L, what fraction have local false nearest neighbors. The methods we describe are insensitive to R_T' as long as $R_T'>2$. We either used $R_T'=2.41$ or set $R_T'\to\infty$. The latter ignores the test

of change in the additional component relative to the neighborhood extent. This was needed in [1] to distinguish high dimensional "noise," but we are already dealing with a low-dimensional system in this discussion, so we need not really retain the second test to continually check that we do not have a high-dimensional data set.

We begin with examples where we know the answer for d_L and d_E and examine what this test indicates as we vary the tolerances involved in the false neighbor test and the number of neighbors N_B.

The first example we consider is the familiar Lorenz attractor [8]:

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

with parameter values $\sigma=16$, $b=4$, and $r=45.92$. In Figs. 1–4 we have taken $d_E=5$ and analyzed 23 000 points of $x(t)$ data taken at a numerical time step of $\tau_s=0.01$. The time lag suggested by the mutual information criterion is $T=0.1=10\tau_s$, and the time to make a circuit around the attractor is approximately 0.5 in these units. So this amount of data represents about 450 visits of the orbit around the attractor. This is probably more than is needed.

In Fig. 1 we see the percentage of local false nearest neighbors as a function of N_B and R_T. Clearly, $d_L=1$ will not do for the Lorenz system. In Fig. 2, the situation has much improved, while in Fig. 3, where $d_L=3$, a careful examination of the vertical scale shows yet further improvement. With a sufficient choice of d_L, one should see a wide plateau as a function of R_T and N_B, as in the global false neighbors calculation [1]. Selecting an arbitrary value of $R_T=30.27$, well within the range where the number of local false nearest neighbors is rather insensitive to this parameter, we see in Fig. 5 the percentage of local false nearest neighbors in $d_L=2$, 3, and 4 as a function of N_B. At $d_L=3$ the method computes a very small proportion of local false neighbors, independent of N_B. It would be quite safe, as we already knew in this example, to choose $d_L=3$. Since the global false nearest neighbors indicates $d_E=3$ [1], we have only confirmed that the global dimension is also the local dimension of the dynamics, also equal to the number of differential equations in the original system.

Next we turn to the Ikeda [9,10] map of the plane to itself:

$$z(n+1)=p+Bz(n)\exp\{i\kappa-i\alpha/[1+|z(n)|^2]\}\ , \tag{5}$$

where $p=1.0$, $B=0.9$, $\kappa=0.4$, $\alpha=6.0$, and $z(n)$ is complex. The dimension of the attractor associated with this map is $d_A\approx 1.8$, and this is a dynamical system in two variables, the real and imaginary parts of $z(n)$. At these parameter values, the data set composed of the real part of $z(n)$, embedded with time step $T=1$, requires a minimum embedding dimension of $d_E=4$. This d_E is also the sufficient dimension as determined from the embedding theorem [2,3]. In Fig. 6 we show the analysis of 43 000 points of the real part of $z(n)$ from the Ikeda map with an embedding dimension $d_E=d_W=5$. In these calculations we again choose $R_T=30.27$.

By virtue of the high proportion of false neighbors, combined with a significant dependence on N_B, $d_L=1$ is excluded. In Fig. 7 we omit the $d_L=1$ results on the graph and concentrate on the remainder, and see that the percentage of local false nearest neighbors for dimensions 2 through 4 is quite small, less than 0.25% for the range $15\le N_B\le 69$ that we considered. While one might con-

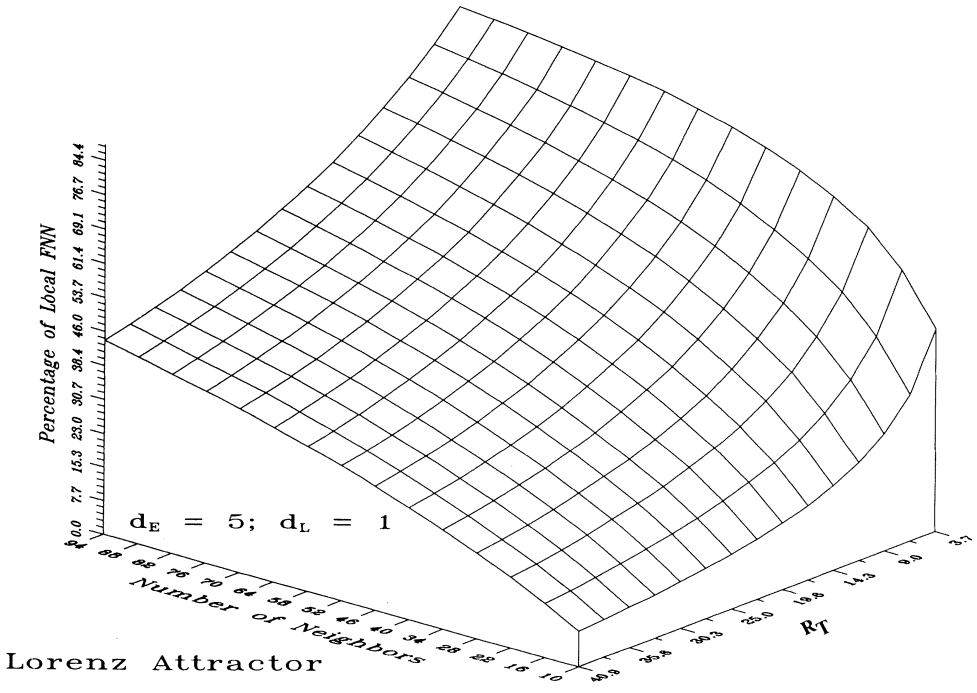

FIG. 1. From the simple approach to local false neighbors, the percentage of local false nearest neighbors (FNN's) for 23 000 points from the $x(t)$ component of the Ref. [8] Lorenz model. The working dimension $d_W=d_E=5$ and $d_L=1$. The percentage of local FNN is exhibited as a function of the number of neighbors and of the tolerance R_T in determining the threshold for change in the distance of nearest neighbors in going from dimension d_L to d_L+1. $d_L=1$ is clearly not chosen as the local dynamical dimension.

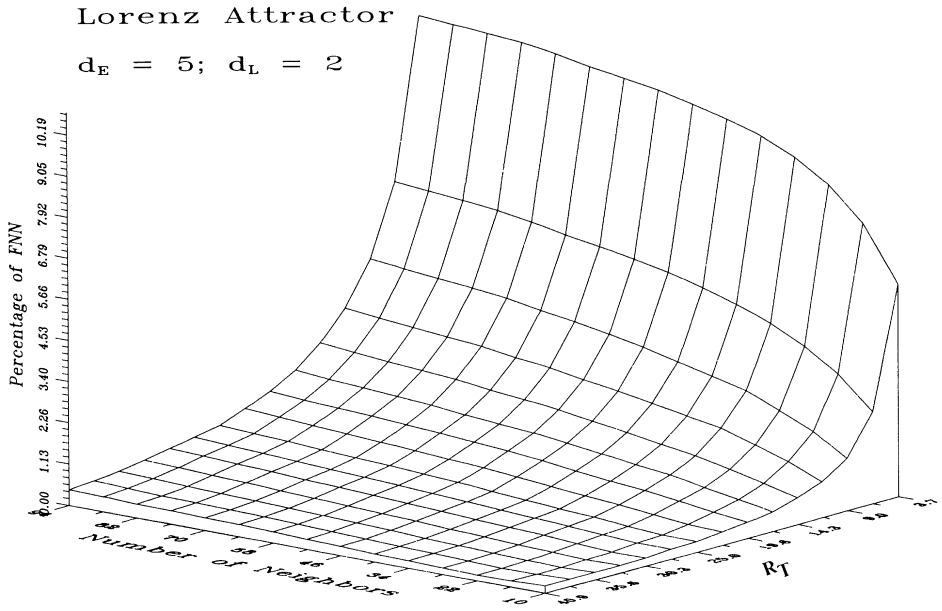

FIG. 2. From the simple approach to local false neighbors, the percentage of local FNN's for 23 000 points from the $x(t)$ component of the Ref. [8] Lorenz model. The working dimension $d_W = d_E = 5$ and $d_L = 2$. The percentage of local FNN's is exhibited as a function of the number of neighbors and of the tolerance R_T in determining the threshold for change in the distance of nearest neighbors in going from dimension d_L to $d_L + 1$. The percentage of local FNN's becomes rather independent of N_B and R_T at this d_L, so one might choose $d_L = 2$ as the local dynamical dimension for this system.

clude that $d_L = 2$ is thus quite acceptable for this data set, looking back at Fig. 5 we see we might have concluded $d_L = 2$ with the same force for the Lorenz system. (Since $d_A \approx 2.06$ for the Lorenz attractor, this is not unreasonable to find, but nonetheless wrong.) Only in $d_L = 3$ for the Ikeda map would we conclude that the number of false local nearest neighbors has become independent of N_B and thus feel safe about adopting this value. There is some achievement here since $3 < d_E = 4$, but we shall be able to do better below by enlarging our considerations.

We have just shown that this simple local false nearest neighbor test works moderately well. Unfortunately, it seems not to be a sufficiently sharp tool to discriminate $d_L = 2$ and 3 for the Ikeda map where we know the true $d_L = 2$.

One potential problem with this test is the method we employ for deciding which $d_L \leq d_W$ components of the original N_B d_W-dimensional vectors to use in choosing neighbors in dimension d_L. Suppose, for example, that the points in the data set lie on a plane, so $d_L = 2$, but the

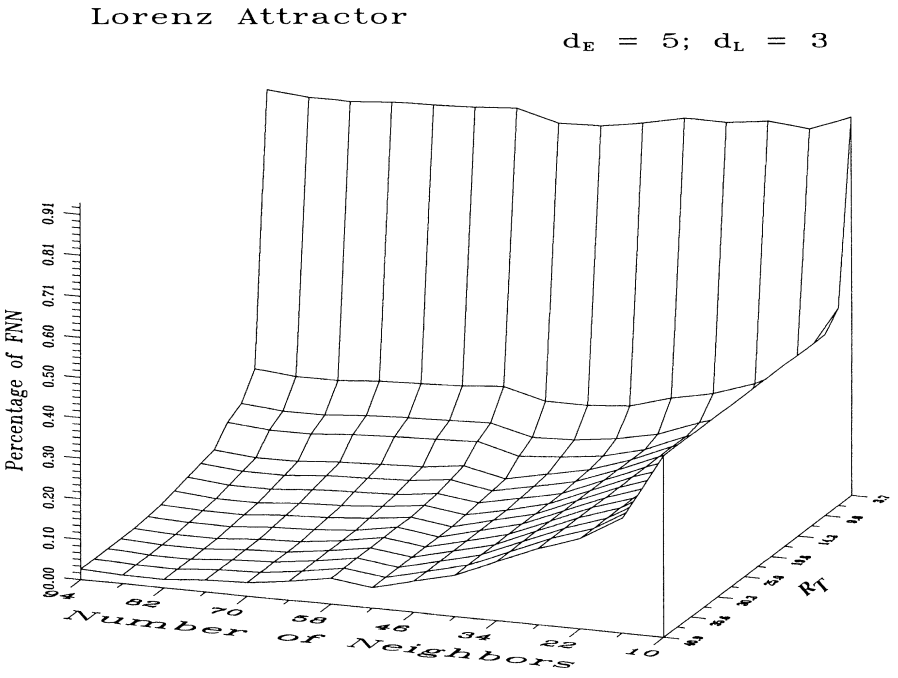

FIG. 3. From the simple approach to local false neighbors, the percentage of local FNN's for 23 000 points from the $x(t)$ component of the Ref. [8] Lorenz model. The working dimension $d_W = d_E = 5$ and $d_L = 3$. The percentage of local FNN's is exhibited as a function of the number of neighbors and of the tolerance R_T in determining the threshold for change in the distance of nearest neighbors in going from dimension d_L to $d_L + 1$. Note the change in scale relative to Figs. 1 and 2. $d_L = 3$ would be quite acceptable as the local dimension.

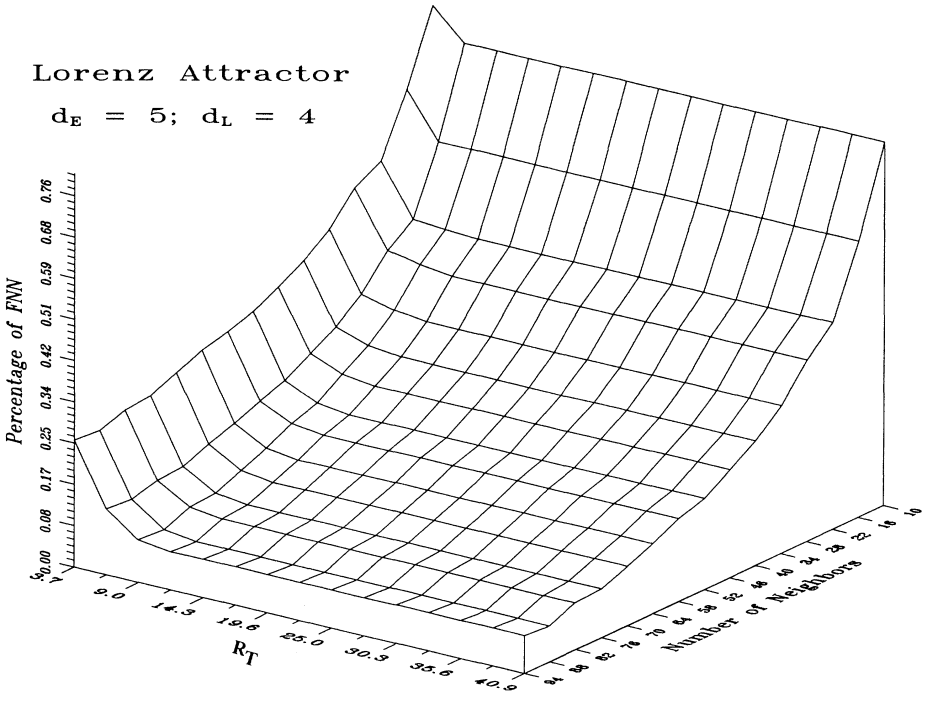

FIG. 4. From the simple approach to local false neighbors, the percentage of local FNN's for 23 000 points from the $x(t)$ component of the Ref. [8] Lorenz model. The working dimension $d_W = d_E = 5$ and $d_L = 4$. The percentage of local FNN's is exhibited as a function of the number of neighbors and of the tolerance R_T in determining the threshold for change in the distance of nearest neighbors in going from dimension d_L to $d_L + 1$. Note the change is scale relative to Figs. 1–3.

orientation of this plane twists and folds in the time delay coordinate system so an embedding dimension $d_E > 2$ is required to globally unfold the attractor. Then it is possible that just taking the first two components of each vector will occasionally result in the plane where the dynamics lies being oriented nearly normal to the (y_1, y_2) plane in time delay space. In this circumstance we would encounter sets of points which are artificially projected near to each other, in the sense of Euclidean distance, and thus identify substantially more false nearest neighbors than is correct. We suspect this is why the method slightly rejects $d_L = 2$ for data from the Ikeda map.

B. A refinement using dynamics

First we address the matter of the coordinate system in which to ask the question of local false neighbors. The issue is that the d_L-dimensional space must be chosen by a local projection from $\mathbb{R}^{d_E} \to \mathbb{R}^{d_L}$. The directions in Euclidean $\mathbb{R}^{d_E}$ space established by the time delay coordinate choice $x(k + jT)$, $j = 0, 1, \ldots, d_E - 1$ may be inconveniently oriented so points which are far from each oth-

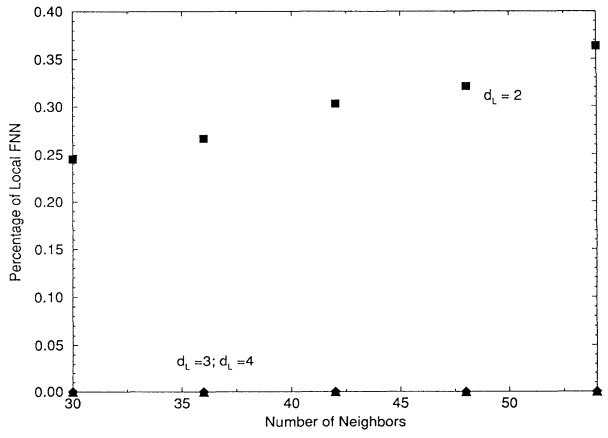

FIG. 5. At $R_T = 40.27$, the percentage of local FNN's in $d_L = 2, 3, 4$ as a function of N_B for the Ref. [8] Lorenz model. $d_W = d_E = 5$ here, and 33 000 data points are used. The time delay is $T = 10$. From this data we might conclude that $d_L = 2$ is not a good choice since the local FNN percentage is not independent of N_B, while at $d_L = 3$ it has become independent.

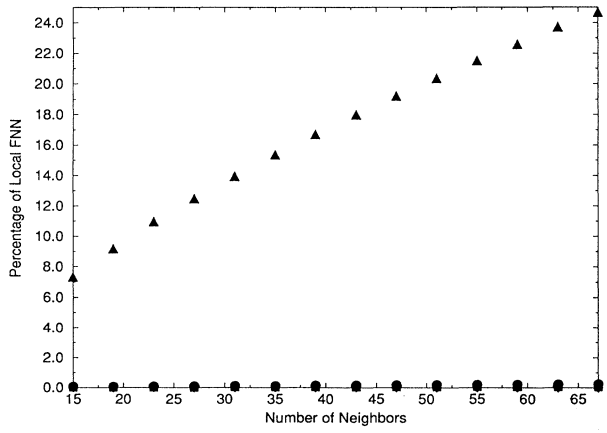

FIG. 6. Percentage of local FNN as a function of N_B for data from the Ikeda map at $d_W = d_E = 5$ for $d_L = 1, 2, 3, 4$. $d_L = 1$ is shown in triangles. 43 000 points of the real part of $z(n)$ are used. From this figure we would conclude that $d_L = 1$ is not a good local dimension.

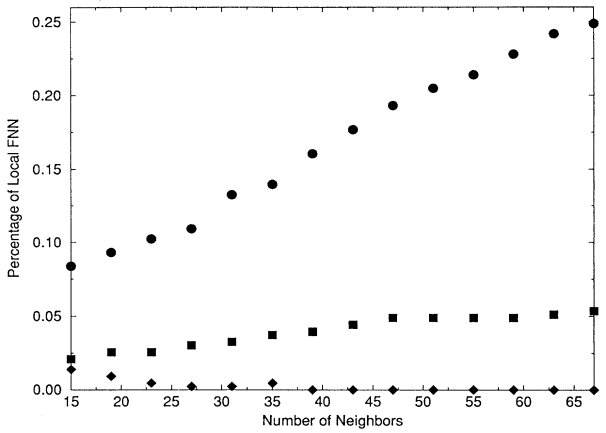

FIG. 7. Percentage of local FNN as a function of N_B for data from the Ikeda map at $d_W=d_E=5$ for $d_L=2,3,4$. $d_L=2$ is in solid circles; $d_L=3$ is in solid squares; $d_L=4$ is in solid diamonds. 43 000 points of the real part of $z(n)$ are used. From this figure we would probably conclude that $d_L=2$ is not a good local dimension, and be inclined to choose $d_L=3$ since independence of N_B is achieved in this dimension.

er within the neighborhood in $\mathbb{R}^{d_E}$ of the N_B true neighbors are inappropriately projected near to each other. To avoid this we choose a local coordinate system in which to select d_L vectors by looking at the *principal components* of the data set comprised of the $\mathbf{y}^{(r)}(k)$. The eigendirections of the local sample covariance matrix relative to $\mathbf{y}(k)$:

$$R_{\alpha\beta}=\frac{1}{N_B}\sum_{r=1}^{N_B}[y^{(r)}(k)-y(k)]_\alpha[y^{(r)}(k)-y(k)]_\beta\ , \tag{6}$$

namely the $\mathbf{e}(i)$, $i=1,2,\ldots,d_E$ satisfying

$$\begin{aligned}&\mathbf{R}(k)\cdot\mathbf{e}(i)=\sigma(i)\mathbf{e}(i)\ ,\\ &\mathbf{e}(i)\cdot\mathbf{e}(j)=\delta_{ij}\ ,\quad \sigma(i)\geq\sigma(i+1)\ ,\end{aligned} \tag{7}$$

provides a local coordinate basis. We choose this set of local coordinates for the sole purpose of selecting directions where most of the data are located, that is, those eigendirections with the largest eigenvalues $\sigma(i)$, and selecting directions "pointing out of the data," namely those directions corresponding to the smallest eigenvalues. This is clearly an heuristic choice, and one could use other criteria than capturing the most variance in the data set, which is the intent of projecting on to the first d_L principal components. For example, one might wish to minimize the entropy of the distribution of points projected along the first direction (as then it looks the least like noise), then choose the second direction to minimize the entropy of vectors projected orthogonal to the first direction, etc. All of these choices have merit. Except for principal components, they require a large amount of computation. Our goal in this paper is really quite modest with regard to the use of principal component analysis. We only wish to avoid accidental projection of points into the neighborhood of $\mathbf{y}(k)$'s in any given d_L. Principal component analysis using the eigendirections of $\mathbf{R}(k)$ to determine a set of basis vectors in each d_L will suffice since it meets our requirements and is computationally undemanding.

In the principal component coordinate system a vector $\mathbf{y}^{(r)}(k)$ has components

$$(\mathbf{y}^{(r)}(k)\cdot\mathbf{e}(1),\mathbf{y}^{(r)}(k)\cdot\mathbf{e}(2),\ldots,\mathbf{y}^{(r)}(k)\cdot\mathbf{e}(d_E))\ , \tag{8}$$

and we select d_L-dimensional vectors $\mathbf{w}^{(r)}(d_L)$ from among these components

$$\mathbf{w}^{(r)}(d_L)=(y^{(r)}(k)\cdot\mathbf{e}(1),\mathbf{y}^{(r)}(k)\cdot\mathbf{e}(2),\ldots,\mathbf{y}^{(r)}(k)\cdot\mathbf{e}(d_L))\ , \tag{9}$$

for each of the N_B true neighbors in d_E. Using these vectors in $\mathbb{R}^{d_L}$ we evaluate the distances $|\mathbf{w}^{(r)}(d_L)-\mathbf{w}^{(0)}(d_L)|^2$ and select the nearest neighbor $\mathbf{w}^{\mathrm{NN}}(d_L)$ of $\mathbf{w}^{(0)}(d_L)$. Now we have identified that d_E-dimensional vector $\mathbf{y}^{\mathrm{NN}}(d_L)=\mathbf{y}(m)$, which is the nearest neighbor of $\mathbf{y}(k)$ with nearness determined in dimension d_L. The "time" $t_0+m\tau_s$ associated with this vector has no special relation to the time $t_0+k\tau_s$ of $\mathbf{y}(k)$ since its nearness is a property of state space location, not temporal sequence. At this stage we are finished with the computations in the d_L-dimensional rotated coordinate system, and we return to the original d_E-dimensional space. The choice of trial local dimension affects only the identity of that point we call a nearest neighbor to $\mathbf{y}(k)$, i.e., the value of m.

We now ask what happens to these two nearest neighbors as we move forward in time. In particular, we consider the value of the difference

$$|x(k+(d_E-1)T+\Delta)-x(m+(d_E-1)T+\Delta)|\ , \tag{10}$$

as a function of the number of steps $\Delta\tau_s$ forward in time from the present reference vector. For every neighborhood surrounding $\mathbf{y}(k)$ and for each trial local dimension d_L, we compute the smallest Δ for which this difference grows larger than βR_A, a significant fraction of the size of whole attractor R_A, and call it $K(k;d_L)$. (If we run off the end of the data set prematurely, we invalidate this neighborhood and remove it from consideration.) We choose R_A as

$$\begin{aligned}&R_A=\frac{1}{N}\sum_{k=1}^{N}|x(k)-\bar{x}|\ ,\\ &\bar{x}=\frac{1}{N}\sum_{k=1}^{N}x(k)\ .\end{aligned} \tag{11}$$

Other choices for R_A, such as the rms size of the attractor, do not change our conclusions. We have looked at β in the range $0.05\leq\beta\leq0.2$ and find no dependence on our results on this fraction. Of course specific quantities may depend on β, but the value determined for d_L should not depend on β over a reasonable range. In fact, this is a useful check to ensure a correct d_L has been computed. Note that now the question of local false neighbors is

posed in dynamical terms, by looking at the future evolution of two local neighbors, rather than in purely geometrical language, as examining what happens as dimension increases. However, in our context of time delay embedding, these are clearly equivalent, as the embedding theorem does not require one to take evenly spaced samples in time for each component, though it is certainly common practice.

We collect the values of $K(k;d_L)$ for each point on the attractor, and ask for the proportion of $K(k;d_L)$ less than T, the information theoretic time delay used in the time delay reconstruction itself, relative to the total number of neighborhoods surveyed. (We use T simply as a useful characteristic time scale of the data. As with the case for β, the exact value does not influence the resulting determination of d_L within a broad range.) This leads us to define

$$P_K(d_L)=\frac{[\text{number of } K(k;d_L)\le T]}{(\text{number of neighborhoods})} . \tag{12}$$

If $K(k;d_L)$ is less than T, it indicates a lack of predictability beyond this time. We expect false neighbors to be less predictable than true neighbors, because they are near each other for geometrical rather than dynamical reasons and, therefore, we will see more neighbors with poor prediction times. As we increase d_L, $P_K(d_L)$ will decline, and once the correct d_L occurs, the variation of $P_K(d_L)$ with d_L will cease as we encounter the natural predictability of the dynamics. Furthermore, $P_K(d_L)$ ceases to vary substantially with N_B at or above the correct d_L, but choosing an insufficient d_L causes it to change with N_B, because fewer or more points can be spuriously projected down and show up as false nearest neighbors as a function of the neighborhood size.

In Fig. 8 we show our results of this procedure for data from the real part of $z(n)$ of the Ikeda map. Here, and in all data sets which follow, we employ 20 000 scalar data points. The $P_K(d_L)$ where the indicated spatial difference grows larger than the fraction $\beta=0.1$ of R_A, is shown for various d_L and various N_B.

Starting at $d_L=2$ we see a clear decrease to a very flat plateau in $P_K(d_L)$, at a level independent of N_B. The interpretation is that the correct $d_L=2$. We have taken $10\le N_B\le 100$, which is quite a large range for the number of neighbors, and global embedding dimension $d_W=d_E+1=5$. In embedding dimensions 4, 6, and 7 we have essentially identical results. This is an indication that $d_L=2$ is reliable.

Next we consider the three-dimensional model proposed by Lorenz as a "tiny" representation of the circulation of the atmosphere in a hemisphere [11]:

$$\begin{aligned}\dot{x}&=-y^2-z^2-a(x-F) ,\\ \dot{y}&=xy-bxz-y+G ,\\ \dot{z}&=bxy+xz-z .\end{aligned} \tag{13}$$

We use the values $a=0.25$, $b=4.0$, $F=8.0$, and $G=1.0$ where Lorenz points out irregular behavior is encountered. The attractor has a dimension d_A slightly greater than 2.5. Global false nearest neighbors suggests that at $d_E=4$ one has less than 0.5% false nearest neighbors, and thus this would be a good value for d_E. At $d_E=5$ this number is less than 0.05%, so a cautious person might choose $d_E=5$. The sufficient condition from the embedding theorem is $d_E=6$. We have chosen to work with $d_W=6$, and in Fig. 9 we show $P_K(d_L)$ for this model as a function of d_L and with varying N_B. From average mutual information considerations we have $T=23\tau_s$. This was our criterion for a good or poor prediction time. Clearly the method chooses $d_L=3$. Again, changes in embedding dimension made no difference.

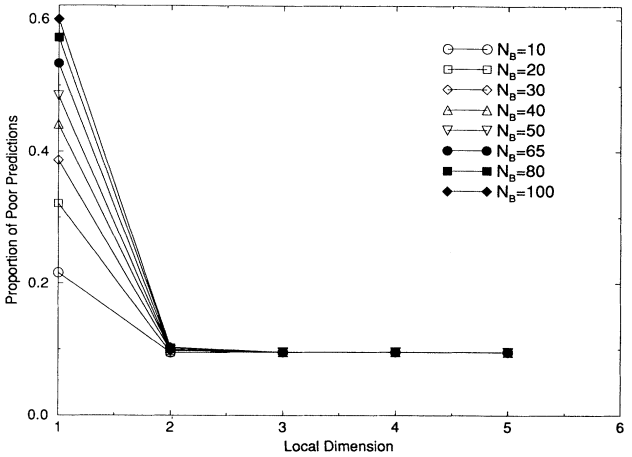

FIG. 8. $P_K(d_L)$ for 20 000 points of the real part of $z(n)$ from the Ikeda map for various $10\le N_B\le 100$; $d_W=d_E=5$. $d_L=2$ is clearly chosen as the dimension where independence of d_L and N_B occurs for $P_K(d_L)$.

Returning to the Lorenz [8] three-dimensional model of convection, we show in Fig. 10 the proportion of poor predictions as a function of d_L for various N_B and in $d_W=5$. Depending on how much one expands the verti-

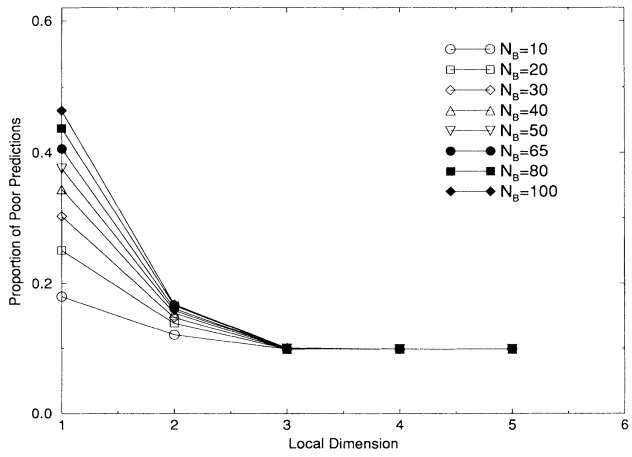

FIG. 9. $P_K(d_L)$ for 20 000 points of $x(t)$ from the Ref. [11] Lorenz model for various $10\le N_B\le 100$; $d_W=d_E=6$. $d_L=3$ is clearly chosen as the dimension where independence of d_L and N_B occurs for $P_K(d_L)$.

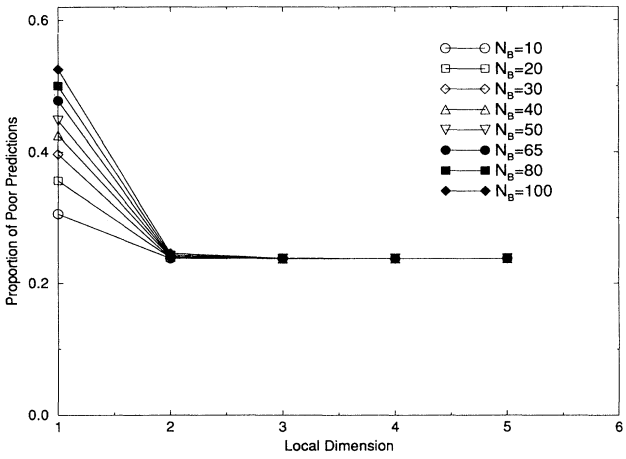

FIG. 10. $P_K(d_L)$ for 20 000 points of $x(t)$ from the Ref. [8] Lorenz model for various $10 \leq N_B \leq 100$; $d_W = d_E = 5$. From this figure, which uses data read forward in time, one might choose $d_L = 2$.

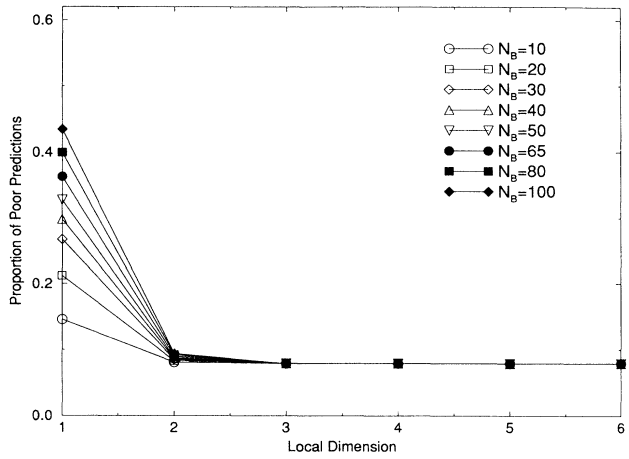

FIG. 12. $P_K(d_L)$ for 20 000 points of a voltage from the hysteretic circuit of Pecora and Carroll [13] for various $10 \leq N_B \leq 100$; $d_W = d_E = 6$. $d_L = 3$ is clearly chosen as the dimension where independence of d_L and N_B occurs for $P_K(d_L)$. This data set is read forward in time.

cal scale, this figure might indicate choosing $d_L = 2$ contrary to what the simple method above would have suggested. This is resolved by looking at precisely the same data set *backward* in time, as shown in Fig. 11. Since the predictability forward and backward in the correct d_L should be the same, we argue that the apparent indication of $d_L = 2$ for the forward data is due to the large magnitude of smallest Lyapunov exponent for this model $\lambda \approx -22.5$ [4]. This is also responsible for the nearness of d_A to 2. This suggests looking at data both forward and backward, as in [6,12], to firmly establish d_L. When this forward-backward method is applied to data from the Ikeda map for the Ref. [11] Lorenz model, no change in d_L from the forward results occurs.

As a final example we examine data from the hysteretic circuit of Pecora and Carroll [13]. The sampling time for this data was $\tau_s = 0.1$ ms, and from average mutual information we determine $T = 6\tau_s$. Global false nearest neighbors suggests an embedding dimension of $d_E = 5$, and we use here $d_W = 6$. In Figs. 12 and 13 we see the local false-nearest-neighbors determination of d_L both looking at the data forwards and backwards in time. It is clear in each case that $d_L = 3$ is selected without ambiguity. This is consistent with the models of the nonlinear circuit discussed by Carroll and Pecora. This example, in addition to being an analysis of laboratory data, is also our first discussion of contaminated data. There is a residual few percent noise level in the hysteretic circuit data, as one

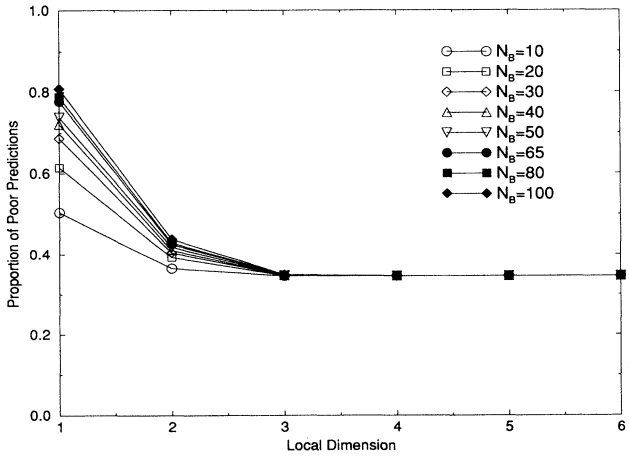

FIG. 11. $P_K(d_L)$ for 20 000 points of $x(t)$ from the Ref. [8] Lorenz model for various $10 \leq N_B \leq 100$; $d_W = d_E = 5$. From this figure, which uses data read backward in time, one clearly would select $d_L = 3$.

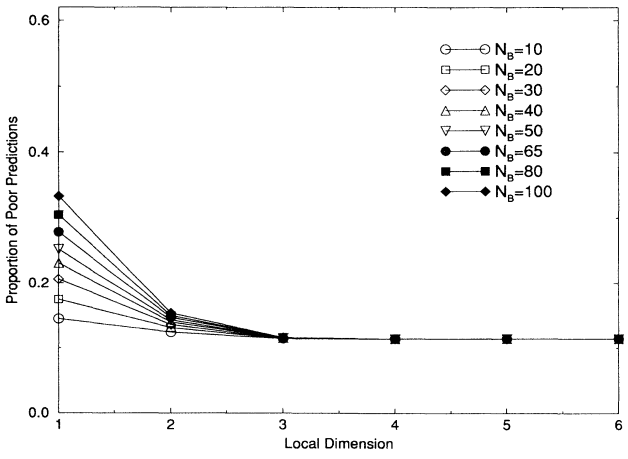

FIG. 13. $P_K(d_L)$ for 20 000 points of a voltage from the hysteretic circuit of Pecora and Carroll [13] for various $10 \leq N_B \leq 100$; $d_W = d_E = 6$. $d_L = 3$ is clearly chosen as the dimension where independence of d_L and N_B occurs for $P_K(d_L)$. This data set is read backward in time.

can observe by looking carefully at the global false-nearest-neighbor curves.

C. Contaminated data

We began our study of local false nearest neighbors because other methods [6,4,12] were so sensitive to contamination of the signal of interest, we felt we needed another tool. We have tested the local false neighbor method with varying amounts of noise in both the Lorenz model [8] and the Ikeda map. Of course, the hysteretic data comes automatically contaminated, as it is real laboratory data and not from a digital simulation.

With Lorenz model [8] data, contaminated with additive Gaussian noise with variance 1% that of the clean signal, we see in Fig. 14 $P_K(d_L)$ for various N_B. Again, $d_L=2$ is indicated from the data read forward in time, but looking at the data backward in time in Fig. 15, we see that $d_L=3$ is selected. The same level of contamination is shown for an analysis of Lorenz [11] data in Fig. 16. The choice of $d_L=3$ is not changed for this model at this noise level.

Using Ikeda map data contaminated by 1% uniform additive noise we have the results in Fig. 17. Once again the value $d_L=2$ is selected, showing the robustness of the method to noise. Other methods [6,12,4] fail at much lower noise levels, typically 0.01%. By the way, 1% global noise, which is what we are discussing here, can be 100% noise in one of our neighborhoods.

To investigate the deterioration of the method in the presence of substantial amounts of contamination, we show in Figs. 18–20 the effect of 2% noise on Ikeda map data ($d_L=2$ is still indicated), and the effect of noise levels up to 50% on Ikeda map data with two different choices of N_B. The unambiguous signal that $d_L=2$ remains until the noise level is 10%, but for noisier signals than that, we cannot determine a local dimension. This is actually a very good sign since at that level of *global* noise, we are swamping the local dynamics with unwanted contamination. Nonetheless, we conclude that for quite substantial noise levels, the local false-nearest-neighbor method which combines geometric ideas with a bit of dynamics is quite robust.

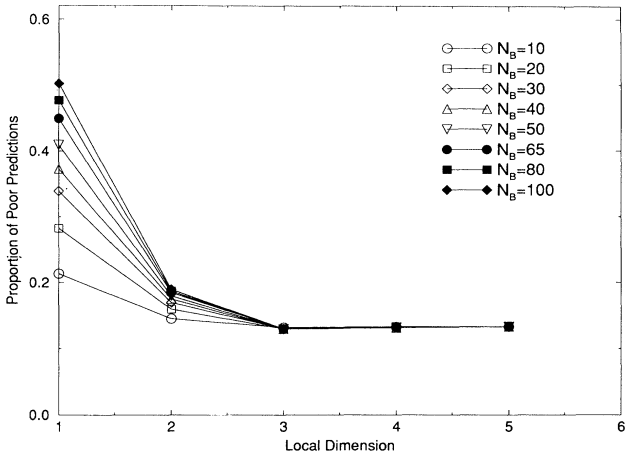

FIG. 15. $P_K(d_L)$ for 20 000 points of $x(t)$ from the Ref. [8] Lorenz model for various $10 \le N_B \le 100$; $d_W=d_E=5$. These data are contaminated with 1% noise relative to the global size of the attractor. From this figure, which uses data read backward in time, one clearly would select $d_L=3$.

We experimented using the proportion of poor predictions as an alternative criterion for finding the minimum embedding dimension, by embedding in successively larger dimensions, searching for the nearest neighbor (over all data) and computing the time for the two points to separate a "large" distance. Unlike the situation for local dimension, we did not see a definitively flat plateau once the correct dimension was used. Often we found a distinct local minimum in the proportion of poor predic-

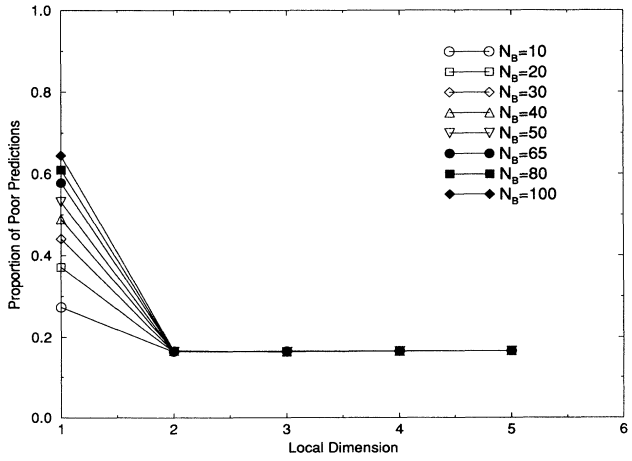

FIG. 14. $P_K(d_L)$ for 20 000 points of $x(t)$ from the Ref. [8] Lorenz model, for various $10 \le N_B \le 100$; $d_W=d_E=5$. These data are contaminated with 1% noise relative to the global size of the attractor. From this figure, which uses data read forward in time, one might choose $d_L=2$.

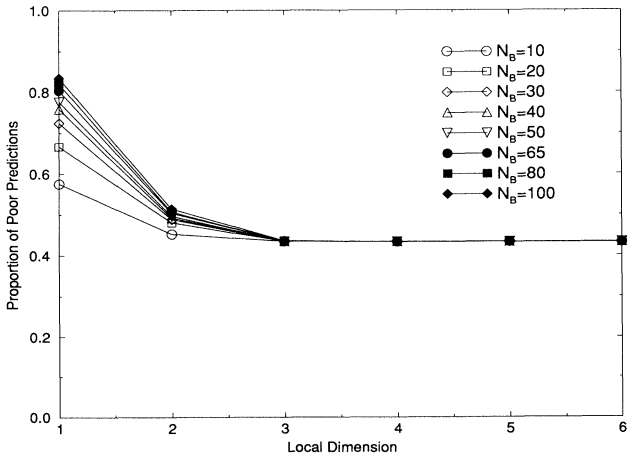

FIG. 16. $P_K(d_L)$ for 20 000 points of $x(t)$ from the Ref. [11] Lorenz model for various $10 \le N_B \le 100$; $d_W=d_E=6$. These data are contaminated with 1% noise relative to the global size of the attractor. $d_L=3$ is clearly chosen as the dimension where independence of d_L and N_B occurs for $P_K(d_L)$.

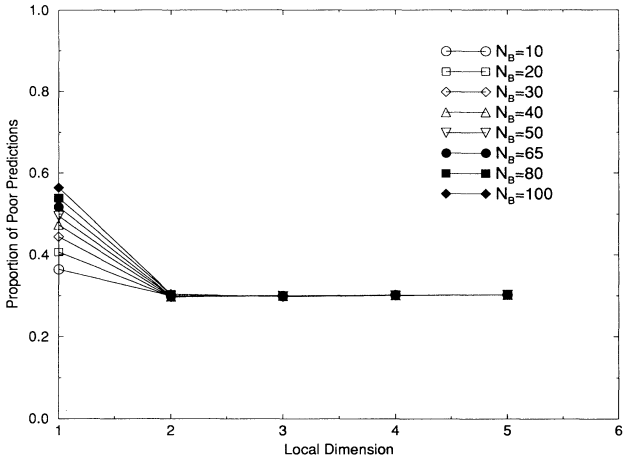

FIG. 17. $P_K(d_L)$ for 20 000 points of the real part of $z(n)$ from the Ikeda map for various $10 \leq N_B \leq 100$; $d_W = d_E = 5$. These data are contaminated with noise of size 1% relative to the global size of the attractor. $d_L = 2$ is clearly chosen as the dimension where independence of d_L and N_B occurs for $P_K(d_L)$.

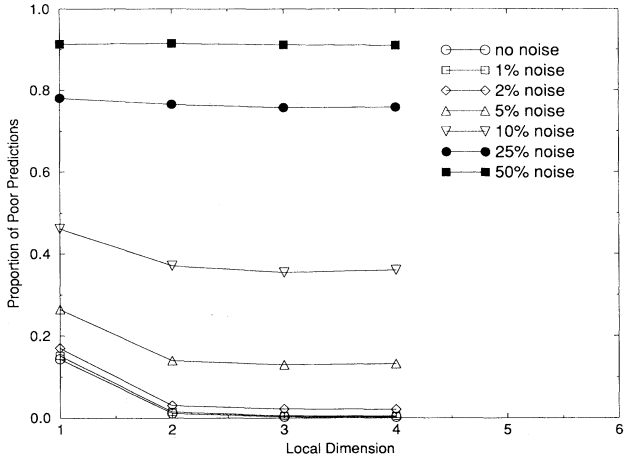

FIG. 19. $P_K(d_L)$ for 20 000 points of the real part of $z(n)$ from the Ikeda map with $d_W = d_E = 4$, $N_B = 100$, and varying amounts of noise added to the data. The percentage of noise is given relative to the global size of the attractor and ranges from 0% to 50%. When the noise level is 10% or less, $d_L = 2$ is clearly chosen. After that the method of local false nearest neighbors is overwhelmed by the contamination. In this figure $N_B = 100$ and $\beta = 0.5$.

tions occurring at reasonable values for embedding dimension, but in other cases, such as with the Ikeda data set, a distinct minimum occurred at a dimension known to be insufficient.

III. CONCLUSIONS

In this paper we examine two methods for establishing the local or dynamical dimension d_L of a system which we have observed solely through a scalar variable and whose state space we have reconstructed using time delay embedding. The sufficient dimension d_E required for that embedding is governed by the global, topological theorem of Mañé and Takens [14,15], which reveals that $d_E > 2d_A$, where d_A is the box-counting [3] dimension of the system attractor. The dimension d_E we use here is the necessary embedding dimension as established by the method of false nearest neighbors [1] directly from the data. The necessary d_E is less than or equal to the sufficient dimension and both are greater than or equal to the dimension of the local dynamics d_L we have sought in

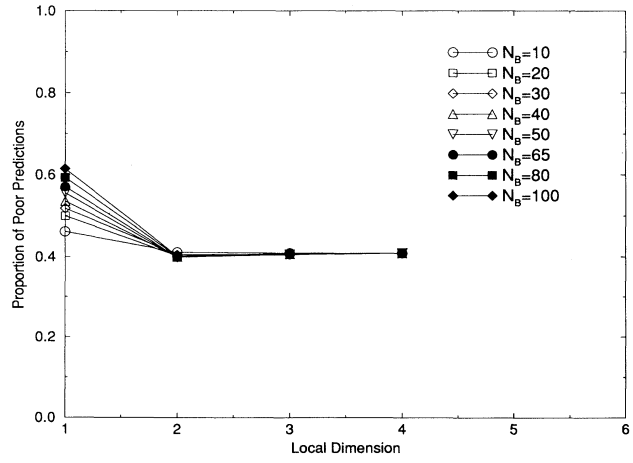

FIG. 18. $P_K(d_L)$ for 20 000 points of the real part of $z(n)$ from the Ikeda map for various $10 \leq N_B \leq 100$; $d_W = d_E = 5$. These data are contaminated with noise of size 2% relative to the global size of the attractor. $d_L = 2$ is clearly chosen as the dimension where independence of d_L and N_B occurs for $P_K(d_L)$.

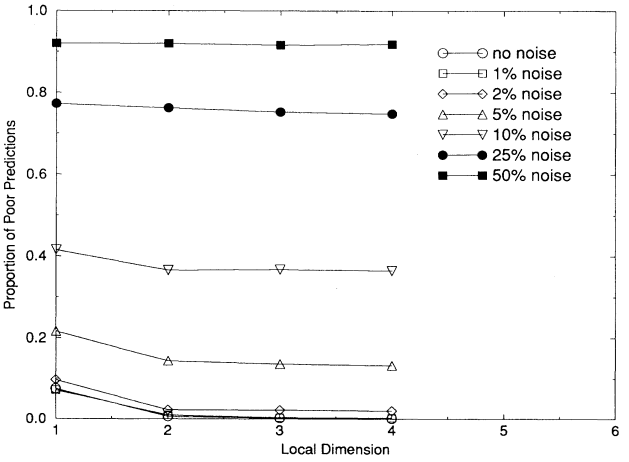

FIG. 20. $P_K(d_L)$ for 20 000 points of the real part of $z(n)$ from the Ikeda map with $d_W = d_E = 4$, $N_B = 40$, and varying amounts of noise added to the data. The percentage of noise is given relative to the global size of the attractor and ranges from 0% to 50%. When the noise level is 10% or less, $d_L = 2$ is clearly chosen. After that the method of local false nearest neighbors is overwhelmed by the contamination. In this figure $N_B = 40$ and $\beta = 0.5$.

this paper.

We know of a few earlier (related) methods to determine d_L. First is the work of Parlitz [6], who implemented a suggestion (not known to him) of Eckmann and Ruelle [2]. He looks at the global Lyapunov exponents of the system forward and backward in time. True Lyapunov exponents, of which there are d_L, will reverse in sign under this operation of time reversal. Fake or spurious exponents due solely to the choice $d_E > d_L$ will do something else. Abarbanel and Sushchik [12] examined Parlitz's work by looking at local Lyapunov exponents [16,17], which give a repeated look at the effect of time reversal and allow one to establish whether the zero exponent present for a flow is true or spurious. No local Lyapunov exponent for a flow is necessarily exactly zero, except in the limit as the number of observation steps goes to infinity (recovering the global Lyapunov exponent), so one can determine when an exponent slightly different from zero flips sign under time reversal, removing the puzzle of identifying the sign reversal of the global zero exponent.

Another method described in Ref. [12] focuses on the more fundamental property being examined in the time reversal of Lyapunov exponents: determinism of the data set forward and backward in time. This method asks when local maps from neighborhood to neighborhood forward and backward in time in the reconstructed state space are accurately inverses of one another. The actual calculation uses polynomial local maps, but any parametric form would do.

In the examples presented in these papers each method is remarkably successful when applied to clean, contamination free data. Since much of the data are simulated, it can be made very clean indeed. Unfortunately, as soon as contamination is added to the data, even at the very modest level of 0.1% of the global size of the attractor, the method gives spurious results. The origin of these spurious results can be traced to the necessity in each case of determining distances between points on the attractor with great accuracy. Small amounts of noise will spoil one's ability to establish these distances, and since the computations typically involve properties of rather ill-conditioned matrices, reliable results depend sensitively on the accurate determination of all quantities entering these matrices.

The present hunt for d_L avoids these methods which require the accurate measurement of local distances, and time-consuming or potentially unreliable local parametric fitting. We looked at two approaches in this paper: (1) a simple local version of the global false neighbor technique, and (2) an improved variant using as estimator of the predictability of the observed dynamics forward and backward in time. Both are robust against noise contamination, but the second, by probing how neighborhoods evolve in dimensions less than or equal to d_E is more accurate, but even more importantly, gives a quite unambiguous indication of the correct value of d_L. We applied the methods to several examples: the Ikeda map of the plane to itself, to two three-dimensional systems of ordinary differential equations suggested by Lorenz as models of various properties of the atmosphere, and to experimental data from observations on a nonlinear circuit by Carroll and Pecora. In three of the examples (Ikeda map, Lorenz model in Ref. [11], and the experimental data) $d_E > d_L$, so the challenge in identifying the active degrees of freedom is precisely within the scope of this paper. In the case of the Lorenz model in Ref. [8], $d_E = d_L$, but this too is produced by the method. Furthermore, and central to the goal of this work, the method works in a clear fashion even up to the order of 10% contamination of the data by noise. This is a level of global noise, so the contamination at the local level is substantially more than that. The concepts presented in the paper are simple, and it would seem that a variety of statistical tests should work well. In fact, we found it surprisingly difficult to find tests that reliably produced clear and correct answers in our model examples (achieving $d_L = 2$ for Ikeda map data, and $d_L = 3$ for Ref. [8] Lorenz data is challenging) with both clean and noisy signals.

We know of some earlier attempts to establish d_E and d_L using essentially geometric ideas. The earliest, as far as we can ascertain, is due to Froehling, *et al.* [18], who determine a quantity close to what we call d_L in this paper. They identify neighborhoods of each phase space point $\mathbf{y}(n)$ in d_E and then seek the best local hyperplane on which the points in the neighborhood of $\mathbf{y}(n)$ lie. The criterion of "best" is that of a local least-squares fit of the data to the hyperplane and then an evaluation of the residuals in the data relative to all points lying in the hyperplane. They begin with "test" hyperplanes of dimension zero, namely points, then of dimension one, namely lines, etc. These residuals (called χ^2 in their paper) decrease quite rapidly once the dimension of the local hyperplane has exceeded what we call d_L here. The method is quite reminiscent of the first approach we describe in this paper.

The method proposed by Broomhead, Jones, and King [19] attempts to discern d_L by examining the scaling of the singular values of a $N_B \times d_E$ matrix consisting of the vector differences from the reference point, as a function of the radius ϵ considered in identifying neighborhoods. These singular values are the square roots of the eigenvalues of the local sample covariance matrix that we use in computing a local coordinate system. Essentially the method of Ref. [19] examines how well the data fit a local d_L-dimensional hyperplane. For fractal attractors, identifying the correct local dimension from the plots of $\sigma(i)^{1/2}$ vs ϵ [they choose $|\mathbf{y}^{(r)} - \mathbf{y}(k)| < \epsilon$, and vary ϵ, not N_B] does not always appear to be clear-cut, and it appears that noise could spoof this method on a local scale. The method only examines geometrical properties of the data, and combining the scaling results for different neighborhoods into a single value for d_L is not automated.

It is amusing to note that the paper of Froehling *et al.* [18] just cited contains, almost as a throwaway comment, the basic idea of the singular value method. See their Sec. 4 [18].

The work of Kaplan and Glass [20] is another attempt to identify properties of the dynamics when only scalar observations are available which does not require precise

computation of distances on the attractor. They examine the effective local vector field of the dynamics seeking a dimension, which we would call d_E, where crossings of orbits as determined by the direction of the local vector field in reconstructed state space cease. There is probably a local version of their method which would establish d_L, but we have not pursued that.

There are two final considerations we wish to make. First, we want to touch on the importance of knowing d_L, and second, we want to muse on the meaning of d_L when the underlying system producing the observations is very high dimensional. Clearly this kind of qualitative issue has been the main motivation for the work presented here. d_L is the number of degrees of freedom of the underlying dynamical system which are activated in the determination of the orbits on the observed attractor. If one evaluates properties of the system which depend on the dimension of the state space in which the reconstruction occurs, the Lyapunov exponents probably being the most important physically because of their connection with the predictability of the system, then in dimensions larger than d_L, false exponents will appear. There are examples [4] where one of the false exponents can be positive and can even be the largest positive exponent, so the misinterpretation of the significance of a large positive Lyapunov exponent is possible. If we have established d_L, then all exponents, whether we know them very accurately or not, as might be the situation in the presence of noise, are real. Further if one wants to build models of the observed system, then working in the minimal number of dimensions, namely d_L, will reduce the amount of work in establishing and verifying that model as well as reduce its sensitivity on noise which would occupy without any dynamical significance any dimension offered it.

In the case of analysis of a continuum system such as fluid dynamics or a very-high-dimension system such as represented by general circulation models of the climate or weather, the number of variables in the basic differential equations may be in the millions or, in principle, actually infinite for partial differential equations. A finite hopefully small d_L then represents the number of degrees of freedom active on the attractor, namely the number which one must model to produce an accurate representation of the orbits of the system for prediction or possible for control purposes. In Ref. [12] for very clean data this was demonstrated on a delay differential equation where an infinite number of degrees of freedom are in the basic equation. In some qualitative sense one can think of the "other" degrees of freedom as frozen out of the observations by their association with large, negative Lyapunov exponents which drive them onto the attractor and thus, since the data are on the attractor alone, make them in some practical sense unobservable. While a more accurate quantitative characterization of this view is not known to us, it may suffice in a pragmatic sense to establish d_L and proceed.

ACKNOWLEDGMENTS

We thank the members of INLS for numerous discussions on this subject; the comments and assistance of J. J. Sidorowich were especially helpful. We are most appreciative to Tom Carroll and Lou Pecora of the U.S. Naval Research Laboratory for allowing us to use their data on the hysteretic circuit and for numerous discussions about that data. This work was supported in part by the U.S. Department of Energy, Office of Basic Energy Sciences, Division of Engineering and Geosciences, under Contract No. DE-FG03-90ER14138, and in part by the Army Research Office (Contract No. DAAL03-91-C-052), and by the Office of Naval Research (Contract No. N00014-91-C-0125), under subcontract to the Lockheed/Sanders Corporation.

[1] Matthew B. Kennel, R. Brown, and H. D. I. Abarbanel, Phys. Rev. A **45**, 3403 (1992).
[2] J.-P. Eckmann and D. Ruelle, Rev. Mod. Phys. **57**, 617 (1985).
[3] M. Casdagli, T. Sauer, and J. A. Yorke, J. Stat. Phys. **65**, 579 (1991).
[4] R. Brown, P. Bryant, and H. D. I. Abarbanel, Phys. Rev. A **43**, 2787 (1991); Phys. Rev. Lett. **65**, 1523 (1990).
[5] K. Briggs, Phys. Lett. A **151**, 27 (1990).
[6] U. Parlitz, Int. J. Bif. Chaos **2**, 155 (1992).
[7] A. M. Fraser and H. L. Swinney, Phys. Rev. A **33**, 1134 (1986); A. M. Fraser, IEEE Trans. Inf. Theory **35**, 245 (1989); A. M. Fraser, Physica D **34**, 391 (1989).
[8] E. N. Lorenz, J. Atmos. Sci. **20**, 130 (1963).
[9] K. Ikeda, Opt. Commun. **30**, 257 (1979).
[10] S. M. Hammel, C. K. R. T. Jones, and J. V. Moloney, J. Opt. Soc. Am. B **2**, 552 (1985).
[11] E. N. Lorenz, Tellus A **36**, 98 (1984).
[12] H. D. I. Abarbanel and M. M. Sushchik, Int. J. Bif. Chaos (to be published).
[13] T. L. Carroll and L. M. Pecora, IEEE Trans. Circ. Syst. **38**, 453 (1991).
[14] R. Mañé, in *Dynamical Systems and Turbulence, Warwick, 1980,* edited by D. Rand and L. S. Young, Lecture Notes in Mathematics Vol. 898 (Springer, Berlin), (1981), p. 230.
[15] F. Takens, in *Dynamical Systems and Turbulence, Warwick, 1980* (Ref. [14]), 366.
[16] H. D. I. Abarbanel, R. Brown, and Matthew B. Kennel, J. Nonlin. Sci. **1**, 175 (1991).
[17] H. D. I. Abarbanel, R. Brown, and Matthew B. Kennel, J. Nonlin. Sci. **2**, 343 (1992).
[18] H. Froehling, J. P. Crutchfield, J. D. Farmer, N. H. Packard, and R. Shaw, Physica D **3**, 605 (1981).
[19] D. S. Broomhead, R. Jones, and G. P. King, J. Phys. A **20** L563 (1987).
[20] D. T. Kaplan and L. Glass, Phys. Rev. Lett. **68**, 427 (1992).