Extracting dynamical structure from unstable periodic orbits

Kevin T. Dolan
Center for Neurodynamics, University of Missouri at St. Louis, St. Louis, Missouri 63121
(Received 13 December 2000; published 19 July 2001)

The topological recurrence algorithm provides a fast and robust method for detecting the presence of unstable periodic orbits (UPO's) in short, noisy experimental data files. We present here a technique for improving this method by using a matrix fitting algorithm to extract dynamical information about the system from these UPO's. This method greatly increases the sensitivity of the algorithm, and also provides a method for identifying false positive results.

DOI: 10.1103/PhysRevE.64.026213 PACS number(s): 05.45.−a, 05.10.−a

I. INTRODUCTION

Identifying the presence of unstable periodic orbits (UPO's) in experimental data has become a powerful tool for the analysis of noisy physical and biological systems [1–3]. A simple statistical method for finding UPO's known as the topological recurrence method [4] has been developed for the analysis of noisy time series data. This method has been demonstrated on a noisy, periodically forced Van der Pol oscillator [4], as well as on biological data from the crayfish caudal photoreceptor system [5,6], and thermally sensitive neurons [7,8]. The topological recurrence method works by specifying and searching for a pattern that is indicative of an encounter with a UPO. These possible encounters with UPO's are counted and compared with surrogate data files in order to assess the statistical probability that the number of encounters found could be the result of random chance.

The algorithm discussed here operates on a discrete time series, such as time intervals between the firings of a neuron [9,10], a Poincaré section embedding from a continuous flow, or even discrete maps such as the Hénon or logistic map. The description of this algorithm will be limited to period-1 fixed points, although the algorithm can be applied to higher period orbits as well. First, the data are presented as a return map (T_n vs T_{n+1}). This return map is a two-dimensional projection of an embedding of the time series data. An example of data from the rat facial cold receptor [8] is shown in Fig. 1.

The 45° line of periodicity shows where consecutive points are nearly equal to each other. In the vicinity of an unstable fixed point the system can often be approximated as a two-dimensional mapping [11]. The pattern we search for is defined as follows: Three consecutive points approach the line of periodicity with sequentially decreasing perpendicular distance, then three points diverge from the line of periodicity with sequentially increasing perpendicular distances. The third point is shared, making a total of five points on the return map, or six data points. An example of such an encounter with a UPO is shown in Fig. 2.

The number of such encounters found is counted. To determine whether the number of encounters found is larger than what would be expected due to chance, surrogate data files are generated and the number of encounters found in them is also counted. The statistical significance is determined with the following measure,

$$K = \frac{N - \overline{N}_S}{\sigma_S}, \qquad (1)$$

where N is the number of encounters in the original data file, $\overline{N}_S$ is the average number of encounters found in the surrogates, and σ_S is the standard deviation of that average. Assuming Gaussian statistics, which should be valid for values of $\overline{N}_S$ greater than about 20 [12], $K > 2$ indicates about a 95% confidence and $K > 3$ indicates about a 99% confidence [13]. If shuffled surrogates are used, we can reject the null hypothesis that the data is an uncorrelated process. Using surrogates that preserve the power spectrum of the original data, such as the amplitude adjusted Fourier transformed (AAFT) surrogates [14], allows the rejection of the stronger null hypothesis that the data set being analyzed is a linear stochastic process passed through a static nonlinear transformation [14,15]. Analysis of several types of colored noise [12], using both shuffled surrogates and AAFT surrogates has demonstrated that the number of false encounters typically found in data sets in which no real periodic orbits are

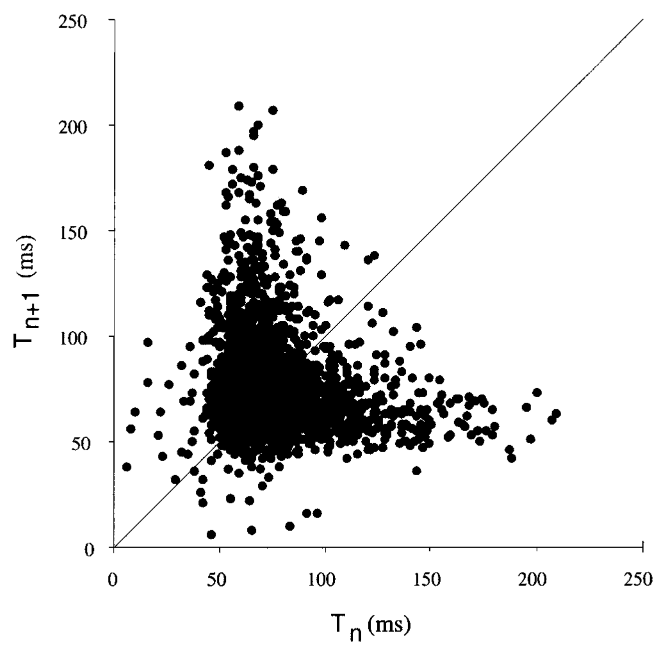

FIG. 1. Return map of interspike time intervals from the rat facial cold receptor [8].

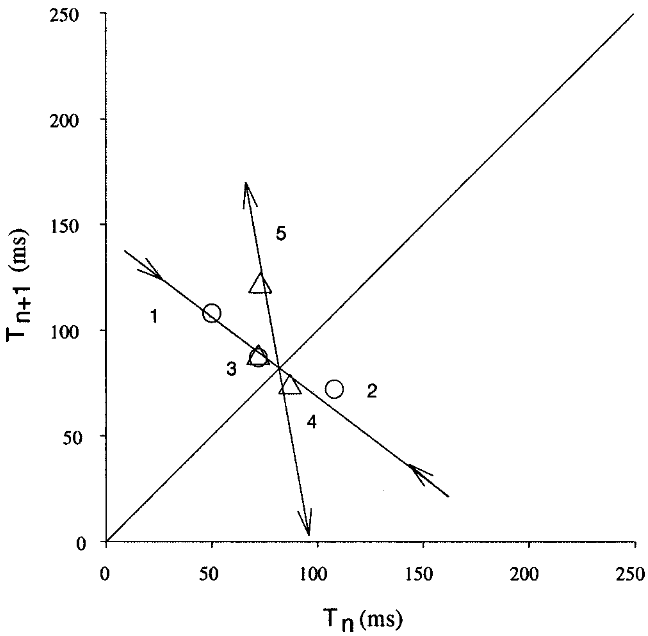

FIG. 2. A typical encounter from the data file shown in Fig. 1. Three points converge towards the line of periodicity (circles), followed by three points diverging (triangles). The third point is shared.

present typically does not depend on the power spectrum of the data. Shuffled surrogates are thus often used for this type of analysis, but it should be noted that technically rejection of this stronger hypothesis requires that either power spectrum preserving surrogates be used, or that the power spectrum is explicitly shown not to have an effect in that particular case.

Since the pattern we have searched for is what we would expect to see in the presence of UPO's, a rejection of the null hypothesis is also a very strong indication that real UPO's are present in the data. In principle it is possible that other types of nonlinear correlations could influence the probability of finding false encounters, thus allowing for the possibility of a positive statistic when no UPO's are present.

At this point two relevant questions present themselves. First, how can we improve the ability of this method to detect UPO's in situations where the surrogate test described above is inconclusive, and second, how can we assure ourselves that we are seeing real UPO's when we do find a significantly positive statistic?

II. EXTRACTING STRUCTURE FROM THE ENCOUNTERS

One approach to answering these questions is to look more closely at the encounters that have been found. The data points that make up the individual encounters have, in principle, all of the information needed to determine the location of the UPO, as well as the eigenvalues and eigenvectors associated with it. This information can be used to improve the ability of the method to detect UPO's by rejecting encounters that do not match certain criteria that a real UPO encounter would be expected to satisfy. The distributions of these fixed points and eigenvalues can be used to establish that the encounters found in the original data really are encounters with UPO's and not the result of some other type of correlations.

In the vicinity of a UPO the system will often behave approximately as a two-dimensional linear mapping [11],

$$T_{n+1}=A_1T_n+A_2T_{n-1}+B. \tag{2}$$

The coefficients A_1, A_2, and B can be determined by doing a linear least-squares fit on the points making up all of the UPO encounters from the file. The location and eigenvalues of the UPO are then given by

$$\lambda_1=\frac{1}{2}(A_1+\sqrt{A_1^2+4A_2}),$$

$$\lambda_2=\frac{1}{2}(A_1-\sqrt{A_1^2+4A_2}),$$

$$T_F=\frac{B}{1-A_1-A_2}, \tag{3}$$

where λ_1 and λ_2 are the eigenvalues of the mapping, and T_F is the location of the fixed point. This method has been used in the control of chaos [11,16,17], where data near the vicinity of a UPO is used to determine the location and eigenvalues of the orbit. In this case the topological recurrence algorithm simply provides a method for determining the data to which the fit should be applied.

This linear map fitting algorithm provides an accurate and reliable method for determining the location and eigenvalues of UPO's, but it does so by using all of the data from every encounter in the file and finding the best fit possible. This is very useful in situations where the system is known to be a dynamical system with only one UPO of a given period, but

TABLE I. Comparison of methods on the Hénon map.

	N	$\overline{N}_S$	σ_S	K
Standard TR method	838	299.73	13.74	39.18
Matrix fit	838	147.22	11.21	61.62

TABLE II. Comparison of methods on the Rössler system.

	N	$\overline{N}_S$	σ_S	K
Standard TR method	419	278.45	15.15	9.28
Matrix fit	379	135.81	10.07	24.14

TABLE III. Comparison of methods on a short data file.

	N	$\overline{N}_S$	σ_S	K
Standard TR method	68	56.05	5.51	2.17
Matrix fit	62	26.40	4.57	7.79

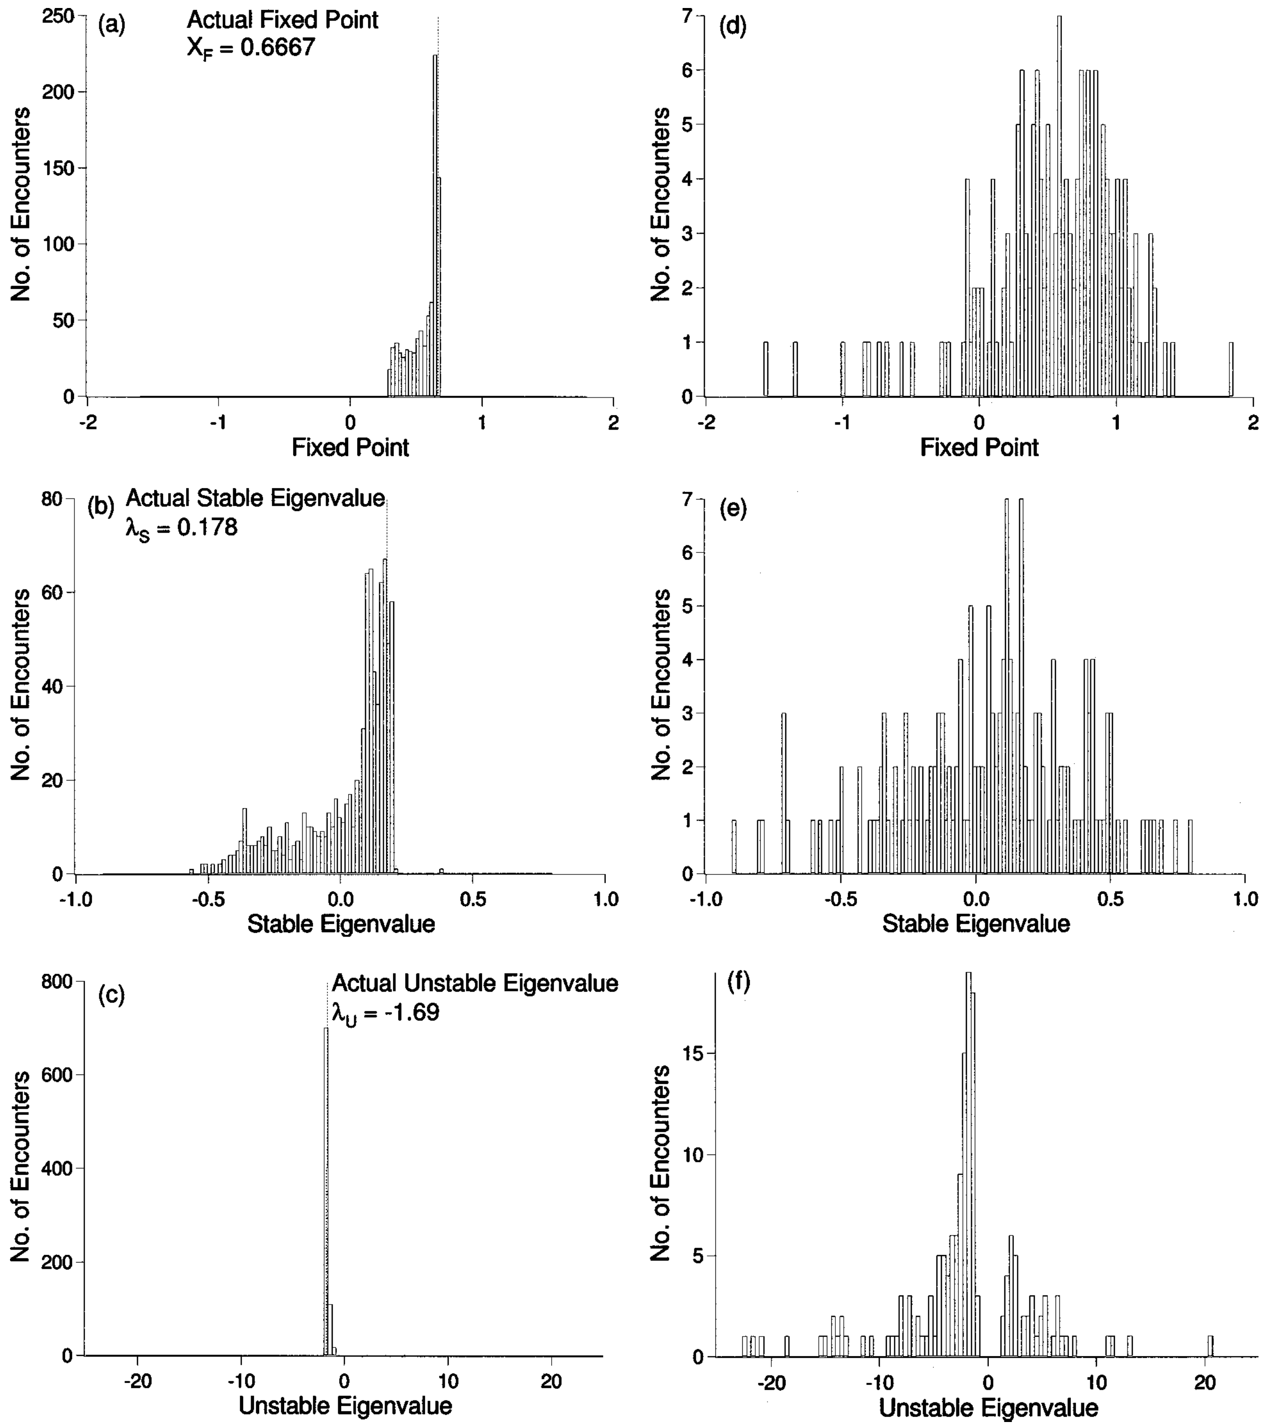

FIG. 3. Distribution of fixed points and eigenvalues of the encounters from the Hénon map in the chaotic regime (a)–(c), and one of its surrogates (d)–(f). The analytically calculated values are shown with vertical dotted lines in (a)–(c).

is not very helpful in establishing whether low-dimensional dynamics are present in a data file in the first place.

One possibility for improving the statistical significance with which UPO's can be detected would be to calculate the location and eigenvalues of each encounter found and then use this information to decide whether or not the encounter represents a real UPO. A good first step is to recognize that

TABLE IV. TR Analysis of the Hénon map (stable period 1).

	N	$\overline{N_S}$	σ_S	K
Standard TR method	419	303.43	13.24	8.73
Matrix fit	270	151.92	10.96	10.78

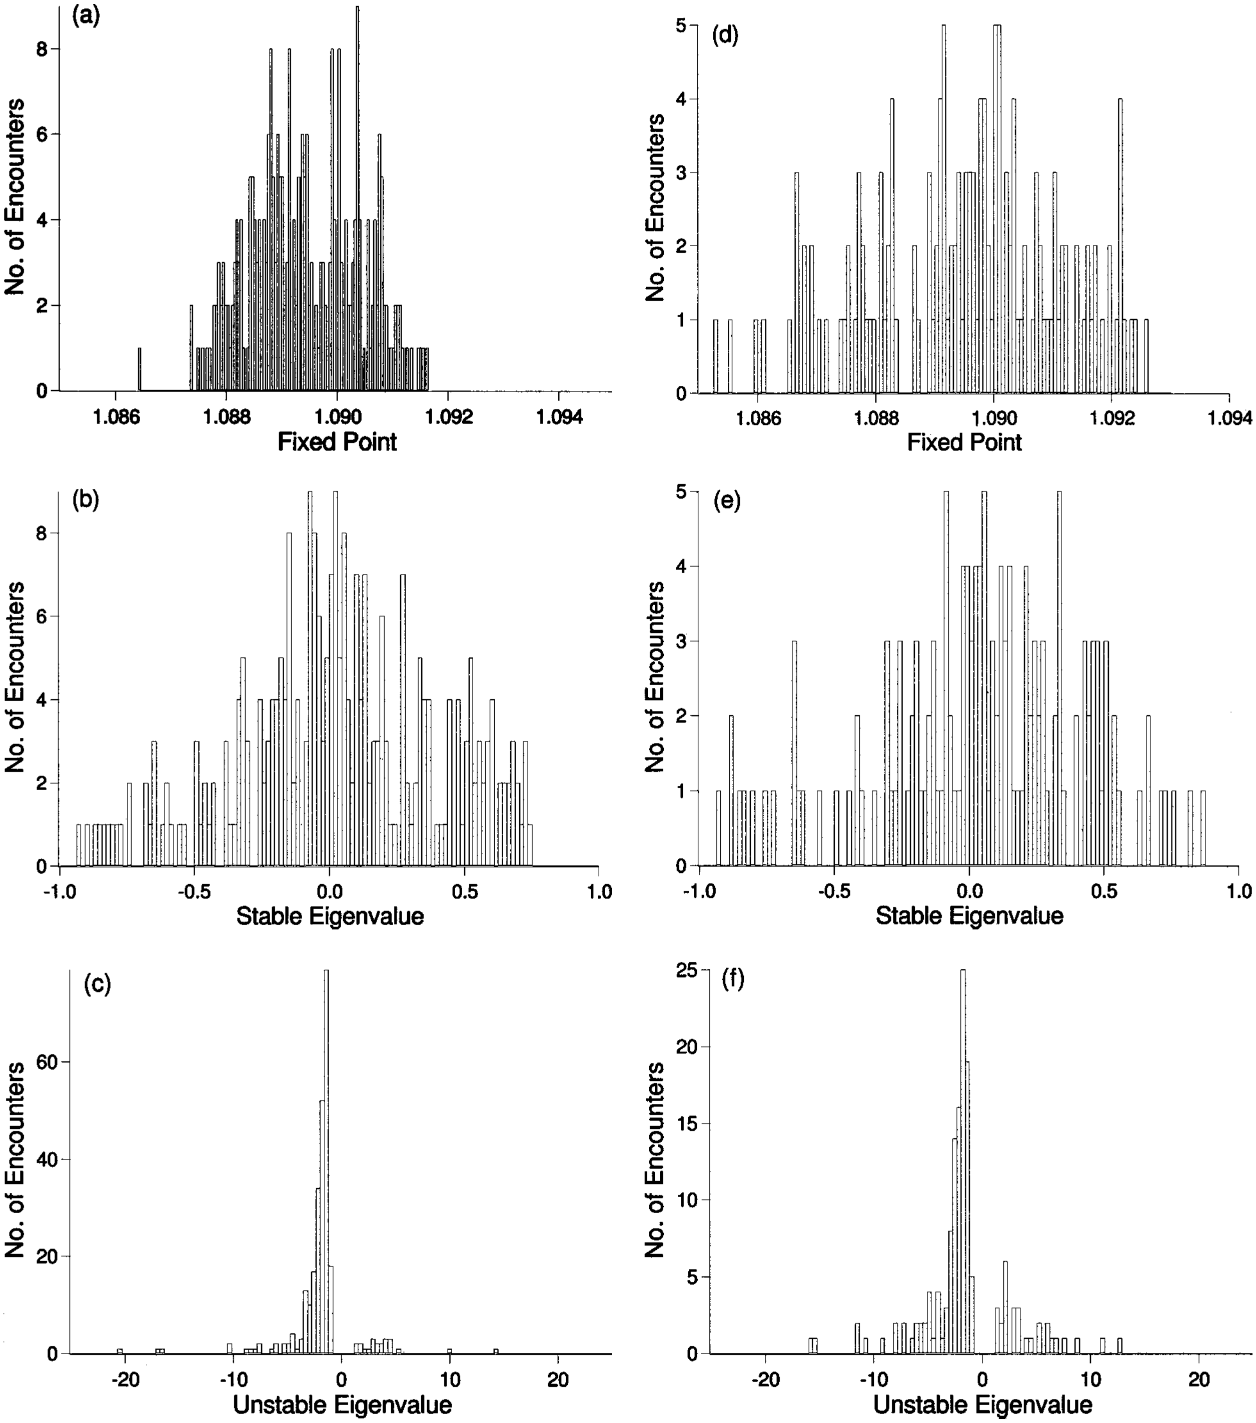

FIG. 4. Distribution of fixed points and eigenvalues of the encounters from the Hénon map in the stable period-1 regime (a)–(c), and one of its surrogates (d)–(f).

although the encounter implies convergence towards and divergence from the line of periodicity, it does not guarantee that the linear fit will come up with a stable and unstable eigenvalue. In fact, we find that approximately half of the false encounters from surrogate data files produce either two stable or two unstable eigenvalues. These encounters clearly do not represent real UPO's and can be rejected. We will also reject any encounters in which the location of the fixed point estimated by the linear fit does not lie near the points that actually make up the encounter.

The algorithm used to do this can be summarized as follows: First the data file is searched for sequences matching the criteria listed in Sec. I. For each encounter, the fixed point and both eigenvalues are calculated. Recall that in Eq. (2) there are three unknowns; A_1, A_2, and B. This means that we need three equations. Our encounter is made up of four

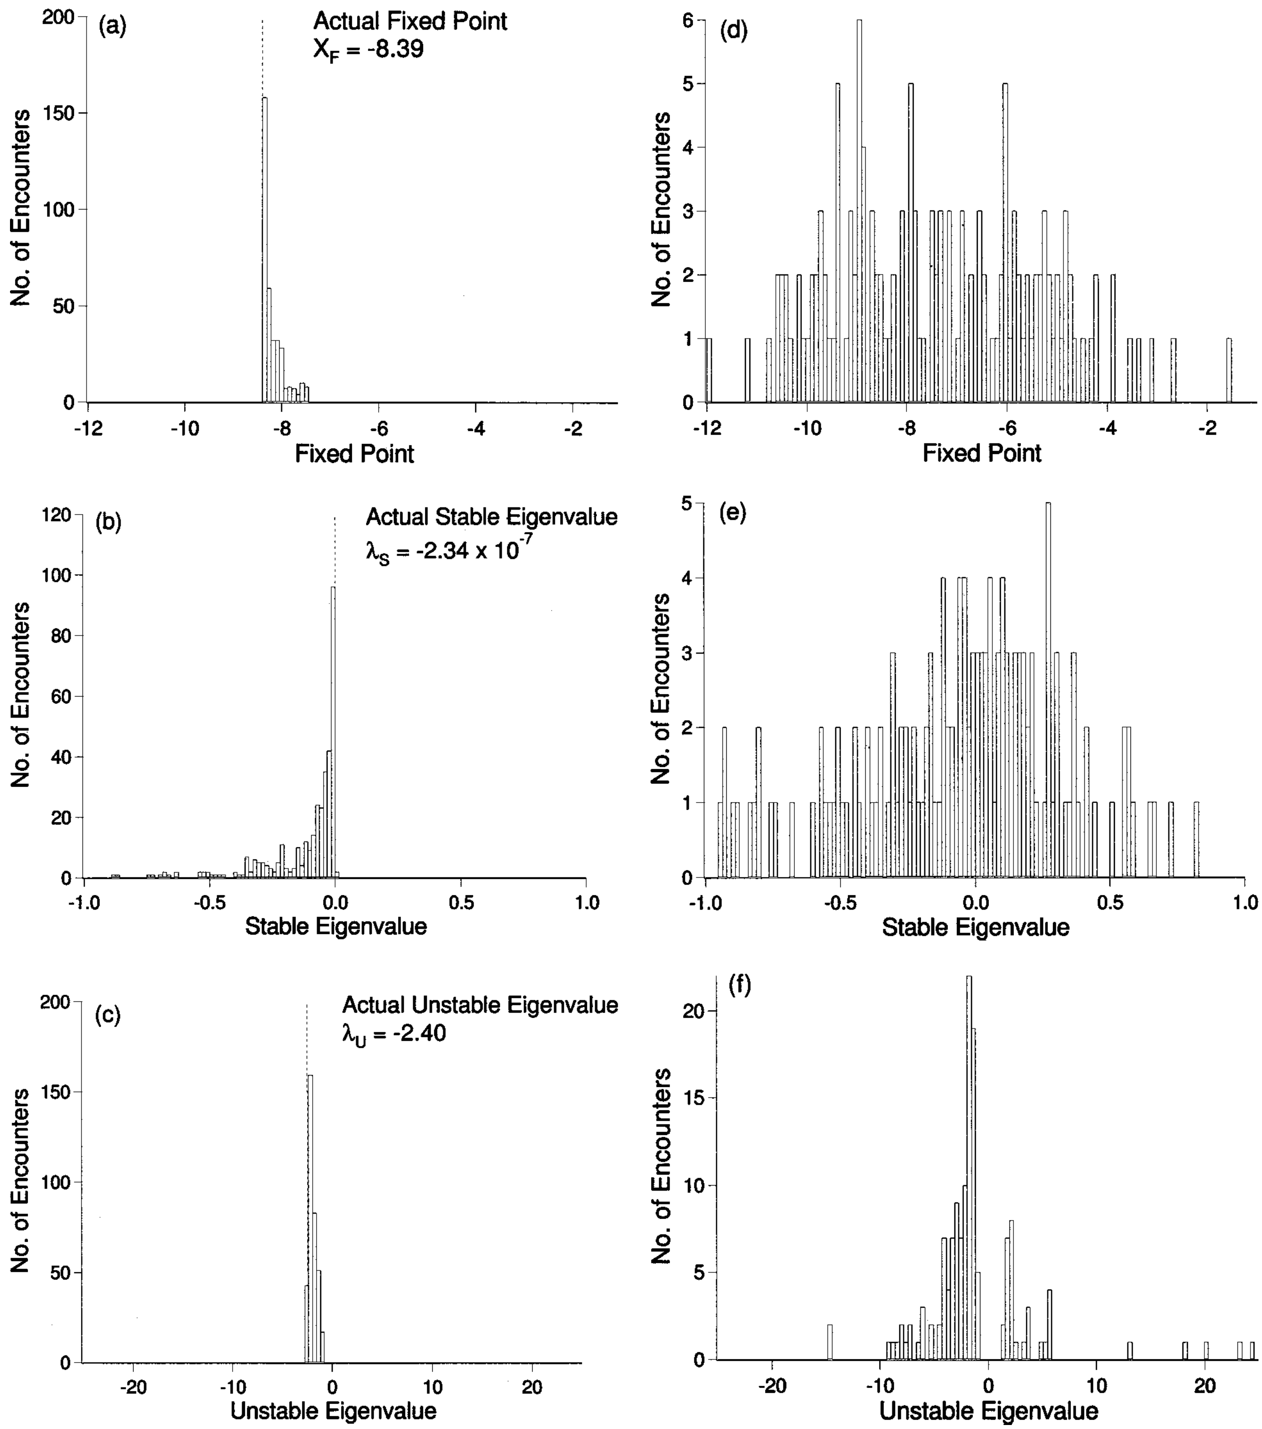

FIG. 5. Distribution of fixed points and eigenvalues of the encounters from the Rössler system in the chaotic regime (a)–(c), and one of its surrogates (d)–(f). The analytically calculated values are shown with vertical dotted lines in (a)–(c).

iterations of Eq. (2), but we only need three to solve for the unknowns. Rather than doing a least-squares fit to determine the unknowns from all four iterations, we choose instead to just use the last three iterations. The reason for this choice is that in practice the system will typically converge towards the UPO very quickly, and then diverge more gradually. The first iteration is thus often not well represented by the local linear fit anyway. The three equations used to perform our matrix fit are thus

$$T_{i+3}=A_1T_{i+2}+A_2T_{i+1}+B,$$

$$T_{i+4}=A_1T_{i+3}+A_2T_{i+2}+B,$$

$$T_{i+5}=A_1T_{i+4}+A_2T_{i+3}+B, \tag{4}$$

where T_i is the first point in the encounter. Note that the first transition in the encounter is from T_{i+1} to T_{i+2}, and is not used here. This also means that the point T_i is not used in the

TABLE V. TR Analysis of the Rössler system (stable period 1).

	N	$\overline{N}_S$	σ_S	K
Standard TR method	466	306.75	14.36	11.09
Matrix fit	279	152.03	10.98	11.57

fit at all. Once A_1, A_2, and B have been determined the fixed point and eigenvalues can be found using Eq. (3). If the encounter has two stable eigenvalues ($|\lambda_j|<1$) or two unstable eigenvalues ($|\lambda_j|>1$), then we reject it. Furthermore, if the location of the fixed point T_F is too far from the point of nearest approach (T_2,T_3) we also reject it. Specifically, if

$$\left|T_F-\frac{T_2+T_3}{2}\right|>\sigma_T, \tag{5}$$

where σ_T is the standard deviation of the distribution of the original time series. This is to eliminate encounters where some quirk of the linear fit causes the estimated fixed point to be some very large, and meaningless number.

If real UPO's are present in the original data, then we would expect most of the encounters to satisfy these requirements. For the surrogate data, however, all encounters found are false encounters. By reducing the number of encounters found in the surrogate data without eliminating any real encounters from the original data, we should greatly increase the value of the K statistic from Eq. (1). To demonstrate this improvement we present the noisy Hénon map,

$$X_{n+1}=bY_n-aX_n^2+1,$$

$$Y_{n+1}=X_n+\epsilon\xi_n, \tag{6}$$

where ξ_n is a uniform random deviate from -1 to $+1$, and ϵ is the intensity of the noise. With the parameters $b=0.30$ and $a=1.2$ this map exhibits chaotic behavior. A 5000 point data file was generated with $\epsilon=0.01$. The value of the variable X was recorded at each iteration, and the resulting file was analyzed using 100 shuffled surrogate files. Table I shows the results for both the initial criteria of convergence and divergence, as well as for the new requirements of the linear fit.

We see that although all of the encounters from the original data satisfy the requirements imposed by the linear fit, about half of the encounters found in the surrogate files are rejected. The statistic K is therefore much higher for the linear fit algorithm than for the original TR algorithm. This effect can also be easily seen when analyzing the Rössler system [18],

$$\dot{x}=-(y+z),$$

$$\dot{y}=x+ay,$$

$$\dot{z}=b+z(x-c)+\sqrt{2D}\xi(t), \tag{7}$$

where $\xi(t)$ is Gaussian white noise with zero mean and unit variance. The parameter D is the amplitude of this dynamical noise. For $a=b=0.2$, $c=5.7$ this system is chaotic. The topological recurrence algorithm can only operate on time series data, so a Poincaré section is made by slicing the three-dimensional phase space with the $x=0$ plane and recording the values of y and z every time the flow crosses this plane in the positive direction. A time series with $D=0.001$ of 5000 points was made by recording the values of the variable y from the Poincaré section. This file was analyzed in the same manner as the Hénon map, see Table II.

As with the Hénon map, almost all of the encounters from the original data satisfy the requirements of the linear fit algorithm, and more than half of the surrogate encounters are rejected. In both of these examples the statistical significance is well above the 99% level even without the improvement of the linear fit algorithm. In order to determine if this method is able to produce statistically significant results even when the standard TR method fails, another data file from the Rössler system was generated with the same parameters as before, but this time with only 1000 points and much more noise ($D=0.01$). As we can see in Table III, the standard TR method gives a statistical significance well below 99%, but the linear fit method easily breaks the 99% confidence level.

III. REJECTING FALSE POSITIVE RESULTS

An important question that occurs when doing this kind of analysis is whether or not a statistically significant number of encounters in a data file indicates that a UPO is really present. It is possible that correlations in the data could influence the probability of the pattern described above occurring due to chance. Ideally the surrogate data that is used in the topological recurrence test should preserve any statistical properties of the original data that could affect such probabilities, but in practice this simply is not possible. This means that although a statistically significant positive result allows us to reject the null hypothesis that the data are simply noise, it does not prove that a UPO is really present.

Rejection of a null hypothesis only allows us to prove that a data file is not accurately described by a particular model. In principle it is not possible to prove that a model is correct, nor is it necessary to do so. All that is required for a model to be useful is that it accurately describe the observable properties of the data. A rejection of the null hypothesis using the TR method provides an indication that a nonlinear dynamical system with UPO's may be an accurate model for the data, but it does not guarantee that the encounters found are from real UPO's, since it is possible that statistical correlations in the data could bias the results.

Fortunately the linear fit algorithm presented here also provides a method for determining whether the encounters found in the data represent a real UPO, or are simply the result of statistical correlations. We will demonstrate this technique by looking at the distribution of the eigenvalues and fixed points of the linear mapping that was fit to each encounter. Figure 3 shows histograms of the fixed points and eigenvalues of the encounters from the Hénon map data presented in Sec. II, as well as from one of its surrogates.

It is clear from these figures that the majority of the encounters found in the data do represent the real unstable fixed point in the Hénon map at the parameter values used. The actual fixed point and eigenvalues can be calculated analytically for the Hénon map, and are $X_F \approx 0.667$, $\lambda_S \approx 0.178$, and $\lambda_U \approx -1.69$. We see that the peaks in these distributions correspond quite well to the actual values. In contrast, the surrogate data produce very broad distributions. The peak in the fixed point distribution lies at the mean value of the surrogate data, and in fact the distribution of fixed points simply mimics that of the data itself. The distribution of stable eigenvalues is peaked at zero, but is too broad to give a reasonable estimate of an actual stable eigenvalue, and the unstable eigenvalue distribution is not only very broad, but also contains large numbers of encounters with both positive and negative unstable eigenvalues. Even if this surrogate file had been analyzed by itself and had given a significantly positive statistic, we could not conclude that it contains real UPO's.

We propose the following method for quantifying how sharp the peaks in these distributions are, and ultimately whether the model of the data as a dynamical system with UPO's is a reasonable one. First we take as our estimate of the actual location and eigenvalues of the UPO to be the mode of the distributions. This is because the mean and variance of the distributions are usually not well defined. The unstable eigenvalue distribution in particular exhibits a power law type of behavior for false encounters, as can be seen in Fig. 3(f). This means that both the mean and variance of the distribution of unstable eigenvalues (assuming that the null hypothesis is true) would be expected to grow rapidly with the length of the time series. Additionally, a time series from a dynamical system with real UPO's would still be expected to have some false encounters due to random chance. These false encounters can take on a wide range of values of both λ_U and X_F. Nonparametric statistics are thus needed for this analysis, and the mode works well as an estimator of the fixed-point location and eigenvalues as long as the number of false encounters is small compared to the number of real UPO encounters.

The mode can usually be determined fairly accurately from a histogram like the ones in Fig. 3, but in general any binning artifacts can be avoided by using an inhomogeneous Poisson process waiting time rate estimate [19], as described in [20]. If a real UPO is present, then we would expect the majority of the encounters to have fixed-point and eigenvalue estimates very close to the actual values. We can therefore quantify how reliable an estimate the mode is from the percentage of encounters whose individual estimates lie within some small distance of the mode. This small distance is selected somewhat arbitrarily, but certainly it should be smaller than the typical distances we see in surrogate data, where the encounters vary over a wide range of eigenvalues and locations.

We suggest the following procedure. Let the modes of the distributions be denoted as $E(X_F)$, $E(\lambda_S)$, and $E(\lambda_U)$ for the fixed-point, stable eigenvalue, and unstable eigenvalue distributions, respectively. We now count the number of encounters for which

$$|X_F - E(X_F)| < \Delta_F,$$

$$|\lambda_S - E(\lambda_S)| < \Delta_S,$$

$$|\lambda_U - E(\lambda_U)| < \Delta_U. \tag{8}$$

The parameters Δ_F, Δ_S, and Δ_U are determined as follows:

$$\Delta_F = \sigma_X,$$

$$\Delta_S = 0.3,$$

$$\Delta_U = |E(\lambda_U)| - 1, \tag{9}$$

where σ_X^2 is the variance of the data distribution. The rationale behind these somewhat arbitrary choices is that they confine the acceptable region to a very small portion of the possible range of values that a given encounter could produce. We will refer to the percentage of encounters that match all three criteria from Eq. (8) as P.

Of course a significant portion of the encounters from a data set with no real UPO's will give fixed-point and eigenvalue estimates that lie very close to the peaks of their distributions, but only a very small percentage of these encounters will meet all three criteria simultaneously. The percentage P is not only an estimate of the sharpness of the distributions, but also a measure of the correlation between the three estimated quantities. If the encounters represent a real UPO, we would expect that when the system converges towards the fixed point on a trajectory very close to the stable manifold that it would get very close to the fixed point and diverge along a trajectory very close to the unstable manifold. If the encounter does not represent a real UPO, then we would have no such expectation. A large value of P is thus a very strong indication that the encounters describe a real UPO. However, it must be emphasized again that P is not a formal confidence level, and cannot be used to reject a null hypothesis.

One type of system that is known to give false positive results when analyzed with the TR method is a noisy system in which there is a stable periodic orbit with a negative Floquet multiplier [21]. A stable fixed point will produce the same type of converging behavior when perturbed by noise that an unstable fixed point produces when the system falls onto the stable manifold. This means that the first half of the pattern searched for is strongly favored when a stable fixed point is present. The second half of the pattern will then happen due to chance with some probability. The result of this is that the overall probability of the pattern occurring may be either favored or disfavored, depending primarily on the sign of the Floquet multiplier. If the pattern is disfavored, then the K statistic will be negative, an effect which has been observed many times in stable periodic systems [4]. If the pattern is favored, then the K statistic will be positive, giving a false positive result. This effect has been used to formulate a technique similar to the TR method to detect the noisy precursors to a bifurcation [22,23]. This effect can be easily seen in both the Hénon map and the Rössler system in their period-1 regions when you get close to the bifurcation to period 2. This presents a real problem when analyzing noisy experimental systems where it is not always easy to identify stable fixed points.

To demonstrate how to identify false positive results like those described above, we refer again to the Hénon map, this time with $a=0.2$ and $b=0.3$. For these parameters the Hénon map has a stable fixed point. The results of analysis with the TR method are shown in Table IV.

We notice here that the fraction of encounters rejected by the matrix fit is almost as high for the original data as for the surrogates. This type of behavior is already a good indication that the encounters may not be with real UPO's. Looking at the distributions of fixed points and eigenvalues in Fig. 4 confirms this suspicion.

Note that the stable eigenvalue distribution is almost identical to that of the surrogates, and the unstable eigenvalue distribution differs only in that there are very few encounters with positive unstable eigenvalues. This is simply due to the strong negative correlations present in the data that are not present in the surrogate. The fixed-point distribution looks much narrower than that of the surrogates, but that alone is not sufficient for us to claim that a dynamical system with UPO's is an accurate model for the data. Applying the method described above we find $P=3.35$, indicating that only 3.35% of the encounters give fixed-point and eigenvalue estimates close to those predicted by the distribution peaks. This is not nearly as high as for the chaotic Hénon map example in Sec. II, for which we get $P=80.7$.

The same analysis can be applied to the Rössler system. Figure 5 shows the distribution of fixed points and eigenvalues for the Rössler data analyzed above, as well as one of its surrogates.

The results are qualitatively very similar to those in Fig. 3, with very sharp peaks near the numerically determined values of $X_F \approx -8.39$, $\lambda_S \approx -2.34\times 10^{-7}$, and $\lambda_U \approx -2.40$ [24]. An interesting point is that although the stable eigenvalue is nearly zero, it is negative, and the distribution of stable eigenvalues in Fig. 5(b) reflects this in that the distribution has a sharp peak at zero, but is asymmetric with almost no positive values at all. As with the chaotic example from the Hénon map, the vast majority of the encounters found in this file have fixed-point and eigenvalue estimates very close to the peaks in the distributions ($P=79.8$).

The Rössler system is also capable of stable periodic behavior and can give false positive results when analyzed with the topological recurrence method. A 5000 point file was generated with the same procedure as in Sec. II, this time with $a=b=0.2$, $c=2.5$, and $D=0.001$. The results of the TR analysis on this file are shown in Table V.

Analysis of the distribution of fixed points and eigenvalues gives $P=26.0$, once again far too low to claim that a real UPO is present.

IV. DISCUSSION

The ability of the topological recurrence method to detect UPO's in very short, noisy data files has made it a very important tool in the analysis of experimental biological and physical systems. We have shown that by rejecting encounters that do not provide a realistic fit to a linear mapping, we can greatly improve the sensitivity of the TR algorithm. We have also provided a method for determining if the encounters found in the data represent a real UPO, or some other type of correlation. Specifically, we have shown that this method can distinguish between UPO's and noisy stable periodic orbits, which until now could not be done reliably.

The modifications to the topological recurrence method presented here also add very little in the way of computational complexity. One of the important strengths of the TR algorithm is that it can be applied in real time to an experiment. The linear fit algorithm and distribution analysis presented above can be applied in real time as well.

ACKNOWLEDGMENTS

The author would like to thank Frank Moss, Mark Spano, and Alexander Neiman for their discussions and insight. This work was supported by the U.S Office of Naval Research, Physical Science Division.

[1] P. Cvitanovic, Phys. Rev. Lett. **61**, 2729 (1988).
[2] D. P. Lathrop and E. J. Kostelich, Phys. Rev. A **40**, 4028 (1989).
[3] G. Mindlin and R. Gilmore, Physica D **58**, 229 (1992).
[4] D. Pierson and F. Moss, Phys. Rev. Lett. **75**, 2124 (1995).
[5] X. Pei and F. Moss, Nature (London) **379**, 618 (1996).
[6] X. Pei and F. Moss, Int. J. Neural Syst. **7**, 429 (1996).
[7] H. A. Braun, K. Schäfer, K. Voigt, R. Peters, F. Bretschneider, X. Pei, L. Wilkens, and F. Moss, J. Comput. Neurosci. **4**, 335 (1997).
[8] H. A. Braun, M. Dewald, K. Schäfer, K. Voigt, X. Pei, and F. Moss, J. Comput. Neurosci. **7**, 17 (1999).
[9] R. Shaw, *The Dripping Faucet as a Model Chaotic System* (Aerial Press, Santa Cruz, 1984).
[10] T. Sauer, Phys. Rev. Lett. **72**, 3811 (1994).
[11] E. Ott, C. Grebogi, and J. A. Yorke, Phys. Rev. Lett. **64**, 1196 (1990).
[12] K. Dolan, A. Witt, M. L. Spano, A. Neiman, and F. Moss, Phys. Rev. E **59**, 5235 (1999).
[13] R. V. Hogg and E. A. Tanis, *Probability and Statistical Inference* (Macmillan Publishing Company, New York, 1988).
[14] J. Theiler, S. Eubank, A. Longtin, B. Galdrikain, and J. D. Farmer, Physica D **58**, 77 (1992).
[15] J. Theiler and D. Prichard, Physica D **94**, 221 (1996).
[16] W. L. Ditto, S. N. Rauseo, and M. L. Spano, Phys. Rev. Lett. **65**, 3211 (1990).
[17] W. L. Ditto, M. L. Spano, and J. F. Lindner, Physica D **86**, 198 (1995).
[18] O. E. Rössler, Phys. Lett. A **57**, 397 (1976).
[19] Emanual Parzen, *Stochastic Processes* (Holden Day, San Francisco, 1962).
[20] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes in C, The Art of Scientific Computing*

(Cambridge University Press, Cambridge, 1992), p. 479.

[21] S. H. Strogatz, *Nonlinear Dynamics and Chaos* (Addison-Wesley, New York, 1994).

[22] L. Omberg, K. Dolan, A. Neiman, and F. Moss, Phys. Rev. E **61**, 4848 (2000).

[23] K. Wiesenfeld, J. Stat. Phys. **38**, 1071 (1985).

[24] Y. Kuznetsov, *Elements of Applied Bifurcation Theory* (Springer-Verlag, New York, 1995); CONTENT—integrated environment for analysis of dynamical systems (available at: ftp.cwi.nl).