LA-UR -91-3345

LA-UR--91-3345

DE92 002461



TITLE: DETECTING NONLINEAR STRUCTURE IN TIME SERIES

AUTHOR(S): James Theiler

SUBMITTED TO: Proceedings of the 1st Experimental Chaos Conference, held 1-3 October 1991, in Arlington VA

# DETECTING NONLINEAR STRUCTURE IN TIME SERIES

JAMES THEILER,[a,b] BRYAN GALDRIKIAN,[a,b] ANDRÉ LONGTIN,[a,b]
STEPHEN EUBANK,[a,b,c] *and* J. DOYNE FARMER[a,c]

(a) *Theoretical Division, and* (b) *Center for Nonlinear Studies,*
*Los Alamos National Laboratory, Los Alamos, NM 87545, USA;* and
(c) *The Santa Fe Institute, 1660 Old Pecos Trail, Santa Fe, NM 87501, USA.*

ABSTRACT

We describe an approach for evaluating the statistical significance of evidence for nonlinearity in a time series. The formal application of our method requires the careful statement of a null hypothesis which characterizes a candidate linear process, the generation of an ensemble of "surrogate" data sets which are similar to the original time series but consistent with the null hypothesis, and the computation of a discriminating statistic for the original and for each of the surrogate data sets. The idea is to test the original time series *against* the null hypothesis by checking whether the discriminating statistic computed for the original time series differs significantly from the statistics computed for each of the surrogate sets. While some data sets very cleanly exhibit low-dimensional chaos, there are many cases where the evidence is sketchy and difficult to evaluate. We hope to provide a framework within which such claims of nonlinearity can be evaluated.

*Destiny is not a matter of chance; it's a matter of choice.*

— William Jennings Bryan

*I returned, and saw under the sun that the race is not to the swift, nor the battle to the strong, neither yet bread to the wise, nor yet riches to men of understanding, nor yet favor to men of skill; but time and chance happen to them all.*

— Ecclesiastes

## 1. Is it chaos or is it noise? (that is not the question)

The extent to which experimental data is deterministic or random has always elicited controversy. But, in asking whether a given time series is chaos or noise, it is important to remember that these are not categories but extremes. Rarely are systems encountered outside of the numerical laboratory entirely deterministic or purely random. For real experiments, it's always a mix.

Confronted with the erratic fluctuations of an experimental time series, one would ultimately like to disentangle the various influences of chaos, nonchaotic but still nonlinear determinism, linear correlations, noise in the dynamics, and noise in the measurement apparatus. Our aim in this report is more modest. Formally, we seek only to *detect* nonlinear structure in the time series, by asking the question in the language of statistical hypothesis testing. Informally, we hope to provide a control experiment which will aid in the interpretation of traditional tests for chaos in time series.

This report summarizes work that is part of an ongoing project to more reliably characterize nonlinear structure from a time series; further discussion (and a more complete list of references) can be found in our longer paper [1].

## 2. Testing for Significance

Statistical hypothesis testing involves two components: a null hypothesis against which observations are tested, and a discriminating statistic. The null hypothesis is the "too-simple" explanation that we seek to show is inadequate for explaining the data. Our null hypotheses have the form, "the underlying dynamics is *blah*," where *blah* might be, for example, linearly correlated noise. Our aim is to show that the observed data differs significantly from any realization that would be generated under the null hypothesis.

For this purpose, a discriminating statistic is employed, which in general is just a *number* that quantifies some aspect of the time series; it can be anything from a higher moment to an estimated fractal dimension. Basically, we want to show that the value of this statistic is different for the experimental time series than it would have been if the null hypothesis were true.

It is possible in some cases to derive analytically the distribution of a given statistic under a given null hypothesis, and this approach is the basis of many conventional tests for nonlinearity [2]. The method of surrogate data is a direct Monte-Carlo approach for estimating this distribution. An ensemble of surrogate data sets are generated which share given properties of the observed time series (such as mean, variance, and Fourier spectrum) but are otherwise random as specified by the null hypothesis. For each surrogate data set, the discriminating statistic is computed, and from this ensemble of statistics, the distribution is approximated.

While this approach is computationally intensive, it avoids the analytical derivations which can be difficult if not impossible. It basically trades brain for brawn; statistics of this kind are persuasively advocated in Efron [3].

Let $Q_D$ denote the statistic computed for the data given by the observed time series, and $Q_{S_i}$ for the $i$th surrogate data set. Let $Q_S$ and $\sigma_S$ denote the mean and standard deviation of the $Q_{S_i}$'s. We define our measure of "significance" as the difference between the original and the mean surrogate value of the statistic, divided by the standard deviation of the surrogate values.

$$\mathcal{S} \equiv \frac{|Q_D - Q_S|}{\sigma_S} \tag{1}$$

The significance is a dimensionless quantity, but it is common parlance among physicists to call the units of $\mathcal{S}$ "sigmas." Thus, one might speak of a "two sigma effect" as not especially significant, but ten sigmas as very significant. If the distribution of statistic values is gaussian (and numerical experiments indicate that this is often a reasonable approximation), then the $p$-value associated with a significance $\mathcal{S}$ is given by $p = \mathrm{erfc}[\mathcal{S}/\sqrt{2}]$; this is the probability of observing a significance $\mathcal{S}$ or larger if the null hypothesis is true.

### 2.1. *Choice of Null Hypothesis and Discriminating Statistic*

This report will concentrate on the null hypothesis of linearly correlated gaussian noise; that is, all the structure of the time series resides in its Fourier spectrum. Our algorithm for generating a surrogate data set is to take the Fourier transform of the original series, randomize the phases, and then take the inverse Fourier transform. More general null hypotheses have also been considered [1].

Since we are motivated by the possibility that the underlying dynamics may be chaotic, our first choices for discriminating statistics are just the conventional discriminants of nonlinear dynamics: correlation dimension, Lyapunov exponent, and forecasting error. Other choices are also possible [1].

An advantage of brute numerics is that the choice of null hypothesis and discriminating statistic can be made independently; there is considerable flexibility to tailor the test to the situation at hand.

## 3. Numerical Experiments

To properly gauge the utility of the method of surrogate data will eventually require many tests with data from actual experiments. To whet our appetites, however, and to give a sense of how this approach ought to work in practice, we begin with some examples from the numerical laboratory.

First, we note that a time series from a linear process should by construction give a negative result (that is, the null hypothesis should not be rejected); this was checked and found to be the case.

Using time series generated by the Hénon map, we confirmed our expectation that the significance of the test for nonlinearity increases with the length of the data set. By using sums of independent Hénon time series we generated data sets of varying dimension, and found that the significance of the test decreases for more complex systems. That is, the higher the dimension, the more difficult it was to distinguish from linearly correlated noise.

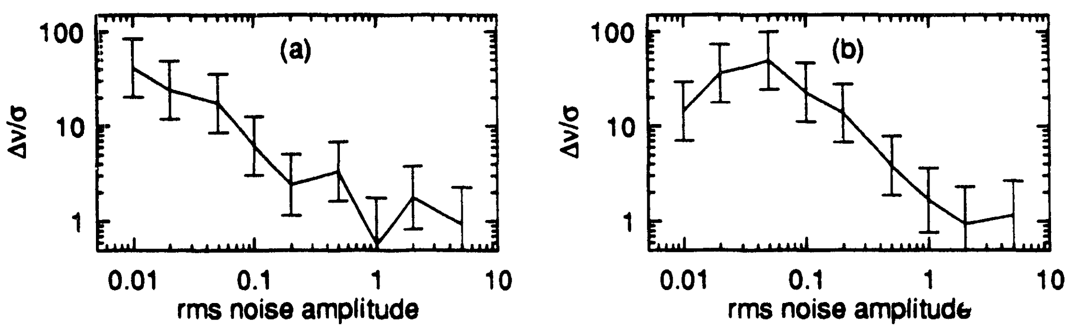


Figure 1: *Effect of noise on significance for a time series of $N = 512$ points, derived from the cosine map with $\lambda = 2.8$: (a) observational noise; (b) dynamical noise.*

### 3.1. *Effect of noise*

To test whether nonlinear determinism can be detected even when it is mixed with noise, we added both dynamical ($\eta$) and observational ($\varepsilon$) noise to the cosine map: $y_t = \lambda \cos(\pi y_{t-1}) + \eta_t$; $x_t = y_t + \varepsilon_t$. In Fig. 1, we plot significance as a function of noise level for both dynamical and observational noise. As expected, significance decreases with increasing noise level, though we remark that the nonlinearity is still observable even with considerable noise. In the absence of noise, the rms amplitude of the signal is 0.36; thus we are able to detect significant nonlinearity even with a signal to noise ratio of one, using a time series of length $N = 512$.

### 3.2. *Don't bleach chaotic data*

A common approach to testing for nonlinearity involves first "subtracting out" the linear correlations [2,4]. Any correlations that are left are then necessarily nonlinear correlations. So the test for nonlinear structure is converted to the conceptually simpler test for any determinism at all. Given a time series $x_t$, the residuals $e_t$ are given by the linear model

$$e_t = x_t - \left[ a_o + \sum_{k=1}^{q} a_k x_{t-k} \right], \tag{2}$$

where $q$ is the order of the model, and the coefficients $a_k$ are chosen to minimize the variance of the residuals. Because the residuals $e_t$ are spectrally white (equal power at all frequencies), the process of determining residuals is sometimes called "pre-whitening" or "bleaching".

While the fit is based on the best auto-regressive (AR) *model*, the linear map that takes $x_t$ to $e_t$ in Eq. (2) is a moving-average (MA) *filter*, and will not formally change the structure of the attractor for finite $q$. For example, if $x_t$ lies on a low-dimensional attractor, then $e_t$ will lie on an attractor of the same dimension. However, in practice, the distortion induced by the linear map can drastically affect the *appearance* of the attractor and can likewise affect *estimates* of its dimension.

In an experiment with a time series from the Hénon map, significance was computed for the raw time series and for time series obtained by bleaching with ever larger values of $q$; the significance was found to decrease with increasing $q$. We have contrived some cases where low-order ($q \leq 3$) filtering enhances the significance, but for larger values of $q$, the significance again decreases. In general, we do not recommend statistical tests for nonlinearity that are based on best estimates of the residuals, as these usually require higher order filtering.

## 4. Experimental chaos: two examples from real laboratories

### 4.1. *Superfluid convection*

Data from a superfluid convection cell provides an example where the evidence for low-dimensional chaos is quite clear. Using discriminating statistics of dimension and forecasting error, we obtain fifteen and forty sigmas of significance, respectively.

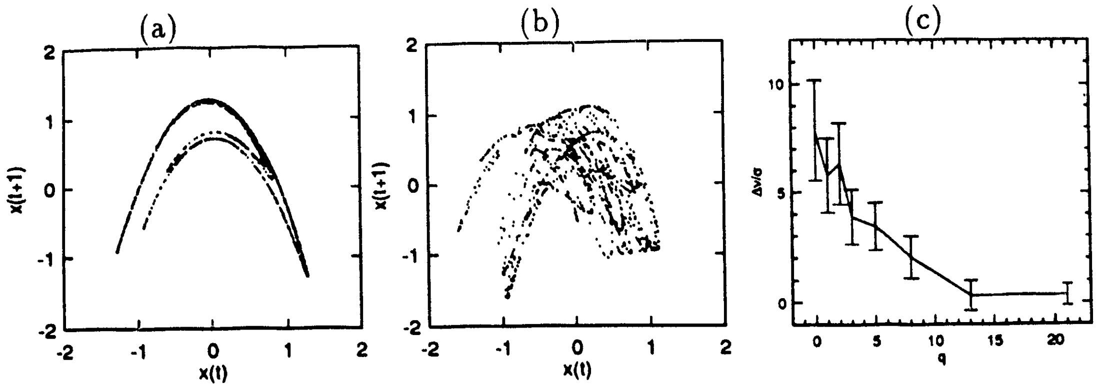


Figure 2: *Effect of bleaching on significance for a time series derived from the Hénon map. For $q = 0$, the raw data set is used. For $q > 0$, the q-th order residuals (as computed by Eq. (2)) is used. In (a), the unfiltered $q = 0$ Hénon attractor is shown. In (b), a $q = 5$ filter distorts the attractor considerably, and hides the determinism that is evident in the raw data. In (c), the significance is computed using the dimension statistic on a time series of length $N = 128$. Significance decreases with increasing q, and is essentially insignificant for $q > 8$. It is apparent that attempting to "subtract out" the linear component only decreases the power of the test.*

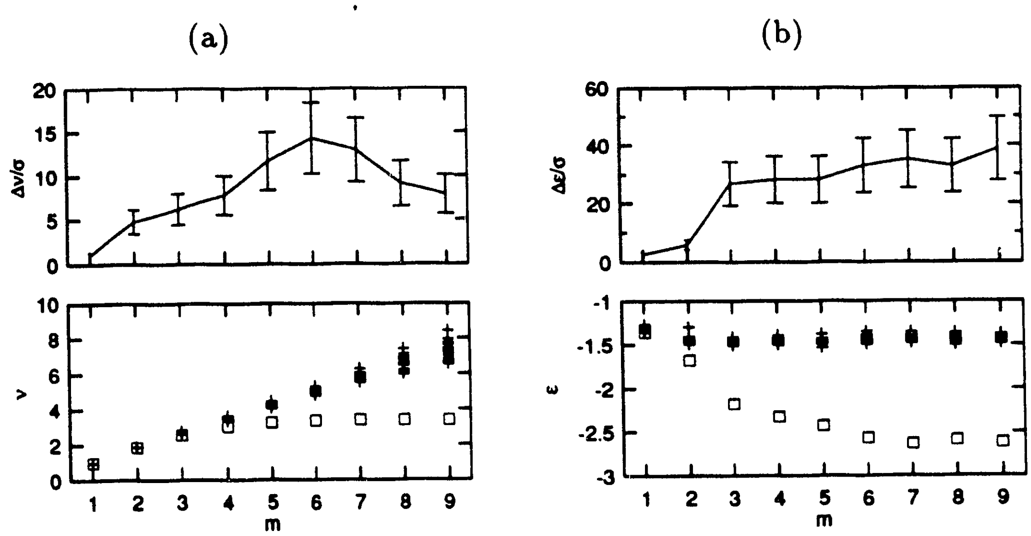


Figure 3: *Data from a fluid convection experiment exhibits very significant nonlinear structure, using (a) dimension, and (b) forecasting error. The top panel in these figures show the significance, measured in "sigmas," and the bottom panel shows the values of the statistics, with squares (□) for the original data and pluses (+) for the surrogates. Both panels plot these statistics against the embedding dimension m. Not only is the evidence for nonlinear structure statistically significant, but the estimated dimension of about $\nu = 3.8$ suggests that the underlying dynamics is in fact low-dimensional chaos.*

This data was also analyzed by Farmer and Sidorowich [5], who found sizable increases in predictability using nonlinear rather than linear predictors.

*4.2. Electroencephalogram (EEG)*

That the brain should exhibit chaos is an idea that some authors have been unable to resist. Our own investigations so far have been mixed; some data exhibit nonlinear structure and some do not. A more systematic survey is clearly in order. In the meantime, we present the "best" result, the one time series we looked at which exhibited the most statistically significant evidence for nonlinear structure. The time series was obtained from the brain of an individual who was staring at a spot on a wall. Here we see almost ten sigmas of significance with the dimension algorithm (around five sigmas with the forecasting algorithm), but it is worth pointing out that we do not see any evidence that the time series is in fact low-dimensional (the correlation dimension $\nu$ does not converge with increasing embedding dimension $m$). We are only able to reject the null hypothesis that the data arise from a linear stochastic process, we cannot specify the nature of the nonlinear structure that we are evidently seeing.

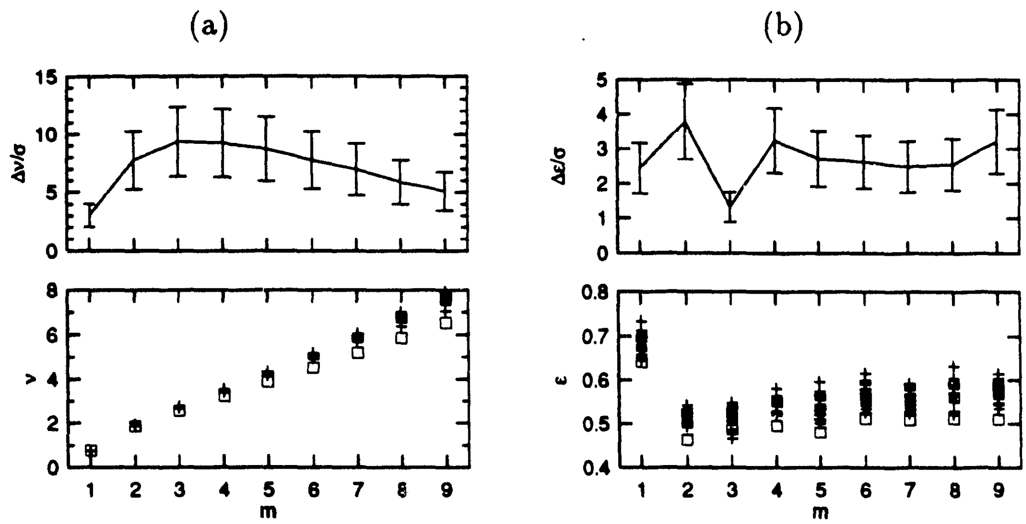


Figure 4: *Data from an electroencephalogram time series exhibits significant nonlinear structure, using both (a) the dimension statistic, and (b) the forecasting error. The evidence for low-dimensional chaos, however, is weak, since the estimated dimension increases almost as rapidly with embedding dimension for the original time series as it does for the surrogates.*

## 5. Conclusion

In this report, we have provided a framework for evaluating the statistical significance of evidence for nonlinearity in a stationary time series. The test properly fails to find nonlinear structure in linear stochastic systems, and is shown to correctly identify nonlinearity in several well-known examples of low-dimensional chaotic time series, even when contaminated with considerable noise.

Our experiments with chaotic data found that using linear pre-processing to "bleach" out the linear correlations decreases the power of the test to detect nonlinearity. Consequently, we advocate a direct application of the method of surrogate data to the raw time series, instead of to a time series of residuals.

Finally, we tested several experimental data sets for nonlinear structure.

Much work remains to make the method of surrogate data a powerful and flexible and idiot-proof tool for nonlinear analysis. In particular, we hope to expand the hierarchy of null hypotheses and to broaden the battery of discriminating statistics. The algorithms should also be extended to deal with multivariate time series and input-output systems. Further investigation of the effectiveness of various statistics for different null hypotheses and in different situations will be valuable not only for increasing our ability to reject null hypotheses, but also for the more qualitative task of characterizing the nature of the nonlinearity that might be evidenced by one statistic but not another.

## 6. Acknowledgements

We are grateful to Martin Casdagli, Xiangdong He, William Brock, Blake LeBaron, Peter Grassberger, Mark Berge, and Bette Korber for useful discussions. We also thank Bob Ecke for providing convection data, and Paul Nuñez and Arden Nelson for providing EEG data. This work was partially supported by the National Institute for Mental Health under grant 1-R01-MH47184-01, and performed under the auspices of the Department of Energy. We urge the reader to use these results for peaceful purposes; in particular, we do not advocate the use of war as a surrogate foreign policy.

## 7. References

1. J. Theiler, B. Galdrikian, A. Longtin, S. Eubank, and J. D. Farmer, "Using surrogate data to detect nonlinearity in time series," in *Nonlinear Prediction and Modeling*, M. Casdagli and S. Eubank, eds. (Addison-Wesley, 1991). To appear.
2. H. Tong, *Non-linear Time Series: A Dynamical System Approach.* (Clarendon Press, Oxford 1990).
3. B. Efron, "Computers and the theory of statistics: thinking the unthinkable," *SIAM Review* **21**, 460–480 (1979).
4. W. A. Brock, W. D. Dechert, and J. Scheinkman, "A test for independence based on the correlation dimension," Technical Report 8702, Social Systems Research Institute, University of Wisconsin, Madison, (1986).
5. J. D. Farmer and J. J. Sidorowich, "Exploiting chaos to predict the future and reduce noise," in *Evolution, Learning and Cognition*, Y. C. Lee, ed. (World Scientific, 1988), pp. 277–330.

# END

# DATE FILMED

12/06/91