ACTA NEUROBIOLOGIAE EXPERIMENTALIS

Acta Neurobiol. Exp. 2000, 60: 109-121

Nonlinearity in human resting, eyes-closed EEG: an in-depth case study

Walter S. Pritchard[1,2,3] and **Cornelis J. Stam**[4]

[1]Psychophysiology Laboratory, Bowman Gray Technical Center 611-12, Research and Development, R.J. Reynolds Tobacco Company, Winston-Salem, NC 27102, USA, Email: pritchw@rjrt.com; [2]Department of Physiology and Pharmacology, Wake Forest University School of Medicine, USA; [3]Department of Psychiatry and Behavioral Science, Mercer University School of Medicine, USA; [4] Department of Clinical Neurophysiology, Leyenburg Hospital, P.O. Box 40551, 2504 LN The Hague, The Netherlands

Abstract. The question of nonlinearity in the human electroencephalogram (EEG) is important, since linear methods of EEG analysis are more well-developed and computationally faster than nonlinear methods. Furthermore, the presence or absence of nonlinearity has important theoretical implications for understanding the nature of the brain's oscillatory activity. Using a linear summary measure as a control, we report a failure to reject the null hypothesis of a (largely) stationary linear-Gaussian process for normal, resting, eyes-closed EEG from a single participant. We found significant evidence of nonlinearity at two occipital sites (O1 and O2) where the 8-12.5 Hz alpha rhythm was prominent. However, this element of nonlinear structure appeared trivial, as (1) we found no evidence of time irreversibility at these loci, and (2) best-fitting linear models accounted on-average for over 94% of the variance in the data with nonlinear modeling doing no better. Half of the remaining variance could be accounted for by nonstationarity. While our findings technically apply only to the one individual tested, his EEG was typical of those seen under the conditions that we employed.

Key words: EEG, time series, nonlinearity, surrogate data, chaos

Review

INTRODUCTION

Nonlinear time-series analysis and the method of surrogate data

In recent years, nonlinear time-series analysis has found increasing application in the life sciences. Originally, the goal of such analysis focused on demonstrating the presence of low-dimensional chaos in biological time series, and, indeed, numerous "chaos sightings" were reported, e.g., Pritchard and Duke (1992a). More recently, the emphasis of nonlinear time-series analysis in general has switched from proving the alternative hypothesis (of chaos) to a more restrained approach of rejecting/failing to reject the null hypothesis of a stationary linear-Gaussian process (LGP), that is, a linear-stochastic process with Gaussian inputs[1] (Kaplan and Glass 1995). Rejection of the LGP null hypothesis for any dynamical process has important practical implications, given that nonlinear methods of time-series analysis are in-general computationally more intensive and historically less well-developed than linear methods (Harvey 1989, Pritchard and Duke 1992b). Furthermore, testing the LGP null hypothesis has important theoretical implications regarding the fundamental nature of oscillatory brain activity.

Testing the LGP null hypothesis involves generating suitable control, or surrogate, data that are LGPs having the same linear properties (e.g., mean, variance, and power spectrum/autocorrelation function (ACF)[2]) as the original data (Kennel and Isabelle 1992, Theiler et al. 1992). One selects a summary measure that reflects nonlinear aspects of the data (e.g., dimensions, Lyapunov exponents, entropies, nonlinear forecast errors) and applies it to the original and the surrogate data - if the original data are in-fact an LGP, then the original and average-surrogate values will not differ significantly.

Phase-angle randomization (PAR)

The leading method of producing LGP surrogates is to fast-Fourier transform (FFT) the original data, randomize on the interval $(0, 2\pi)$ all the phase angles without changing power at any frequency, and then apply an inverse FFT (see Appendix 1 for further technical details). Surrogate-data testing based on phase-angle randomization (PAR) has indicated that the EEG, although not low-dimensional chaos, does have an element of nonlinear structure (e.g., Pritchard et al. 1995, Rombouts et al. 1995, Ehlers et al. 1998).

However, PAR-based surrogate-data testing has two drawbacks. First, it can spuriously indicate nonlinearity when applied to perfect linear-periodic processes, i.e., sine waves. When a sine wave is FFTed and its frequency does not match one of the FFT frequency steps (which is the typical case), both power and phase are spread across a range of frequencies. When the phases are then randomized, an inverse FFT produces not the original sine wave, but a superimposed set of sine waves beating against each other. This results in a time series having altered nonlinear-dynamical properties (e.g., it is more complex) that in turn results in altered nonlinear summary measures (Stam et al. 1998a, Pritchard and Stam 2000, in press). It also has altered linear properties (e.g., it is not, in-fact, isospectral with the original data). To the degree that human alpha EEG can be quite sinusoidal in appearance with a very sustained, oscillating autocorrelation function, one may suspect that it might to some degree be susceptible to this FFT frequency mismatch artifact.[3]

---

[1]The inputs are also variously termed "innovations", "shocks", "disturbances", or "errors". They come from the normally distributed stochastic parameter(s) that drive(s) the dynamics of an LGP (the two major classes of LGPs, moving averages (MAs) and autoregressive models (ARs) are discussed in Appendices 1 and 2, respectively).

[2]By the Wiener-Khintchine theorem, the power spectrum and the ACF provide equivalent information, the former being a frequency-domain representation and the latter a time-domain representation.

[3]There is one out to the PAR FFT-mismatch artifact for sine waves, namely, if the sine wave in question were to have perfect end-point wrap-around (the last portion of the time series fits exactly with the first portion, so that, if the series were plotted on a sheet of paper and the paper were curled into a cylinder, there would be a perfect sine wave going "endlessly" around the cylinders circumference). However, the FFT requires data that are a power of 2, and power-of-2 length sine waves "naturally" having perfect wrap-around are almost never encountered with experimental data. One solution might be to systematically shorten the time series until wrap-around is obtained. However, this would result in a loss of data and would also require use of either the computationally slower DFT, which does not require data length to be a power of 2 (for today's computers, the relative slowness of the DFT is really not much of an issue), or a multiple-radix FFT (we are grateful to Danny Kaplan for pointing out the latter technique to us).

The second problem with the PAR surrogate technique is that it cannot distinguish dynamic from static nonlinearity. That is, it will flag an LGP that has undergone a static (time-independent), nonlinear transform of some sort (e.g., due to the conditions of measurement) as being nonlinear. Both these problems can be addressed by examining time irreversibility.

Time irreversibility

A stationary time series is time reversible if its probabilistic structure going forward in time is identical to that going backwards in time (Ramsey and Rothman 1996; the topic of stationarity is discussed in the next section). Weiss (1975) showed that all stationary LGPs are time reversible and cannot be made irreversible by a static, nonlinear transform. Thus, if a stationary time series having nonlinear structure (as determined using PAR surrogates) can be demonstrated to be time-irreversible as well, then that structure is most likely dynamic in nature. Conversely, if no evidence of time irreversibility is found, then the dynamic nature of any nonlinearity detected using the PAR technique becomes suspect.

We note here that assessing time irreversibility does not (necessarily) involve the generation of surrogate data, and thus avoids the FFT frequency-mismatch artifact (Stam et al. 1998a). The major weakness of the application of time irreversibility tests is that they are best applied inferentially to multiple-participant data sets, since for a single time series they either provide no estimate of significance (Stam et al. 1998a) or must bootstrap one using computationally intensive numerical techniques (Diks et al. 1995, Rothman 1996). It should also be noted that time irreversibility by itself does not necessarily imply nonlinearity, since stationary linear-stochastic processes with non-Gaussian inputs are also time irreversible (Ramsey and Rothman 1996).

Nonstationarity

Assessing nonlinearity and time irreversibility both presume that the data being examined are stationary (to be defined in the following paragraphs). This means that if a time series can be shown to be nonstationary, than the dynamical nature of nonlinearity detected by means of the PAR technique in conjunction with testing for time irreversibility becomes suspect. This is because nonstationary processes are also time irreversible (Ramsey and Rothman 1996). Thus, an assessment of nonstationarity would seem to be recommended in any assessment of nonlinearity.

Informally, a stationary time series is one that in some sense remains the same across time. The question of stationarity is important since many of the methods of both linear and nonlinear time series analysis presume that the data in question are stationary (e.g, for linear time-series analysis, autoregressive (AR) modeling (see Appendix 2); for nonlinear time-series analysis, estimation of the correlation dimension as a nonlinear summary measure (see below)).

In linear time-series analysis, two types of stationarity are defined (Harvey 1988). A weakly stationary process has a mean, variance, and power spectrum/ACF that do not depend on time, i.e., are constant. Since the mean, variance, and power spectrum/ACF are all linear properties of a time series, a weakly stationary process may be thought of as being "linearly stationary". A process is said to be strictly stationary if, in addition to having a constant mean, variance, and power spectrum/ACF, its probabilistic structure is time-independent. Since the probabilistic structure of a process reflects its nonlinear properties, a strictly stationary process may be thought of as being "nonlinearly stationary". Note that a strictly stationary process is also weakly stationary, but not *vice versa*. An exception is an LGP, which is fully described by the mean, variance, and power spectrum/ACF, and hence is strictly stationary also.

In theoretical approaches to the physics of dynamical systems, the equations of a system are deterministic and known *a priori*. In such cases, stationarity is equivalent with constant dynamics wherein the parameters of the governing differential or difference equations do not change. But with complex real world phenomena, the dynamics usually are not known *a priori* and stationarity, like any other dynamical property, has to be assessed from "the data" alone.

Along these lines, it is important to note that stationarity cannot be considered independently of the notion of time scale. For example, intermittent chaotic process alternate between periods of laminar behavior and periods of highly irregular behavior. Assessed on a small enough time scale, a realization of such a process[4] would appear

[4]A note on terminology: in time-series analysis, "process" refers to an equation or set of equations for generating time-series data. A particular time series generated by a given process is referred to as a "realization".

to be stationary (e.g, during either the laminar or irregular phases the variance would appear largely constant). However, on a longer time scale, a realization of an intermittent-chaotic process would appear highly nonstationary (e.g., the variance is dramatically larger during the irregular phases relative to the laminar phases). But viewed on a still longer time scale, the realization is seen (correctly, since the dynamics remain constant) to be stationary.

For the EEG, there is a general consensus that at a short enough scale it can be considered stationary (ignoring for the moment smaller-amplitude, stimulus-induced transients (so-called event-related potentials, or ERPs) embedded in the larger-amplitude, ongoing activity). But how short is "short enough"? The work of Lehmann and colleagues (e.g., Lehmann 1981) indicates that EEG topographic microstates last only on the order of tenths of seconds, and this is true of alpha activity as well (Cantero et al.). But this clearly is not long enough to do any meaningful sort of dynamical analysis, either linear or nonlinear! The trick is to find a segment long enough to permit dynamical analysis but short enough to be relatively stationary.

Linear vs. nonlinear prediction

A final approach to the question of nonlinearity involves comparing linear vs. nonlinear modeling of the data. For example, one could compare goodness of fit of a best-fitting (e.g., lowest-AIC; see Appendix 2) linear model with that of a nonlinear modeling algorithm. Since the latter is able to take advantage of both linear and nonlinear predictability in the time series, its goodness of fit should be better than that of a linear model if the time series is a realization of a nonlinear process (this assumes that the data are not overfit, that is, noise peculiar to the time series in question is being fit along with any systematic dynamics).

It should be noted that a good linear fit, while a necessary condition for inferring a linear-stochastic process, is not by itself a sufficient condition. Linear modeling relies upon certain criteria (e.g., the AIC; see Appendix 2) for determining a best fit. However, these criteria can fail when faced with, for example, chaotic time series. In unpublished simulations we have found that a best-fitting autoregressive (AR) model (see Appendix 2) fit to time series from the Lorenz system in the chaotic regime typically produces an adjusted R^2 (see Appendix 2) of over 0.98 (in theory, any nonlinear process can be in-sample modeled to any degree of goodness desired by a linear model of sufficient power).

Previous studies of alpha nonlinearity

Brandt et al. (2000) examined a set of human alpha EEGs, comparing K-nearest-neighbor, local-linear (KNNLL) forecasting (Casdagli 1989) with forecasting based upon AR modeling. They found almost no difference between the two methods, suggesting that human alpha activity is a largely a linear-stochastic process. Stam et al. (1999a) used a new surrogate-data procedure developed by Schreiber and Schmitz (1996) that iterates until an exact match between the original-data and surrogate-data amplitude distributions is achieved in conjunction with a close match between power spectra. They reported that for alpha data from a large number of participants (60), only 1.25% were classified as "nonlinear" (it should be noted that this percentages is based on the conservative Bonferroni correction when multiple *t*-tests are employed).

Pritchard and Stam (2000) used still another type of surrogate data based on the Wold decomposition and designed specifically to avoid the FFT-mismatch artifact (subtract a best-fitting sine wave from the data, PAR the residuals, reverse FFT, then add the sine wave back in). They reported significant evidence of nonlinearity in 10% of the alpha EEGs examined, although inferential statistical analysis indicated that the mean level of nonlinearity across EEGs differed significantly from zero.

Goals

The present study intensively investigates alpha EEG segments from a single participant. The goals were as follows:

1. Assess the nonlinearity of EEG data using the PAR technique. Based on previous findings, we expected variations in nonlinearity from segment to segment with perhaps a weak element of overall nonlinear structure at some loci. In assessing nonlinearity, we took an approach similar to that of Palus (1996): apply the PAR technique to both a nonlinear and a linear summary measure. If the former flags the time series as nonlinear, but the results for the latter indicate that the linear properties of the data have been significantly altered as well, then the detected nonlinearity is suspect. Simulations FFT frequency-mismatching sine waves indicated that

this procedure substantially attenuated but did not completely eliminate the FFT frequency mismatch artifact.
2. Assess the time irreversibility of the data. If EEG nonlinear structure as determined in (1) is positively associated with time-irreversibility, then one may propose that the nonlinear structure may be dynamic in nature if nonlinear structure is not positively related to nonstationarity (again, in theory, nonstationary processes are time irreversible and may spuriously appear nonlinear to methods assuming stationary data).
3. Assess the stationarity of the data. If nonlinear structure is positively associated with time-irreversibility but not nonstationarity, then one may propose that the nonlinear structure may be dynamic in nature.
4. Compare linear and nonlinear modeling of the data. If nonlinear structure is positively associated with time-irreversibility and better nonlinear modeling but not nonstationarity, then that nonlinear structure is most likely dynamic in nature.

PROCEDURES FOR ASSESSING NONLINEARITY, TIME IRREVERSIBILITY, AND NONSTATIONARITY

EEG data

We examined a set of EEGs recorded in 1997 from a normal, then-46-year old male (the first author) under resting, eyes-closed conditions (loci O1, O2, Pz, Cz, and Fz of the International 10-20 system; nasal reference; forehead ground; digitization rate 128 Hz; on-line band pass 0.3-30 Hz (-12 dB/octave rolloff)). The recordings were 64-s in length and extremely clean in terms of a lack of EOG, EMG, ECG, and movement artifacts. The recordings at O1, O2, and Pz featured typical examples of the prominent 10-12.5 Hz alpha rhythm (the data from O2 are presented in Fig. 1), although there was notable alpha power in all channels (Fig. 2).

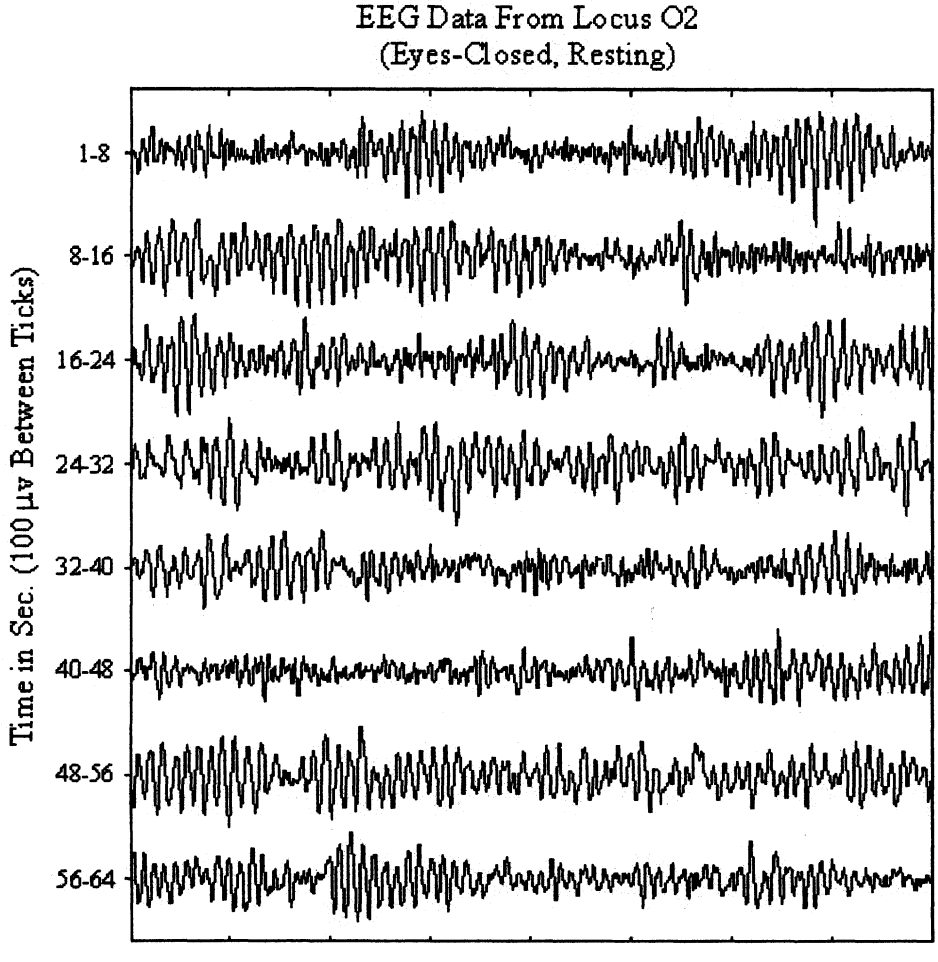

Fig. 1. EEG recorded from locus O2 (right occipital). This EEG is a typical example of the 8-12.5 Hz alpha rhythm best seen in normal persons under resting, eyes-closed conditions.

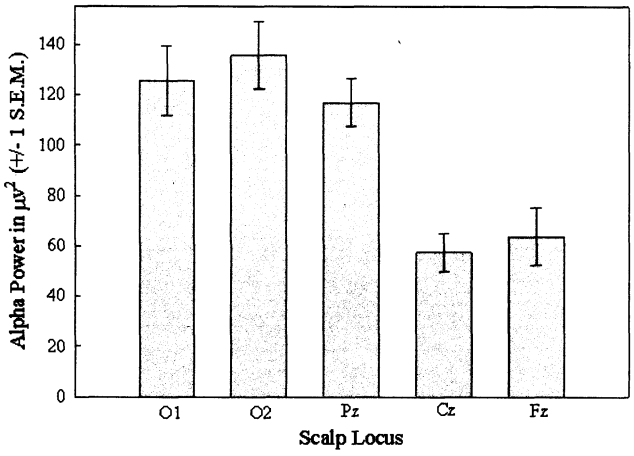

Fig. 2. Total alpha power as a function of scalp locus. SEM, Standard Error of the Mean.

Assessing nonlinearity

NONLINEAR SUMMARY MEASURE

For PAR surrogate-data testing, we employed estimated correlation dimension (D_2) as a nonlinear summary measure. For each EEG, a state-space trajectory was reconstructed using the method of delays as reviewed in Pritchard and Duke (1992b). The embedding dimension was selected using the method of false nearest neighbors (Kennel et al. 1992), and the lag was set by a geometric method based upon the rate of expansion of the reconstructed state-space vectors from the main state-space diagonal (Rosenstein et al. 1994). D_2 was then estimated using a maximum-likelihood method proposed by Takens (1985) and extended by Ellner (1988), excluding vector pairs closer than 3 (autocorrelation time). For further details regarding the Takens--Ellner method, see Pritchard and Duke (1992b). Within the general context of EEG research, estimated D_2 typically is employed not "as intended" (i.e., a measure of a stationary, deterministic system's degrees of freedom), but rather as an (admittedly vague) atheoretical measure of the complexity of the reconstructed EEG state-space trajectory.

We chose estimated D_2 as a nonlinear summary measure for two reasons. First, it appears to be a conservative method for assessing nonlinearity relative to other nonlinear summary measures (Ehlers et al. 1998). Second, it "suffers considerably" from nonstationarity (Kantz and Schreiber 1997, p. 17). For example, Timmer (1998) generated realizations of second-order autoregressive (AR(2)) processes (see Appendix 2) have varying amounts of either cyclic amplitude or period modulation (so-called cyclostationary processes). He found that, using the correlation dimension as a nonlinear summary measure, the surrogate-data technique spuriously indicated nonlinearity at nonstationarity parameter settings producing time series indistinguishable "by eye" from a stationary AR(2) time series (cf. Timmer, Fig. 1). Estimated D_2 thus seems attractive for assessing the question of the impact of nonstationarity on the assessment of nonlinearity in EEG data.

Linear summary measure

We employed the Ljung-Box Q-statistic (Q_{LB}) (Ljung and Box, 1978) as a linear summary measure. For a time series of length N, letting r_k symbolize the autocorrelation at lag k, this statistic is given by the following equation:

$$Q_{LB} = (N)(N+2)[\sum_{k=1}^{L} (r_k^2)/(N-k)]$$

Here, L is a parameter to be set. Following the recommendation of Diebold (1997), we set L equal to $N^{1/2}$ (that is, 32 lags). Under the null hypothesis that all $r_k = 0$ (that is, that the series is a white noise), Q_{LB} distributes as χ^2 with L df. In effect, Q_{LB} summarizes how sustained the ACF is and thus is solely a function of the linear properties of the data.

THE "FINAL" NONLINEAR SUMMARY STATISTIC

Each of the five EEG channels was divided into 15 eight-s segments overlapping by four s. For each segment, D_2 was estimated and Q_{LB} computed for the original segment and 1,000 PAR surrogates. Either surrogate distribution was rejected if there was significant evidence of non-normality (Jarque-Bera statistic (Jarque and Bera 1980; see Appendix 2) >5.99; $P<0.05$), and a new set of surrogates was generated. This process was iterated until the normality criteria were achieved. For both D_2 and Q_{LB}, a measure termed NSigmas was then computed as ((original-data statistic) - (average surrogate statistic))/(surrogate standard deviation). (These were equivalent to z-scores, since the surrogate distributions were essentially normal). A "final" estimate of nonlinearity was then taken as (D_2 NSigmas) plus (Q_{LB}

NSigmas). Thus, for example, if D_2 NSigmas was +4 (indicating greater nonlinearity) but Q_{LB} NSigmas was -4 (indicating corresponding reduced linearity, i.e., a less sustained ACF), then the final nonlinear summary measure would be 0, indicating that the D_2-based nonlinearity was spurious.

Assessing time ireversibility

We employed two tests of time irreversibility, the TIR statistic and the TR test. The TIR statistic is based on nonlinear cross prediction using a KNNLL nonlinear-prediction algorithm. First, the data are nonlinearly modeled. That model is then applied to a time-reversed copy of the data. We employed the procedure outlined in (Stam et al. 1998a) with the exception that TIR was defined as the normalized average prediction error forward in time minus that backwards in time. If the null hypothesis of time reversibility is true, the expected value of TIR is zero. Increasingly positive TIR values indicate increasing evidence of time irreversibility.

The TR test (Ramsey and Rothman 1996) examines the temporal symmetry of the bicovariance function of a time series and its residuals for a number of lags (we employed 3; for a given model of a time series x, the residuals are defined as (predicted x_t) minus (actual x_t)). Our procedure was modeled after Rothman (1996) with the exception that we employed a lowest-AIC AR model (see Appendix 2) rather than an ARMA (autoregressive moving average) model (the former has an exact solution rather than requiring numerical estimation). Details regarding our implementation of the TR test may be found in Appendix 3.

The TR test attempts to distinguish between two types of time irreversibility. So-called Type I irreversibility can (but doesn't necessarily have to) result from a nonlinear stochastic process with Gaussian inputs. In contrast, Type II irreversibility results from a linear-stochastic process with non-Gaussian inputs. For further details, see Appendix 3.

Assessing nonstationarity

The first approach to assessing nonstationarity was straightforward: in each segment, the variance of the first four statistical moments across that segment was examined. For this purpose, each EEG segment was divided into 16 half-second sub-segments, and the mean, variance, skew, and kurtosis was computed for each sub-segment. The variance across each set of 16 was then computed for each of these four statistical moments.

Three other measures of nonstationarity from traditional linear time-series analysis were also computed: (1) Autocorrelation in the squared residuals (Harvey 1989; see Appendix 2). This measure tests the null hypothesis of a stationary variance vs. the alternative hypothesis that the time series has a form of variance nonstationarity termed ARCH (autoregressive conditional heteroskedasticity). For a time series with ARCH, the variance is an AR process in the variance; (2) The ARCH LM (Lagrange-Multiplier) test (Harvey 1989; see Appendix 2), which also examines the question of ARCH by regressing squared residuals on lagged, squared residuals (we chose the number of lags based upon a lowest-AIC criterion and employed ARCH LM adjusted R^2 as a statistic independent of the order of the best AR model); (3) The White test (White 1980) of the null hypothesis of a stationary variance against an alternative hypothesis of variance nonstationarity in-general (this test also involves a regression procedure performed on the residuals - we employed the *P*-value of the White *F*-ratio as a statistic independent of order of the best AR model).

Linear vs. nonlinear modeling

The data were linearly modeled using best-fitting AR models according to the AIC (see Appendix 2). The nonlinear predictability of the data was assessed using the locally constant nonlinear prediction (LCNP) statistic. This statistic is the output of a KNN nonlinear-forecasting algorithm. We followed the procedure outlined in (Kaplan and Glass 1995), with K = 30. If a time series has no predictability (either linear or nonlinear), the LCNP statistic has an expected value of unity. To the degree that a time series is predictable the LCNP statistic will be less than one, reaching a value of zero if the data are perfectly predictable (thus, LCNP is like adjusted R^2, but lower values indicate greater predictability). As a "final" LCNP value for each EEG segment, we chose the lowest of the 30 computed.

APPLICATION OF PROCEDURES TO EEG DATA

Nonlinearity

Treating the set of EEGs at each locus as a population of one (although, again, this set was typical of such EEGs

in general), we tested for nonlinearity using *t*-tests vs. a "population" mean of zero. There was significant nonlinearity in the occipital EEGs (for O2, $t(14) = 2.27$, $P=0.038$; for O1, $t(14) = 2.38$, $P=0.031$).

Time irreversibility

Neither the TIR statistic nor the TR test produced significant evidence of time irreversibility in the EEGs at any locus (for the TR test, this was true both for the time series itself and the residuals from the lowest-AIC (see Appendix 2) AR model).

Nonstationarity

We next sought to determine if the nonlinearity at the occipital loci was associated with nonstationarity. For the nonlinearity data from O1 and O2, there was a significant (Pearson product-moment) correlation with variance of the variance ($r_{xy} = 0.607$, $P<0.001$), variance of the skew ($r_{xy} = 0.437$, $P=0.016$), and White *P*-value ($r_{xy} = -0.431$, $P=0.017$; this correlation is negative because lower *P*-values mean greater variance nonstationarity). Thus, the nonlinearity at O1 and O2 was significantly associated with nonstationarity, in particular, variance nonstationarity (we note here that some of our EEG's did have a non-Gaussian amplitude distribution as determined by the J-B test (see Appendix 2); however, these EEGs had distributions that still were highly "mound-shaped", and degree of non-normality (J-B p-value) did not correlate with either nonlinearity or nonstationarity).

The significant relation at O1 and O2 between nonlinearity and nonstationarity might seem puzzling at first, since, again, nonstationary time series are time irreversible also (Ramsey and Rothman 1996). We believe, however, that this is due to the fact that, as outlined in the next section, the EEGs at these loci were largely stationary linear-stochastic processes.

Linear vs. nonlinear modeling

The EEGs at O1 and O2 were extremely well-modeled linearly: the average adjusted R^2 of the lowest-AIC AR models (see Appendix 2) of these EEG segments was 0.943 (SEM 0.004), leaving only (1 - average adjusted R^2 =) 5.7 % of the variance unaccounted for (as outlined in Appendix 2, adjusted R^2 is a somewhat more conservative estimate of variance accounted for relative to R^2, since a slight penalty for AR model order is imposed). Further, all our AR models passed two standard diagnostic tests for structural stability, CUSUM (cumulative sum of residuals) and CUSUM of squares (cumulative sum of squared residuals), both of which examine the temporal stability of the coefficients of an AR model (Harvey 1989).

Additional evidence that our EEGs were linear-stochastic processes came from the fact that nonlinear modeling did no better than linear. The average LCNP statistic at the occipital loci was 0.140 (SEM 0.007), a value comparable to the average adjusted R^2 of 0.943. Most notably, the LCNP statistic correlated -0.939 with adjusted R^2.

A multiple regression of nonlinearity at the occipital loci on the three significant nonstationarity correlates (variance of the variance, variance of the skew, and White *P*-value) produced a multiple adjusted R^2 of 0.479, indicating that nonstationarity accounted roughly half the variance in nonlinearity. This translates into circa (5.7% x 0.479 =) 2.7% of the variance left over from a stationary linear model. Thus, it seems no surprise that neither the TR test nor the TIR statistic could detect significant time irreversibility in the occipital EEGs, given their dynamically trivial degree of nonstationarity and hence nonlinearity. The small amount of remaining unaccounted for variance in nonlinearity presumably represents either a static, nonlinear transform (recall that some of our EEGs were slightly non-Gaussian) or perhaps less-than-perfect correction for the FFT frequency-mismatch artifact.

DISCUSSION

We employed a systematic strategy toward assessing EEG nonlinearity outlined in the "Goals" section of the Introduction and found little evidence for rejecting the LGP null hypothesis with regard to a typical example of normal, resting, eyes-closed EEG. Of course, our findings would need to be extended beyond a single participant in order to be considered applicable to such EEGs in-general, although previous results with multiple participants indicate at best only weak nonlinear structure in the alpha rhythm (Stam et al. 1999a, Brandt et al. 2000, Pritchard and Stam 2000).

Do our findings suggest that nonlinear EEG analysis should be abandoned? We believe that the answer is "no". For example, estimated D_2 used as a measure of the dynamical complexity has been shown to improve the

ability to predictively classify normal elderly from probable Alzheimer's patients (Pritchard et al. 1994), although our current results indicate that a linear measure of complexity may serve just as well. Further, there is growing evidence that the EEG in many disease states is nonlinear, including periodic lateralized epileptiform discharges (Stam et al. 1998b), terminal-stage Creutzfeldt-Jakob EEG (Stam et al. 1997), periodic complexes (Stam et al. 1999b), and frontal intermittent rhythmic delta activity (FIRDA; Stam and Pritchard 1999b). In this context, then, it would seem that EEG linearity is normal with nonlinearity indicating pathology, leading to the interesting possibility that subtle increases in EEG nonlinearity could serve as "leading indicators" of future disease states.

APPENDIECES

Appendix 1: phase-angle randomization (PAR)

The Wold Decomposition Theorem proves that any stationary, discrete-time process (see footnote 4 for a definition of a "process") can be split into two sub-processes, one deterministic and the other non-deterministic, or linear-stochastic (stationary is defined subsequently in the text). The term "linear" is used because the non-deterministic process can be represented as linearly filtered white noise, the latter being defined as a process that has no autocorrelation at any lag zero. That is, white noise is completely unpredictable from a linear standpoint (but would be nonlinearly predictable if the observations, although uncorrelated, were not independent). If the white noise has a Gaussian (normal) distribution, then it is termed Gaussian white noise (GWN). GWN is neither linearly nor nonlinearly predictable, since, as discussed in the text, normally distributed data are completely characterized by their linear properties (and hence the observations are independent as well as uncorrelated).

Let y_t represent a zero-mean, stationary linear-stochastic process (any process having a non-zero mean can be made zero-mean by subtracting the mean from each observation). The Wold Theorem states that such a process can be written as linearly filtered GWN as follows:

$$y_t = \sum_{k=0}^{\infty} b_k \varepsilon_{1-k} \qquad (1)$$

where $b_0 = 1$ and $\varepsilon_t \sim GWN(0, \sigma^2)$, the latter expression meaning the ε_t's have a Gaussian distribution and are zero-mean with variance σ^2. (1) is termed an infinite moving average, and the ε_t's are referred to as inputs (see footnote 1 for alternate terminology). They are termed inputs because, when input to the linear filter represented by the b_k's , the result is the y_t time series (finite moving averages (MAs) constitute one major class of linear-stochastic models; the other major class, autoregressive (AR) models, is discussed in Appendix 2; combined AR and MA (ARMA) models are also possible).

Recall that the Discrete Fourier Transform (DFT) allows any discrete time series to be decomposed into a sum of sine waves. That is, the DFT allows the time series to be transformed from the time domain to the frequency domain (and back again) with no loss of information. Instead of amplitude as a function of time for N data points, following the DFT we have amplitude and phase angle as a function of frequency for (N/2) +1 frequencies ranging from zero to the Nyquist frequency. (The latter is half the sampling rate (number of data points per unit time); if, as is customary, time is in seconds, then frequency will be in cycles/second, or Hertz (Hz)). Given that the Fourier transform employs sine waves as it representation, or basis, the phase angle at each frequency designates where (in terms of radians) in the cycle (of the sine wave of that frequency) the oscillation begins. Theoretically, the Fourier transform of GWN has equal amplitude at each frequency and phase angles that are a set of random numbers on the interval $(0,2\pi)$.

In the frequency domain, any linear filter (such as that represented by the b's in (1) can be represented by a transfer function consisting of two parts, the transfer gain $G(\omega)$ and the transfer phase-shift $\phi(\omega)$ (ω symbolizes angular frequency, which can be conveted to Hz by dividing by 2π (one complete cycle of a sine wave = 2π radians)). If $A_{input}(\omega)$ is the amplitude of the input signal at angular frequency ω, then the amplitude of the output signal $A_{output} = G(\omega)A_{input}$. Similarly, if $\phi i_{nput}(\omega)$ is the phase angle of the input signal at angular frequency ω, then the phase angle of the output signal $\phi_{output} = \phi(\omega)\phi_{input}$.

Again, we know from a theoretical standpoint that GWN has equal amplitude at all frequencies and random phase angles. We also know from the Wold Theorem that any linear-stochastic process can be produced by linearly filtering GWN using (1). Finally, we know that a linear filter in the time domain has an equivalent repre-

sentation in the frequency domain termed a linear transfer function. So any stationary stochastic process can be represented in the frequency domain by applying a linear transfer function to GWN. Since the phase angles of GWN are random, any linear transform performed on them (such as $\phi(\omega)$) will also be random. This means that for a linear-stochastic process, all the frequency domain information is in the amplitudes, or, as it is more customary to say, all the information is in the power spectrum (in the frequency domain, power is amplitude squared). This in-turn means that if a time series has its phase angles randomized while keeping it power spectrum constant, it has by definition been transformed into a linear-stochastic process.

Appendix 2: topics from linear time-series analysis

AUTOREGRESSIVE (AR) MODELING

An AR model fits the data based on a weighted fit of p previous values, were p is model order. Irregularities in the dynamics of an AR model are driven by a stochastic (random) parameter. For example, a second-order autoregressive (AR(2)) process is given by

$$x_t = a_0 + a_1 x_{t-1} + a_2 x_{t-2} + \varepsilon_t$$

where the x_i are the individual data points and the stochastic parameter ε_t is usually Gaussian white noise. For our AR modeling, each EEG was fit using OLS (ordinary least squares, i.e., solving the Yule-Walker equations (see Weigend 1996)) for p varying from 1 to 35. The best-fitting model was selected as the one having the lowest AIC (see immediately below). For a time series x_1, x_2, , x_N, the residuals from an AR model are equal to x_t predicted by the model minus actual x_t. They are usually symbolzied e_t and, if the realization being modeled is in fact produced by a linear procees, are IID (independently identically distributed) white noise.

AKAIKE INFORMATION CRITERION (AIC)

The AIC (Akaike 1974) imposes a penalty for increasing model order that is asymptotically efficient, that is, given a set of models, the model with the lowest AIC best approximates the true data-generating process in terms of one-step-ahead prediction error variance (Diebold 1997). This prevents "overfitting" the data, i.e., memorizing a particular realization to the point that the ability to predict other realizations of the same process suffers. The AIC is computed as follows:

$$AIC = [exp\ (2p/N)\ ]\ \{[\sum_{t=1}^{N} e_t^2]/N\}$$

where e_t are the residuals from the model and p is the model order. The AIC is the gold standard of best-fit criteria in linear time-series analysis.

Adjusted R^2

The adjusted R^2 of a best-fitting AR model gives a slightly conservative estimate of variance accounted for by the model and is computed as follows:

$$R^2 adj. = 1 - \{[\sum_{t=1}^{N} e_t^2/N - p]/[\sum_{t=1}^{N} (x_t - M)^2/N - 1]\}$$

where M is the mean value of the time series. It is "slightly conservative" since it imposes a small penalty on model order given by the quantity inside the left set of brackets (the sum of squared residuals (see AIC above) over N-p).

THE JARQUE-BERA (J-B) STATISTIC

Symbolizing the skew of a size-N distribution as S and the kurtosis as K, the J-B statistic is computed as follows:

$$J\text{-}B = [N/6][S^2 + (1/4)(K-3)^2]$$

For a normal distribution, S = 0 (the distribution is symmetric) and K = 3. Under the null hypothesis of normality, J-B is distributed as χ^2 with 2 df.

AUTOCORRELATION OF SQUARED RESIDUALS

This test computes the Q_{LB} statistic for the squared residuals from a best-fitting linear model. In this case, the df are equal to the m parameter of the Q_{LB} formula given in the text minus the order of the linear model.

ARCH LM TEST

This procedure is based on the regression of squared residuals from a best-fitting linear model on lagged,

squared residuals. The number of lagged residuals needs to be specified. For two lags, the regression equation would be:

$$e_t^2 = a_1 + a_2 e_{t-1}^2 + a_3 e_{t-2}^2$$

The null hypothesis is that the coefficients a_i of the lagged, squared residuals are all zero (no ARCH). The procedure is a Lagrange multiplier (LM) test (for mathematical details, see Harvey 1989) resulting in a χ^2 having df equal to the number of lagged, squared residuals.

WHITE'S TEST

This test is performed on the residuals of a best-fitting linear model. If the autoregression is, for example, second order:

$$x_t = a_1 x_{t-1} + a_2 x_{t-2}$$

then White's test is based on the augmented regression:

$$e_t^2 = a_1 + a_2 x_{t-1} + a_3 x_{t-2} + a_4 x_{t-1}^2 + a_5 x_{t-2}^2 + a_6 x_{t-1} x_{t-2}$$

where e_t symbolizes the residuals. The output of the test is a statistic that has an asymptotic χ^2 distribution with df equal to the number of regressors in the test regression. The statistic tests the null hypothesis that the coefficients a_i of all the variables in the augmented regression are zero. White offered the procedure as a general test of linear-model misspecification, since the null hypothesis will be rejected if the residuals are heteroskedastic (the variance of the residuals is changing over time), if the residuals and the regressors are not independent, and/or an incorrect linear model has been specified.

Appendix 3: The TR test

First, each EEG segment is transformed to zero mean and unit variance. The sample bicovariances for a range of k are then computed as:

$$B_{2,1}(k) = [(N-k)^{-1}] \left[\sum_{t=k+1}^{N} (x_t^2)(x_{t-k})\right] \qquad (1a)$$

$$B_{1,2}(k) = [(N-k)^{-1}] \left[\sum_{t=k+1}^{N} (x_t)(x_{t-k^2})\right] \qquad (1b)$$

Sample gamma statistics are then computed as:

$$g_{2,1}(k) = B_{2,1}(k) - B_{1,2}(k). \qquad (1c)$$

A value for gamma by itself is meaningless without a standard-deviation estimate. If x_t is an IID random variable, then it can be shown that:

$$Var[g_{2,1}(k) = \{2[(\mu_4)(\mu_2) - \mu_3^2]/(N-k)\} - \\ - 2[(\mu_2^3)(N-2k)/[(N-k)^2] \qquad (2)$$

where

$$\mu_2 = E[x_t^2],\ \mu_3 = E[x_t^3], \text{and} \mu_4 = \mu_4 = E[x_t^4]$$

The variance formula given by (2) will not be accurate for time series that are not IID, in which case the variance must be estimated by Monte Carlo. Following Ramsey and Rothman, this can be done by (a) finding a best-fitting (lowest-AIC) AR model of each EEG (see Appendix 1), (b) estimating the variance of the stochastic inputs to the model, (c) generating a time series that is a realization of the model having stochastic inputs that are $N(0,s^2)$, with s^2 equal to the variance estimated in step b, and (d) computing the gamma statistics of the realization using equation set (1). Steps c and d were repeated 1,000 times to obtain good estimates of the gamma-statistic standard deviations for k = 1 to 30. Using these standard deviations, the 30 gamma statistics were standardized. Two test statistics were then computed, a portmanteau statistic representing the sum of square of the 30 standardized gammas, and a gmax statistic representing the maximum gamma.

For each value of k, p-values assessing the time reversibility null hypothesis for the EEG in question were estimated using a second 1000-iteration Monte Carlo. The first two steps were the same as (c) and (d) above. In step (c) the 30 gamma statistics of the realizations were estimated using the 30 standard deviations obtained above, and the portmanteau and gmax statistics computed. For each value of k, a count Np of the number of times that original-data portmanteau exceeded realization portmanteau was made. A similar counted Ng was made for gmax. For each statistic, a *P*-value was estimated as 1 - ((final count)/1000). A final *P*-value was take as the average of the two. The procedure is based on the fact that the realizations are known to be time-re-

versible LGPs. Thus, if the original data are time irreversible, one would expect its portmanteau and gmax to be greater than the corresponding statistic for the realizations a significant number of times (at least 950 out of 1,000).

A third Monte Carlo yielded a *P*-value for residuals from the best linear model. First, portmanteau and gmax statistics were computed for the original-data residuals using formula (2) (it was assumed the residuals would be "IID-enough" to use this formula). Then, across 1,000 iterations, counts Np and Ng were made as above for residuals, not from the original data, but from the best linear model of a realization of the original-data best linear model. A "final" *P*-value was again taken as the average of the portmanteau and gmax *P*-values.

The procedure distinguishes Type I from Type II irreversibility as follows: If a given time series is a linear--stochastic process, then the time series should be well-modeled linearly and the residuals should be IID (independently, identically distributed), and therefore time reversible. Thus a finding that a given time series is time irreversible but its residuals are time reversible would indicate Type II time irreversibility (linear-stochastic process with non-Gaussian inputs). In contrast, if a given time series is a nonlinear stochastic process, then a linear model will do less well, and time irreversibility in the residuals should still be detectable with probability greater than the size of the test. Thus, a finding that both a time series and its (best-linear model) residuals are time irreversible would indicate Type I time irreversibility (nonlinear-stochastic process with Gaussian inputs). Empirically, in simulations we have found that lowest-AIC AR models of chaotic realizations of the Lorenz system result in an excellent fit (see main text). When the TR test is applied to such a time series, although the time series is labeled as reversible, the residuals are labeled as highly irreversible. This is because, although the AR model provides an excellent fit, the residuals still reflect the chaotic nature of the system (unlike the residuals from an AR model applied to a linear-stochastic process, which are IID white noise).

ACKNOWLEDGEMENTS

The authors thank Dr. Philip Rothman for his advice in implementing the TR test. The term "chaos sightings" was borrowed from Dr. James Theiler. Based on our results, we propose that the similar term "nonlinearity sightings" may be applicable to the EEG also.

REFERENCES

Akaike H.A. (1974) A new look at the statistical model identification. IEEE Trans. Auto. Control 19: 716-723.

Brandt M.E., Ademoglu A., Pritchard W.S. (2000) Nonlinear prediction and complexity of alpha EEG activity. Int. J. Bifurc. Chaos (in press)

Cantero J.L., Atienza A., Salas R.M., Gómez C.M. (1999) Brain spatial microstates of human spontaneous alpha activity in relaxed wakefulness, drowsiness period, and REM sleep. Brain Topog. 11: 257-263.

Casdagli M. (1989) Nonlinear prediction of chaotic time series. Physica D 35: 335-356.

Diebold F.X. (1997) Elements of Forecasting. South-Western College Publishing, Cincinnati.

Diks C., van Houwelingen J.C., Takens F., DeGoede J. (1995) Reversibility as a criterion for discriminating time series. Phys. Let. A 201: 221-228.

Ehlers C.L.,Ehlers C.L., Havstad J., Prichard D., Theiler J. (1998) Low doses of ethanol reduce evidence for nonlinear structure in brain activity. J. Neurosci. 18: 7474-7486.

Ellner S. (1988) Estimating attractor dimensions from limited data: A new method, with error estimates. Phys. Let. A 133: 128-133.

Harvey A.C. (1989) Forecasting, structural time series models and the Kalman Filter. Cambridge University Press, Cambridge, 554 p.

Jarque C.M., Bera A.K. (1980) Efficient tests for normality, homoscedasticity and serial independence of regression residuals. Econom. Let. 6: 255-259.

Kaplan D., Glass L. (1995) Understanding nonlinear dynamics. Springer-Verlag, New York, 480 p.

Kantz H., Schreiber T. (1997) Nonlinear time series analysis. Cambridge University Press, Cambridge, p. 304.

Kennel M.B., Brown R., Abarbanel H.D.I. (1992) Determining minimum embedding dimension using a geometrical construct. Phys. Rev. A 45: 3403-3411.

Kennel M.B., Isabelle S. (1992) Method to distinguish possible chaos from colored noise and to determine embedding parameters. Phys. Rev. A 46: 3111-3118.

Lehmann D (1981) Spatial analysis of evoked and spontaneous EEG potential fields. In: Recent advances in EEG and EMG data processing (Eds. N. Yamaguchi and K. Fujisawa). Elsevier, Amsterdam.

Ljung G.M., Box G.E.P. (1978) On a measure of lack of fit in time series models. Biometrika 66: 67-72.

Palus M. (1996) Nonlinearity in normal human EEG: Cycles, temporal asymmetry, nonstationarity and randomness, not chaos. Biol. Cybern. 75: 389-396.

Pritchard W.S., Duke D.W. (1992a) Dimensional analysis of no-task human EEG using the Grassberger-Procaccia method. Psychophysiology 29: 182-192.

Pritchard W.S., Duke D.W. (1992b) Measuring chaos in the brain: A tutorial review of nonlinear dynamical EEG analysis. Int. J. Neurosci. 67: 31-80.

Pritchard W.S., Duke D.W., Coburn K.L., Moore N.C., Tucker K.A., Jann M.W., Hostetler R. (1994) EEG-based, neural-net predictive classification of Alzheimer's disease versus control subjects is augmented by nonlinear EEG measures. Electroencephalogr. Clin. Neurophysiol. 19: 118-130.

Pritchard W.S., Duke D.W., Krieble K.K. (1995) Dimensional analysis of resting human EEG II: Surrogate-data testing indicates nonlinearity but not low-dimensional chaos. Psychophysiology 32: 486-491.

Pritchard W.S., Stam C.J. (2000) A new type of surrogate data for highly periodic time series: application to human alpha EEG comparing inferential with non-inferential statistical analysis. In: Nonlinear dynamics and brain functioning (Eds. N. Pradhan, P.E. Rapp and R. Sreenivasan). Nova Science, Huntington, NY. (in press)

Ramsey J.B., Rothman P. (1996) Time irreversibility and business cycle asymmetry. J. Money Credit Banking 28: 1-21.

Rombouts S.A.R.B., Keunen R.W.M., Stam C.J. (1995) Investigation of non-linear structure in multichannel EEG. Phys. Lett. A 202: 352-358.

Rothman P. (1996) FORTAN programs for running the TR test: A guide and examples. East Carolina University, Department of Economics, available on the internet at http://ecuvax.cis.ecu.edu/~ecrothma/rothman.htm.

Rosenstein M.T., Collins J.J., De Luca C.J. (1994) Reconstruction expansion as a geometry-based framework for choosing proper delay times. Physica D 73: 82-98.

Schreiber T., Schmitz A. (1996) Improved surrogate data for nonlinearity tests. Phys. Rev. E 77: 635-638.

Stam C.J., Nicolai J., Keunen R.W.M. (1998b) Nonlinear dynamical analysis of periodic lateralized epileptifirm discharges. Clin. Electroencephalogr. 29: 101-105.

Stam C.J., Pijn J.P.M., Pritchard W.S. (1998a) Reliable detection of non-linearity in experimental time series with strong periodic components. Physica D 112: 361-380.

Stam C.J., Pijn J.P.M., Suffczynski P., Lopes da Silva F.H. (1999a) Dynamics of the human alpha rhythm: Evidence for non-linearity?. Clin. Neurophysiol. 110: 1801-1813.

Stam C.J., Pritchard W.S. (1999) Dynamics underlying rhythmic and non-rhythmic variants of abnormal, waking delta activity. Int. J. Psychophysiol. 34: 5-20.

Stam C.J., van Woerkom T.C.A.M., Keunen R.W.M. (1997) Non-linear analysis of the electroencephalogram in Creutzfeldt-Jakob disease. Biol. Cybern. 77: 247-256.

Stam C.J., Vliegen J.H.R., Nicolai J. (1999b) Investigation of the dynamics underlying periodic complexes in the EEG. Biol. Cybern. 80: 57-69.

Takens F. (1985) On the numerical determination of the dimension of an attractor. Lect. Notes Math. 1125: 99-115.

Theiler J., Eubank S., Longtin A., Galdrikian B., Farmer J.D. (1992) Testing for nonlinearity in time series: the method of surrogate data. Physica D 58: 77-94.

Timmer J. (1998) Power of surrogate data testing with respect to nonstationarity. Phys. Rev. E 58 :5153-5156.

Weigend A.S. (1996) Time series analysis and prediction. In: Time series analysis and prediction (Eds. P. Smolensky, M.C. Mozer and D.E. Rumelhart). Lawrence Erlbaum, Hillsdale, NJ, p. 1-55.

Weiss G. (1975) Time-reversibility of linear stochastic processes. J. Appl. Prob. 12: 831-836.

White H. (1980) A heteroskedasticity-consistent covariance estimator and a direct test for heteroskedasticity. Econometrica 48: 817-838.