Available online at www.sciencedirect.com
SCIENCE DIRECT®
Physica D 190 (2004) 167–176
PHYSICA D
www.elsevier.com/locate/physd

The delay vector variance method for detecting determinism and nonlinearity in time series

Temujin Gautama [a,*], Danilo P. Mandic [b], Marc M. Van Hulle [a]

[a] *Laboratorium voor Neuro en Psychofysiologie, K.U. Leuven, Campus Gasthuisberg, Herestraat 49, B-3000 Leuven, Belgium*
[b] *Department of Electrical and Electronic Engineering, Imperial College of Science, Technology and Medicine, Exhibition Road, SW7 2BT London, UK*

Received 4 April 2002; accepted 6 November 2003
Communicated by R.E. Goldstein

Abstract

A novel 'delay vector variance' (DVV) method for detecting the presence of determinism and nonlinearity in a time series is introduced. The method is based upon the examination of local predictability of a signal. Additionally, it spans the complete range of local linear models due to the standardisation to the distribution of pairwise distances between delay vectors. This provides consistent and easy-to-interpret diagrams, which convey information about the nature of a time series. In order to gain further insight into the technique, a DVV scatter diagram is introduced, which plots the DVV curve against that for a globally linear model (surrogate data). This way, the deviation from the bisector line represents a qualitative measure of nonlinearity, which can be used additionally for constructing a quantitative measure for statistical testing. The proposed method is compared to existing methods on synthetic, as well as standard real-world signals, and is shown to provide more consistent results overall, compared to other, established nonlinearity analysis methods.

PACS: 05.45.−a; 05.45.Tp

Keywords: Nonlinearity analysis; Surrogate data; Deterministic; Stochastic

1. Introduction

The analysis of the predictability of a time series has received ample attention in the signal processing community, since it is a necessary (yet insufficient) condition for indicating the presence of deterministic chaos. However, since many nonlinearity analysis techniques rest upon chaos theory, determinism and nonlinearity have been confounded in a single analysis, and, subsequently, methods analysing *only* determinism in a time series are not widespread.

To this cause, the methods introduced by Kaplan [4] and Kennel et al. [6], both rest upon the examination of this predictability of a time series. Furthermore, Kaplan's δ–ϵ method has been used in combination with the surrogate data strategy for examining the linear or nonlinear nature of a time series [5]. However, a unified analysis

* Corresponding author. Tel.: +32-16-345960; fax: +32-16-345960.
E-mail addresses: temu@neuro.kuleuven.ac.be (T. Gautama), d.mandic@ic.ac.uk (D.P. Mandic), marc@neuro.kuleuven.ac.be (M.M. Van Hulle).

0167-2789/$ – see front matter © 2003 Elsevier B.V. All rights reserved.
doi:10.1016/j.physd.2003.11.001

of the predictability and the degree of nonlinearity remains an open issue. It is natural to ask, for instance, whether predictability implies the presence of nonlinearities. Therefore, there is a need for a comprehensive study to provide further insight into this question.

We propose a novel analysis method, the 'delay vector variance' (DVV) method, which follows an approach similar to Kaplan's, namely that of characterising a time series based upon its predictability, and comparing the result to those obtained for linearised versions of the signal (surrogates). Due to the standardisation within the algorithm, the method is robust to the presence of noise, and further allows for both an easy-to-interpret qualitative visualisation, and a quantification of the analysis results. To further illustrate the benefits of the proposed approach, we investigate the similarities and differences between the proposed approach and that described in [4], and perform a comparative study including other widespread nonlinearity analysis methods.

2. Nonlinearity methods

Several established nonlinearity analysis methods, which are briefly described in this section, are based on the time-delay embedding representation of a time series, $\mathbf{X} = \{\mathbf{x}(k)|k = 1, \ldots, N\}$, i.e., based upon a set of N delay vectors (DVs) of a given embedding dimension m, denoted by $\mathbf{x}(k) = [x_{k-m}, \ldots, x_{k-1}]$, a vector containing m consecutive time samples. Every DV $\mathbf{x}(k)$ has a corresponding *target*, namely the following sample, x_k.

Nonlinearity analyses can be performed by computing a test statistic for the original time series and for its 99 so-called 'surrogates' (using a one-tailed rank test, this allows for a statistical test at the level of 0.01). These are realisations of a composite null hypothesis, namely that the time series is generated from a linear and stationary process driven by a Gaussian white noise input, the output of which is amplitude transformed by a zero-memory observation function. Many methods exist for the generation of such surrogate time series (for an overview, see [11]). We have opted for the iterative amplitude adjusted Fourier transform (iAAFT) method, which yields time series with amplitude spectra identical to that of the original time series, and approximately identical signal distributions.[1] In the sequel, we briefly tackle several key issues on time series characterisation, such as deterministic versus stochastic, and linear versus nonlinear behaviour.

2.1. Deterministic versus stochastic plots

The method introduced by Casdagli [3] examines the (robust) average prediction error $E(n)$ for local linear models of a given embedding dimension m. The degree of locality is controlled by the number of nearest DVs (in the Euclidean sense in the m-dimensional space), n, that are used for estimating the model parameters. The complete set of DVs is divided into a test set, $\mathbf{V}_{\text{test}}$, and a training set $\mathbf{V}_{\text{train}}$. For every DV $\mathbf{x}(k)$ in the test set, a subset Ω_k is generated by grouping the n DVs in the training set that are nearest to $\mathbf{x}(k)$. The prediction error, $E(n)$, is computed as the mean (robust) prediction error over $\mathbf{V}_{\text{test}}$, namely

$$E(n) = \langle |x_k - \hat{x}_k| \rangle_{\mathbf{x}(k) \in \mathbf{V}_{\text{test}}}, \tag{1}$$

where $\hat{x}_k$ is the output generated by a linear model, when provided with $\mathbf{x}(k)$ as its input, and x_k the target for $\mathbf{x}(k)$. The model parameters are determined using the set Ω_k ($\Omega_k \subset \mathbf{V}_{\text{train}}$, thus $\mathbf{x}(k) \notin \Omega_k$). The mean prediction error $E(n)$ is then computed as a function of the degree of locality (number n of DVs in the sets Ω_k). The resulting plots, representing E as a function of the number of nearest neighbours n, are referred to as 'DVS plots'. The number of DVs in the sets Ω_k yielding the lowest mean prediction error, n_{opt}, i.e., the position of the minimum in the DVS

[1] We apply the end-to-end compensation to alleviate the bias towards a flatter amplitude spectrum (for a detailed discussion, see [7,11]).

plot, is used as an indicator of the nature of the time series under examination. A minimum on the left-hand side (near the 'local linear extreme') indicates a deterministic nature, and a minimum on the right-hand side (near the 'global linear extreme') indicates a linear and stochastic nature. Minima occurring in between the two extremes, for increasing values of n_{opt}, correspond to gradually fitting 'more linear' and 'more stochastic' models, and are an indication of nonlinearity [3]. However, the DVS method does not allow for a quantitative analysis.

2.2. Traditional nonlinearity metrics

To undertake the performance comparison between the proposed DVV method and other nonlinearity analysis methods, we have implemented two traditional measures of nonlinearity, which have also been used in [9], namely the third-order autocovariance (C3):

$$t^{\text{C3}}(\tau) = \langle x_k x_{k-\tau} x_{k-2\tau} \rangle, \tag{2}$$

and a measure of the deviation due to time reversibility (REV):

$$t^{\text{REV}}(\tau) = \langle (x_k - x_{k-\tau})^3 \rangle, \tag{3}$$

where τ is a time lag for which simplicity and convenient comparison is set to unity in all simulations. In combination with the surrogate data strategy, both measures yield two-tailed tests for nonlinearity.

2.3. The δ–ϵ method

The method proposed by Kaplan [4] was initially used for examining the degree of predictability of a time series without constructing its model, assuming a continuous function that maps the DVs onto their corresponding targets. The analysis can be summarised as follows:

- The pairwise (Euclidean) distances between DVs $\mathbf{x}(i)$ and $\mathbf{x}(j)$ are computed and denoted by $\delta_{i,j}$. The distance between corresponding targets is denoted by $\epsilon_{i,j}$.
- The ϵ-values are averaged, conditional to δ, i.e., $\epsilon(r) = \bar{\epsilon}_{j,k}$, for $r \leq \delta_{j,k} < r + \Delta r$, where Δr denotes the width of the 'bins' used for averaging $\epsilon_{j,k}$.
- The smallest value for $\epsilon(r)$ is denoted by $E = \lim_{r \to 0} \epsilon(r)$, and is a measure for the predictability of the time series.

Thus, the heuristic for determining E is the Y-intercept of the linear regression of the N_δ (δ, ϵ)-pairs with smallest δ. This value can be used as a test statistic for a left-tailed nonlinearity test using surrogate data.[2] In our simulations, we have set $N_\delta = 500$.

3. Proposed method

We introduce a novel analysis of a time series which examines the predictability of a time series by virtue of the observation of the variability of the targets. The approach is somewhat related to the false nearest neighbours [6] and the δ–ϵ methods described earlier.

For a given embedding dimension m, the mean target variance, σ^{*2}, is computed over all sets Ω_k. A set Ω_k is generated by grouping those DVs that are within a certain distance to $\mathbf{x}(k)$, which is varied in a manner standardised

[2] A Matlab implementation of the δ–ϵ method and the test statistic is publicly available from http://www.macalester.edu/~kaplan/Software/.

with respect to the distribution of pairwise distances between DVs. This way, the threshold scales automatically with the embedding dimension m, as well as with the dynamical range of the time series at hand, and thus, the complete range of pairwise distances is examined.[3] The proposed DVV method can be summarised as follows for a given embedding dimension m:

- The mean, μ_d, and standard deviation, σ_d, are computed over all pairwise distances between DVs, $\|\mathbf{x}(i) - \mathbf{x}(j)\|$ $(i \neq j)$.
- The sets Ω_k are generated such that $\Omega_k = \{\mathbf{x}(i) | \, \|\mathbf{x}(k) - \mathbf{x}(i)\| \leq \tau_d\}$, i.e., sets which consist of all DVs that lie closer to $\mathbf{x}(k)$ than a certain distance τ_d, taken from the interval $[\min\{0, \mu_d - n_d\sigma_d\}; \mu_d + n_d\sigma_d]$, e.g., uniformly spaced, where n_d is a parameter controlling the span over which to perform the DVV analysis.
- For every set Ω_k, the variance of the corresponding targets, σ_k^2, is computed. The average over all sets Ω_k, normalised by the variance of the time series, σ_x^2, yields the measure of unpredictability, σ^{*2}:

$$\sigma^{*2} = \frac{(1/N)\sum_{k=1}^{N} \sigma_k^2}{\sigma_x^2}. \tag{4}$$

We only consider a variance measurement *valid*, if the set Ω_k contains at least 30 DVs.

As a result of the standardisation of the distance axis, the resulting 'DVV plots' are easy-to-interpret. The presence of a strong deterministic component will lead to small target variances for small spans. At the extreme right, the DVV plots smoothly converge to unity, since for maximum spans, *all* DVs belong to the same set, and the variance of the targets is equal to the variance of the time series. If this is not the case, the span parameter, n_d, should be increased.

In the following step, the linear or nonlinear nature of the time series is examined by performing DVV analyses on both the original and a number of surrogate time series, using the optimal embedding dimension of the original time series. Due to the standardisation of the distance axis, these plots can be conveniently combined in a *scatter diagram*, where the horizontal axis corresponds to the DVV plot of the original time series, and the vertical to that of the surrogate time series. If the surrogate time series yield DVV plots similar to that of the original time series, the 'DVV scatter diagram' coincides with the bisector line, and the original time series is likely to be linear. The deviation from the bisector line is, thus, an indication of nonlinearity, and can be quantified by the root mean square error (RMSE) between the σ^{*2}'s of the original time series and the σ^{*2}'s averaged over the DVV plots of the surrogate time series (note that while computing this average, as well as with computing the RMSE, only the valid measurements are taken into account). In this way, a single test statistic is obtained, and traditional (right-tailed) surrogate testing can be performed (the deviation from the average is computed for the original, and surrogate time series).

4. Simulations

To verify the proposed analysis, a number of time series are generated, for which we can control the predictability and the degree of nonlinearity.

4.1. Signals

A unit-variance deterministic signal (sum of three sine waves, scaled to unit variance) is contaminated with uniformly distributed white noise with standard deviation σ_n. After standardising to unit variance, the resulting

[3] For computational reasons, we restrict the analysis by computing the pairwise distances between a subset of 500 DVs and the complete set of DVs, selected by subsampling the DVs. Note that this is not equivalent to subsampling the time series.

signal, n_k, is passed through a second-order nonlinear nonlinear system, described by

$$x_k = \arctan\left(\gamma_{\mathrm{nl}}\mathbf{C}^{\mathrm{T}}\mathbf{x}(k)\right) + \tfrac{1}{2}n_k,$$

where γ_{nl} controls the degree of nonlinearity, $\mathbf{C} = [0.2, -0.5]$, and $\mathbf{x}(k)$ are DVs of embedding dimension $m = 2$. This is a benchmark nonlinear system referred to as model II in [8]. In this way, the predictability is influenced by σ_{n}, whereas the degree of nonlinearity is controlled by γ_{nl}. In total, we generate nine time series, with $\sigma_{\mathrm{n}} \in \{0, 0.25, 0.5\}$ and $\gamma_{\mathrm{nl}} \in \{0, 0.5, 1.0\}$. We refer this set of signals as the 'tile' set, since, in a way, it tiles the space mentioned in [10], formed by a deterministic/stochastic and a linear/nonlinear axis.

The algorithms are further tested on five time series that have been used for benchmarking in [1]: Model 1 (deterministic chaotic), Model 2 (generalised autoregressive conditional heteroscedastic, or GARCH process), Model 3 (nonlinear moving average, NLMA process), Model 4 (autoregressive conditional heteroscedastic, ARCH process) and Model 5 (autoregressive moving average, ARMA model), each of which consist of 2000 samples. Thus, only the Model 5 series is linear (for a more detailed description, see [1]). Finally, to complete the comparison, we include four standard time series which have been analysed frequently in the context of nonlinearity, namely the Sunspots series (280 samples), the Laser data from the Santa Fe Competition (1000 samples), the first coordinate of a realisation of the Lorenz series (1000 samples), and a realisation of a Hénon series (1000 samples).

4.2. Nonlinearity analyses

For the tile set, we used an embedding dimension of $m = 2$, and all analyses for the benchmark set were performed with $m = 3$. For the standard time series, we arbitrarily set m to 2. This convenience did not influence the generality of our results. The results of the rank tests for the tile set are shown in Table 1 (significant rejections at the level of 0.1 of the null hypothesis, i.e., an underlying Gaussian linear stochastic model, the output of which is amplitude transformed, are shown in italics). Note that the DVS method is not included in this table, since it does not allow for a quantitative analysis. In the absence of noise ($\sigma_{\mathrm{n}} = 0$), only the δ–ϵ and DVV method detected nonlinearities for slopes γ_{nl} exceeding 1.5. When noise was added to the driving signals, the time reversal metric (REV) was able to detect the nonlinear nature for high slopes, but the δ–ϵ method failed. The third-order cumulant (C3) was unable to detect nonlinearities in this type of signals. Only the DVV method consistently detected nonlinear behaviour for $\gamma_{\mathrm{nl}} \geq 2$, for all noise levels.

Table 1
Results of the rank tests for the tile set[a]

γ_{nl}	σ_{n}	δ–ϵ	REV	C3	DVV
0.0	0.0	99	38	31	22
1.5	0.0	*6*	59	45	*100*
2.0	0.0	*2*	73	54	*100*
2.5	0.0	*2*	52	65	*100*
0.0	0.5	56	81	36	52
1.5	0.5	73	82	52	*98*
2.0	0.5	94	*100*	54	*100*
2.5	0.5	95	87	43	*100*
0.0	1.0	52	82	34	28
1.5	1.0	52	89	57	82
2.0	1.0	11	24	71	*100*
2.5	1.0	76	41	38	*100*

[a] Significant rejections of the null hypothesis at the level of 0.1 are given in italics.

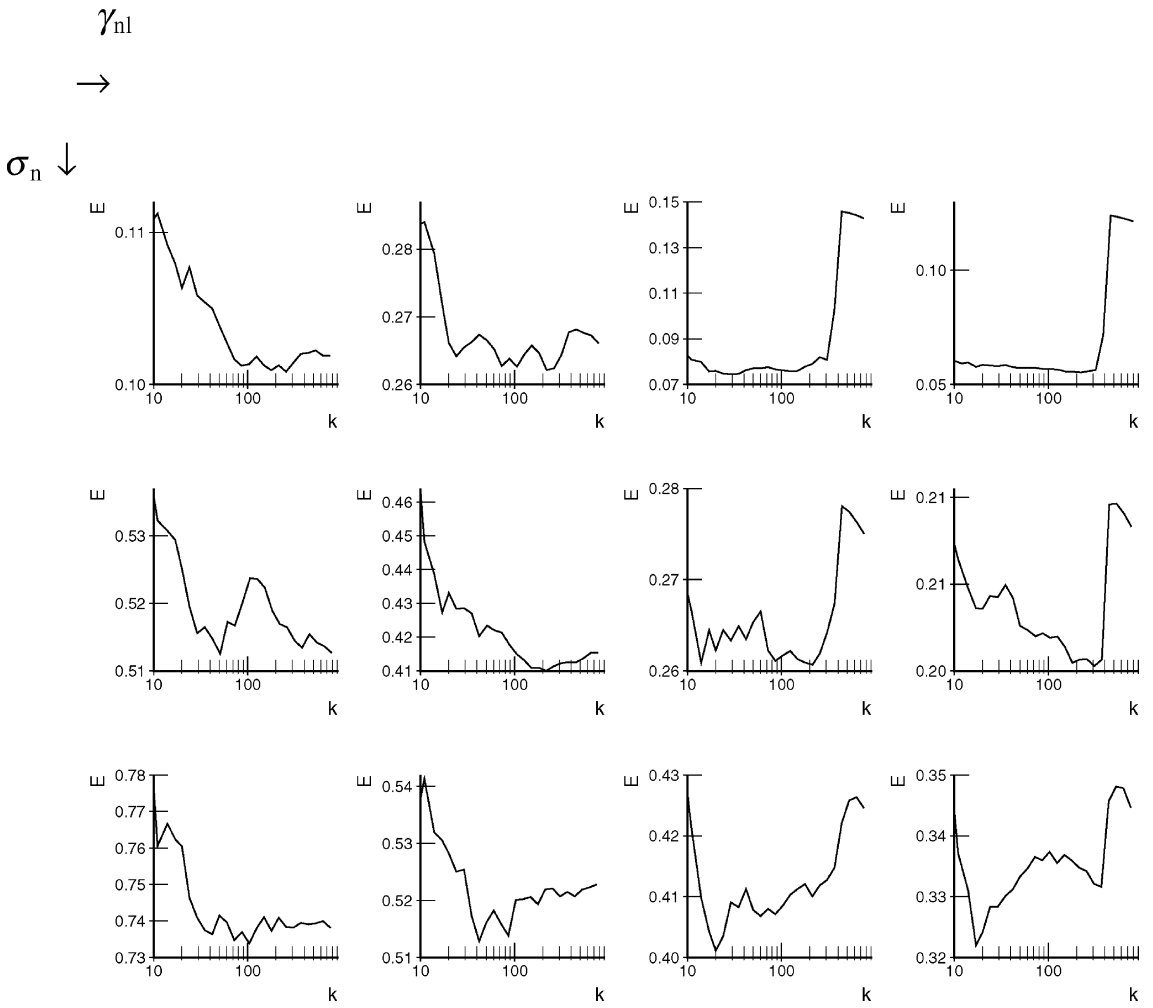

Fig. 1. DVS results for the tile set. The degree of nonlinearity increases from left to right, the noise level from top to bottom.

The results for the DVS and the DVV analyses are shown visually in Figs. 1 and 2, respectively. The degree of nonlinearity, γ_{nl}, increases from left to right, and the noise level, σ_n, increases from top to bottom. The DVS plots in Fig. 1 show that, as γ_{nl} increases, the error discrepancy between the best local linear model and the global linear model becomes larger, indicating, indeed, a higher degree of nonlinearity. In the DVV scatter diagrams (Fig. 2), the effect of increasing nonlinearity as described above, corresponds to a stronger deviation from the bisector line (dashed line). The effect of increasing σ_n in the DVS plots is a higher error value at the optimal degree of locality. The span on the horizontal axes of the DVV scatter diagrams becomes smaller as σ_n increases. Both methods are in agreement and show a gradual change as a function of the degree of nonlinearity and the noise level. Thus, for instance in the first columns of the tile figures, the lowest error increases (Fig. 1), and the horizontal range spanned by the DVV scatter diagrams decreases (Fig. 2) from top to bottom, i.e., for increasing noise levels. Conversely, considering the first line in the tile figures, from left to right, i.e., for increasing degrees of nonlinearity, the minimum becomes more pronounced in Fig. 1, and the deviation from the bisector line grows stronger in Fig. 2.

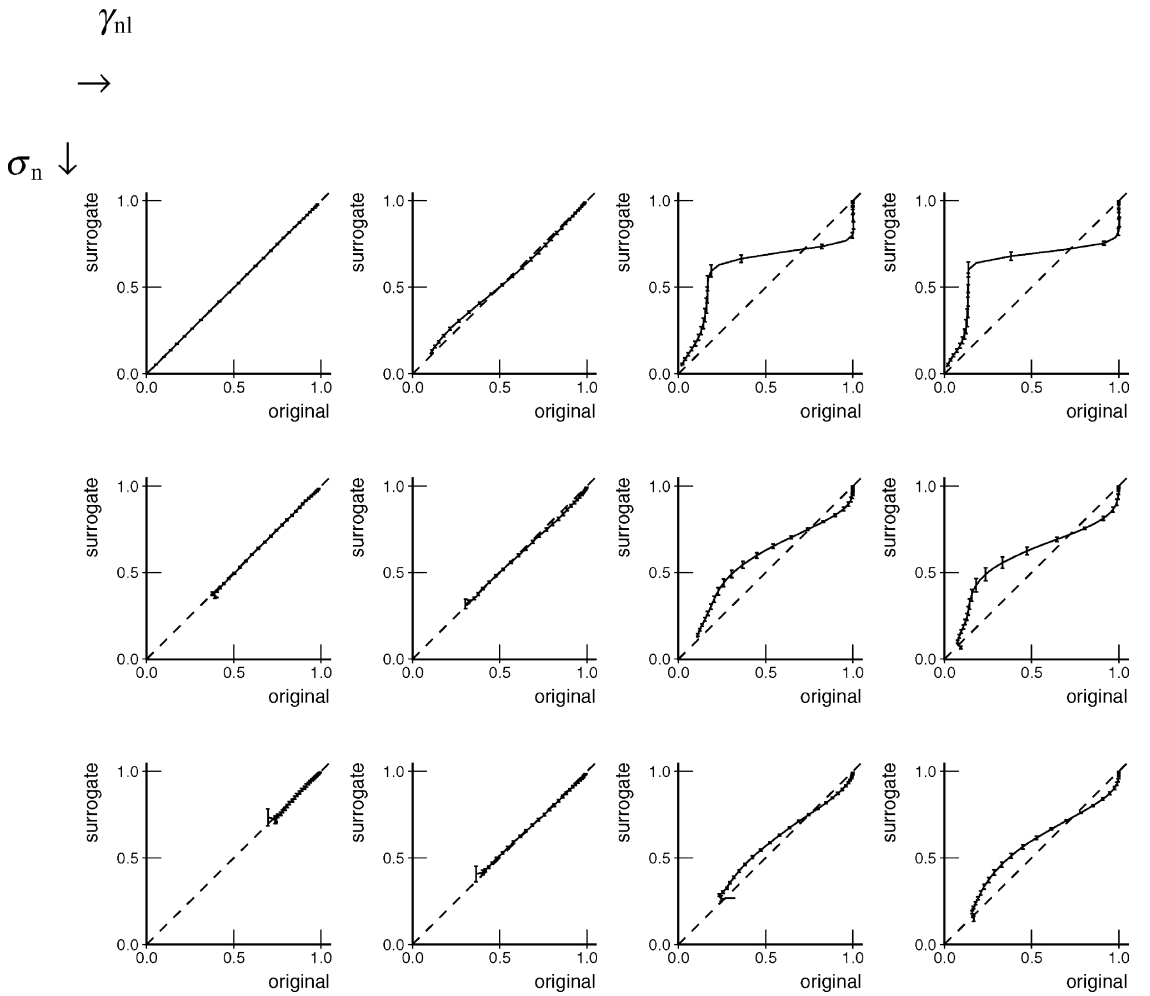

Fig. 2. DVV results for the tile set. The degree of nonlinearity increases from left to right, the noise level from top to bottom. The error bars indicate one standard deviation from the mean of σ^{*2}.

The results for the remaining time series are shown in Table 2. Since the DVS plots do not allow for a quantitative analysis, they are shown in Figs. 3 and 4 for the benchmark and standard sets, respectively. The corresponding DVV scatter diagrams are visualised in Figs. 5 and 6 for comparison. It is clear that the different methods yield different results. All methods detect nonlinearities in the Hénon and Model 4 time series, and the δ–ϵ, DVV and DVS methods consistently detect nonlinearities for all other chaotic series (Laser, Lorenz and Model 1). Nonlinear behaviour is detected in the Sunspots time series by DVS and REV. The DVV method is the only one that rejects the null hypothesis for all nonlinear signals described in [1] (Models 1–4).[4] None of the methods detect nonlinearities in the linear time series, Model 5.

[4] Note that in [1], the δ–ϵ method also rejects the linearity hypothesis for Model 2, whereas in our simulations, the null hypothesis is accepted, albeit marginally so. This can be due to our choice of the time lag τ, which has been optimised for in [1], but has been set to unity in all our simulations, or to a different approach to the estimation of E.

Table 2
Results of the rank tests for the standard and benchmark time series[a]

Signal	δ–ϵ	REV	C3	DVV
Model 1	*3*	*1*	13	*100*
Model 2	12	*97*	84	*100*
Model 3	*8*	73	*100*	*100*
Model 4	*1*	*97*	*100*	*100*
Model 5	57	41	32	16
Sunspots	30	*100*	8	43
Laser	*1*	45	*1*	*100*
Lorenz	*1*	*1*	19	*100*
Hénon	*1*	*1*	*100*	*100*

[a] Significant rejections of the null hypothesis at the level of 0.1 are given in italics.

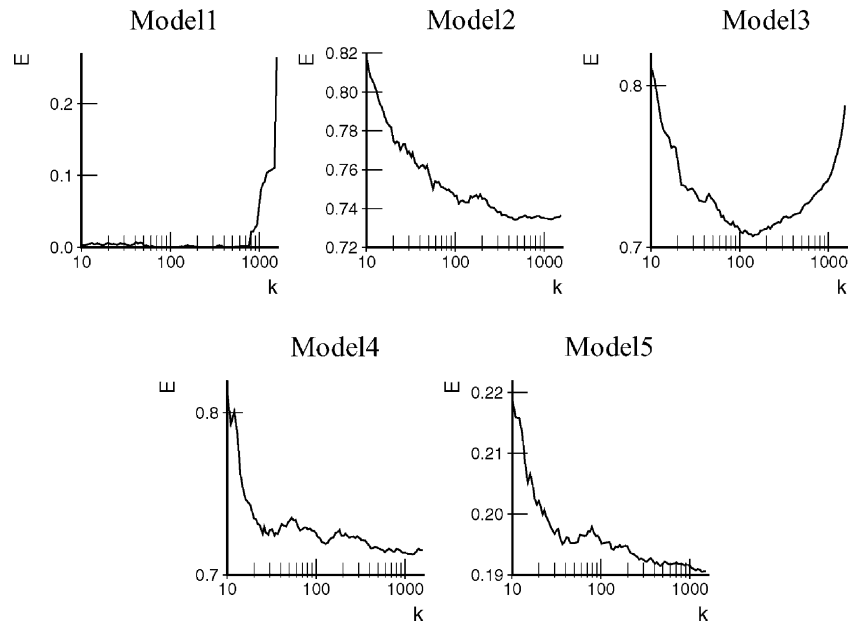

Fig. 3. DVS plots of the five benchmarks signals used in [1].

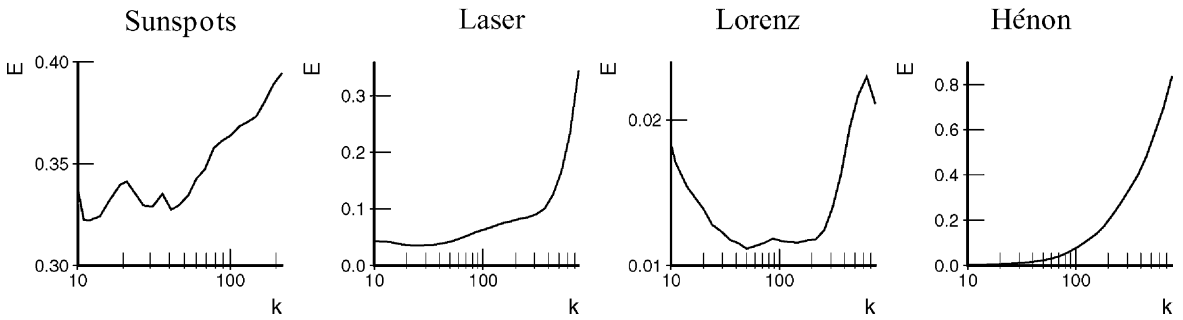

Fig. 4. DVS plots of the four standard signals.

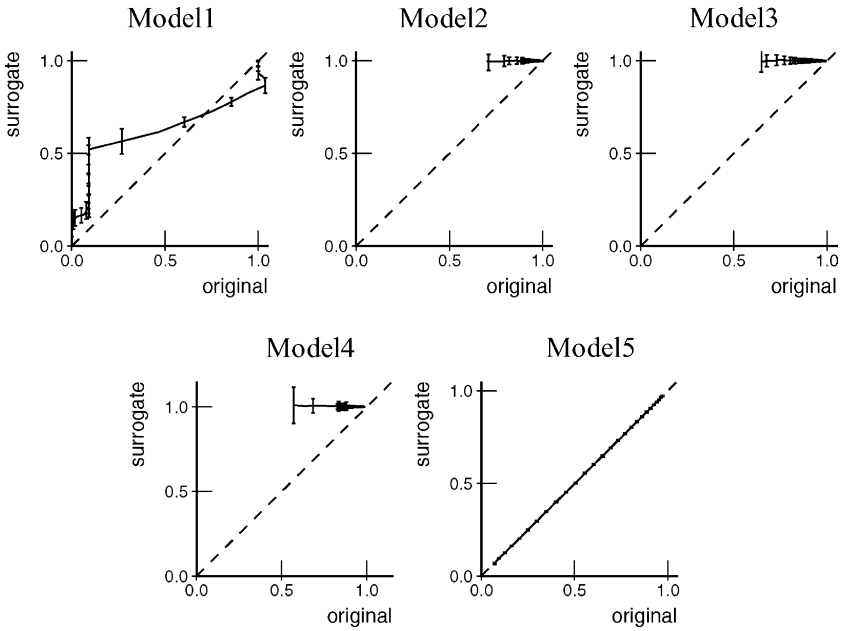

Fig. 5. DVV scatter diagrams of the five benchmarks signals used in [1].

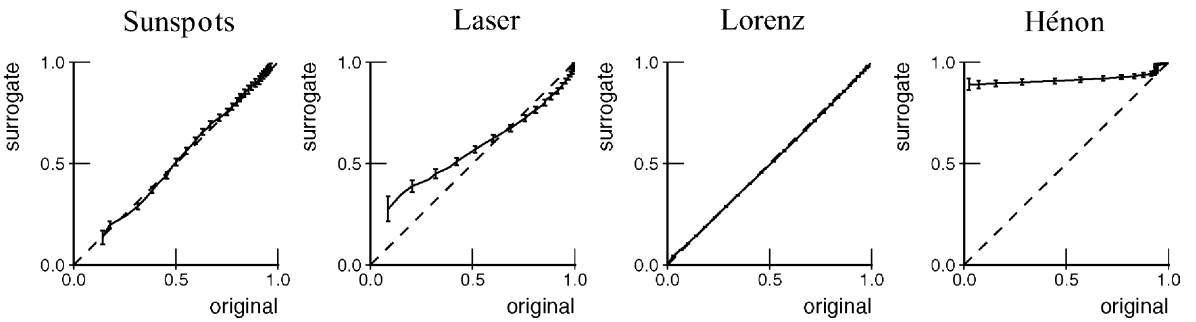

Fig. 6. DVV scatter diagrams of four standard signals.

5. Discussion

We have introduced a novel nonlinearity analysis technique, the DVV method, which examines the local predictability of a time series, given m previous samples. It is related to the false nearest neighbour approach [6] and the δ–ϵ technique [4]. A comparison has been made between the proposed DVV method and several established nonlinearity analysis techniques, namely the 'deterministic versus stochastic' (DVS) plots, a time reversal (REV) technique, a third-order cumulant (C3) technique, and the δ–ϵ method. All the methods except the DVS plots compute a test statistic on the original time series and on 99 surrogates, which implement a composite null hypothesis of a linear, stationary and Gaussian process, the output of which is amplitude transformed by a zero-memory nonlinearity. The DVS plot yields a single plot, the position of the local minimum of which reflects the degree of nonlinearity that is present in the time series.

In order to gain further insight, we have generated a set of synthetic time series by gradually increasing the degree of nonlinearity and the noise level. This way, a space formed by the degree of nonlinearity and the degree of

determinism in a time series is tiled, and the performance of the different analysis methods can be systematically examined. Each of the 12 time series in the set is driven by a stationary, Gaussian process. Thus, rejections of the null hypothesis can be attributed to the presence of nonlinearities. Only the proposed DVV and the DVS method consistently detect nonlinearities when present in the time series in our simulations.[5] Furthermore, analysis results have been compared for the benchmark time series described in [1], which consist of four nonlinear and one linear time series, showing that the DVV method is able to reject the linearity hypothesis for these four types of nonlinearity as well, whereas the other methods fail to detect the nonlinearities is one or more of these four time series, albeit that the δ–ϵ method only marginally fails to reject one of the nonlinear cases. A final comparison has been made of the results for four time series, which have been frequently addressed in the context of nonlinearity analysis, namely the Sunspots, Laser, Lorenz and Hénon map time series. The DVS method indicates the presence of nonlinearity in all four of these time series and δ–ϵ, REV and DVV in only three.

Overall, we have found that the proposed DVV method yields more consistent results compared to the other methods, and that the results are straightforward to interpret. We have restricted our analyses to a convenient choice of the time lag, τ, in all simulations, and, in the case of the DVS and δ–ϵ method, the embedding dimension, m, has been arbitrarily chosen. For the standard time series, we have also determined m using Cao's method [2], yielding $m = 12, 9, 7, 6$ for the Sunspots, Laser, Lorenz and Hénon series. The results for the DVV method remained the same, but the δ–ϵ method failed on the Lorenz series. A proper choice of these parameters can greatly influence analysis results, but is by no means evident. Therefore, the relative insensitivity of the DVV method to the parameter choice is a desirable property for a robust analysis method.

Acknowledgements

Temujin Gautama is supported by a scholarship from the Flemish Regional Ministry of Education (GOA 2000/11). Danilo P. Mandic is supported by QLG3-CT-2000-30161, while at the Laboratorium voor Neuro and Psychofysiologie, K.U. Leuven, Belgium. Marc M. Van Hulle is supported by research grants received from the Fund for Scientific Research (G.0185.96N), the National Lottery (Belgium) (9.0185.96), the Flemish Regional Ministry of Education (Belgium) (GOA 95/99-06; 2000/11), the Flemish Ministry for Science and Technology (VIS/98/012), and the European Commission, fifth framework programme (QLG3-CT-2000-30161 and IST-2001-32114).

References

[1] W.A. Barnett, A.R. Gallant, M.J. Hinich, J.A. Jungeilges, D.T. Kaplan, M.J. Jensen, A single-blind controlled competition among tests for nonlinearity and chaos, J. Econ. 77 (1997) 297–302.
[2] L. Cao, Practical method for determining the minimum embedding dimension of a scalar time series, Physica D 110 (1997) 43–50.
[3] M. Casdagli, Chaos and deterministic versus stochastic non-linear modelling, J. R. Statist. Soc. B 54 (2) (1991) 303–328.
[4] D.T. Kaplan, Exceptional events as evidence for determinism, Physica D 73 (1994) 38–48.
[5] D.T. Kaplan, Nonlinearity and nonstationarity: the use of surrogate data in interpreting fluctuations in heart rate, in: Proceedings of the Third Annual Workshop on Computer Applications of Blood Pressure and Heart Rate Signals, Florence, Italy, 1995.
[6] M.B. Kennel, R. Brown, H.D.I. Abarbanel, Determining embedding dimension for phase-space reconstruction using a geometrical construction, Phys. Rev. A 45 (1992) 3403–3411.
[7] D. Kugiumtzis, Test your surrogate data before you test for nonlinearity, Phys. Rev. E 60 (1999) 2808–2816.
[8] K.S. Narendra, K. Parthasarathy, Identification and control of dynamical systems using neural networks, IEEE Trans. Neural Networks 1 (1) (1990) 4–27.
[9] T. Schreiber, A. Schmitz, On the discrimination power of measures for nonlinearity in a time series, Phys. Rev. E 55 (1997) 5443–5447.
[10] T. Schreiber, Interdisciplinary application of nonlinear time series methods, Phys. Rep. 308 (1999) 1–64.
[11] T. Schreiber, A. Schmitz, Surrogate time series, Physica D 142 (2000) 346–382.

[5] Note that this set of simulations is limited to one specific type of nonlinearity.