Linear and non-linear methods for automatic seizure detection in scalp electro-encephalogram recordings

P. E. McSharry[1,2]   T. He[1]   L. A. Smith[2,3]   L. Tarassenko[1]

[1]Department of Engineering Science, University of Oxford, Oxford, UK
[2]Mathematical Institute, University of Oxford, Oxford, UK
[3]Centre for the Analysis of Time Series, London School of Economics, London, UK

Abstract—*The electro-encephalogram is a time-varying signal that measures electrical activity in the brain. A conceptually intuitive non-linear technique, multi-dimensional probability evolution (MDPE), is introduced. It is based on the time evolution of the probability density function within a multi-dimensional state space. A synthetic recording is employed to illustrate why MDPE is capable of detecting changes in the underlying dynamics that are invisible to linear statistics. If a non-linear statistic cannot outperform a simple linear statistic such as variance, then there is no reason to advocate its use. Both variance and MDPE were able to detect the seizure in each of the ten scalp EEG recordings investigated. Although MDPE produced fewer false positives, there is no firm evidence to suggest that MDPE, or any other non-linear statistic considered, outperforms variance-based methods at identifying seizures.*

1 Introduction

EPILEPSY IS the most common serious neurological disorder, with 4–5% of the population suffering at some time in their life. If it were possible to develop reliable and robust indicators of a seizure ahead of its onset, this would have considerable impact on the quality of life of a very large number of sufferers.

The electro-encephalogram (EEG) is a time-varying signal generated by electrical activity within the brain and is recorded either from intracranial electrodes inside the brain or scalp electrodes on the surface of the head. This paper focuses on the detection of epileptic seizures from scalp EEG recordings. Analysis of the EEG involves investigating recordings of long duration, for which the underlying dynamics associated with different brain states, such as non-seizure, pre-seizure, seizure and post-seizure, are typically obscured by noise and artifacts. The amount of time taken to examine such recordings visually suggests that an automatic seizure detection system would alleviate the work of EEG technicians who continue to score multi-channel records manually.

Previous research has demonstrated that the dynamic processes underlying EEG signals are likely to be non-linear (THEILER, 1995; SCHREIBER, 2000) but found no evidence of low-dimensional chaos. There has been much debate concerning whether or not the processes underlying epileptic seizures are low-dimensional or display deterministic chaos (MARTINERIE *et al.*, 1998; LEHNERTZ and ELGER, 1998). One proposed mechanism for an epileptic seizure is that neurons in a particular region of the brain become synchronised (LARTER and SPEELMAN, 1999), leading to a reduction in the complexity of the observed electrical activity. Some evidence of this synchronisation has been found by investigating EEG signals from neighbouring channels using intracranial (ARNHOLD *et al.*, 1999) and scalp (ALBANO *et al.*, 2000) electrodes.

Although understanding the dynamics underlying the EEG is an important goal, the operational questions here are whether or not the fundamental non-linearities are sufficiently robust that they may be detected and whether or not they are related to the epileptic seizures. If both of these are the case, then non-linear statistics will be able to outperform traditional linear statistics. As linear statistics have the advantage of being easy both to understand and to implement, this paper aims to investigate exactly how much is to be gained from using non-linear statistics.

2 Methods

2.1 *Synthetic data*

As the dynamics underlying the EEG can never be known exactly, it is useful to test different analysis techniques on a system for which the dynamical equations are known. The following system is defined such that a single parameter controls the amount of non-linearity in the dynamics at any given instant in time. Consider a random process with linear temporal
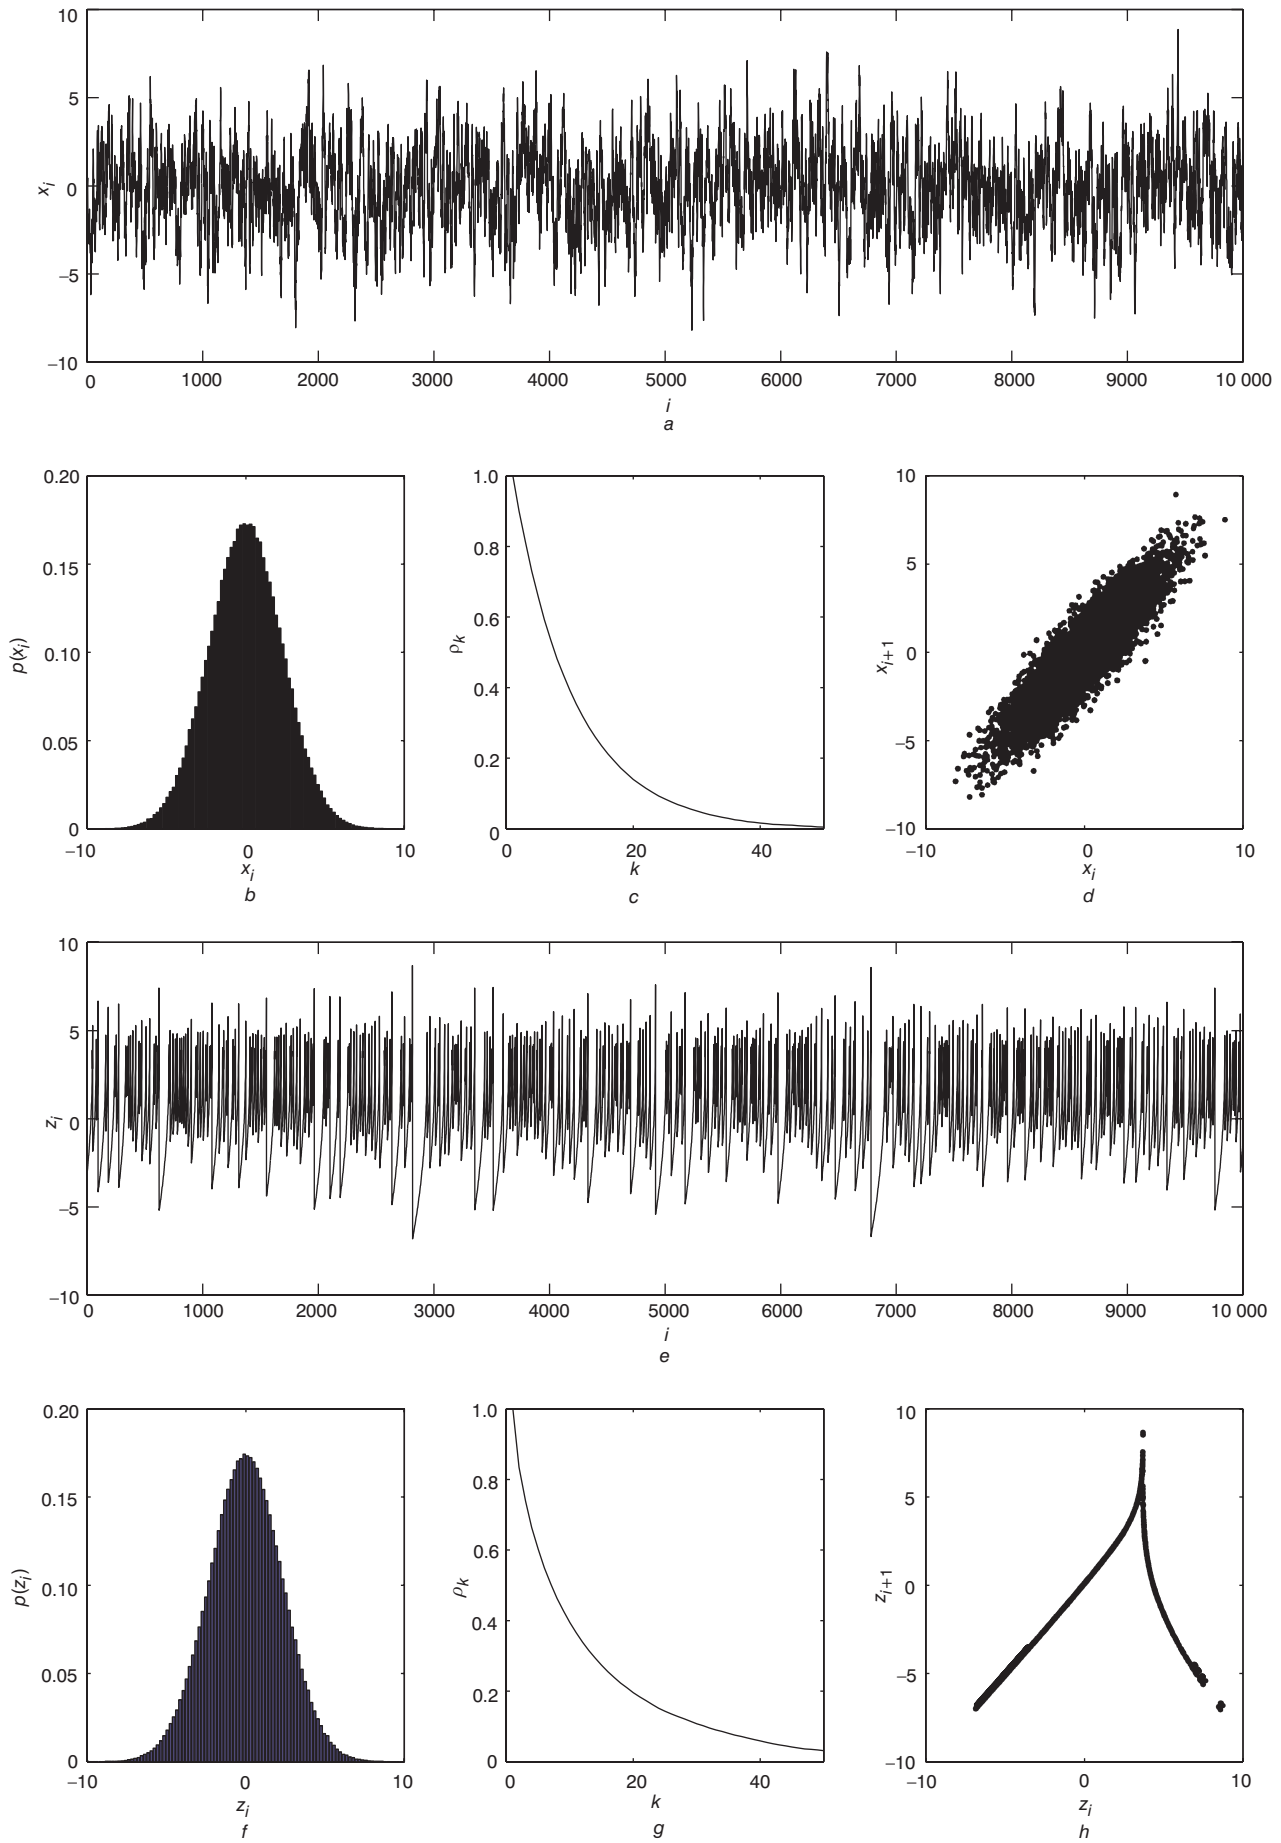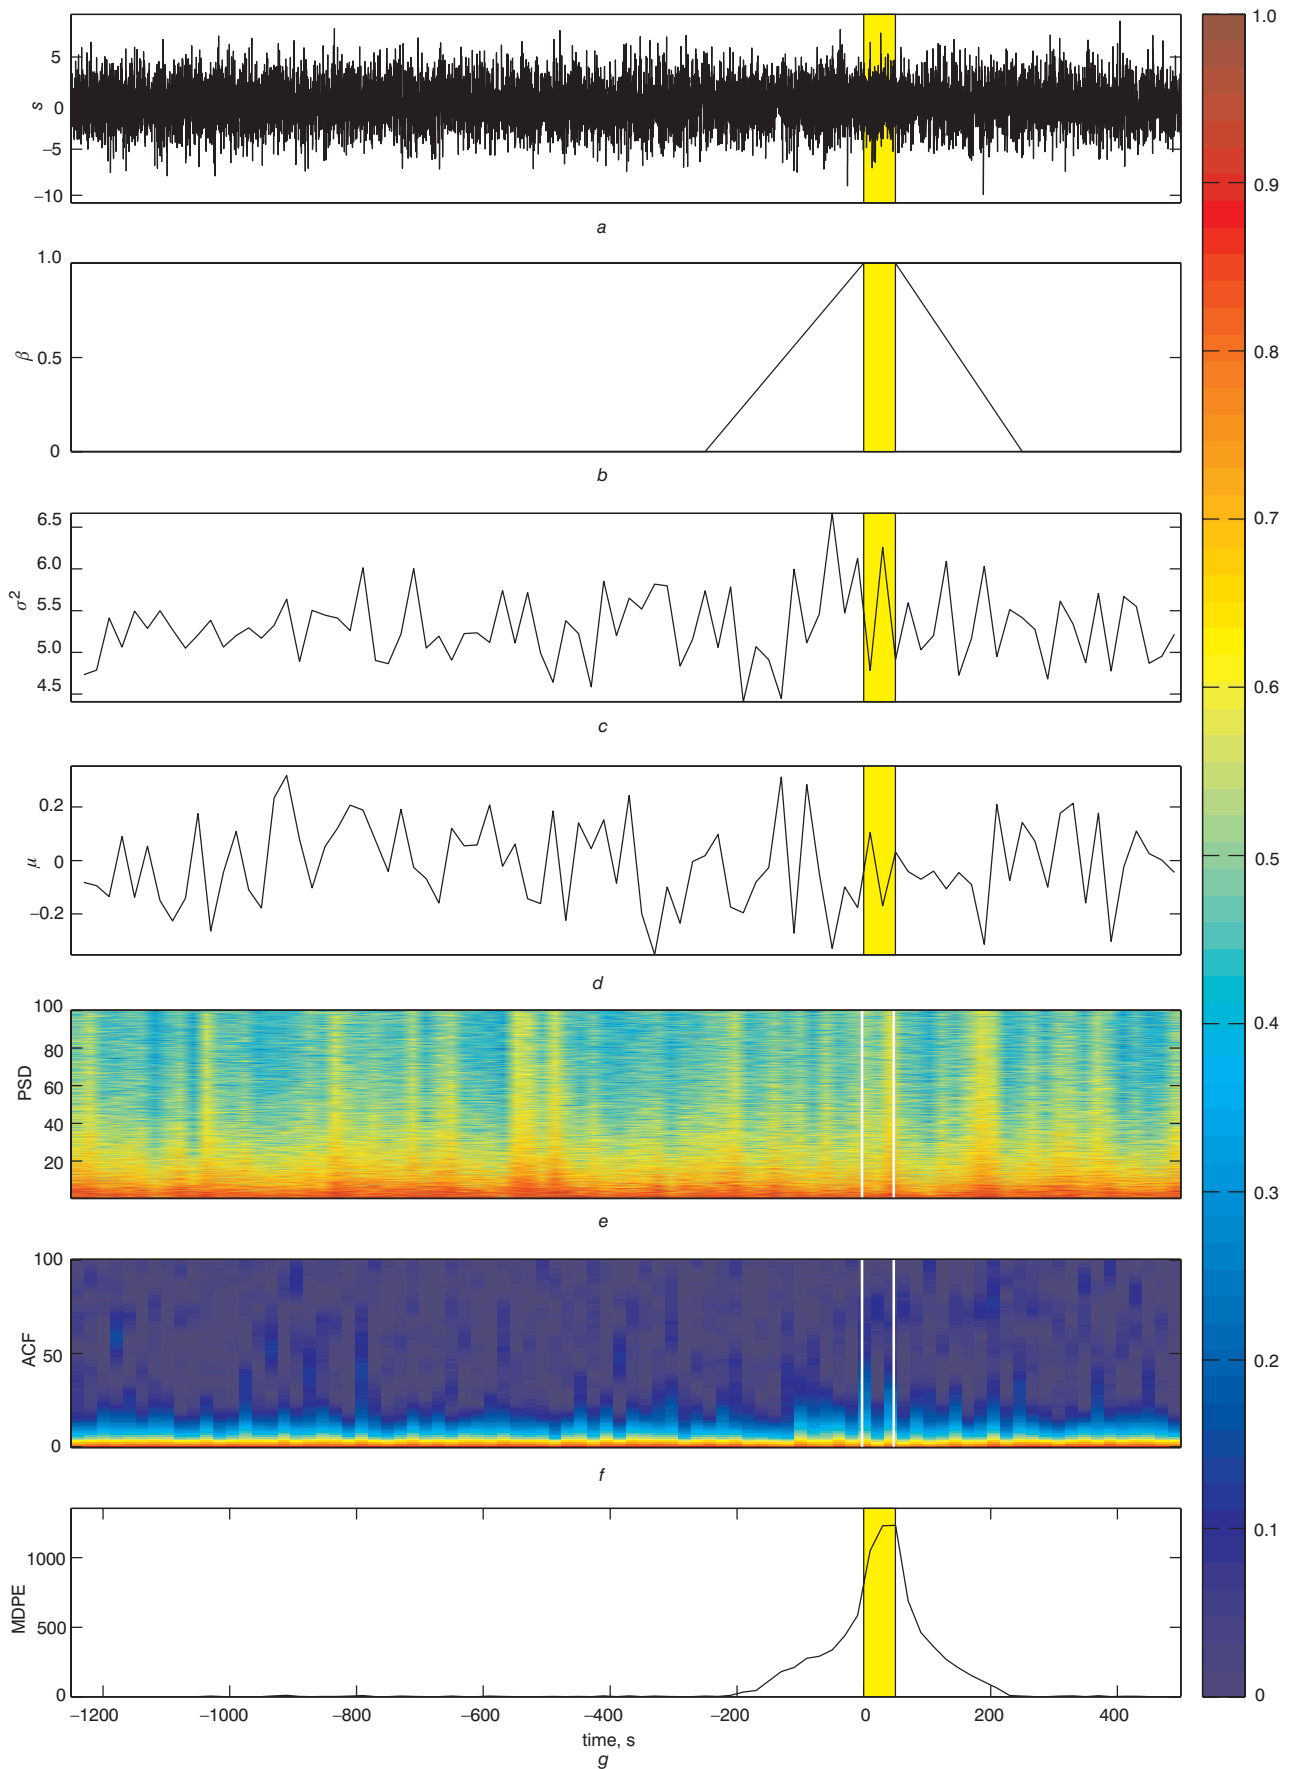
Fig. 1 *Comparison of AR(1) process x_i given by (1) with non-linear deterministic process z_i given by (2) for a = 0.95: (a) time series of x_i; (b) PDF of x_i; (c) auto-correlation function of x_i; (d) return map of x_i; (e) time series of z_i; (f) PDF of z_i; (g) auto-correlation function of z_i; and (h) return map of z_i*

Fig. 2 *(a) Example of synthetic signal obtained by combining AR(1) random process (1) and non-linear deterministic process (2) using (b) control parameter β. Non-overlapping 20 s window analysis of (c) mean μ, (d) variance σ, (e) power spectral density (arbitrary units), and (f) auto-correlation function demonstrates that linear statistics fail to detect non-linear deterministic process when β > 0. In contrast, (g) MDPE technique with learning period of 200 s clearly identifies dynamical changes and also reflects strength of non-linear activity*

correlations, such as the auto-regressive process of order one AR(1) (CHATFIELD, 1989)

$$x_{i+1} = \alpha x_i + \epsilon_i \tag{1}$$

where $-1 < \alpha < 1$, and ϵ_i is a normally distributed random variable with mean zero and standard deviation one. The standard deviation of this process is $\sigma_x = 1/\sqrt{1-\alpha^2}$, and its auto-correlation ρ_k is given by $\rho_k = \alpha^{|k|}$, where k is the time delay.

The non-linear deterministic system known as the skew-tent map, given by

$$y_{i+1} = \begin{cases} \dfrac{1}{a} y_i & 0 \leqslant y_i \leqslant a \\ \dfrac{1-y_i}{1-a} & a \leqslant y_i \leqslant 1 \end{cases} \tag{2}$$

is a non-invertible transformation of the unit interval into itself, with the parameter a chosen to satisfy $0 < a < 1$, and its invariant measure is uniform on the unit interval. This dynamical system is chaotic, with Lyapunov exponent $\Lambda = -a \log_2(a) - (1-a) \log_2(1-a)$.

If the parameter values of these two systems are chosen such that $\alpha = 2a - 1$, then both systems will have identical power spectra (SMITH, 1997). To enhance the similarity between these two systems, a measurement function $z_i = h(y_i)$ can be used to transform the output of the skew-tent map y_i, so that the probability density function (PDF) of z_i is also normally distributed with mean zero and standard deviation σ_x, as in the case of the AR(1) process. The required measurement function is

$$z_i = \frac{\sqrt{2}}{\sigma_x} \Phi\left[2\left(y_i - \frac{1}{2}\right)\right] \tag{3}$$

where Φ is the inverse error function (PRESS *et al.*, 1992). This means, of course, that the auto-correlation functions of x_i and z_i will differ somewhat, although those of x_i and y_i are identical.

Fig. 1 contrasts the statistical properties of x_i and z_i for $a = 0.95$. Although the two time series look different to the eye, their PDFs are identical, and their auto-correlation functions are similar (some disparity is introduced as the measurement function (3) is non-linear). The difference in the underlying dynamical equations is best seen by investigating their return maps (also known as delay reconstructions), as shown in Figs 1d and h. These two-dimensional delay reconstructions are able to resolve the dynamics underlying both systems given by (1) and (2). The difference between the multi-dimensional distributions displayed in the state space (see Figs 1d and f) is utilised to develop a technique for detecting dynamical changes in Section 2.5. Note that linear statistics, such as variance, the auto-correlation function and indeed any quantity defined with respect to the power spectrum, fail to distinguish between the two systems illustrated in Fig. 1.

A synthetic signal s_i is defined as the linear combination of the random process x_i and the non-linear deterministic process z_i

$$s_i = \sqrt{\beta_i} z_i + \sqrt{1-\beta_i} x_i \tag{4}$$

where β_i controls the amount of non-linear deterministic structure in s_i. This combination ensures that, if both x_i and z_i have standard deviation σ_x, then that of s_i is also σ_x. Furthermore, as both x_i and z_i are mean zero, s_i is also mean zero. Note that $\beta_i = 1$ gives the purely non-linear deterministic signal z_i, whereas $\beta_i = 0$ yields the random linear process x_i, which is devoid of any non-linear structure.

By analogy to the claims (LEHNERTZ and ELGER, 1998; MARTINERIE *et al.*, 1998) that epileptic seizures are associated with low-dimensional non-linear processes, $\beta_i = 1$ is chosen to correspond to the seizure period of the synthetic signal. A non-seizure state is defined by $\beta_i = 0$, and the pre-seizure and post-seizure states are both given by $\beta_i \in [0, 1]$. Figs 2a and b illustrate the synthetic signal s_i obtained for a given control parameter β_i.

2.2 EEG recordings

The multi-electrode scalp EEG signals analysed in this paper were recorded at the National Hospital for Neurology and Neurosurgery, London (MCGROGAN, 2001). A sample rate of 200 Hz was used to record 20 channels with positions shown in Fig. 3. The data were recorded with 12-bit resolution. A modified 10–20 system was used that omits the P_z position but includes two additional electrodes. Seizure onset times were marked by experienced technicians, and the seizures lasted for at least 40 s, apart from recording F/1, for which the seizure lasted for 24 s. All seizures are classified as generalised seizures, implying that activity related to the seizure should be present in every channel.
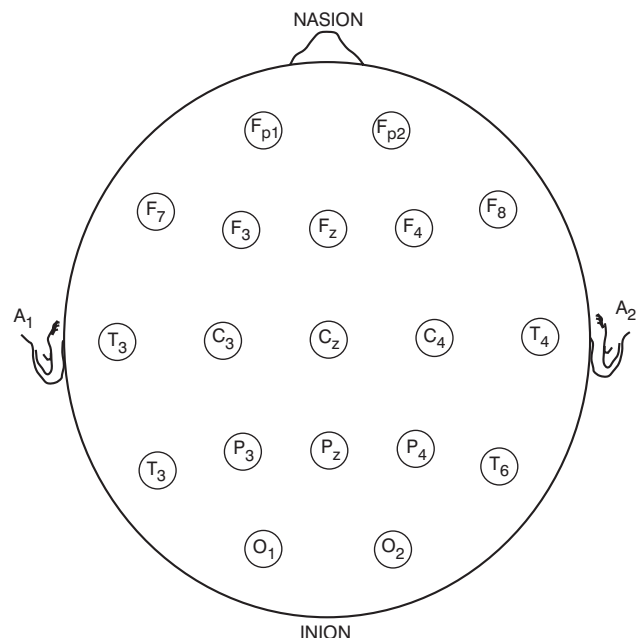

Fig. 3 *Spatial geometry of 10–20 electrode system*

Table 1 *Channel numbering and naming convention of 10–20 system: prefrontal (F_p), frontal (F), central (C), temporal (T), parietal (P), occipital (O) and ear or mastoid (A)*

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
F_{p2}	F_8	T_4	T_6	O_2	F_{p1}	F_7	T_3	T_5	O_1	F_4	F_z	F_3	C_4	C_z	C_3	P_4	P_3	A_1	A_2

The database contains eight seizures recorded from six different patients. Table 1 shows the correspondence between the electrodes and the channel numbering convention of the 10–20 system.

2.3 Linear analysis

2.3.1 *Variance*: One of the simplest linear statistics that can be used for investigating the dynamics underlying the EEG is the variance of the signal calculated in consecutive non-overlapping windows. Let s_i denote the EEG signal at time i. The variance of this EEG signal is given by

$$\sigma^2 = \langle [s_i - \mu]^2 \rangle \tag{5}$$

where the mean is $\mu = \langle s_i \rangle$, and $\langle \cdot \rangle$ is the average taken over the time interval being considered.

ESTELLER *et al.* (1999) suggest measuring the energy (simply, the sum of s_i^2 without removal of the mean) of the signal in consecutive windows of the EEG signal. For the recordings investigated in this paper, the means do not vary significantly, and therefore the results obtained using variance are equivalent to those using energy.

The F-test (PRESS *et al.*, 1992) provides a statistical test of the hypothesis that two given data sets have different variances. The F statistic is the ratio of one variance to the other, so that $F \gg 1$ and $F \ll 1$ both indicate significant differences. The probability that F would be as large as it is if the first data set's underlying distribution actually has smaller variance than the second is given by $p = Q(F|v_1, v_2)$ where v_1 and v_2 are the number of degrees of freedom in the first and second data sets, respectively. This probability can be computed using $Q(F|v_1, v_2) = I_{v_2/(v_2+v_1 F)}(v_2/2, v_1/2)$, where $I_x(a, b)$ is the incomplete beta function. Note that the negative logarithm of the probability $\gamma = -\log_{10} p$ was calculated using log co-ordinates to compute $I_x(a, b)$, to avoid problems with computer precision.

2.3.2 *Power spectrum*: A popular approach for investigating the EEG signal is to utilise its power spectrum (BLANCO *et al.*, 1995). There are a number of different statistics that aim to summarise the information contained in the power spectrum. These include calculation of the total integral of the power spectrum over all non-zero frequencies (note that this equals the variance of the signal) and the median frequency, which estimates the 'typical' frequency present in the signal (WIDMAN *et al.*, 2000). It has also been postulated that rhythmic behaviour, characterised by a peak in the power spectrum at a specific frequency, can be used to identify epileptic seizures (MURRO, 1991).

2.3.3 *Auto-correlation function*: The auto-correlation function ρ_k of a process s_i is given by (CHATFIELD, 1989)

$$\rho_k = \frac{\langle [s_i - \mu][s_{i+k} - \mu] \rangle}{\sigma^2} \tag{6}$$

where k is the time lag. ρ_k quantifies the amount of linear correlation between the signal and itself shifted by a time lag k. This function satisfies $\rho_0 = 1$; values of $\rho_k \sim 1$ reflect strong linear correlations; $\rho_k \sim -1$ implies strong linear anti-correlations; and $\rho_k \sim 0$ indicates that no linear correlations exist.

2.4 Non-linear analysis

The non-linear analysis of data recorded from an experimental system usually begins with a state space reconstruction. An advantage of obtaining a multi-dimensional state space is that it may reveal the underlying dynamics, as demonstrated in Figs 1*d* and *h*. One method for obtaining such a reconstruction is to employ delays of the EEG signal recorded at a single electrode.

A delay vector reconstruction (PACKARD *et al.*, 1980; TAKENS, 1981; SAUER *et al.*, 1991) of the signal s_i is defined by

$$\boldsymbol{x}_i = [s_{i-(m-1)\tau}, \ldots, s_{i-\tau}, s_i] \tag{7}$$

where m is the reconstruction dimension, and τ is the time delay. In the results presented here, $m = 2$, and τ was chosen using the geometrical approach introduced in ROSENSTEIN *et al.* (1994). Note that, in the case of the processes given by (1) and (2), these equations imply that $m = 2$ and $\tau = 1$ yield a suitable reconstruction of the state space.

Over the last decade, a plethora of papers have been published claiming that low-dimensional dynamics underly various complex systems. The correlation dimension D_2 (GRASSBERGER and PROCACCIA, 1983) provides a measure of the complexity of the fractal geometry sculpted by the observations in the reconstructed state space. A number of papers have warned of the difficulties in estimating D_2 and have demonstrated that vast amounts of data are required for even the simplest dynamical systems (SMITH, 1988; ECKMANN and RUELLE, 1992). This difficulty is exacerbated by the fact that real-world systems are invariably contaminated by measurement noise and that the processes underlying the EEG may be non-stationary.

The requirement for long data sets implies that the dimension estimators will not converge to a meaningful estimate of D_2 for short windows of EEG data. Hence, D_2 is a rather blunt instrument for detecting changes in the dynamics over time. LEHNERTZ and ELGER (1998) have calculated an effective correlation dimension in a non-overlapping window of 20 s and claim that a significant decrease in this dimension occurs prior to and during seizure for most subjects. Note that D_2 may still be a useful statistic, just not a good dimension estimate.

A number of other non-linear statistics have been used to investigate changes in the EEG: correlation density (LERNER, 1996; MARTINERIE *et al.*, 1998), cross-correlation integral (LE VAN QUYEN *et al.*, 1999; 2000; 2001), Lyapunov exponents (IASEMIDIS and SACKELLARES, 1996; IASEMIDIS *et al.*, 1999; SACKELLARES *et al.*, 2000), similarity measures (HIVELY, 1999) and non-linear predictability (BLINOWSKA and MALINOWSKI, 1991; HERNÁNDEZ *et al.*, 1995; CASDAGLI *et al.*, 1996; 1997; CASDAGLI, 1997). The fundamental link between all these non-linear statistics is that they are defined with respect to a particular state space. If the distribution of points in this state space varies, then these statistics are likely to change. The new technique presented in the following Section directly quantifies changes within this state space.

2.5 Multi-dimensional PDF evolution

The multi-dimensional probability evolution (MDPE) technique provides a method for detecting changes in the trajectories woven by the EEG signal as it evolves throughout the reconstructed state space. The idea is simply to measure how often different regions of state space are visited while the EEG displays non-seizure activity. By monitoring how often the

trajectory visits different parts of the state space, the MDPE technique is capable of detecting changes in the underlying dynamics.

A reference set $\mathcal{A}$, representing non-seizure activity, is constructed from the state vectors recorded during the learning period. By choosing N_c centres $\boldsymbol{\xi}_i$ ($i = 1, \ldots, N_c$) in the state space, it is possible to define a partition of the state space such that each point $\boldsymbol{x}$ in the reference set $\mathcal{A}$ is a member of partition $\mathcal{B}_i$ if

$$\|\boldsymbol{x} - \xi_i\| < \min_{j \neq i} \|\boldsymbol{x} - \xi_j\| \qquad (8)$$
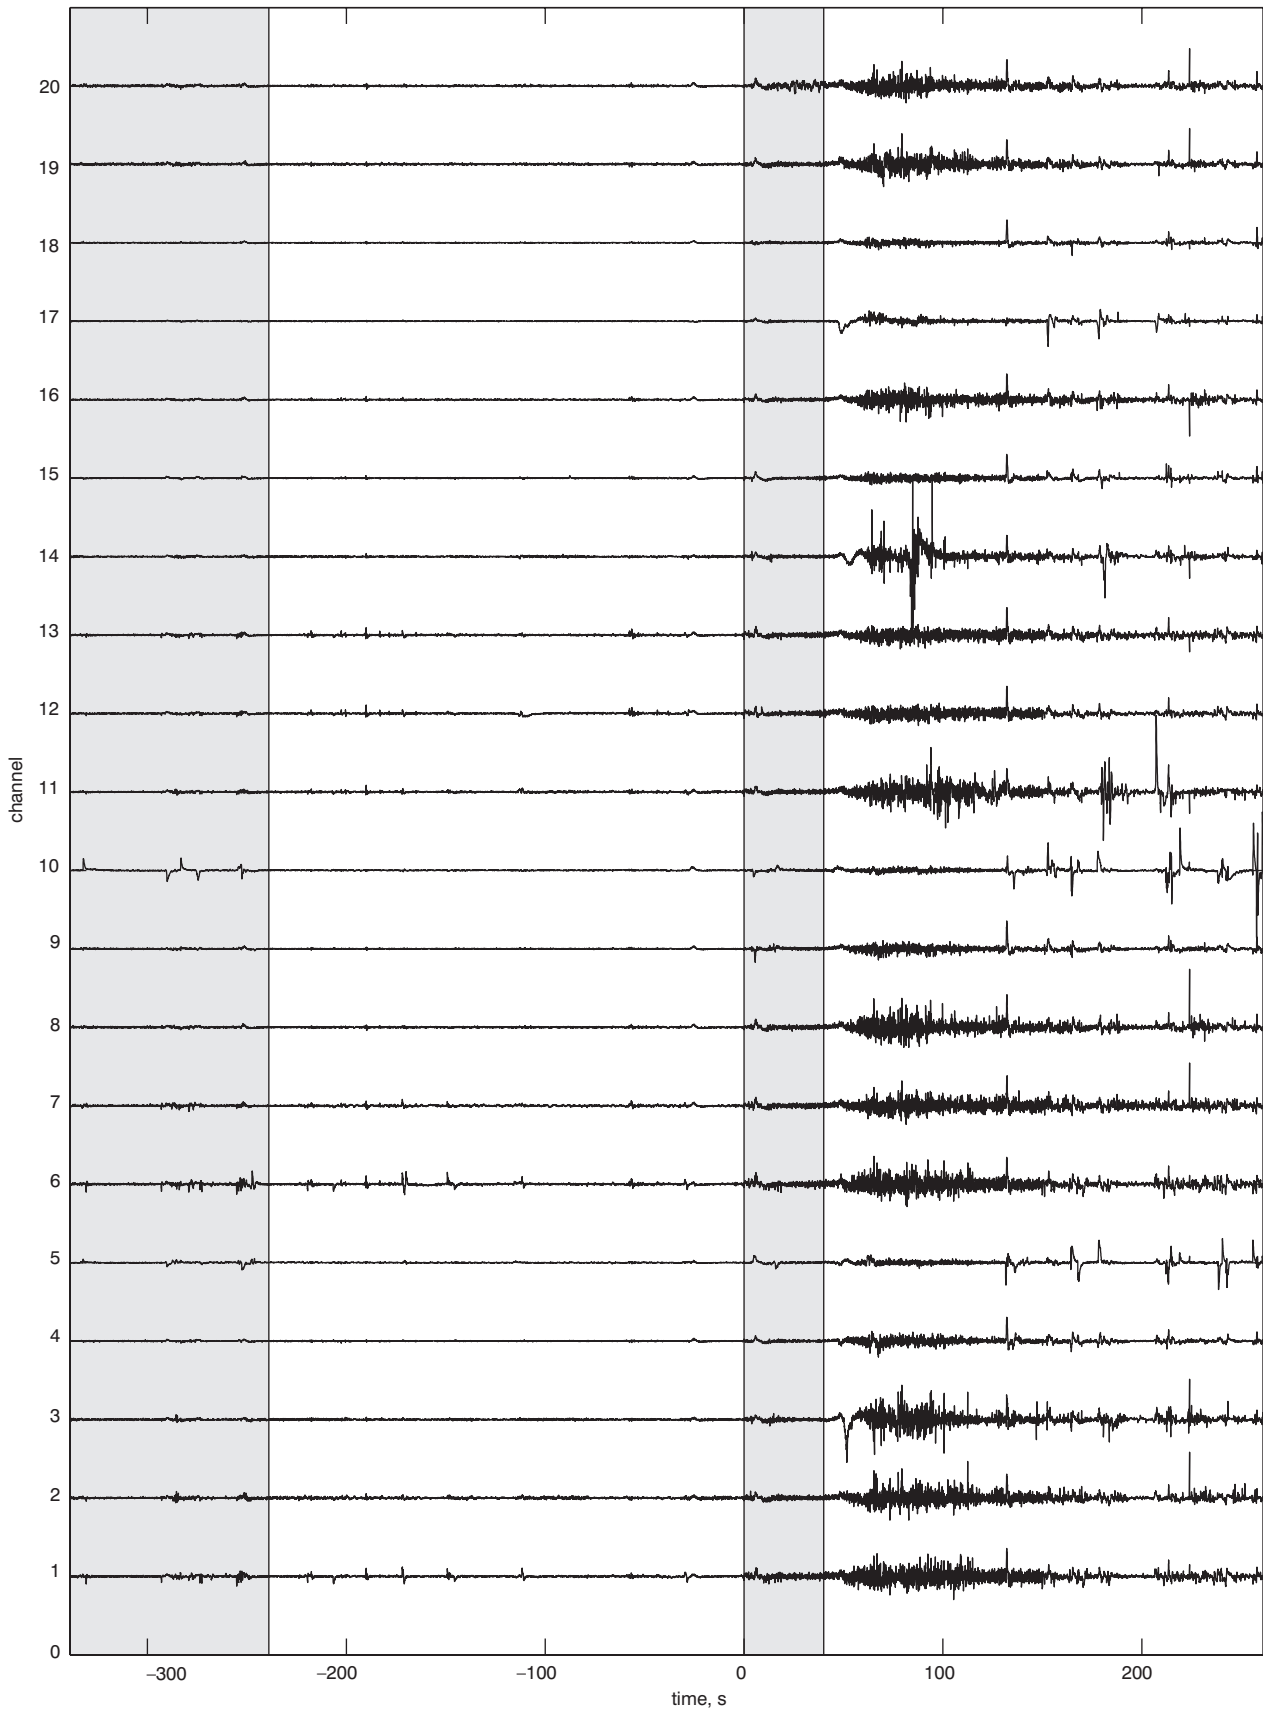
Fig. 4 *Multi-channel scalp recordings for A/1. Leftmost shaded region shows non-seizure state used as reference dynamics, whereas seizure is indicated by rightmost shaded region. All voltage amplitudes are scaled by an equal amount*
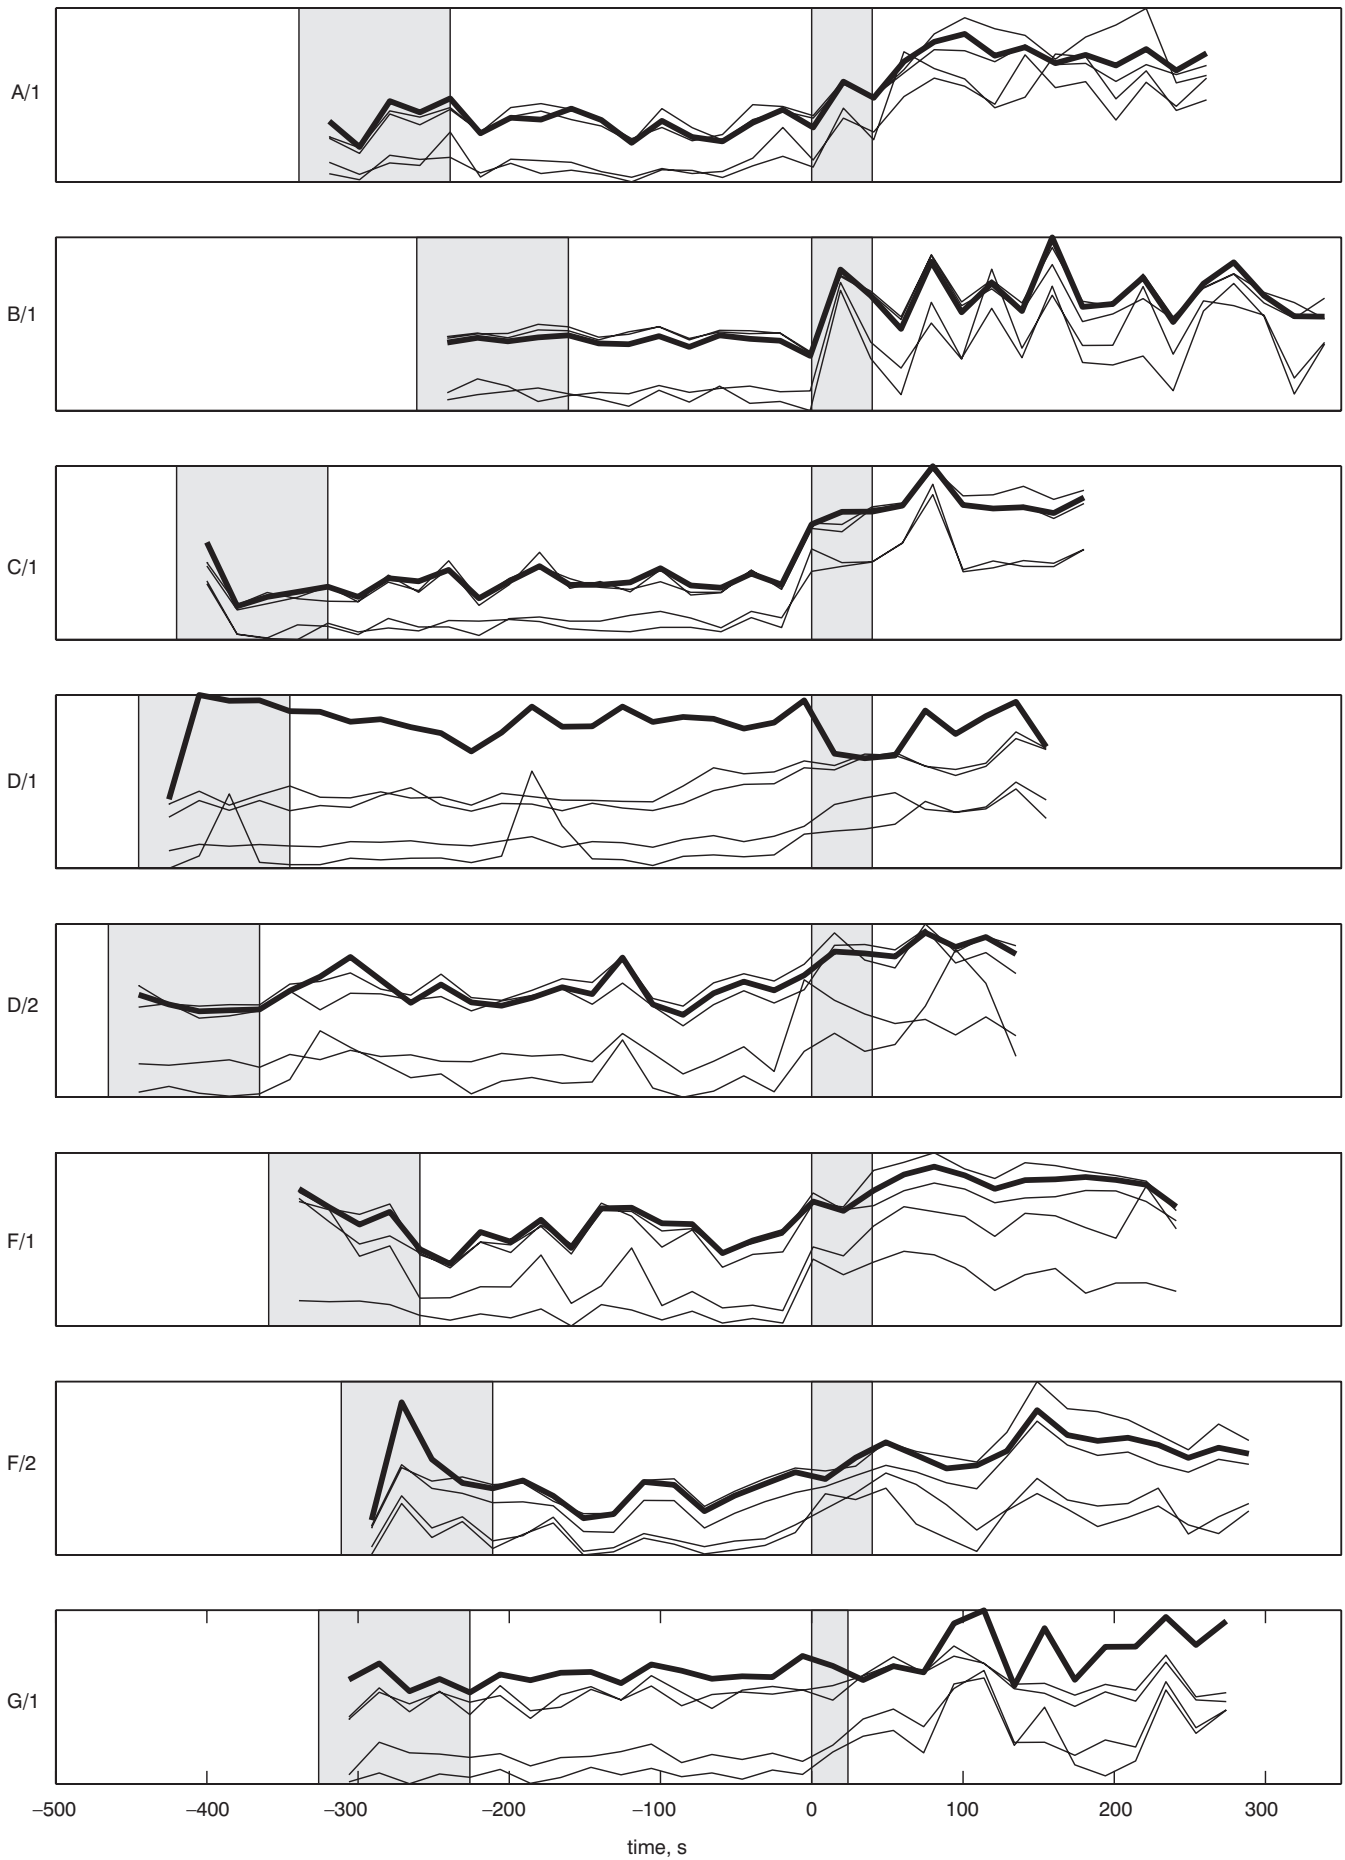
Fig. 5 *(– – –) Running variance for each of four (F₃, F₄, P₃ and P₄) scalp electrodes and (——) mean variance averaged over all 20 electrodes. Each label specifies patient (A–G) and recording number for that patient. Learning and seizure periods are illustrated by leftmost and rightmost shaded regions, respectively. Vertical axes have logarithmic scale ranging between minimum and maximum variances*

The centres were chosen uniformly from the distribution obtained by the trajectories in the learning period (SMITH, 1992). A value of $N_c = 100$ was employed in the results presented in Section 3. Counting the number of points n_i^0 in each of the N_c partitions yields a discrete distribution for the reference set $\mathcal{A}$. Similarly, for any given window of the recording, it is possible to calculate its distribution n_i in the state space.

A χ^2-test (PRESS *et al.*, 1992) can be used to compare the distribution n_i with that of the reference set n_i^0. Suppose that the total number of points in the reference set and the window are N^0 and N, respectively. The χ^2-statistic is

$$\chi^2 = \sum_{i=1}^{N} \frac{[rn_i^0 - (1/r)n_i]^2}{n_i^0 + n_i} \qquad (9)$$

where $r = \sqrt{N/N^0}$. Note that χ^2 is zero if both distributions have equal numbers in each partition. In the case of both time series having different numbers of points, the number of degrees of freedom is $v = N_c$. The χ^2 probability distribution $p = Q(\chi^2|v)$, an incomplete gamma function, gives the probability of observing values greater than χ^2 under the null hypothesis that the data sets are drawn from the same distribution (PRESS *et al.*, 1992). In other words, a small value of p indicates that there is a significant difference between the two distributions. The negative logarithm of the probability, $\gamma = -\log_{10} p$, was calculated to avoid problems with computer precision.

3 Results

The variation in the statistics is now investigated for both the synthetic signal and the real scalp EEG recordings.

3.1 *Synthetic data*

The invisibility of the dynamical changes induced by blending in the non-linear deterministic process in the synthetic signal are demonstrated for the linear statistics in this Section. Figs 2*a* and *b* show the synthetic signal s_i and the associated control parameter β_i. A moving window analysis of the mean, variance, power spectrum and auto-correlation function are shown in Figs 2*c–f*; none of these linear statistics is capable of detecting the changes due to the non-linear correlations. In contrast, Fig. 2*g* demonstrates the probability that the multi-dimensional PDF has not changed from that computed within the learning period (first 100 s). Low probabilities (high values of γ) reflect excursions into new or rarely visited regions of state space. Such excursions reflect abnormal behaviour with respect to the learning data set, and the shape of the MDPE parameter γ between -200 and 200 s reveals (Fig. 2*g*) very clearly the dynamical changes introduced by the non-linear deterministic process when $\beta_i \neq 0$ (Fig. 2*b*).

3.2 *EEG recordings*

The scalp EEG recordings contain noise and artifacts, unlike the clean synthetic signal, and this suggests that, even if non-linearities were active in the underlying dynamics associated with the seizure, it may be more difficult to distinguish them from these artifacts. In addition, the duration

of the EEG recording that is available before the onset of the seizure ranges between 250 and 450 s. This limits the use of out-of-sample tests. Without being further away from the onset of the seizure, it is not obvious whether any part of the recording available before the seizure represents non-seizure activity.

Doubling the duration of the learning set did not significantly improve the results. The duration of the learning set was fixed at 100 s because, as the size of the learning set is increased, more pre-seizure data are introduced, thus possibly including activity relating to the seizure onset. Also note that the size of the learning set should be limited by the need to assess the technique on out-of-sample data prior to the seizure.

Fig. 4 shows the multi-electrode activity for patient A/1. The voltage amplitudes are scaled by an equal amount, so that the amount of activity recorded at each electrode can be visualised. There is a large increase in the variance at each electrode during and after the seizure. Almost any statistic, including 'non-linear statistics', is likely to detect this change in variance. The important question is whether or not there is anything other than an increase in variance taking place in this recording around the time of the seizure.

Fig. 5 illustrates the spatio-temporal changes in variance for each of the eight scalp recordings from six patients. The first number of the label is the patient index, and the second is the recording index. In each panel, the lines denote the variance, computed over non-overlapping 20 s windows for four of the electrodes, F_3, F_4, P_3 and P_4, which are evenly distributed about the centre of the skull (broken line), and the variance averaged over all 20 channels (solid line). The rightmost shaded region on each panel denotes the time period that clinicians have marked as containing the epileptic seizure for that particular recording. The leftmost shaded region on each panel, containing the first 100 s of each recording, is used as a learning set. These results demonstrate that the variance usually increases just after the onset of the seizure. Variance provides a benchmark with which other statistics can be compared. If a non-linear statistic cannot outperform a simple linear statistic such as variance, then there is no reason to advocate complex non-linear statistics.

This increase in variance is associated with a number of different changes in the power spectrum (and thus the auto-correlation function), as shown in more detail for patient A/1 in Fig. 6. In particular, note the rhythmicity (periodic activity) associated with an increase in power concentrated at a frequency that varies between 2 and 4 Hz during the 150 s that follow the seizure. This rhythmicity was displayed in only five out of the 20 channels. Some traces of rhythmicity were also found in patients B/1, E/1 and E/2. Without having access to a large population of patients, it is difficult to say whether or not rhythmicity is a hallmark of epileptic seizures. Furthermore, rhythmicity may be associated with some classes of seizure and not with others.

Fig. 7 shows the results of applying the *F*-test to the learning data and 20 s windows of data throughout the recording. This analysis was carried out separately for each electrode for each of the eight EEG recordings. The grey-scale colour bar indicates the significance as defined in Section 2.3. The results for four particular electrodes (F_3, F_4, P_3 and P_4) are shown separately in Fig. 8. Variance seems to be capable of detecting the seizures to within 100 s of the markings.

A leave-one-out approach was used to obtain five out-of-sample values for γ in the learning set. These values provide information about the amount of variability in γ due to non-seizure activity. Given that there are only five values of γ representing the learning set, under ideal conditions and the
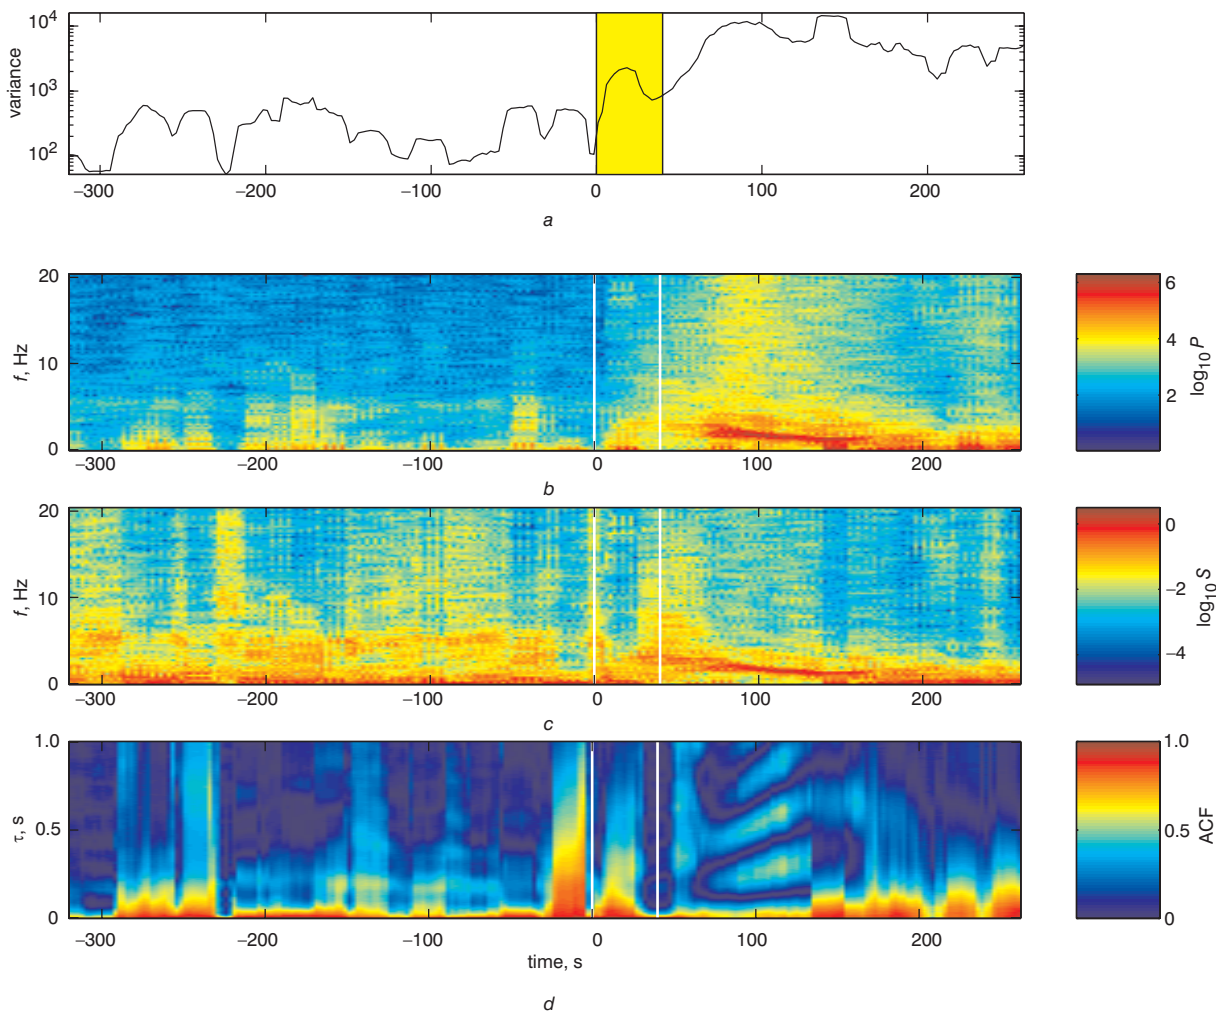
Fig. 6 *Analysis of EEG signal from channel 13 or F_3 of recording A/1: (a) variance; (b) power spectral density P; (c) normalised power spectral density S (total power is same in each window); and (d) auto-correlation function. Two vertical lines indicate duration of seizure. Windows of 20 s with overlap of 17.5 s were used*

assumption that these values are independent, the probability of observing a value greater than the largest of these is approximately 20%. A broken line is used to indicate the threshold given by the maximum value of γ for the learning set in Fig. 8.

The MDPE technique was also applied to non-overlapping 20 s windows of the EEG recordings, as shown in Figs 9 and 10. A broken line indicates the threshold corresponding to the maximum value of γ for the learning set in Fig. 10. Delay reconstructions given by (7), with $m = 2$ and τ calculated using ROSENSTEIN *et al.* (1994) (see Table 2) were employed. Note that, although there is little variation between the delays chosen for different electrodes, the delays tend to vary quite a lot from one recording to the next. The overall results are very similar to those yielded by the variance analysis in Figs 7 and 8. Note, however, that MDPE provides a more accurate identification of the seizure onset in recordings E/1, E/2 and F/1. In addition, MDPE produces fewer large values of γ (possibly leading to false positives) in the times between the learning period and the onset of the seizure. A comparison of Figs 8 and 10 demonstrates that variance yields possible false warnings in F_4 of A/1, P_3 of D/1, F_1, F_4, P_4 of E/1, and F_1, F_4, P_3 of E/2. The MDPE non-linear statistic may therefore be more robust than variance, but without investigating more EEG recordings with longer non-seizure periods, it is too early to say whether this is generally the case.

4 Discussion

The occurrence of non-seizure dynamics that were not observed during the learning period is likely to cause false positives when the subsequent EEG signals are analysed. This is true for any technique that aims to construct a model of normal activity from a fixed learning data set. In previous publications, where databases of epileptic seizure recordings have been investigated, learning periods of 250 s (LE VAN QUYEN *et al.*, 1999), 300 s (LE VAN QUYEN *et al.*, 2000; 2001) and 400 s (MARTINERIE *et al.*, 1998) have been used. It was assumed that these learning periods captured adequate representations of non-seizure activity: 'this reference window was chosen during a state as 'normal' as possible, i.e. far from the seizure, free of artefacts and containing common features of interictal activity, e.g. isolated spikes' (LE VAN QUYEN *et al.*, 2000).

It is useful to expand on this assumption, as it implies that the learning set contains

 (i)   no seizure activity
(ii)   a complete representation of numerous types of EEG activity that are non-seizure.

In particular, a suitable learning set should include the dynamics associated with non-seizure activity, such as muscle activity and movement artifacts, (when these are detected independently and excluded from the analysis).
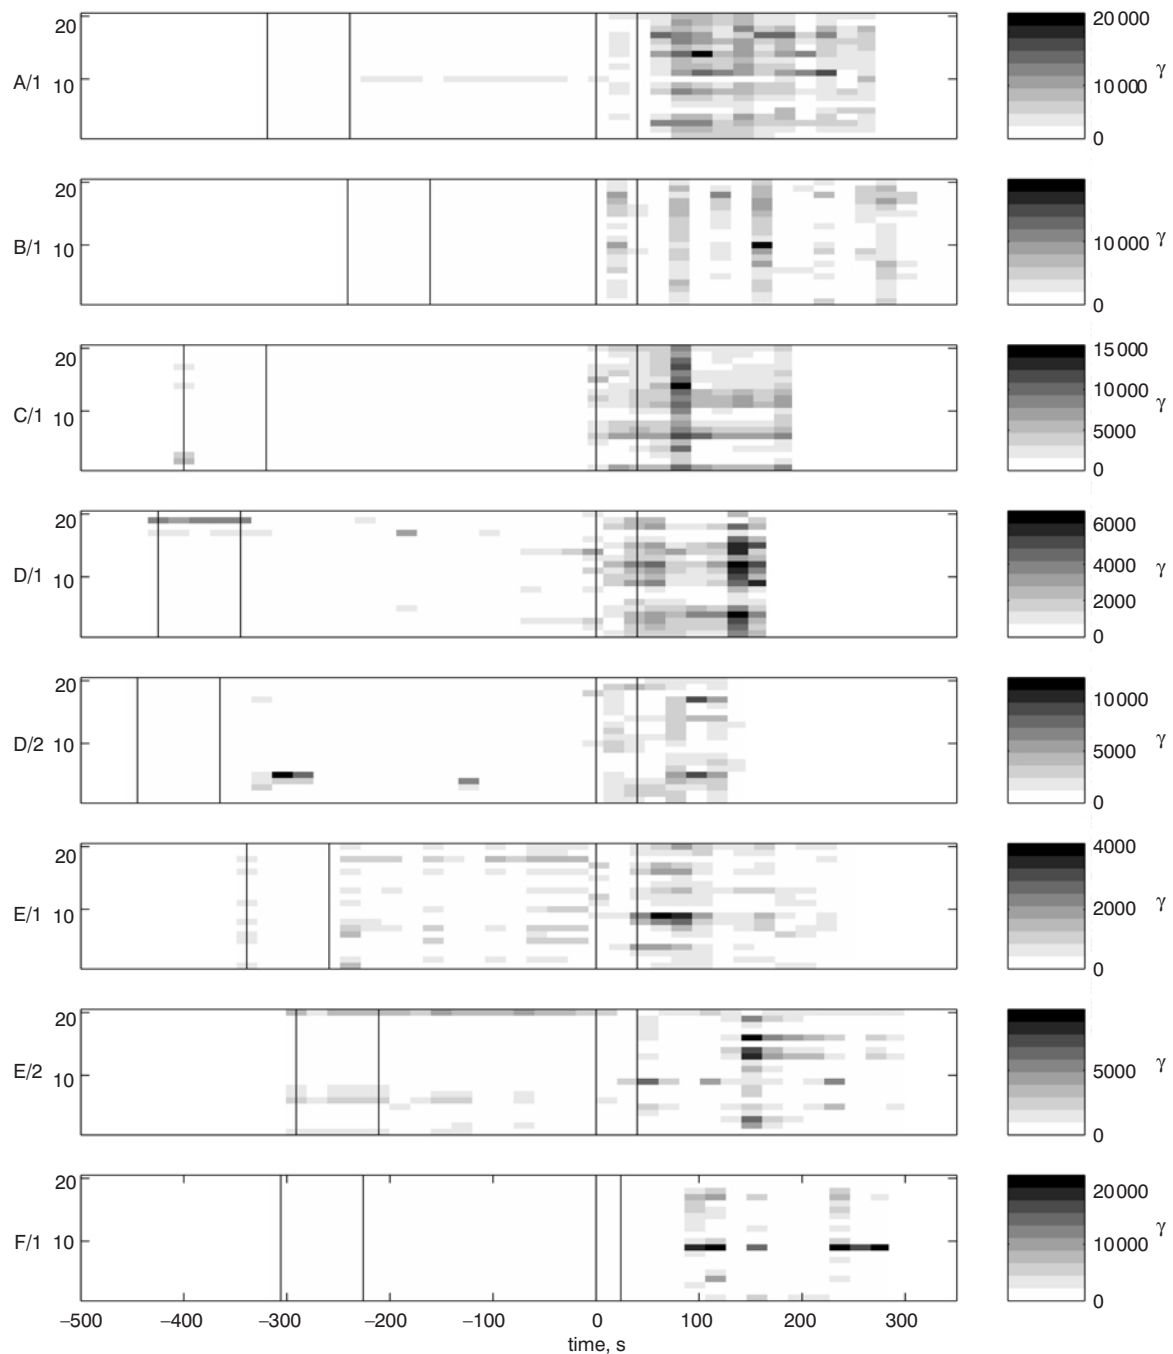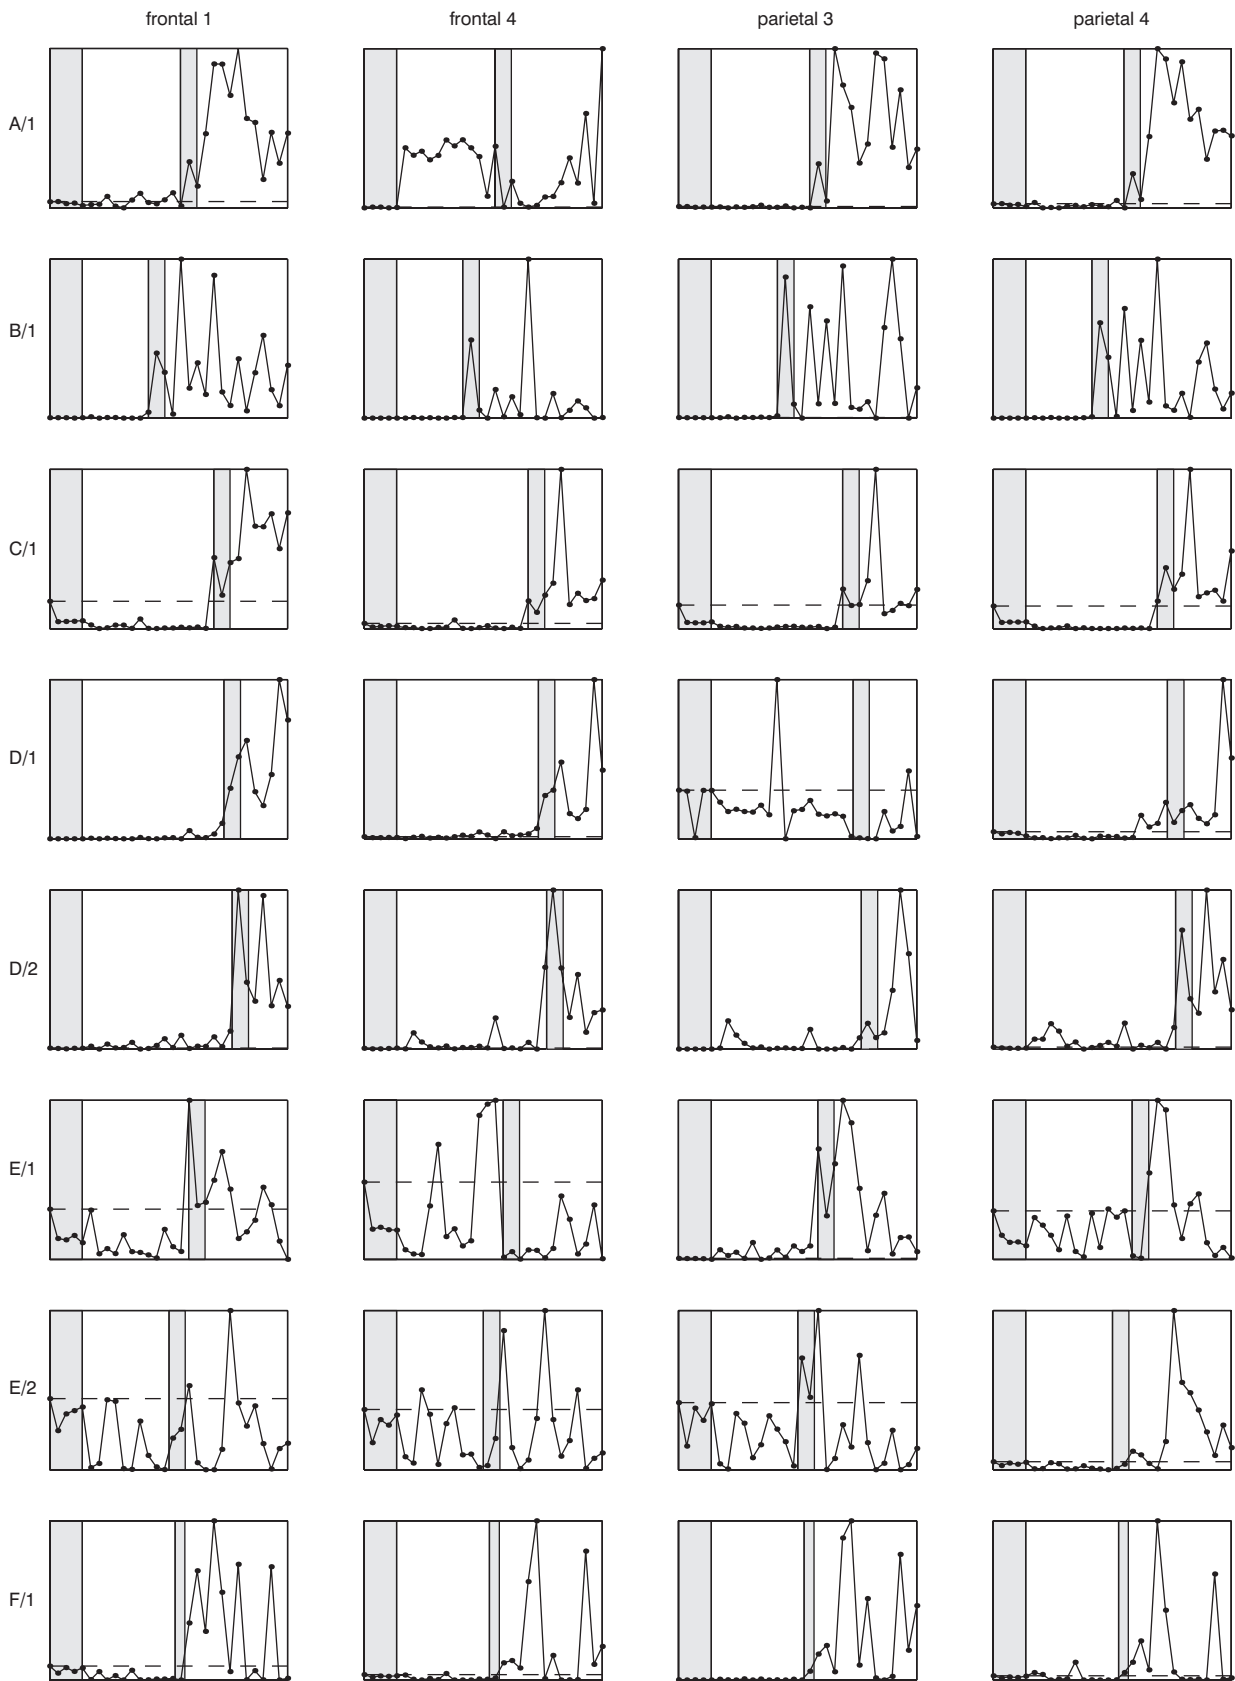
Fig. 7 *Variance analysis of non-overlapping 20 s windows for EEG recordings where vertical axis corresponds to channel number. Grey-scale bar shows significance of F-test given by γ. First pair of vertical lines indicates learning period, and the second pair indicates duration of seizure*

Many different statistics, both linear and non-linear have been suggested in the literature; the evaluation of these is non-trivial given the limited data available. If enough statistics are constructed on a given database and evaluated on the same database, then eventually one statistic will demonstrate apparent success, and yet it is unlikely that it will generalise to new recordings. A statistic that requires the setting of a few parameters could be tuned to provide remarkable seizure predictions on a given database, but such in-sample over-fitting yields a misleading view of the likely success of the statistic in an operational setting.

Ideally, a test statistic should be able to distinguish between detecting a 'rare' event and detecting an 'abnormal' event. An evaluation of the robustness of a statistic against non-seizure activity is a fundamental requirement for any practical method that might be used for detecting or predicting epileptic seizures. A practical technique for seizure detection must be applied to a wide variety of non-seizure data to quantify the number of false positives. Without this test, published results fail to provide sufficient information for clinicians.

Rather than blind application of non-linear statistics, it makes sense to ask whether or not there is any conceptual basis for these techniques. If a complicated non-linear statistic yields superior results and yet there is no explanation for its success, it is likely that a simpler statistic exists that is just as good.

Fig. 8 *Variance analysis of non-overlapping 20 s windows for four scalp electrodes (F_3, F_4, P_3 and P_4). Leftmost shaded region indicates learning period, whereas rightmost reflects duration of seizure. Time is displayed on horizontal axis, and γ from F-test is shown on vertical axis. (– – –) Maximum γ obtained in learning set. Note that vertical axis ranges from 0 to maximum value of γ for that particular panel*

For the database investigated here, there was no evidence of any non-linearities in the dynamics underlying the EEG associated with the epileptic seizures. The variance displayed clear changes around the time of the seizure, and these may provide an accurate seizure detection system, although both muscle activity and movement artifact also increase variance. It
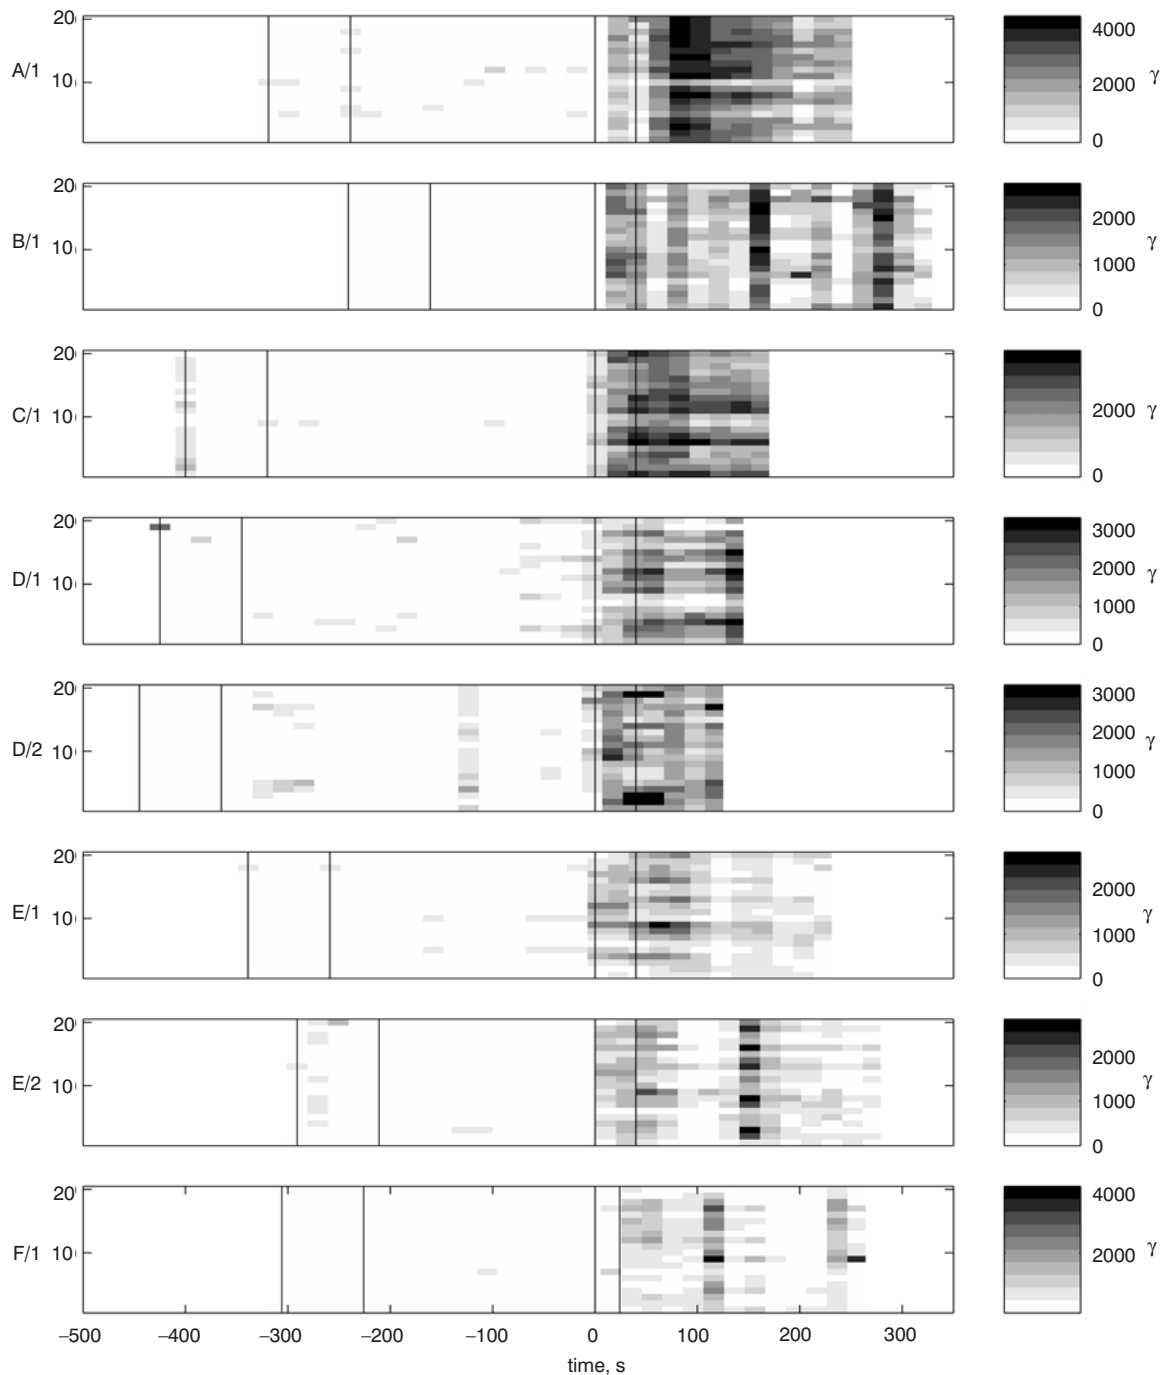
Fig. 9 *MDPE analysis of non-overlapping 20 s windows for EEG recordings where vertical axis corresponds to channel number. Grey-scale bar shows significance of χ^2-test given by γ. First pair of vertical lines indicates learning period, and second pair indicates duration of seizure*

remains to be seen whether the occurrence of suspected false warnings is due to pre-seizure activity or simply due to the fact that the learning period is not long enough to obtain an accurate model of the dynamics associated with normal activity.

5 Conclusions

A synthetic signal was constructed by combining a linear random process and a non-linear deterministic process. As expected, no linear statistic could be found that was able to detect transitions between these two processes, despite the difference between their underlying dynamics. In contrast, the MDPE statistic introduced in this paper is capable of detecting subtle changes in the underlying state space that are associated with changes in the dynamical equations used to generate the synthetic signal.

An *F*-test was used to compute the significance of the observed difference between the variances of the recording during the learning period and that during the testing window. Similarly, a χ^2-test was used to obtain the significance of the observed differences between the multi-dimensional distributions observed in the state space during these periods.
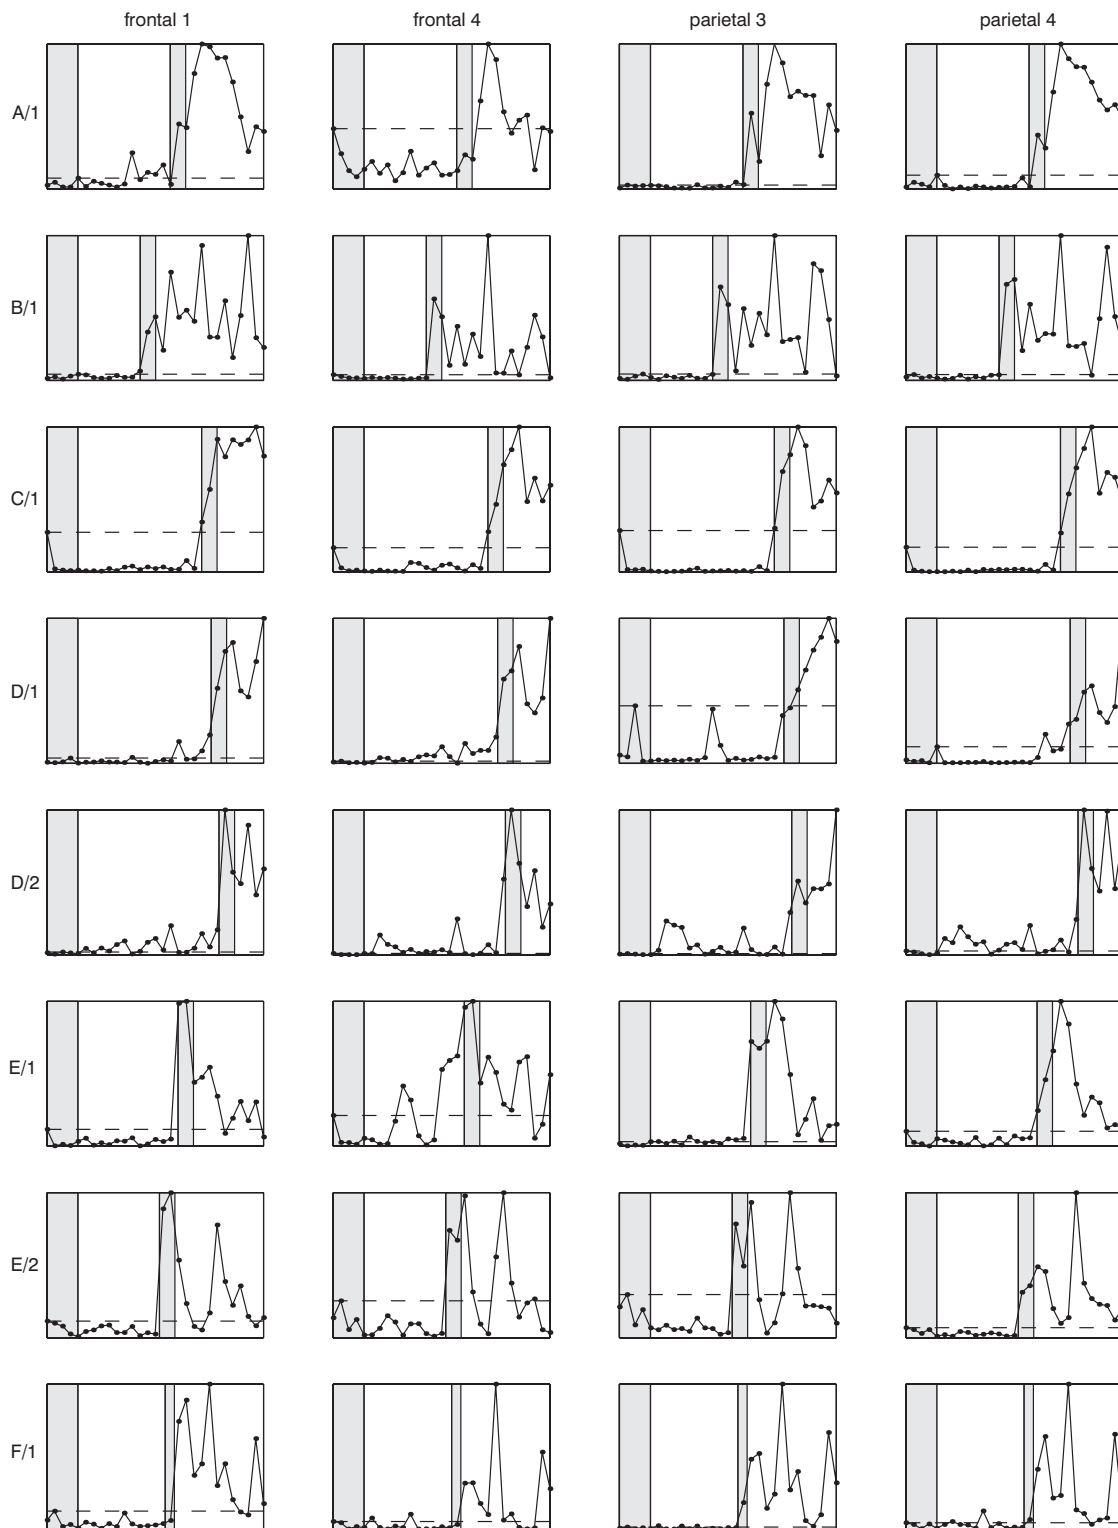
Fig. 10 *MDPE analysis of non-overlapping 20 s windows for four scalp electrodes (F_3, F_4, P_3 and P_4). Leftmost shaded region indicates learning period, whereas rightmost reflects duration of seizure. Time is displayed on horizontal axis, and γ from χ^2-test is shown on vertical axis. (— — —) Maximum γ obtained in learning set. Note that vertical axis, ranges from 0 to maximum value of γ for that particular panel*

A database of scalp EEG recordings was also analysed using linear statistics and the MDPE statistic. Both MDPE and variance were able to detect all of the seizures, but the MDPE provided more accurate detection of the seizure onset in recordings E/1, E/2 and F/1. There is no clear justification for preferring the MDPE statistic over variance, except that the former may generate fewer false positives. Evaluating this claim requires tests on larger databases with longer non-seizure periods. The availability of more learning data would also allow the MDPE to build up a more accurate model of non-seizure dynamics. Non-linear statistics greatly increase the scope of automatic detection, but their use must be justified on a case-by-case basis.

Table 2 *Time delays for different channels of EEG recorded from each patient*

Channel/patient	A/1	B/1	C/1	D/1	D/2	E/1	E/2	F/1
1	3	11	5	3	7	4	4	6
2	3	11	5	3	7	4	4	5
3	3	11	3	3	4	3	3	3
4	3	11	6	5	7	4	3	3
5	3	10	6	4	6	3	3	3
6	3	11	5	6	9	4	4	7
7	3	11	5	3	7	4	4	13
8	3	11	6	3	6	3	3	3
9	3	10	6	5	7	4	3	3
10	4	10	6	4	6	3	3	3
11	3	10	5	3	7	4	4	4
12	5	10	5	7	8	4	4	5
13	3	10	5	3	7	3	3	5
14	3	10	5	3	7	4	4	4
15	5	9	5	6	7	4	5	5
16	3	10	5	3	7	4	4	4
17	3	10	5	5	7	4	4	4
18	3	10	4	5	7	4	4	5
19	3	11	6	3	7	4	4	7
20	3	11	5	5	7	3	3	8

References

ALBANO, A. M., CELLUCCI, C. J., HARNER, R. N., and RAPP, P. E. (2000): 'Optimization of embedding parameters for prediction of seizure onset with mutual information', in: MEES, A. I. (Ed.): 'Nonlinear dynamics and statistics: proceedings of an Issac Newton Institute' (Birkhauser, Boston, 2000), pp. 435–451

ARNHOLD, J., GRASSBERGER, P., LEHNERTZ, K., and ELGER, C. (1999): 'A robust method for detecting interdependences: application to intracranially recorded EEG', *Physica D*, **134**, pp. 419–430

BLANCO, S., QUIAN QUIROGA, R., ROSSO, O. A., and KOCHEN, S. (1995): 'Time-frequency analysis of electroencephalogram series', *Phys. Rev. E*, **51**, pp. 2624–2631

BLINOWSKA, K. J., and MALINOWSKI, M. (1991): 'Non-linear and linear forecasting of the EEG time series', *Biol. Cybern.*, **66**, pp. 159–165

CASDAGLI, M., IASEMIDIS, L. D., SAVIT, R. S., GILMORE, R. L., ROPER, S. N., and SACKELLARES, J. C. (1996): 'Characterizing nonlinearity in invasive EEG recordings from temporal lobe epilepsy', *Physica D*, **99**, pp. 381–399

CASDAGLI, M. C. (1997): 'Characterizing non-linearity in weather and epilepsy data', *Fields Inst. Commun.*, **11**, pp. 201–222

CASDAGLI, M. C., IASEMIDIS, L. D., SAVIT, R. S., GILMORE, R. L., ROPER, S. N., and SACKELLARES, J. C. (1997): 'Non-linearity in invasive EEG recordings from patients temporal lobe epilepsy', *Electroencephalogr. Clin. Neurophys.*, **102**, pp. 98–105

CHATFIELD, C. (1989): 'The analysis of time series, 4th edn' (Chapman and Hall, London, New York, 1989)

ECKMANN, J. P., and RUELLE, D. (1992): 'Fundamental limitations for estimating dimensions and Liapunov exponents in dynamical systems', *Physica D*, **56**, pp. 185–187

ESTELLER, R., VACHTSEVANOS, G., ECHAUZ, J., D'ALESSANDRO, M., HENRY, T., PENNELL, P., EPSTEIN, C., BAKAY, R., BOWEN, C., SHOR, R., and LITT, B. (1999): 'Accumulated energy is a state-dependent predictor of seizures in mesial temporal lobe epilepsy', *Epilepsia*, **40**, pp. 173–173

GRASSBERGER, P., and PROCACCIA, I. (1983): 'Characterization of strange attractors', *Phys. Rev. Lett.*, **50**, pp. 346–349

HERNÁNDEZ, J. L., VALDÉS, J. L., BISCAY, R., JIMÉNEZ, J. C., and VALDÉS, P. (1995): 'EEG predictability: adequacy of non-linear forecasting methods', *Int. J. Biomed. Comput.*, **38**, pp. 197–206

HIVELY, L. M. (1999): 'Detecting dynamical change in nonlinear time series', *Phys. Lett. A.*, **258**, pp. 103–114

IASEMIDIS, L. D., and SACKELLARES, J. C. (1996): 'Chaos theory and epilepsy', *The Neuroscientist*, **2**, pp. 118–126

IASEMIDIS, L. D., PRINCIPE, J. C., and SACKELLARES, J. C. (1999): 'Measurement and quantification of spatio-temporal dynamics of human epileptic seizures' in AKAY, M. (Ed.): 'Nonlinear signal processing' (IEEE Press, 1999)

LARTER, R., and SPEELMAN, B. (1999): 'A coupled ordinary differential equation lattice model for the simulation of epileptic seizures', *Chaos*, **9**, pp. 795–804

LE VAN QUYEN, M., MARTINERIE, J., BAULAC, M., and VARELA, F. (1999): 'Anticipating epileptic seizures in real time by a non-linear analysis of similarity between EEG recordings', *NeuroReport*, **10**, pp. 2149–2155

LE VAN QUYEN, M., MARTINERIE, J., BAULAC, M., and VARELA, F. (2000): 'Spatio-temporal characterizations of non-linear changes in intracranial activities prior to human temporal lobe seizures', *Eur. J. Neurosci.*, **12**, pp. 2124–2134

LE VAN QUYEN, M., MARTINERIE, J., NAVARRO, V., BOON, P., D'HAVÉ, M., ADAM, C., RENAULT, B., VARELA, F., and BAULAC, M. (2001): 'Anticipation of epileptic seizures from standard EEG recordings', *The Lancet*, **357**, pp. 183–188

LEHNERTZ, K., and ELGER, C. E. (1998): 'Can epileptic seizures be predicted? Evidence from nonlinear time series analysis of brain electrical activity', *Phys. Rev. Lett.*, **80**, pp. 5019–5022

LERNER, D. E. (1996): 'Monitoring changing dynamics with correlation integrals: case study of an epileptic seizure', *Physica D*, **97**, pp. 563–576

MARTINERIE, J., ADAM, C., LE VAN QUYEN, M., BAULAC, M., CLEMENCEAU, S., RENAULT, B., and VARELA, F. J. (1998): 'Epileptic seizures can be anticipated by non-linear analysis', *Nature Med.*, **4**, pp. 1173–1176

McGROGAN, N. (2001): 'Neural network detection of epileptic seizures in the electroencephalogram'. PhD thesis, University of Oxford

MURRO, A. M. (1991): 'Computerized seizure detection of complex partial seizures', *Electroenceph. Clin. Neurophysiol.*, **79**, pp. 330–333

PACKARD, N., CRUTCHFIELD, J., FARMER, J. D., and SHAW, R. (1980): 'Geometry from a time series', *Phys. Rev. Lett.*, **45**, pp. 712–716

PRESS, W. H., FLANNERY, B. P., TEUKOLSKY, S. A., and VETTERLING, W. T. (1992): 'Numerical recipes in C' (CUP, Cambridge, 1992), 2nd edn

ROSENSTEIN, M. T., COLLINS, J. J., and DE LUCA, C. J. (1994): 'Reconstruction expansion as a geometry-based framework for choosing proper time delays', *Physica D*, **73**, pp. 82–98

SACKELLARES, J. C., IASEMIDIS, L. D., SHIAU, D., GILMORE, R. L., and ROPER, S. N. (2000): 'Epilepsy—when chaos fails' in LEHNERTZ, K., and ELGER, C. E. (Eds): 'Chaos in the brain?' (World Scientific, Singapore, 2000)

SAUER, T., YORKE, J. A., and CASDAGLI, M. (1991): 'Embedology', *J. Stats. Phys.*, **65**, pp. 579–616

SCHREIBER, T. (2000): 'Is nonlinearity evident in time series of brain electrical activity?' in LEHNERTZ, K., ELGER, C. E., ARNHOLD, J., and GRASSBERGER, P. (Eds): 'Chaos in brain?' (World Scientific, Singapore, 2000)

SMITH, L. A. (1988): 'Intrinsic limits on dimension calculations', *Phys. Lett. A*, **133**, pp. 283–288

SMITH, L. A. (1992): 'Identification and prediction of low-dimensional dynamics', *Physica D*, **58**, pp. 50–76

SMITH, L. A. (1997): 'The maintenance of uncertainty' in CINI, G. (Ed.): 'Nonlinearity in geophysics and astrophysics', Vol. CXXXIII, International School of Physics 'Enrico Fermi' (Società Italiana di Fisica, Bologna, Italy, 1997), pp. 177–246

TAKENS, F. (1981): 'Detecting strange attractors in fluid turbulence' in RAND, D., and YOUNG, L. S. (Eds): 'Dynamical systems and turbulence, Vol. 898' (Springer-Verlag, New York, 1981), p. 366

THEILER, J. (1995): 'On the evidence for low-dimensional chaos in an epileptic electroencephalogram', *Phys. Lett. A*, **196**, pp. 335–341

WIDMAN, G., SCHREIBER, T., REHBERG, B., HOEFT, A., and ELGER, C. E. (2000): 'Quantification of depth of anesthesia by nonlinear time series analysis of brain electrical activity', *Phys. Rev. E.*, **62**, pp. 4898–4903

Author's biography

PATRICK E. MCSHARRY was born in Leitrim, Ireland, in 1972. He received the BA in Theoretical Physics, in 1993, and the MSc in Electronic & Electrical Engineering, in 1995, from Trinity College, Dublin. He was awarded a Marie Curie Research Fellowship in 1995 and received the DPhil in Mathematics from the University of Oxford, in 1999. He is currently with the Department of Engineering Science and the Oxford Centre for Industrial & Applied Mathematics in the University of Oxford. His research interests include time series analysis, mathematical modelling, forecasting, pattern recognition and biomedical signal processing.