Available online at www.sciencedirect.com
SCIENCE DIRECT®
Physica A 349 (2005) 582–596
PHYSICA A
www.elsevier.com/locate/physa

A new method for change-point detection developed for on-line analysis of the heart beat variability during sleep

M. Staudacher[a], S. Telser[b], A. Amann[c], H. Hinterhuber[b], M. Ritsch-Marte[a,*]

[a]*Institute for Medical Physics, Müllerstr.44, Innsbruck Medical University, A-6020 Innsbruck, Austria*
[b]*Sleep Research Laboratory, Department of Psychiatry, Anichstrasse 35, Innsbruck Medical University, A-6020 Innsbruck, Austria*
[c]*Department of Anaesthesiology and Critical Care Medicine, Anichstrasse 35, Innsbruck Medical University, A-6020 Innsbruck, Austria*

Received 16 August 2004
Available online 10 November 2004

Abstract

We present a novel scaling-dependent measure for times series analysis, the progressive detrended fluctuation analysis (PDFA). Since this method progressively includes and analyzes all data points of the time series, it is suitable for on-line change-point detection: Sudden changes in the statistics of the data points, in the type of correlation or in the statistical variance, or both, are reliably indicated and localized in time. This is first shown for numerous artificially generated data sets of Gaussian random numbers. Also time series with various non-stationarities, such as non-polynomial trends and "spiking", are included as examples. Although generally applicable, our method was specifically developed as a tool for numerical sleep evaluation based on heart rate variability in the ECG-channel of polysomnographic whole night recordings. It is demonstrated that PDFA can detect specific sleep stage transitions, typically ascending transitions involving sympathetic activation as for example

*Corresponding author. Tel.: +43 512 507 3550; fax: +43 512 507 2860.
E-mail address: Monika.Ritsch-Marte@uibk.ac.at (M. Ritsch-Marte).

0378-4371/$ - see front matter
doi:10.1016/j.physa.2004.10.026

short episodes of wakefulness, and that the method is capable to discern between NREM sleep and REM sleep.

PACS: 05.45.Tp; 87.19.Hh; 89.75.Da

Keywords: Time series analysis; Detrended fluctuation analysis; Heart rate variability; Sleep ECG

1. Introduction

With the prevalence of digital data acquisition, time series analysis has become an important tool applicable in fields as diverse as e.g. engineering, financial mathematics, meteorology, and physiology. Also sleep research increasingly relies on numerical methods to analyze and characterize digital recordings from the sleep lab. These so-called polysomnographic recordings include electroencephalographic (EEG) and electrocardiographic (ECG) data, as well as other information such as eye movement, muscle tones, and breathing. Numerical methods, which were developed for time series analysis of heart rate variability (HRV) in healthy probands as well as patients [1–8], have lately also been applied to data from the ECG-channel of the polysomnography; this allowed a characterization of various sleep stages by means of the heart beat rate. It was found that sleep stages (light sleep, deep sleep, REM sleep) significantly differ in the type of correlation in the heart beats. For instance, the so-called Detrended Fluctuation Analysis (DFA), first introduced in Ref. [9] in a different context, has demonstrated that REM-sleep and periods of wakefulness display a long-time memory which is absent in light sleep and deep sleep [10,11].

In this work, we first introduce a new scaling-dependent method suitable for change-point detection. It takes the DFA, which has proven especially successful in the context of analyzing sleep, as a starting point. We call our new method "*progressive detrended fluctuation analysis*" (PDFA), because it analyzes the data points progressively. In our particular example this is a time series comprising the time intervals between the R-wave peaks of two subsequent heart beats (RR-intervals) derived from the sleep ECG. The suggested method allows continuous on-line method processing of the data extracted from an entire whole night recording. Although real-time evaluation may not always be possible, the method is able to follow the changing pattern of successive sleep stages on an intrinsic time-axis. There is no restriction to analyzing a *single* sleep stage that was specifically isolated beforehand. DFA, in contrast, at once requires the whole data-set for evaluation or specific segments cut from the data set (for example *single* sleep stage segments that were specifically isolated beforehand).

The intention of the present paper is to demonstrate the capability of the suggested method to detect change-points in the statistics by applying it first to artificially produced data and second to data from our sleep lab. In the latter context we only briefly summarize the answers to some questions that motivated the development of

the presented method, for more details we refer to Ref. [12]. In particular we demonstrate the feasibility to detect transitions between certain sleep stages (as e.g. short episodes of wakefulness and ascending transitions involving a rise in sympathetic activation). Moreover, we suggest a new way to utilize the scaling-dependence itself to differentiate between episodes of non-REM (NREM) sleep from REM sleep and wakefulness. Other on-line methods to detect heterogeneities in time series, involving wavelet or Hilbert transformations, local Hurst exponents and other scaling-dependent and scaling-independent measures [13] have recently been applied in similar contexts [14–16]. For the validation of our new method we have simultaneously applied other numerical methods, for example based on wavelet transform, to the time series.

2. Definition of the progressive detrended fluctuation analysis (PDFA)

For comparison with our new approach, let us briefly outline the DFA algorithm (for details see e.g. Refs. [17–19]), which is a scaling analysis method developed for problems associated with non-stationary time series containing e.g. polynomial trends. For analyzing a time series $\{\tau_k\}$ of total length N (with $k = 1, \ldots, N$), first the average $\langle\tau\rangle = \frac{1}{N}\sum_{k=1}^{l}\tau_k$ is subtracted. Then the integrated time series $y(l) = \sum_{k=1}^{l}(\tau_k - \langle\tau\rangle)$ is divided into non-overlapping segments of equal length n. In first-order-DFA, in each segment a least-squares line is fitted to the data, representing the local linear trend in that segment. The coordinates of the straight line segments are denoted by $y_{trend}(l)$. Finally, the standard deviation of the integrated time series from the linear trends

$$F(n) = \sqrt{\frac{1}{N}\sum_{l=1}^{N}[y(l) - y_{trend}(l,n)]^2} \tag{1}$$

depends on the segment size n and thus represents a scale-dependent measure. This definition can easily be extended to time series with higher polynomial trends: in qth order DFA, trends of polynomial order q in the integrated time series and of order $(q-1)$ in the original time series $\{\tau_k\}$ are eliminated.

Usually one derives a scale-independent coefficient, the fluctuation exponent α as the slope of $\log F(n)$ versus $\log n$ [19]. For segments of uncorrelated data (“white noise”) α is equal to 0.5. If there are only short-time correlations, the initial slope will be different from 0.5, but α will approach 0.5 for large segments. $0.5 < \alpha \leqslant 1$ indicates persistent long-range power-law correlations, meaning that a value larger than the average of the time series is more likely to be followed by another exceedingly large value than by a smaller value, and vice versa. In contrast, $0 < \alpha < 0.5$ indicates a different type of power-law correlation such that large and small values of the time series are more likely to alternate. For $\alpha \geqslant 1$ correlations exist, but cease to be of a power-law form.

DFA was developed as a method to investigate the long-range fluctuation correlations in a given time interval, where it is typically assumed that the type of

correlation is unknown, but does not change during the time interval to be investigated. The DFA method is not an appropriate tool to detect abrupt changes in the statistics occurring at a specific time, although it is possible to derive some information on the type of correlations contributing to a signal comprised of segments with different statistical properties from the crossover behavior [20]. However, since *all* points of the complete data-set (from the beginning to the end of the recording in time) are taken into account in every calculation step, there is no progressive time-axis and, consequently, it is impossible to localize the statistical change-points accurately in time.

PDFA, on the other hand, creates an intrinsic time-axis by *progressively* extending the data-set data point for data point. Analogous to DFA, first a cumulative time series $y(l)$ is generated for the whole data-set. However, in contrast to DFA, not the total time series is divided into segments of length n and analyzed in one step, but partial sums of length p ($p = 1, \ldots, N$), covering an increasing part of the total time series, are created and analyzed *separately*. *Each* partial sum is divided into non-overlapping segments of length n starting from the beginning, and the local trends are eliminated in each segment. Note that in contrast to DFA, where the window size n is varied, here the segments thus are all of *constant* length n (except for the last one containing the remaining points). Fig. 1 visualizes this difference graphically.

In order to eliminate local trends to qth order, in each segment a polynomial of order q is fitted to the partial sum and the deviations from the partial sum and the polynomial are calculated. Introducing the following function of the partial sum length p (i.e., of the length of the evaluated time interval $p = 1, \ldots, N$)

$$P_{[n]}(p) = \sqrt{\sum_{l=1}^{p} [y(l) - y_{trend}(l, n)]^2}\,, \tag{2}$$

we have defined a scale-dependent measure, which (apart from the partial sum index p) also depends on the chosen window length n. As a reminder of this implicit scaling-dependence we tag the PDFA fluctuation function with the subscript $[n]$. Let us remark that there is no prefactor $1/p$ included in the square root (as compared to the factor $1/N$ in Eq. (1)); thus $P_{[n]}(p)$ does not represent a proper fluctuation variance. We have investigated both forms, with and without the factor $1/p$, but find the definition as in Eq. (2) more convenient, since it suppresses the part of the curves

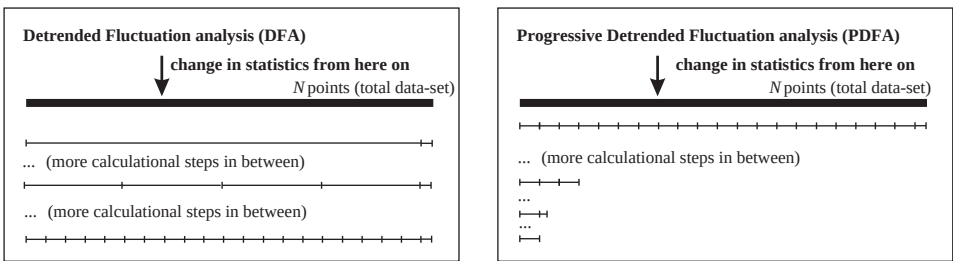

Fig. 1. Schematic diagram visualizing the difference between DFA and PDFA.

belonging to only few points included (with intrinsically poor statistics), and also since it leads to monotonically rising curves which we find easier to read, especially when overlaying curves for several window sizes n. As we shall demonstrate below, repeated analysis of progressively growing partial sums with increasing time interval length allows reliable detection and accurate localization of statistical changes occurring at a particular time. The calculational effort is comparable to DFA.

3. Change-point detection in artificial data-sets

We have thoroughly tested the above defined scaling-dependent measure $P_{[n]}(p)$ on various kinds of numerically generated artificial data-sets. Furthermore, in order to eliminate any residual doubt about systematic errors stemming from the analysis of only "quasi-random numbers", we have also tried the method on segments of true random numbers derived from a physical quantum random number generator accessible on the internet [21]; no different behavior was found. In the following we will show examples that were derived using an algorithm based on Fourier-transformation [22], which can generate sequences of random numbers with selectable *long-range* correlations, i.e., with an autocorrelation function $C(n) \sim n^{-\gamma}$ with $0<\gamma<1$ [19]. The Fourier-approach allows the straightforward generation of segmented "patchwork" time series with predefined change-points where the statistics (i.e., either the correlation exponent γ or the standard deviation σ of the random numbers, or both) change abruptly. Such artificial data-sets represent a perfect test scenario for detecting abrupt changes in correlation.

In Fig. 2 we show the results for the "PDFA change-point segmentation" of a data-set with two built-in change-points, consisting of a patchwork of 3 segments of random numbers of identical Gaussian distribution function with variance $\sigma = 1$: the first and the third contain uncorrelated random numbers (i.e., the fluctuation exponent α introduced in Section 2 equals 0.5), whereas the random numbers in the middle segment are long-range correlated ($\alpha = 0.8$). Using the relation $\gamma = 2(1 - \alpha)$ which holds for the case of long-range correlations (see e.g. Ref. [19]) the corresponding autocorrelation coefficients γ can be calculated from the fluctuation exponents α. In Fig. 2, the numerically evaluated variance and α calculated by first-order DFA are indicated for each segment; they agree well with the chosen input parameters. The PDFA norm $P_{[n]}(p)$, as defined in Eq. (2), is plotted on the left for $n = 50$. The two built-in change-points, which occur at the data points 5000 and 10 000 (marked by arrows), are clearly visible as "kinks" in the curve. To make this even more visible on the right we also show the local slope of the PDFA curves derived by numerical differentiation of the doubly logarithmic representation of the curves. By definition $P_{[n]}(p)$ is a monotonically growing square root function with a slope levelling out towards the end of the data-set. If calculational time is not an important issue, it is possible to process the data-set twice, i.e., backwards and forwards, and average the two; this was done to derive the slopes displayed on the r.h.s of Fig. 2.

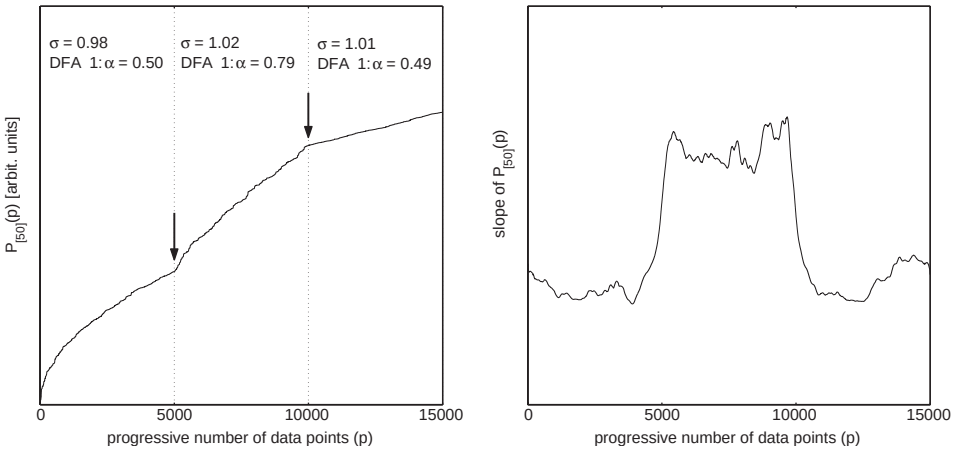

Fig. 2. Change-point detection for artificial data with two built-in change-points. The first and the third of 3 segments contain uncorrelated random numbers, the middle segment long-range correlated random numbers. Left: PDFA function with $P_{[50]}(p)$, right: local slope of the PDFA function obtained by numerical differentiation in a doubly logarithmic scale.

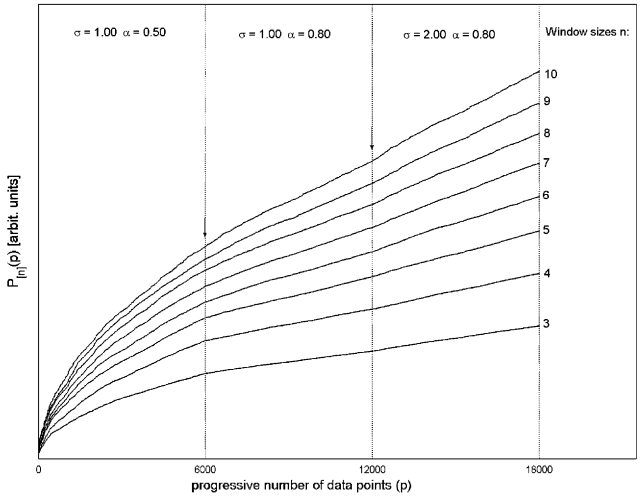

Fig. 3. Two change-points: the first in the DFA fluctuation exponent α (indicating a different autocorrelation coefficient $\gamma = 2(1-\alpha)$) and the second in the standard deviation σ for various window size parameters n from 3 to 10.

In the example discussed in Fig. 2, the 3 segments of Gaussian random numbers have different correlation coefficient γ, but the same standard deviation σ. Another example in Fig. 3 contains a first change-point, where only the correlation is altered (and the statistical variance kept constant) and a second change-point with fixed γ

and altered σ. For the case of a mere change in statistical variance (but unaltered correlations) we find as a general behavior that the slopes of the PDFA curves in double logarithmic scales for uncorrelated or long-range correlated random numbers undergo a change that is *independent* of n, i.e., a homogeneous shift for all short and long window sizes. A similar behavior is well-known for the DFA. For a "γ-change point", on the contrary, one has to distinguish between long and short window size: for small n the slopes of long-range correlated numbers lie below those for uncorrelated numbers, for large n this is reversed.

However, in view of practical applications the important question is "What happens when a change in autocorrelation coefficient is accompanied by a change in the standard deviation of the random numbers?". For many applications, for instance during a transition to REM sleep, the change in correlation is typically accompanied by a change in the standard deviation of the heart beat intervals. It might even be argued that in some time series a change in statistics may accidentally be masked by an accompanying change in the statistical variance. In view of what has just been said above, however, this can never be so for *all* possible window sizes n. This is illustrated in Fig. 4, where the combined change in γ *and* σ is indeed masked for the $n = 8$ curve, but not for the entire set of curves. Whether the two parameter changes add up or act "antagonistically" depends on the window size. This effect can be seen in Fig. 4, where the curves with window sizes $n = 3$–7 fall below the extended "no change-point" curve, whereas the curves with $n = 8$–12 turn steeper. All of this implies that PDFA can in fact discern between changes in correlation and changes in statistical variance, provided several curves for varying window size are calculated.

In practical contexts one frequently encounters time series with strong trends which can only piecewise be approximated by polynomial trends. Sinusoidal trends, for instance due to seasonal or circadian oscillations of a variable, are a common

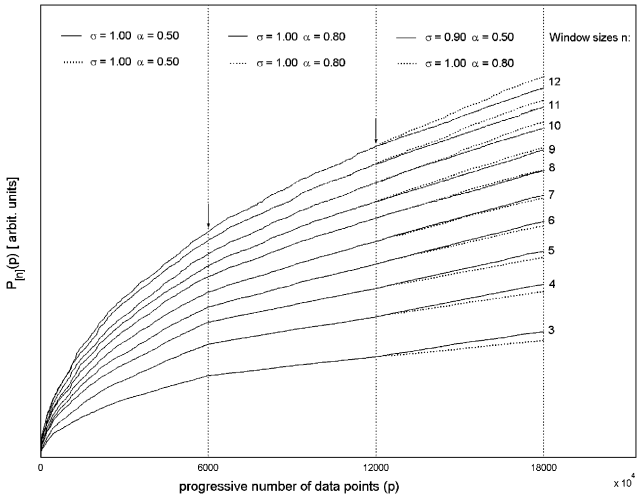

Fig. 4. Simultaneous changes in correlation coefficient and statistical variance may mask the change-point for the curve belonging to a particular window size (here $n = 8$), but never in the entire set of curves.

example. Moreover, the trends themselves may be non-stationary in their nature. As long as these temporal changes stay on a time scale which is small compared to the chosen PDFA window sizes, the PDFA curves will not be affected. As an example we have added a sinusoidal trend of period 3000 to the time series of Fig. 2 and analyzed it with PDFA. Fig. 5(a) shows the original time series, (d) the resulting time series, and (b) and (c) the smoothed PDFA curves of both time series (same seeding for the quasi-random numbers). Obviously the sinusoidal trend did not affect the performance of the PDFA method. As for DFA, it is possible to eliminate trends requiring higher polynomial fits by going to higher order PDFA.

"Spiking", i.e., the presence of narrow peaks which are often of external origin and not related to the intrinsic system dynamics, is another common problem inherent in many real data-sets. Therefore, this issue has been discussed in the literature [20]. In the following we discuss the performance of PDFA for "change-point segmentation" in the presence of spikes. To this end we have generated data-sets with abundant spiking by adding 4% of random (uncorrelated) large-amplitude spikes to a time series of true random numbers (cf. Fig. 6). It can be seen in Fig. 6 that the spikes affect the curves corresponding to small window size n more than it affects the curves belonging to large window-size n, which is only to be expected for

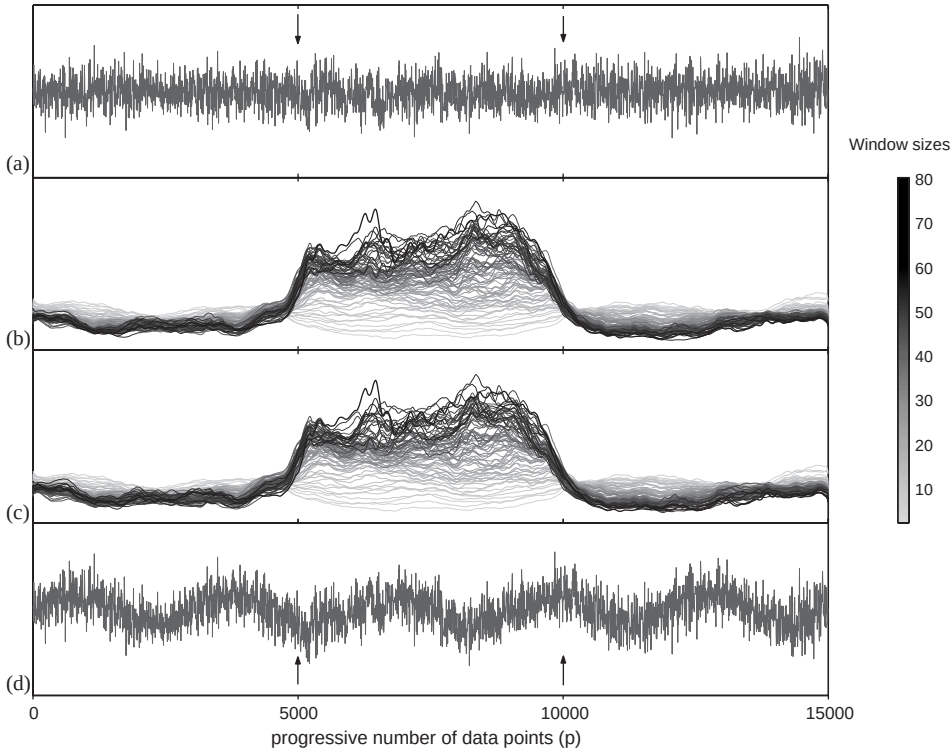

Fig. 5. Sinusoidal trend: (a) original time series without the trend; (b) PDFA curves for the original time series; (c) PDFA curves for the new time series; (d) time series with added sinusoidal trend.

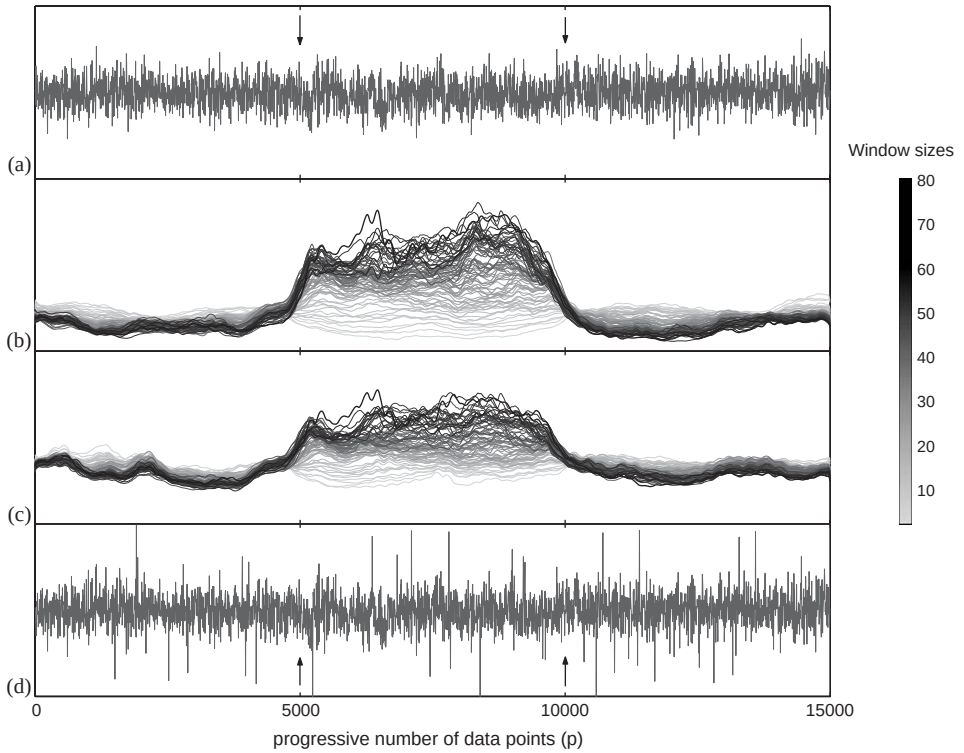

Fig. 6. Impact of spikes on the PDFA analysis: (a) time series without spikes; (b) PDFA curves for original time series; (c) PDFA curves for time series with spikes; (d) time series with 4% of uncorrelated and normally distributed spikes added. It can be seen that PDFA is a very robust method concerning spiking.

such temporally narrow features. The same behavior was also found in sleep datasets with ''spikes'' stemming from movement artifacts during the polysomnographic recording, which could be identified as such (by the numerical analysis) because of their prevalence to the small-n curves.

4. Application to sleep lab data

The starting point for the numerical analysis of sleep ECG was the extraction of the time series of RR-intervals (i.e., the time lapse between two subsequent R-waves) from the digitized ECG-channel of the polysomnographic recording by a home-made computer-program. The program compares the values of the ECG data with a reference curve, keeping only data points exceeding the reference curve. The reference curve is a piecewise constant function of steps with both, the grid interval length and the height of the steps (in comparison to the local maximum in the given grid interval), being freely selectable. Within the remaining points the actual maxima

are obtained by direct comparison of neighboring points. Any of these maxima is accepted as an "R wave peak", only if the distance to the last such peak is larger than a chosen limit. On the other hand, if a reconstructed RR-interval *exceeds* the estimated length by a pre-selected factor, a visual control window is opened that allows to look for R waves that were possibly missed, e.g. due to problems with the recording. For a data-set of length 30 000 s (typically corresponding to a file size of a few 10 MByte) with about 27 000–29 000 R waves, on a Pentium IV PC with 1 GByte RAM our "R finder" algorithm takes 15–70 s, depending on the sampling rate, on the noise level and on the number of artifacts (e.g. movement artifacts) in the recording. For normal, i.e., non-pathological, ECG recordings of reasonable quality the number of cases requiring visual assessment and manual correction remained well below 10 reported incidents per data-set.

None of the data-sets discussed below contained extrasystoles. However, for other recordings with regularly occurring extrasystoles we have checked that it is possible to detect all extrasystoles or to suppress them by choosing a suitable "critical distance parameter". Furthermore, a strong drift in the ECG baseline may be countered by finer time interval length. The fact that this procedure allows *interactive* checking of the resulting RR-intervals is very important for monitoring the quality of the time series data generated. On the other hand, the possibility to change the pre-set sensitivity to deviations from the "typical" pattern avoids unnecessary and cumbersome interruptions of the data processing: dealing with a problematic data-set, for instance, one may start off with initially quite narrow margins for tolerated deviations, and then successively relax these conditions to a level where the number of reported incidents is reduced to the necessary level.

Furthermore, we emphasize that our method does not require any assumptions on the parametric dependence of the underlying process(es) giving rise to the time series. Conversely, our approach might be helpful for gaining insight into the sleep dynamics and for formulating or testing physiological sleep models. The dynamics of sleep–wake transitions, for instance, has been investigated and modelled [23], and stochastic signals have been designed to closely resemble the heart beat dynamics of wake or various sleep stages [24]. One of the questions being investigated involves the feasibility to *specifically* detect vegetative arousals [25] with this method. The intention and the scope of the present paper, however, is restricted to demonstrating the feasibility of the approach by concentrating on examples chosen from typical sleep data-sets exemplarily. Normal night recordings (i.e., following an adaptation night in the sleep lab) from young and healthy male probands were included. These sleep data-sets contained the normal percentage of NREM and REM stages within the physiological sleep cycle pattern. Finally we remark that for the comparison of the numerical results with the manual scoring, exact synchronization of all types of recordings is of paramount importance (no time-leaps, omissions, etc.).

In the following examples we have super-imposed several PDFA curves, corresponding to various selected window sizes, on a color-coded sleep stage diagram derived independently and prior to numerical analysis by manual scoring of the polysomnography according to the most widely accepted Rechtschaffen and Kales criteria [26]. A change in hue of the curves from red to blue indicated an

increase in window size n (cf. color-coded bar attached at the right side in units of heart beats).

In Fig. 7 we have chosen an example from our analyzed data, where the following typical behavior can be observed: transitions from "deeper" sleep stages to "lighter" sleep stages (e.g. $4 \rightarrow 3$) are clearly indicated as pronounced steps in the curves, but not so for descending transitions (e.g. $3 \rightarrow 4$). This is a general effect which we have found to be present in a large number of data-sets (to date about 50) recorded from healthy young probands [12]. A more extended statistical evaluation of this effect is under way. In view of the experience gained from artificial data (cf. Fig. 2), we think that the occurrence of such steps can be attributed to short intervals of arousals *between* the sleep stages in ascending transitions, which are particularly well detected by our algorithm. Note also the different average slope in the REM sleep stage just before 6 h after beginning of the recording in Fig. 7, compared to the sleep stage 2 just after 6 h, as expected due to the larger scaling coefficient α and the larger variance σ typical for the long-time correlated REM sleep. In Fig. 8 we have selected another example from a different whole night recording which demonstrates the capability of our algorithm to detect short embedded incidents of wakefulness, which typically become more frequent closer to awakening in the morning.

Figs. 7 and 8 also illustrate the relevance of the window size parameter n for the trade-off between high detection sensitivity and high time-resolution: Choosing a small window size for the complete analysis leads to good time-resolution of the transition time, however, at the cost of visibility due to a less pronounced pitch of the steps in the curve. On the other hand, for a larger window size the transitions become more conspicuous at the cost of reduced transition time, which is limited by the size of the window.

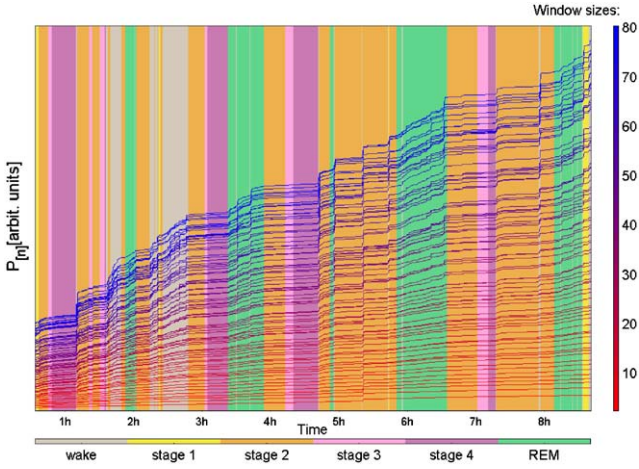

Fig. 7. A typical example for PDFA functions for various window sizes overlaid on the manually scored color-coded sleep stage diagram of a whole night. (For colour reproduction of this figure see on-line publication.)

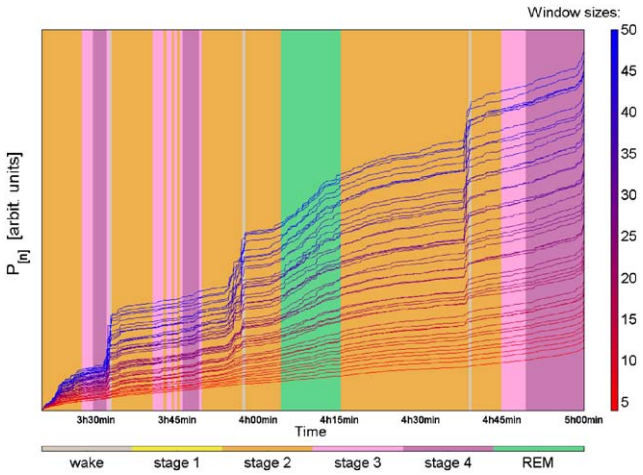

Fig. 8. Detecting short episodes of wakefulness. (For colour reproduction of this figure see on-line publication.)

To distinguish NREM sleep from REM sleep we suggest to utilize the built-in scaling dependence: Scanning the PDFA scaling parameter, i.e., the window size n, across the correlation time of NREM-sleep (which is around 6 heart beats for light sleep [10]), there is a striking difference between periods of REM and NREM sleep. If one superposes several numerically differentiated PDFA curves, episodes of NREM sleep are characterized by a splitting (or "dispersion") of the curves belonging to different window sizes (Fig. 9a), which is absent in REM sleep. Applying the general considerations from Section 3 (in particular the discussion of the second change-point in Fig. 3), we can immediately conclude that NREM sleep cannot be long-range correlated from just viewing Fig. 9, since in NREM sleep the (red) curves belonging to small window sizes are seen to have steeper slopes than the (blue) curves for long window size. In fact, in contrast to REM sleep and wakefulness, NREM sleep is known to be short-range correlated. REM and wakefulness, on the other hand, are known to be long-range correlated and are typically also characterized by a larger variance in the heart beat intervals. However, since REM and wakefulness are long-range correlated, and thus the corresponding correlation times exceed any finite time, one cannot see any significant window-size dependence in the slopes.

For the sake of validation and comparison we have carried out a corresponding calculation on another scale-dependent measure which allows analysis across the data-set. We have chosen the wavelet transform standard deviation $\sigma_{wav}(m)$ described in Ref. [13], where several scale-dependent and scale-independent measures are compared. For our purpose we calculated $\sigma_{wav}(m)$ for a subset of 200 data points (about 3 min) which continuously moved through the whole data set. When the scaling parameter m, which here relates to the width of the chosen wavelet basis, was

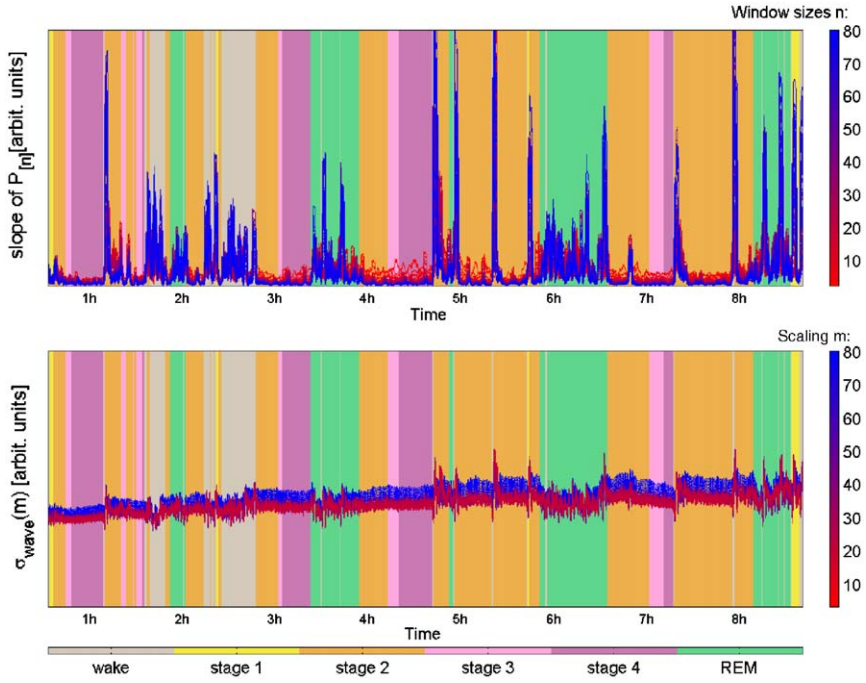

Fig. 9. A "scaling-parameter dispersion" is clearly visible in periods of NREM sleep and absent in REM sleep. Top: window-size dispersion for PDFA, bottom: scaling-parameter dispersion for wavelet transform standard deviation $\sigma_{wav}(m)$. The graphs also demonstrate that the peaks indicating changes in the heart rate statistics coincide for both methods throughout the data-set. (For colour reproduction of this figure see on-line publication.)

scanned in a similar range as n, after having pulled all curves, which typically increase with m, to the same baseline, we saw an analogous pattern of "scaling parameter dispersion" emerging (see Fig. 9b). Furthermore, Fig. 9 convincingly demonstrates that the PDFA measure and the wavelet transform standard deviation measure are sensitive to the same features in the heart rate statistics, as indicated by coinciding peaks for these two independent methods.

The fact that such a "scaling-parameter dispersion", as one might call it, is only present in NREM sleep, can be explained by noting that the correlation time for the long-time correlated sleep stages REM-sleep and wakefulness exceeds any window size. This suggested approach reliably differentiates episodes of NREM and REM sleep visually along the data-set, and thus seems to be a useful tool for "REM versus NREM" staging based on the heart beat statistics alone, complementing other methods. In this context one should mention that a disadvantage of the PDFA method is the longer computation time compared to the wavelet analysis which utilizes a fast algorithm applicable for scales that are integer powers of 2 ($m = 2^j$).

5. Conclusions and outlook

Summarizing, in the present paper our intention was first to introduce an alternative and independent change-point detection method, to demonstrate its capability to detect change points in the statistics of artificially produced non-stationary correlated random numbers and second to discuss its perspective of application to sleep research. On the artificially generated data-sets with built-in change points we conclude that the approach is very powerful for detecting change points in time series—even in the presence of non-polynomial trends and spiking. We have demonstrated that change points which are characterized by an abrupt (or at least non-adiabatical) change in either statistical correlation or statistical variance (or both) of the data points are reliably and accurately localized in time.

Our opinion concerning sleep lab applications, which was gained by analyzing 48 data-sets of whole night recordings so far from 19 healthy male probands with a mean age of about 30 years, is more complex: We have seen a clear indication that the suggested algorithm is capable of detecting *particular* transitions between sleep stages, typically from deeper to lighter sleep (e.g. from sleep stage 4 to stage 3, but not vice versa [12]). We believe that this asymmetry is due to changes in the state of the autonomous nervous system which represent detectable change-points for our method which are only present in markedly ascending transitions. The sleep data results gained by PDFA were validated by a comparison with other independent measures (i.e., a wavelet analysis). It was found that they were sensitive or insensitive, respectively, to exactly the same sleep stage transitions. This fact, of course, rules out any over-ambitious attempts to detect *all* sleep stage transitions occurring during the night by an on-line processing based on this approach.

Nevertheless, it is possible to extract valuable information from the on-line monitoring of the HRV: We suggest to utilize the scale-dependence in the form of a "scaling-parameter dispersion" (for scale-dependent measures such as PDFA or wavelet analysis calculated progressively along the data-set) for the differentiation between NREM and REM sleep. In the future this may develop into a tool for on-line (possibly even real-time) assessment of NREM versus REM sleep. Naturally, the information included in the ECG curve alone is considerably less than the information contained in the total polysomnographic recordings (EEG, EOG, EMG, etc.). Therefore, we view our approach as a supplementary easy-to-implement method, which might serve as a useful first screening tool (e.g. for distinguishing REM from NREM) together with other methods analyzing the heart rate variability.

And finally, motivated by our surprising finding that only *ascending* transitions (in particular into wakefulness) give rise to a signal in our numerical analysis, we intend to assess the usefulness of the method to detect autonomous arousals, which might be the underlying reason for a "change point" detected by our numerical analysis. We intend to investigate their role and their relation to sympathetic activation, which can alternatively be determined by spectral analysis methods [27] and by noninvasive measurements [28].

Acknowledgements

The authors gratefully acknowledge support by the Austrian National Bank (Jubiläumsfondsprojekt No. 9302).

References

[1] P.Ch. Ivanov, et al., Chaos 11 (2001) 641–652.
[2] P.Ch. Ivanov, et al., Nature 399 (1999) 461–465.
[3] C.-K. Peng, S. Havlin, H.E. Stanley, A.L. Goldberger, Chaos 5 (1995) 82–87.
[4] Y. Ashkenazy, et al., Phys. Rev. Lett. 86 (2001) 1900–1903.
[5] P.Ch. Ivanov, et al., Europhys. Lett. 48 (1999) 594–600.
[6] L.A. Amaral, et al., Phys. Rev. Lett. 81 (1998) 2388–2391.
[7] L.A. Amaral, et al., Phys. Rev. Lett. 86 (2001) 6026–6029.
[8] T. Penzel, et al., IEEE Trans. Biomed. Eng. 50 (2003) 1143–1151.
[9] C.-K. Peng, et al., Phys. Rev. E 49 (1994) 1685–1689.
[10] A. Bunde, et al., Phys. Rev. Lett. 85 (2000) 3736–3739.
[11] J.W. Kantelhard, et al., Phys. Rev. E 65 (2002) 051908(1)-041107(6).
[12] S. Telser, M. Staudacher, Y. Ploner, A. Amann, H. Hinterhuber, M. Ritsch-Marte, Somnologie 8 (2004) 33–41.
[13] M.C. Teich, et al., in: M. Akay (Ed.), Nonlinear Biomedical Signal Processing, vol. 2, IEEE Press, New York, 2001.
[14] P.Ch. Ivanov, et al., Physica A 249 (1998) 587–593.
[15] Z.R. Struzik, Fractals 9 (2001) 77–93.
[16] P. Bernaola-Galván, P.Ch. Plamen, L.A. Nunes Amaral, H.E. Stanley, Phys. Rev. Lett. 87 (2001) 168105(1)-168105(4).
[17] K. Hu, et al., Phys. Rev. E 64 (2001) 011114(1)-011114(19).
[18] J.W. Kantelhard, et al., Physica A 316 (2002) 87–114.
[19] J.W. Kantelhard, et al., Physica A 295 (2001) 441–454.
[20] Z. Chen, P.Ch. Ivanov, K. Hu, H.E. Stanley, Phys. Rev. E 65 (2002) 041107(1)-041107(15).
[21] True random numbers generated using a physical quantum random number generator can be downloaded from a website provided by a collaboration between the University of Geneva (N. Gisin) and the company id Quantique http://www.randomnumbers.info/index.jsp.
[22] H.A. Makse, S. Havlin, M. Schwartz, H.E. Stanley, Phys. Rev. E 53 (1996) 5445–5449.
[23] C.-C. Lo, et al., Europhys. Lett. 57 (2001) 625–631.
[24] J.W. Kantelhardt, et al., Europhys. Lett. 62 (2003) 147–153.
[25] M.H. Kryger, T. Roth, W.C. Dement, Principles and Practice of Sleep Medicine, 3rd ed., W.B. Saunders Company, Philadelphia, 2000.
[26] A. Rechtschaffen, A. Kales, A Manual of Standardized Terminology, Techniques and Scoring System for Sleep Stages of Human Subjects, NIH Publ 204, US Gov Print Off, Washington, 1968.
[27] M.H. Bonnet, D.L. Arand, Electroencephalogr. Clin. Neurophysiol. 102 (1997) 390–396.
[28] T. Penzel, U. Brandenburg, R. Fricke, J.-H. Peter, Somnologie 6 (2002) 69–73.