J. theor. Biol. (1994) **170**, 273–281

A Method for Filtering Respiratory Oscillations

A. C. FOWLER†¶, G. KEMBER†, P. JOHNSON‡, S. J. WALTER§, P. FLEMING‖ AND M. CLEMENTS‖

† *Mathematical Institute, Oxford University, 24–29 St Giles', Oxford OX1 3LB,*
‡ *John Radcliffe Hospital, Oxford OX3 9DU,* § *Oxford Brookes University, Headington, Oxford and* ‖ *St Michael's Hospital, Bristol BS2 8EG, U.K.*

(*Received on* 3 *July* 1993, *Accepted in revised form on* 17 *March* 1994)

We present a method based on dynamical systems theory which can be used to filter time series in a way which is superior to classical Fourier decomposition. This method is applied to three data-sets, taken from respiratory measurements of two children in quiet and REM sleep. Our purpose is to filter the several different oscillatory mechanisms which operate, in order to provide clearer signals on which further analysis and diagnosis can be based.

1. Introduction

Normal breathing in humans is an oscillatory phenomenon, but one which has variable irregularities caused by a variety of oscillatory and random influences. We wish to be able to distinguish abnormal from normal irregularities (rhythms), in order to determine the nature of any defect of regulation, and its likelihood of predicting risk for such clinical problems as respiratory illness or even cot death (SIDS).

For example, periodic breathing is common in infants, but if excessive (e.g. more than 5% of sleep time), is considered by some (Kelly & Shannon, 1979) to indicate an infant with increased risk of sudden infant death syndrome (SIDS). Others have also considered this pattern to indicate instability in the chemoreceptor control of breathing, but showed that it is age-dependent (Fleming *et al.*, 1984). However, it is now evident that the most prevalent pattern of periodic breathing in infants has a cycle time of 11–14 seconds and is related primarily to temperature control, since it also persists in lambs after peripheral chemodenervation when in a warm environment (Johnson & Andrews, 1990). Thus, even obvious oscillations in breathing pattern are multifactorial in origin.

Another major (longer term) influence on breathing pattern is sleep state, and especially active or rapid eye movement (REM) sleep, when the regulatory mechanisms (some of which are described above) may be greatly altered. In studying records with a view to understanding the response of the system to varying external conditions, it is therefore necessary to have a method to filter the series so as to reproduce the features of interest.

The classical approach to doing this is to use Fourier analysis, which, for a given segment of a record, gives a decomposition of the signal in terms of its Fourier frequency components. This method, which is 200 years old and designed for linear systems, is inappropriate for data in which trends occur, and where the underlying mechanism is non-linear and the dynamics may be chaotic. Other filtering methods, such as moving averages, suffer similar disadvantages.

In this paper, we apply a method based on dynamical systems theory, which is suitable for providing a non-linear filter for coupled multi-oscillator systems. The method uses a Fourier spectrum to identify significant cyclic components, but the filter does not select any one frequency, or band of frequencies; rather, the data *itself* is used to provide the best

¶ Author to whom correspondence should be addressed.

0022-5193/94/190273 + 09 $08.00/0

"spectral" decomposition functions. In particular, the decomposition yields filtered time series whose periodicity and amplitude may vary, and which in fact may themselves be chaotic. Here, we show how this technique can yield significant insight into respiratory data by applying it to three segments of such data, one taken in quiet sleep of an infant, one from REM sleep of the same infant, and the third from REM sleep of a different infant. The third record, in particular, shows large fluctuations in signal, and is considered to be a severe test for this decomposition method, although the data is considered to be in the normal range.

2. Dynamical Systems Approach

There is now a large (though recent) literature on this topic (e.g. Broomhead & King, 1986; Smith, 1992); here we give only a brief description. Rather than work with a scalar time series $y(t)$, we choose to embed the series as a trajectory in an embedded phase space (Takens, 1981). To do this, we choose an embedding dimension d_E and a time lag Δ, and then we define the d_E-dimensional vector

$$x(t) = (y(t), y(t-\Delta), \ldots y\{t-(d_E-1)\Delta\}). \quad (1)$$

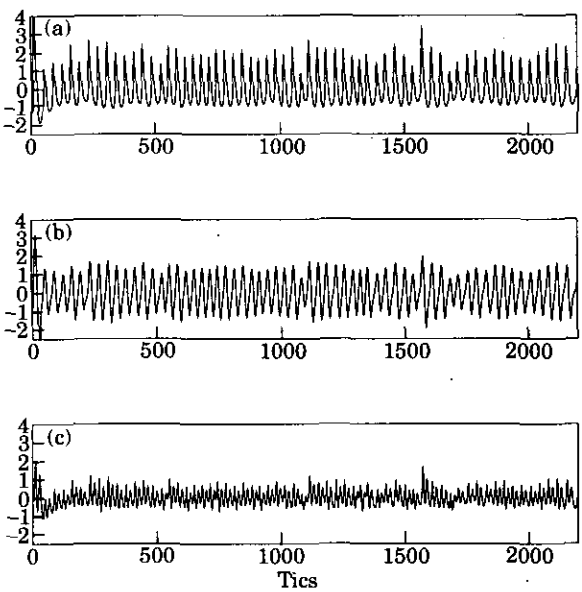

FIG. 1. First infant. (a) Quiet sleep signal y; (b) filtered signal y_F; (c) residual signal y_R. Here and in subsequent figures, the signals are the sum of measurements taken from chest and abdomen receptor bands, and as such are a proxy measure of lung volume. The data is sampled at 20 Hz, and each "tic" is a data point. Thus this sample of about 2300 ticks corresponds to about 115 sec. In this figure and all subsequent ones, the data has been normalized to have zero mean and unit variance. The units of measurement are thus standard deviations.

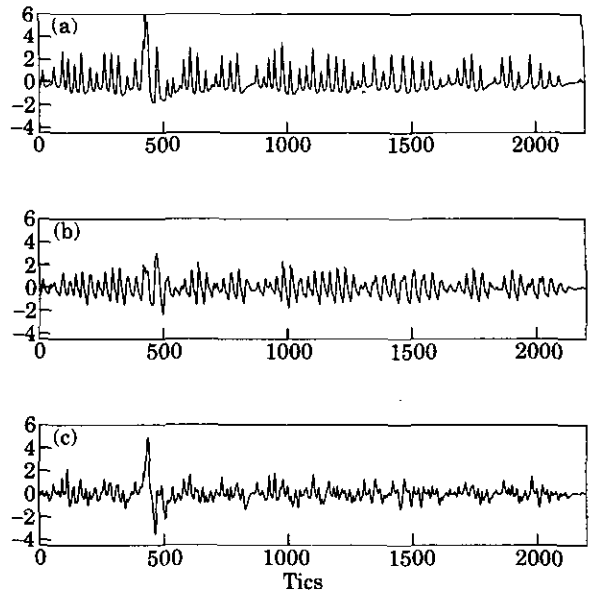

FIG. 2. First infant. (a) REM sleep signal y; (b) filtered y_F; (c) residual y_R.

If the signal $y(t)$ arises from a finite-dimensional deterministic system, then (providing d_E is sufficiently large), the trajectory $x(t)$ gives a genuine image of the dynamics of the signal in its own phase space. To give a simple example, the natural phase-space trajectory of the periodic oscillator $y(t) = \frac{1}{\sqrt{2}} \sin \pi t$ (the amplitude is chosen so that the series has zero mean and unit variance) is $(y, \dot{y}) = \frac{1}{\sqrt{2}}(\sin \pi t, \pi \cos \pi t)$, which is an ellipse. A slightly different choice of co-ordinates $(y, \dot{y}/\pi)$ gives a circle. Similarly, the embedding trajectory $\frac{1}{\sqrt{2}}(\sin \pi t, \sin \pi (t-\Delta))$ is an ellipse, whose eccentricity varies with Δ, and for $\Delta = \frac{1}{2}$ it also is a circle.

More generally, a multi-oscillator trajectory will fill a region in a larger space. For example, a two-oscillator system could have trajectories which lie on or near a torus in a three-dimensional space. By choosing Δ appropriately (using a measure called the singular value fraction; see Kember & Fowler, 1993), we can optimize the spread of the embedded trajectory—effectively we choose Δ so that the basic oscillator is as nearly circular as possible.

Having chosen the embedding in this way, we use a computational algebraic method called *singular value decomposition* (SVD), which has the effect of calculating the principal axes of the "mass" represented by the cloud of points generated in the embedded phase space by the trajectory. In effect, these *singular vectors* give a geometric description of where the trajectory lies, and associated singular values are a measure of the extent of the trajectory "surface" in the corresponding directions.

For example, a doubly periodic two-oscillator signal lying on or near a torus in the embedded phase space would have its two principal singular vectors spanning the plane in which (one hopes) the primary oscillator lies. Our method now projects the embedded trajectory on to this plane, and as a result we obtain a *filtered* (projected) trajectory and a *residual* trajectory (Kember & Fowler, 1994). When these are transformed back as scalar series, we thus have a (non-Fourier) decomposition

$$y(t) = y_F(t) + y_R(t), \tag{2}$$

in which the filtered series $y_F(t)$ represents the primary oscillator, and y_R is the residue. Note that y_F is not constrained to be periodic, nor in general will it be. If y is chaotic, then in general y_F will be also, but less so.

Our procedure continues in this way: if y_F is essentially periodic, then we repeat the process for y_R. However, considered as a time series in its own right, we embed y_R with its own choice of d_E and Δ, chosen to maximize spread and to retain resolution of the particular oscillatory signal we seek to filter. Since the "time window" of width of the embedded vector $x(t)$ is $(d_E - 1)\Delta$, and since the optimal window for "circularizing" a sine wave is a half of the period P, this suggests that (for roughly sinusoidal oscillations) we choose

$$(d_E - 1)\Delta \approx P/2, \tag{3}$$

and then choose Δ by finding the first minimum of the singular value fraction (Kember & Fowler, 1993). In practice, there is flexibility in this choice, and some trial and error is recommended. In particular, if the oscillations are not close to planar in the embedding space, (3) is liable not to be the best choice.

We illustrate the technique using segments of respiratory data taken from an infant in quiet sleep (Fig. 1) and REM sleep (Fig. 2). The signal is the sum of two signals obtained from chest and abdomen movements. Each "tic" is a data point, and represents a time interval of 0·05 sec. It can be seen that the data segment in Fig. 1(a) is fairly regular with variable alterations in amplitude and frequency, whereas in Fig. 2(a), the variation is much more marked, and contains some breathing movements which appear to be biphasic (i.e. at double frequency), as well as greater longer-term oscillations or periodicities. To attempt an effective comparison between two such sleep states in any one infant, as well as between infants, the intrinsic features of the breathing pattern need to be "extracted". Irregular transients, such as are visible in Fig. 2, can cause havoc with spectral methods, linear filters, etc.

We have embedded the signal in Fig. 1 as a trajectory choosing $d_E = 5$, $\Delta = 4$. This dimension is the lowest which can guarantee properly resolving a two-component orbit (i.e. a closed loop in the phase space) without having it intersect itself. A little examination (see Fig. 3) shows that, in the terminology of Kember & Fowler (1993), the singular value

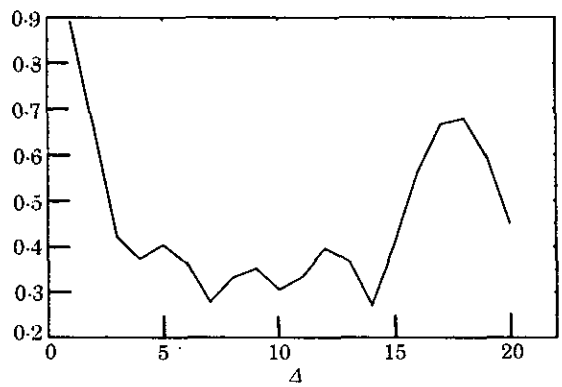

FIG. 3. Singular value fraction (SVF) $f_{SV}(1)$ as a function of lag time Δ. The first minimum is at $\Delta = 4$.

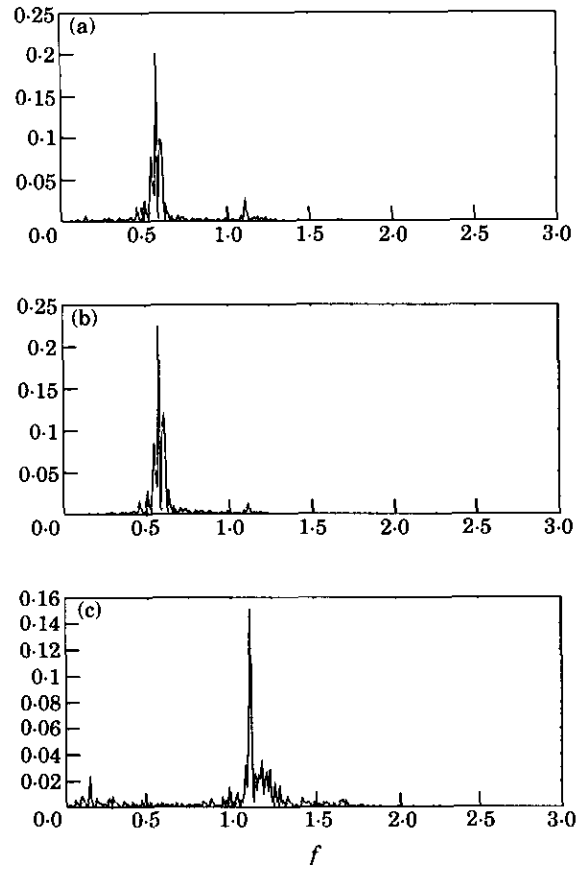

FIG. 4. Spectral power density of series in Fig. 1. (a), (b) and (c) as for Fig. 1.

fraction (SVF) $f_{SV}(1)$ selects a window $\Delta_w \approx 18$, for which the corresponding lag is $\Delta = \Delta_w/(d_E - 1) \approx 4$.

If we project the signal $y(t)$ on to the first two singular vectors, we obtain the filtered signal y_F and the residual signal y_R, shown in Fig. 1. Their respective spectra are shown in Fig. 4. The basic respiratory pattern is represented by y_F, while y_R contains a small second harmonic component. In this case, one pass gives an effectively clean signal, with no other independent oscillators apparently present.

Next we turn to the REM sleep segment shown in Fig. 2. It is filtered as before, and y_F and y_R are also shown in Fig. 2. The respective power spectra are shown in Fig. 5. The primary respiratory frequency is clearly visible at $f \approx 0{\cdot}7$, but there is a lot of noisy periodicity. The filtered signal is still very messy, comprising a band of frequencies round the respiratory frequency.

To further resolve the respiratory signal, we submit y_F to a second pass. We again choose $d_E = 5$ and $\Delta = 4$, and project the resulting embedded trajectory of y_F onto the first two singular vectors. The resultant projected series y_{FF} and the second residual y_{FR} are shown in Fig. 6, and their power spectra in Fig. 7. In this sequential way, we refine the filtering. The interesting feature of this second filter is that, while a linear frame of mind might lead us to expect that the side bands in the spectrum of y_F would be successively

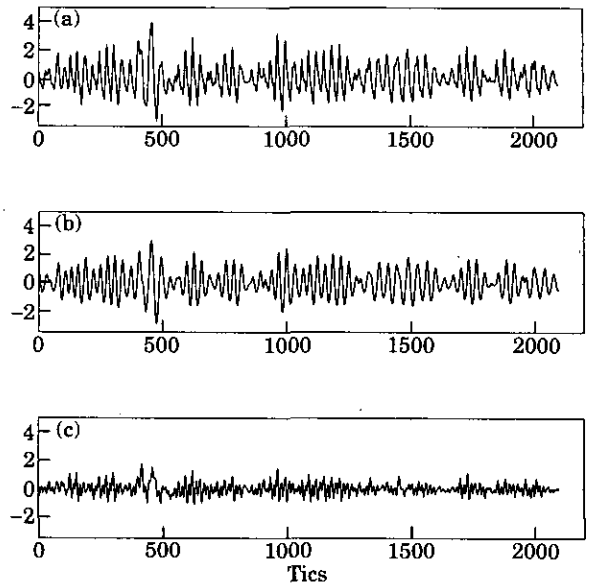

FIG. 6. (a) Filtered series y_F (Fig. 2); (b) second filter y_{FF}; (c) second residual y_{FR}.

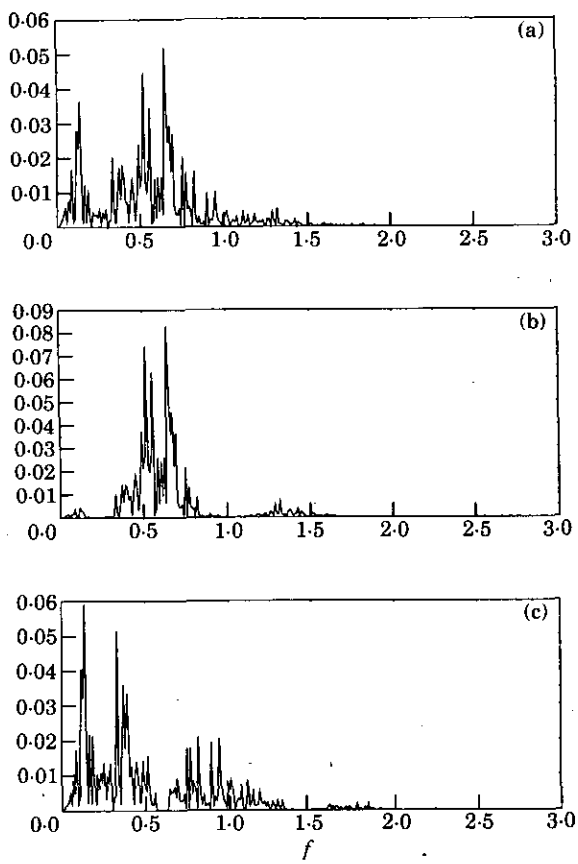

FIG. 5. Spectra of series in Fig. 2. (a), (b) and (c) as for Fig. 2.

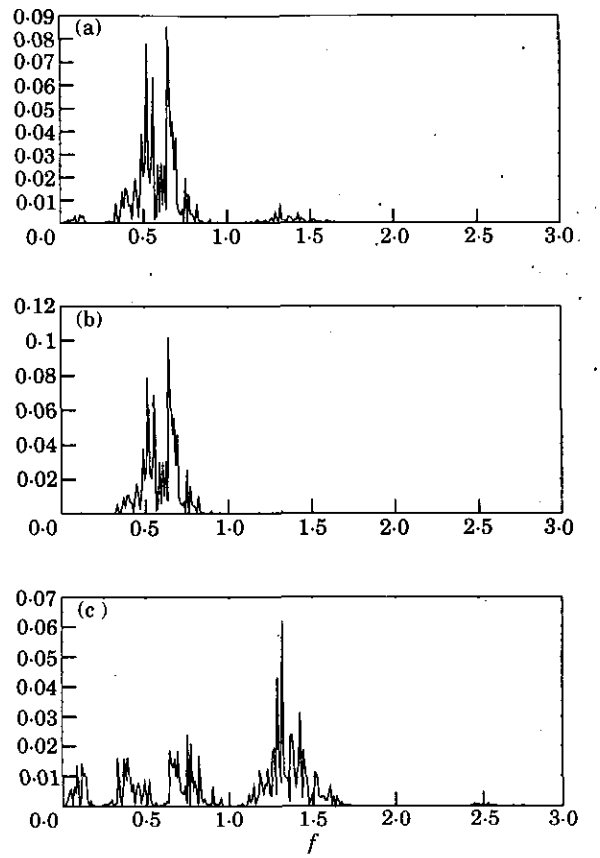

FIG. 7. Spectra of series in Fig. 6. (a), (b) and (c) as for Fig. 6.

cleaned out, this is not the case, and the spectrum of y_{FF} is very similar to that of y_F. This suggests that the iterative filter *converges*, so that in a meaningful sense, y_{FF} in Fig. 6 represents a genuine signal, and the variations in amplitude represent real fluctuations in the depth of breathing. There is no way a linear filter could hope to distinguish this.

In the following section we apply this method to a much longer segment of REM sleep of a different infant, which is a much more noisy-looking series.

3. Application to REM Sleep

We now turn to the systematic application of our filtering method to the time series shown in Fig. 8, which represents the sum of two signals measuring chest and abdomen inflation of a (different) sleeping infant in REM sleep. It is clear that this is an extremely irregular signal, where the basic respiratory pattern is shielded by a slowly varying high-amplitude signal, and as such, it represents a severe test for our technique, which aims to extract the underlying oscillator(s).

FIRST PASS, ACTUAL SIGNAL

In all the embeddings to be described, we choose $d_E = 8$. Some experimentation is necessary, but so long as d_E is large enough to distinguish between different oscillators, the results are robust. We have computed the singular value fraction $f_{SV}(k)$ as a function of Δ for $k = 2$, 4, 6 (Kember & Fowler, 1993). The lowest values reveal no minimum,

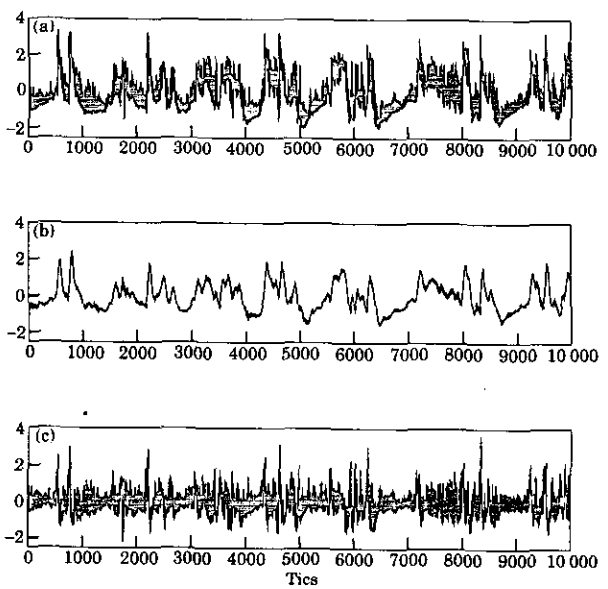

FIG. 8. Second infant. (a) REM sleep signal y; (b) filter y_F; (c) residual y_R.

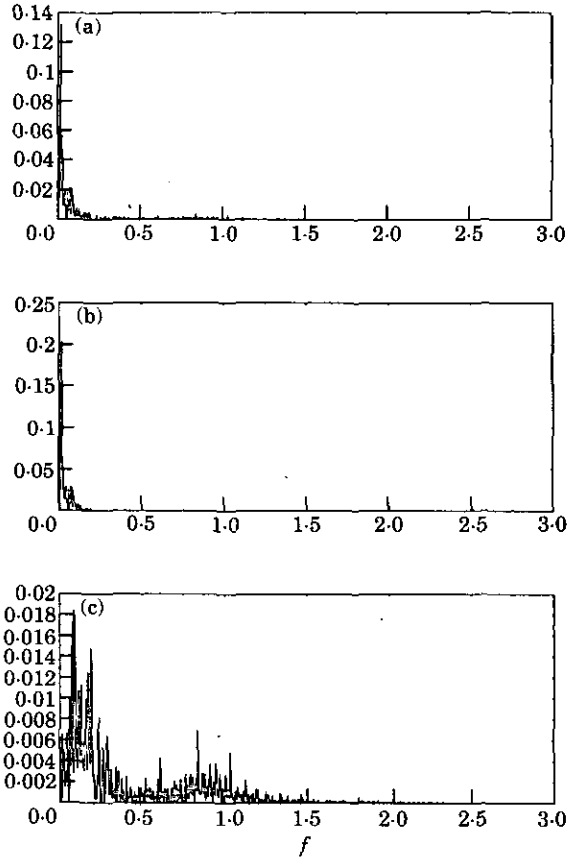

FIG. 9. Spectra of series in Fig. 8, with (a), (b) and (c) as for Fig. 8.

corresponding to the idea that the oscillators lie in the higher singular vectors, but a clear minimum for $k = 6$ is at $\Delta = 8$, and we choose this as the lag time. The resulting singular value decomposition of the embedded series gives the results in Fig. 8, where y_F is the series reconstructed from the projection of the embedded trajectory on to the first singular value, and y_R is the residual signal. Their power spectra (Fig. 9) show that y_F has lost the breathing pattern, and contains most (but not all) of the low-frequency noise, while y_R contains the breathing pattern, but is still contaminated by low-frequency noise.

SECOND PASS, y_F

We now consider the filtered time series y_F obtained from y by projection on to the first singular vector of the preceding embedding. Computation of the SVF, $f_{SV}(k)$, for various k reveals a minimum for higher k near $\Delta = 100$; we choose $\Delta = 94$, and the corresponding embedded trajectory is projected on to the first two singular vectors to obtain the filtered series y_{FF} and residual y_{FR} shown in Fig. 10. We see that it is roughly doubly periodic, consisting of a large-

amplitude slow oscillation with a higher-frequency oscillator riding on it. Though more clearly still noisy, the power spectrum of y_{FF} is roughly doubly periodic, with spectral peaks near $f = 0{\cdot}03$ and $f = 0{\cdot}2$ (see Fig. 11).

THIRD PASS, y_R

We now consider the residual signal from the first pass. Recall that this is the signal which contains all the respiratory pattern. Fig. 12 gives the signal. The SVF seemed to indicate a choice of $\Delta = 8$; however, in practice this was not useful, and so the alternative choice based on a window width $P/2$ of about 20 (at the peak of f_{SV}) suggests $\Delta \approx P/2d_E \approx 3$. With the choice $\Delta = 3$, we find that projection of y_R on to the first singular vector gives the filtered series y_{RF} shown in Fig. 12. Clearly, we have filtered the noisy trend, and the respiratory pattern resides in the second residual, y_{RR}, also shown in Fig. 12. Despite its variability, and after two filters of the original data set, we have a decomposition

$$y = y_F + y_{RF} + y_{RR}, \tag{4}$$

where the respiratory signal is present in y_{RR}, and where in suitably defined spaces, y_F and y_{RR}, and y_{RF} and y_{RR}, are orthogonal. The spectra are shown in Fig. 13.

FOURTH PASS, y_{RR}

We now filter the respiratory signal y_{RR} (Fig. 14). Here, the SVF picks out a value $\Delta = 9$. Projection of the signal on to the first two singular vectors gives the

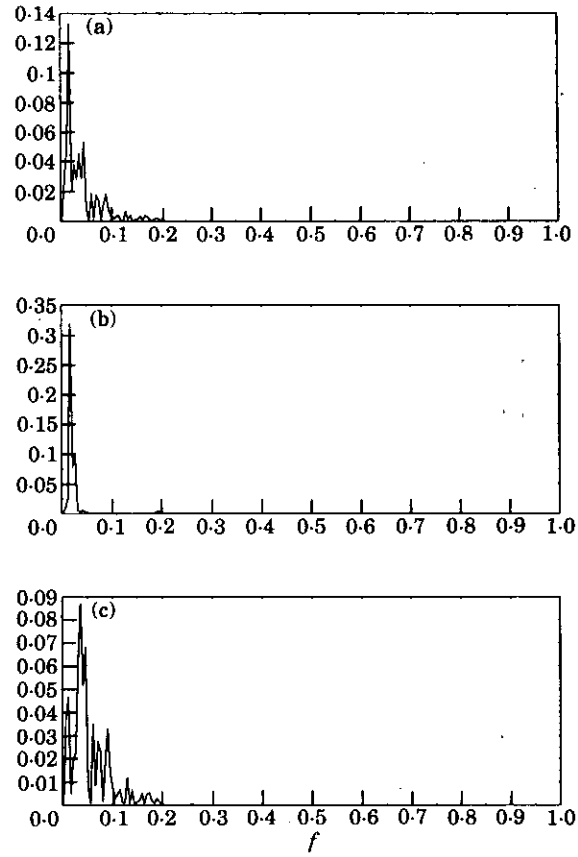

FIG. 11. Spectra of Fig. 10, with (a), (b) and (c) as for Fig. 10.

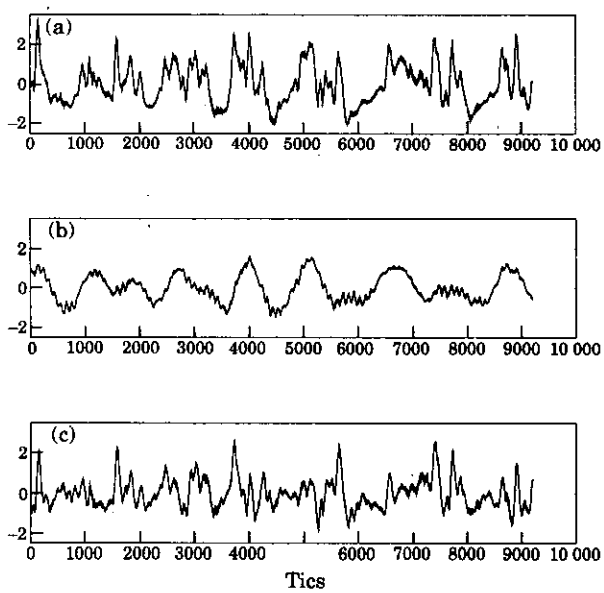

FIG. 10. (a) Filtered series y_F (Fig. 8); (b) second filter y_{FF}; (c) second residual y_{FR}.

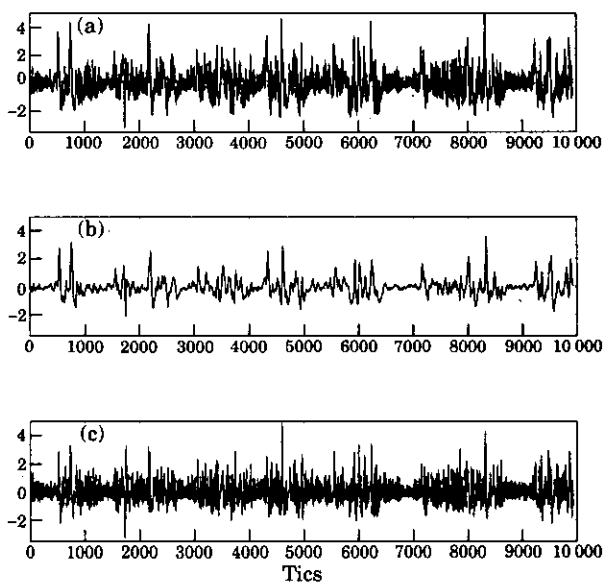

FIG. 12. (a) Residual y_R (Fig. 8); (b) filter y_{RF}; (c) residual y_{R2}.

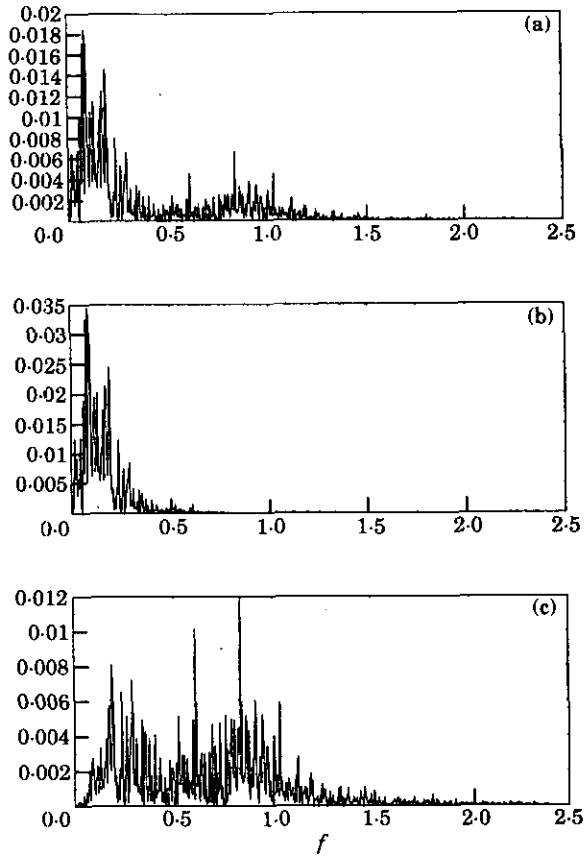

FIG. 13. Spectra of Fig. 12, with (a), (b) and (c) as for Fig. 12.

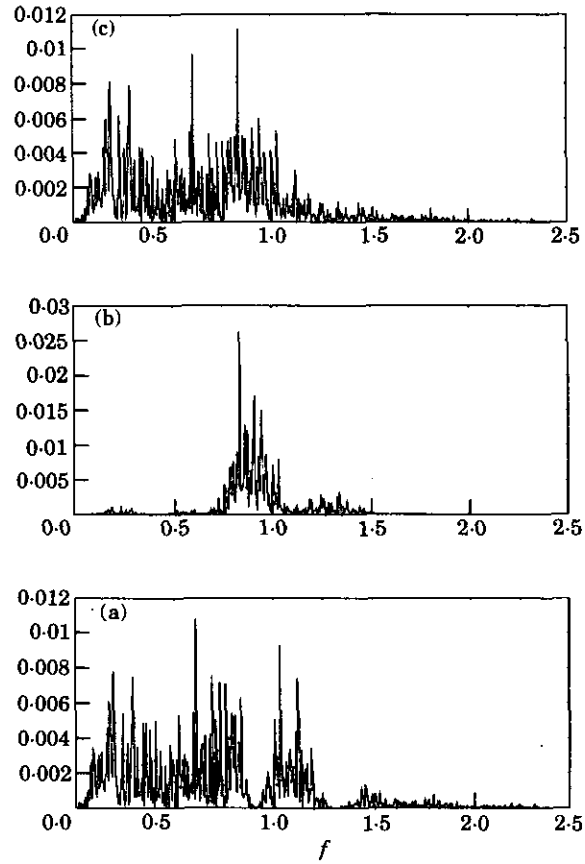

FIG. 15. Spectra of Fig. 14. (a), (b) and (c) as for Fig. 14.

series y_{R^2F}. Although variable, we have now clearly isolated the respiratory signal from the low-frequency noise present in the third residual y_{R^3}. The power spectra are in Fig. 15. A segment of the respiratory projection y_{R^2F} is shown in Fig. 16, and a projected embedding of this segment, $y_{R^2F}(t)$ vs. $y_{R^2F}(t-5)$ (5 being approximately a quarter of the period) is shown in Fig. 17. Although the signal is clearly more than two-dimensional, its cyclic nature is readily apparent.

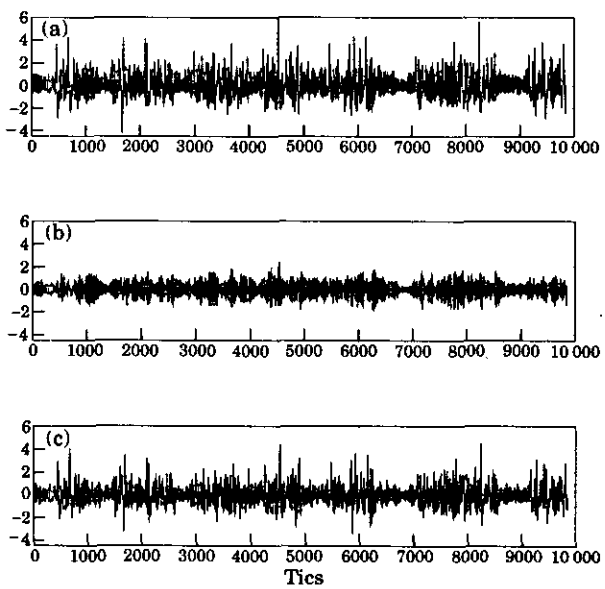

FIG. 14. (a) Second residual y_{R^2} (Fig. 12); (b) filter y_{R^2F}; (c) residual y_{R^3}.

FIFTH PASS, y_{R^3}

Clearly the filtering process can continue *ad infinitum*. In so far as the decomposition isolates individual processes, it may be of interest to pursue one particular branch of the *decomposition tree* shown below, and this we choose to do by analysing the residual respiratory signal y_{R^3} as $y_{R^3F} + y_{R^4}$:

$$
\begin{aligned}
y &\overset{1}{=} y_{F^2} + y_{FR} \\
y_F &\overset{2}{=} y_{RF} + y_{R^2} \\
y_R &\overset{3}{=} y_{R^2F} + y_{R^3} \\
y_{R^2} &\overset{4}{=} y_{R^3F} + y_{R^4} \\
y_{R^3} &\overset{5}{=} \text{respiratory}
\end{aligned}
\tag{5}
$$

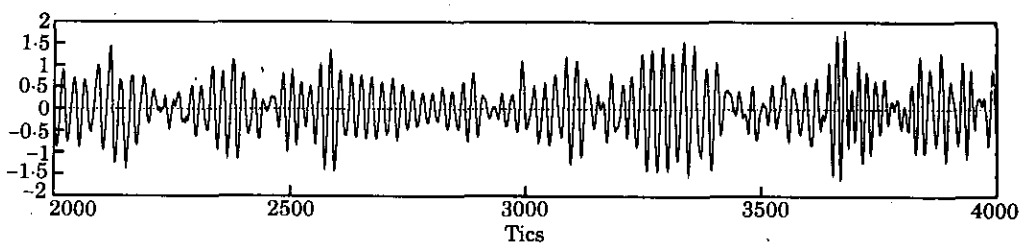

FIG. 16. Segment of y_{R^2F}(SV(1,2)).

In this decomposition, we identify y_{R^2} as "the" respiratory signal, which itself decomposes as the basic respiratory pattern y_{R^2F} and a secondary oscillator y_{R^3} which appears to be forced by the respiratory pattern (at least spectrally). We now filter this secondary oscillator, see Fig. 18. The SVF selects $\Delta = 8$, and the projection of y_{R^3} on to the first two singular values gives the series y_{R^3F} and the residual y_{R^4} in Fig. 18. Although both time series appear equally messy, their respective power spectra in Fig. 19 are strikingly different. In particular, the secondary oscillator is revealed as a chaotic oscillator with base frequency $f \approx 0{\cdot}25$, 0·7 and 1·0, but with power strikingly absent at the even harmonics $f \approx 0{\cdot}45$ and 0·9. The residual y_{R^4} itself has oscillatory characteristics, but with less power and more distributed frequencies than y_{R^2F}.

4. Conclusions

Although intensive and interactive, the method described here appears to be a robust way to filter complicated signals which comprise chaotic oscillations, even where these are swamped by large-amplitude trends. For a clean signal such as is obtained in quiet sleep, the filter easily resolves the primary respiratory pattern. In REM sleep in the first infant, an obviously more variable pattern is similarly resolved.

The much more involved REM sleep segment of the second infant provides a sterner test of the

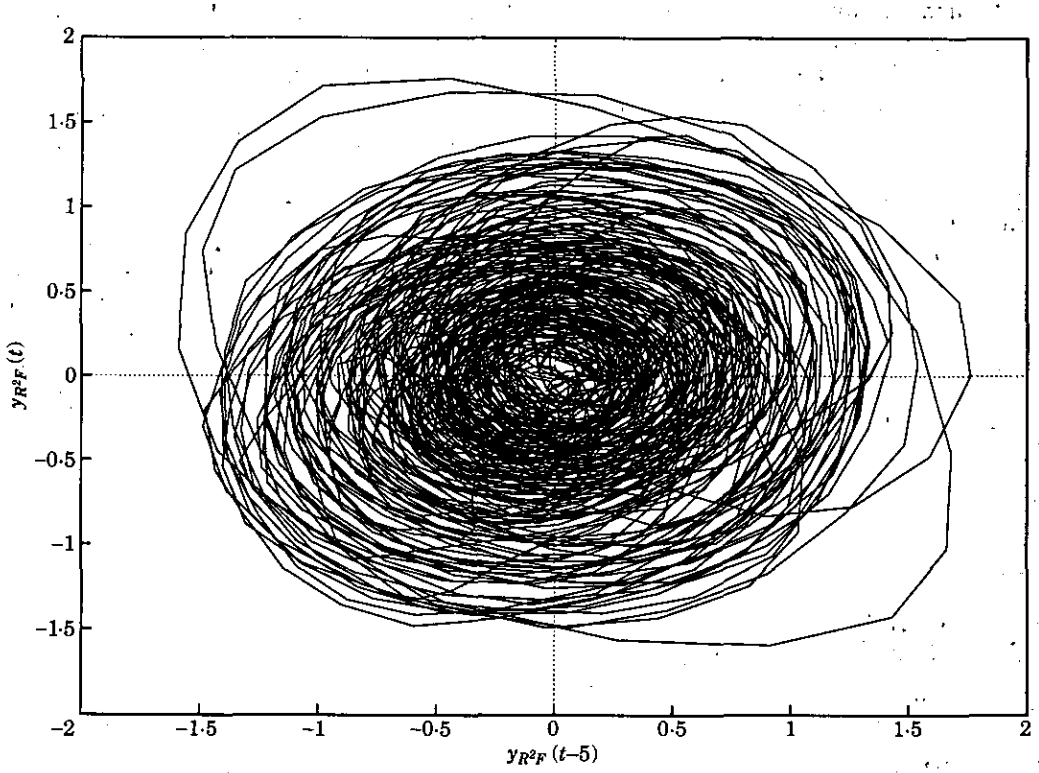

FIG. 17. Two-dimensional projection of y_{R^2F} (Fig. 16).

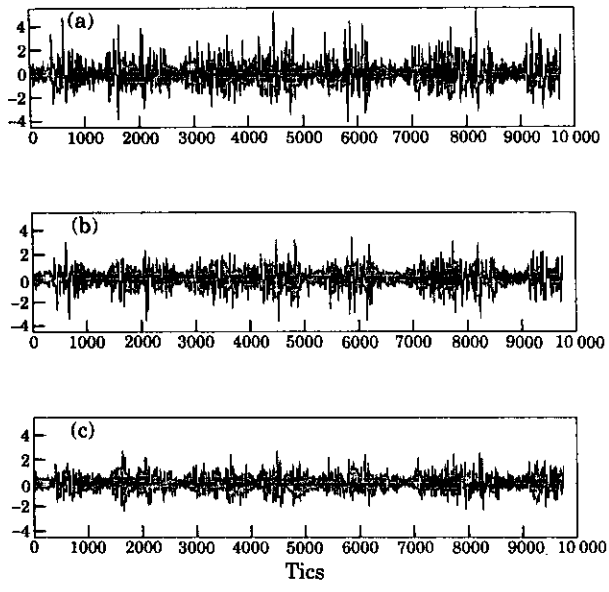

FIG. 18. (a) Third residual y_{R^3}; (b) filter y_{R^3F}; (c) residual y_{R^4}.

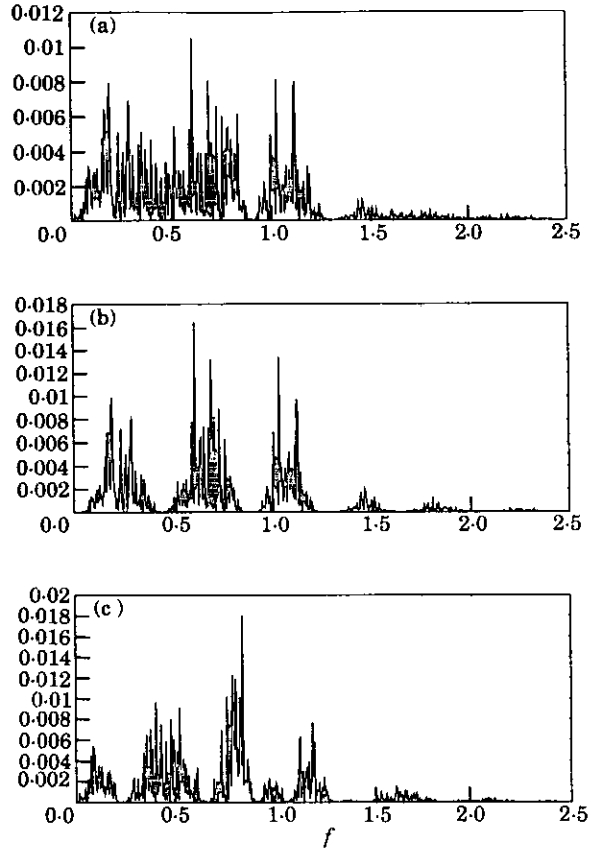

FIG. 19. Spectra of Fig. 18. (a), (b) and (c) as for Fig. 18.

methodology, but the filter appears to be robust. The signal can be decomposed into "independent" series

$$y = y_{FF} + y_{FR} + y_{RF} + y_{R^2F} + y_{R^3F} + y_{R^4}, \qquad (6)$$

in which y_{FF} (Fig. 10) is a doubly periodic oscillator with two disparate frequencies $f \approx 0{\cdot}03$ and $f \approx 0{\cdot}2$; y_{RF} (Fig. 12) is a chaotic signal with broad-band noise in the frequency range $f < 0{\cdot}6$, whose spiky nature we associate with transient "events" owing to spontaneous external disturbances; y_{R^2F} (Fig. 14) is the respiratory signal at a frequency of $f \approx 0{\cdot}88$; y_{R^3F} (Fig. 18) is the secondary oscillator, a chaotic oscillator with a frequency of $f \approx 0{\cdot}25$. Finally y_{R^4} (Fig. 18) is the residual signal, though it still contains some oscillatory components. In practical terms, we would expect that the filtering "tree" is not indefinitely pursued. For example, y_{R^2F} adequately represents "the" breathing signal (and in fact a further filter of it leaves the spectrum almost unchanged).

The filtering technique presented here thus gives a useful way in which complicated respiratory and other signals can be analysed into various distinct oscillatory, chaotic and transient signals.

REFERENCES

BROOMHEAD, D. S. & KING, G. P. (1986). Extracting qualitative dynamics from experimental data. *Physica* **D20,** 217–236.

FLEMING, P. J., GONCALVES A. L., LEVINE, M. R. & WOOLLARD, S. (1984). The development of stability of respiration in human infants: changes in ventilatory responses to spontaneous sighs. *J. Physiol.* **347,** 1–16.

JOHNSON, P. & ANDREWS, D. C. (1990). Hypoxia, temperature control and periodic breathing in postnatal life. In: *Hypoxia: The Adaptations* (Sutton, J. R., Coates, G. & Remmers J. E., eds), pp. 84–87. New York: Dekker.

KELLY, D. H. & SHANNON, D. (1979). Periodic breathing in infants with near-miss sudden infant death syndrome. *Pediatrics* **63,** 355–359.

KEMBER, G. & FOWLER, A. C. (1993). A correlation function for choosing time delays in phase portrait reconstructions. *Phys. Lett.* **A179,** 72–80.

KEMBER, G. & FOWLER, A. C. (1994). A nonlinear filtering technique for multi-oscillator systems. *Comput. math. Applic.*, in press.

SMITH, L. A. (1992). Identification and prediction of low dimensional dynamics. *Physica* **D58,** 50–76.

TAKENS, F. (1981). Detecting strange attractors in fluid turbulence. In: *Dynamical Systems and Turbulence* (Rand, D. & Young, L.-S., eds), pp. 366–381. Berlin: Springer-Verlag.