Pergamon

Computers Math. Applic. Vol. 29, No. 4, pp. 55–67, 1995

0898-1221/95 $9.50 + 0.00

0898-1221(94)00238-X

A Nonlinear Filtering Technique for Multi-Oscillator Systems

A. C. FOWLER AND G. KEMBER
Oxford Centre for Industrial and Applied Mathematics
Mathematical Institute, Oxford University
24-29 St Giles', Oxford, OX1 3LB, England

(Received May 1994; accepted June 1994)

Abstract—The recent application of methods from nonlinear systems theory to the analysis of time series can be used to provide a novel technique for filtering such series, when they represent the coupled output of several distinct oscillators. The method is a nonlinear analogue of Fourier spectral analysis, but is potentially more powerful, in that it allows for local (in time) cycle to cycle variability in oscillator amplitude.

Many natural systems can be thought of as oscillators, and are often subject to the onset of secondary or tertiary oscillations through instability. Many systems combine the effects of different oscillators, and it is well-known that even in the simplest such systems, the resultant signal can be chaotic. A simple example is the forced Duffing equation [1]. The question then arises, given an 'oscillatory' time series, can we distinguish in it the presence of several different oscillators which contribute to the signal?

The classical technique for doing this (or attempting to) is the Fourier power spectrum, a famous use of which was in the analysis of paleoclimatic data [2] to understand ice age cycles. While such techniques are useful where the signal is strictly periodic or quasi-periodic [3], they become less appropriate where the signal is slightly irregular, since the Fourier decomposition is unable to follow trends in a segment of the series; moreover, the spectrum which is obtained depends on the length of the segment which is used. The occurrence of particular 'events' at irregular intervals is another cause of bad results.

To be specific, suppose $X(t)$ has Fourier spectrum $\hat{X}(f)$ given by

$$X(t) = \int_{-\infty}^{\infty} \hat{X}(f) e^{-2\pi i f t}\, df. \tag{1}$$

A band pass filter of width 2ε about a selected frequency f_0 is then

$$X(t) = \int_{f_0-\varepsilon}^{f_0+\varepsilon} \hat{X}(f) e^{-2\pi i f t}\, df, \tag{2}$$

and this can be written in the form

$$X(t) = A(\varepsilon t) e^{-2\pi i f_0 t}, \tag{3}$$

where

$$A(\tau) = \int_{-1}^{1} \tilde{X}(g) e^{-2\pi i g \tau}\, dg,$$

We acknowledge discussions and assistance from P. Johnson, P. Fleming, S. Walter and M. Clements.

$$\tilde{X}(g) = \frac{\hat{X}(f_0 + \varepsilon g)}{\varepsilon}. \tag{4}$$

A narrow band-pass filter thus isolates a particular frequency, at the expense of having a slowly varying amplitude.

In this note, we show how the use of singular systems analysis [4,5] allows a method of filtering to be used which leads to an improved characterisation of different oscillators in a signal. We are particularly interested in cardiorespiratory patterns of young infants, where an understanding of the interaction of cardiac and respiratory signals is required so that (valid) relationships to conditions such as cot death (SIDS) can be properly investigated [6,7]. In practical terms, this often means "extracting" cardiac and respiratory patterns from a phasic signal received from a single sensor. An example is a pulse time series representing an indirect measure of the pulse wave form from the plethysmograph of a pulse oximeter shown in Figure 1. It can be seen that the basic oscillation is subject to a secondary modulation of about four times the period, and our basic task is to find a projection of the series which will separate (filter) the two oscillators. Since one anticipates that they will be (nonlinearly) coupled, it is not obvious that a separation on the basis of Fourier frequencies will be appropriate, although for such a clean spectrum as that for Figure 1, it is relatively easy to do (and gives good results): this is shown in the second two plots in Figure 1, where the actual signal (top) is subjected to a high pass filter and low pass filter with frequencies $f > 1.56\,\mathrm{Hz}$ and $f < 1.56\,\mathrm{Hz}$, respectively. However, spectral filtering techniques would be less useful for more complicated time series.

Given a time series $X(t)$, usually sampled at discrete times, we construct an *embedded phase space* for the trajectory by choosing a time lag Δ and an embedding dimension d_E, and then identifying a point in phase space $\mathbb{R}^{d_E}$ on the trajectory at the current time t by the vector

$$x(t) = (X(t - m_-\Delta), X(t - (m_- - 1)\Delta), \ldots, X(t), \ldots X(t + (m_+ - 1)\Delta), X(t + m_+\Delta)). \tag{5}$$

The embedding dimension is $d_E = m_- + m_+ + 1$, and usually we choose $m_- = m_+$ if d_E is odd, otherwise $|m_+ - m_-| = 1$. Other choices are possible, e.g., $m_- = 0$ or $m_+ = 0$ are both common. Thus, $x(t)$ traces out a trajectory in this artificial phase space, and if the time series arises from a finite d_A-dimensional system, or at least resides on a d_A-dimensional attractor, then a theorem of Takens [8] guarantees that the structure of the attractor will be properly represented, providing $d_E \geq 2d_A + 1$ (as the embedded trajectories will then not intersect).

The uses to which this embedded trajectory can be put are various: prediction, identification of low dimensional dynamics [9], noise reduction [4] and event recognition, with applications such as the determination of sensor failure [10]. The basic tool of dynamical systems methods is *singular value decomposition* (SVD) which is a method for identifying the shape and (crudely) the density of a cluster of points in $\mathbb{R}^{d_E}$. Its use is described in [4].

The continuous limit of the matrix methodology when $m_- = m_+$ has been described by Vautard and Ghil [5], and is useful for purposes of description. Given a continuous time series $X(t)$, its *autocorrelation function* is

$$C_X(\tau) = \lim_{T\to\infty} \frac{1}{2T} \int_{-T}^{T} X(t)X(t - \tau)\, dt, \tag{6}$$

and we define a linear, positive definite operator on $L^2(-\tau, \tau)$, with the usual inner product

$$(f, g) = (2\tau)^{-1} \int_{-\tau}^{\tau} f\bar{g}\, dt,$$

by

$$A\rho = \frac{1}{2\tau} \int_{-\tau}^{\tau} C_X(t - s)\rho(s)\, ds. \tag{7}$$

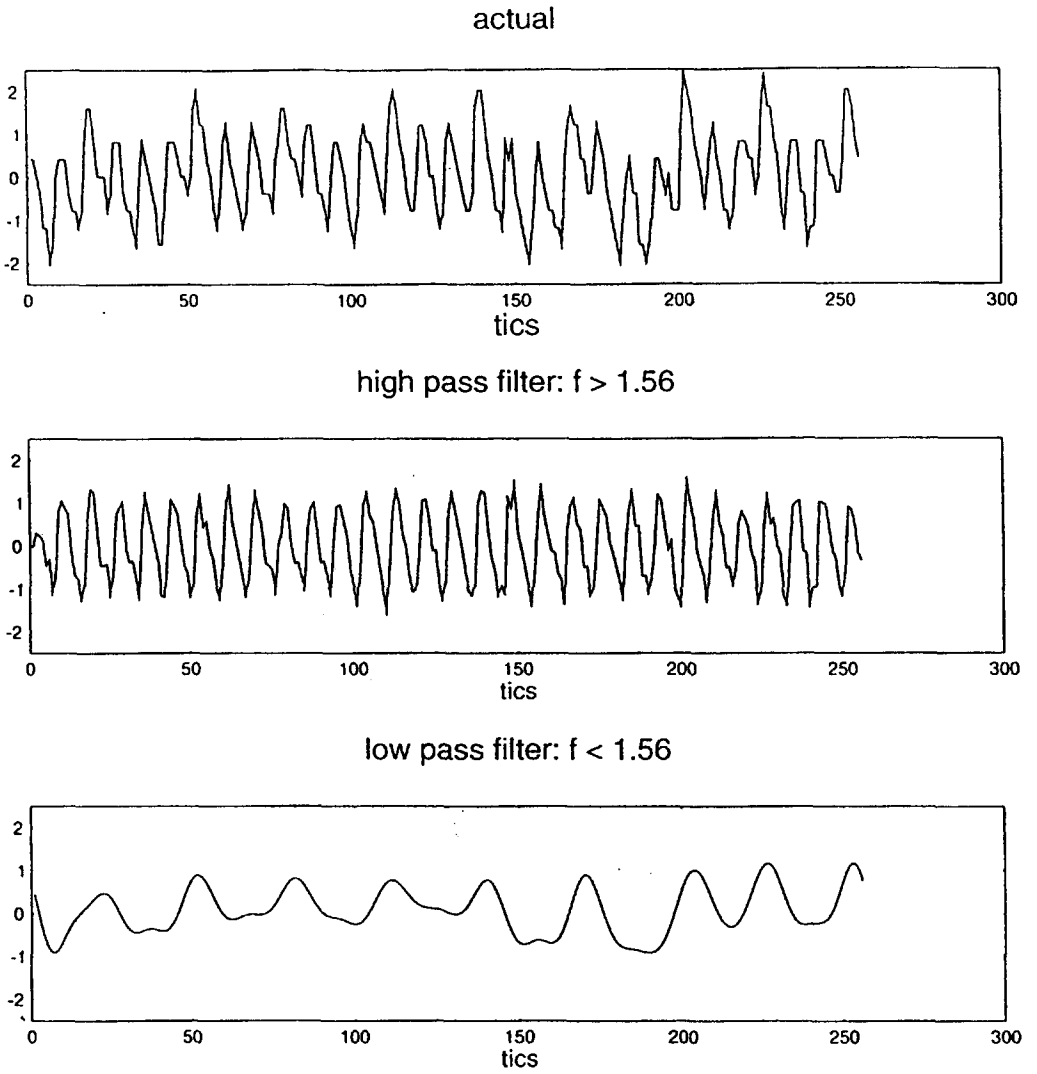

Figure 1. Segment of a plethysmograph time series representing pulse measurement of a young child (top). Filtered series for $f > 1.56$ Hz, $f < 1.56$ Hz. The high pass filter represents the pulse, while the low pass filter represents the respiratory signal.

In SVD terminology, $\tau = \tau_w/2$, τ_w is the window width $(d_E - 1)\Delta$, and $X(t-s)$, $s \in (-\tau, \tau)$ is the continuous analogue, for each t, of the embedded trajectory; thus, $X(t-s)$ corresponds to the trajectory matrix T, and the operator A corresponds to $T^\top T$. The discrete form of this description is given in [11]. Hence the eigenvalues λ of (7), which, by Hilbert-Schmidt theory for positive, symmetric kernels C_X, are a positive decreasing sequence $\lambda_1 \geq \lambda_2 \geq \cdots \geq 0$, correspond to the squares of the singular values σ_i^2 of the trajectory matrix. The (orthonormal) eigenfunctions $\rho(s)$, $s \in (-\tau, \tau)$ are called empirical orthogonal functions (EOFs), and correspond to the singular vectors of the trajectory matrix. In particular, it is important to realise that a vector in the phase space also represents, by its coordinates, a time sequence over the window $(t-\tau, t+\tau)$. One can thus think of a trajectory in the embedding space geometrically as a trajectory in phase space, *or* as a time sequence of functions which are slices of the trajectory over the window $(t-\tau, t+\tau)$.

According to Hilbert-Schmidt theory, we can expand $X(t-s)$ as

$$X(t-s) = \sum_k X_k(t)\rho_k(s), \qquad s \in (-\tau, \tau), \tag{8}$$

and this defines a sequence of principal components $X_k(t)$, given, by inverting (8), as

$$X_k(t) = \frac{1}{2\tau}\int_{-\tau}^{\tau} X(t-s)\bar{\rho}_k(s)\,ds. \tag{9}$$

Equation (9) represents projection of X on to ρ_k, and a filter is then a projection to a set $\{k\}$ of EOFs, so that, from (8),

$$X_F(t) = \sum_{\{k\}} X_k(t)\,\rho_k(0), \tag{10}$$

providing (8) converges at $s = 0$. In the discrete case, this is always so, and is also true here for finite τ. In our calculations below, we carry out these procedures in the discrete version described by SVD. Notice the crucial fact that the filtered signal $X_F(t)$ is determined essentially by values of X in $(t-\tau, t+\tau)$ and is thus *local.*

CHOICE OF TIME LAG

In using a phase space embedding for the purposes described above, there is an arbitrariness available in the choice of the embedding dimension d_E and the lag time Δ. For systems which are (simply) recurrent, that is, the embedded trajectory lies close to a closed loop, the choice of lag time should be such as to maximise the spread, and we have recently suggested [12] that for pseudo-oscillatory systems, this can be done by minimising a quantity called the *singular value fraction* (SVF). We define the SVF as follows. Let $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_{d_E} \geq 0$ be the singular values of the trajectory matrix, normalised so that $\sum_1^{d_E} \sigma_i^2 = 1$. Then the SVF is

$$f_{SV}(k;\Delta) = 1 - \left(\frac{d_E}{d_E - k}\right)\sum_{k+1}^{d_E}\sigma_i^2, \tag{11}$$

and has the following properties: $0 \leq f_{SV} \leq 1$, and if $f_{SV} = 0$, power is equidistributed among all the singular vectors ($\sigma_i^2 = 1/d_E$ for each i), while if $f_{SV} = 1$, then all the power resides in the first k singular vectors. Kember and Fowler [12] showed that for recurrent (i.e., oscillatory) signals, the first minimum Δ_m of f_{SV} plotted for fixed k (usually $k = 1$) versus the embedding lag Δ is the preferred lag, in order to maximise the spread of the signal in the embedding space. The first major maximum Δ_M is half the recurrence time, and we have $\Delta_M \approx d_E\Delta_m$. The window length τ_w is then given by $\tau_w = (d_E - 1)\Delta_m \approx (d_E - 1)P/2d_E$, where P is the recurrence time. At these low values of d_E, the singular vectors resemble Legendre polynomials [13], and we project onto a pair which between them represent the filtered oscillator, with a period $\sim 2\tau_w$.

ITERATIVE SINGULAR SYSTEMS ANALYSIS

The procedure of the method presented here is the following. First, we compute the SVF as a function of Δ. Usually, we select $k = 1$ (sometimes, higher values are used). The embedding dimension is usually chosen at a fairly low value. By observation of Δ_M and/or Δ_m, the SVF indicates a primary recurrence time. We now choose Δ appropriate to this frequency. This has the effect of maximising the spread of the trajectory in the plane of the singular vectors corresponding to this frequency. The second step is, therefore, to project (as per equation (10)) onto the singular vectors. Note that, since the EOFs (or, discretely, singular vectors) $\rho(s)$ depend on the time series (being eigenfunctions of A in (7)), this is in fact a nonlinear filter.

The Fourier transform of (9) is

$$\hat{X}_k(f) = \tilde{\rho}_k(f)\,\hat{X}(f), \tag{12}$$

where the reduced transform $\tilde{\rho}$ is

$$\tilde{\rho}_k(f) = \frac{1}{2\tau}\int_{-\tau}^{\tau} \bar{\rho}_k(t) e^{2\pi i f t}\, dt. \tag{13}$$

Thus, taking the Fourier transform of (10) and substituting from (12),

$$\hat{X}_F = \left[\sum_{\{k\}} \rho_k(0)\tilde{\rho}_k(f)\right] \hat{X}, \tag{14}$$

and the power spectra ($P_X(f) = 2|\hat{X}(f)|^2$) are related by

$$P_{X_F}(f) = |\Sigma_k(f)|^2 P_X(f), \tag{15}$$

where Σ_k denotes the expression in square brackets in (14).

Our third critical step is that we now iterate the filtering process. In general, we find that this iteration 'cleans out' other frequencies from the signal, leaving us with a recurrent signal with a broad band spectrum. Apparently, the map $\hat{X} \to \hat{X}_F$ converges. For a signal consisting of several oscillators, this filtering process can be repeated, *but not necessarily with the same embedding.*

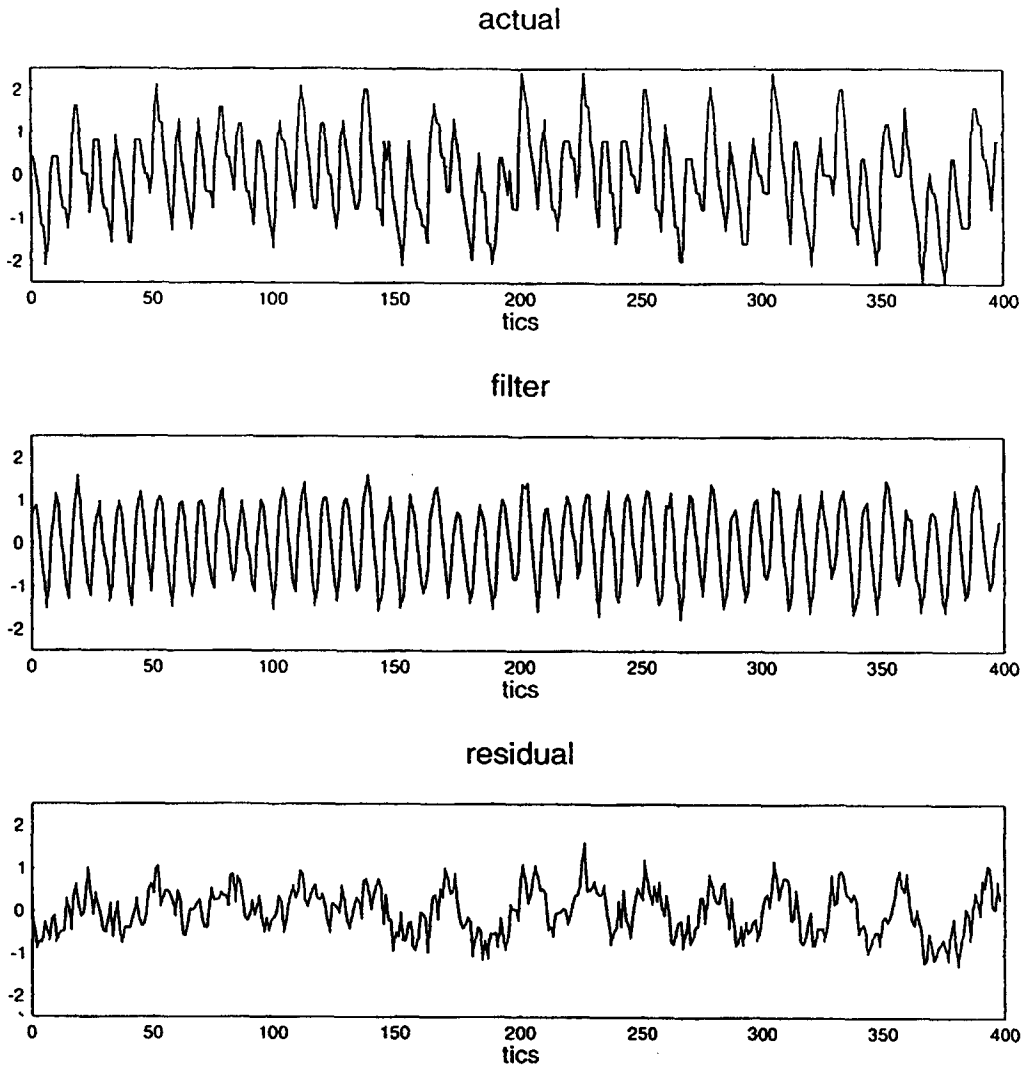

Figure 2. Signal y, filtered signal y_F, and residual signal y_R. y_F is the pulse signal, and is rather cleaner than the high pass filter in Figure 1.

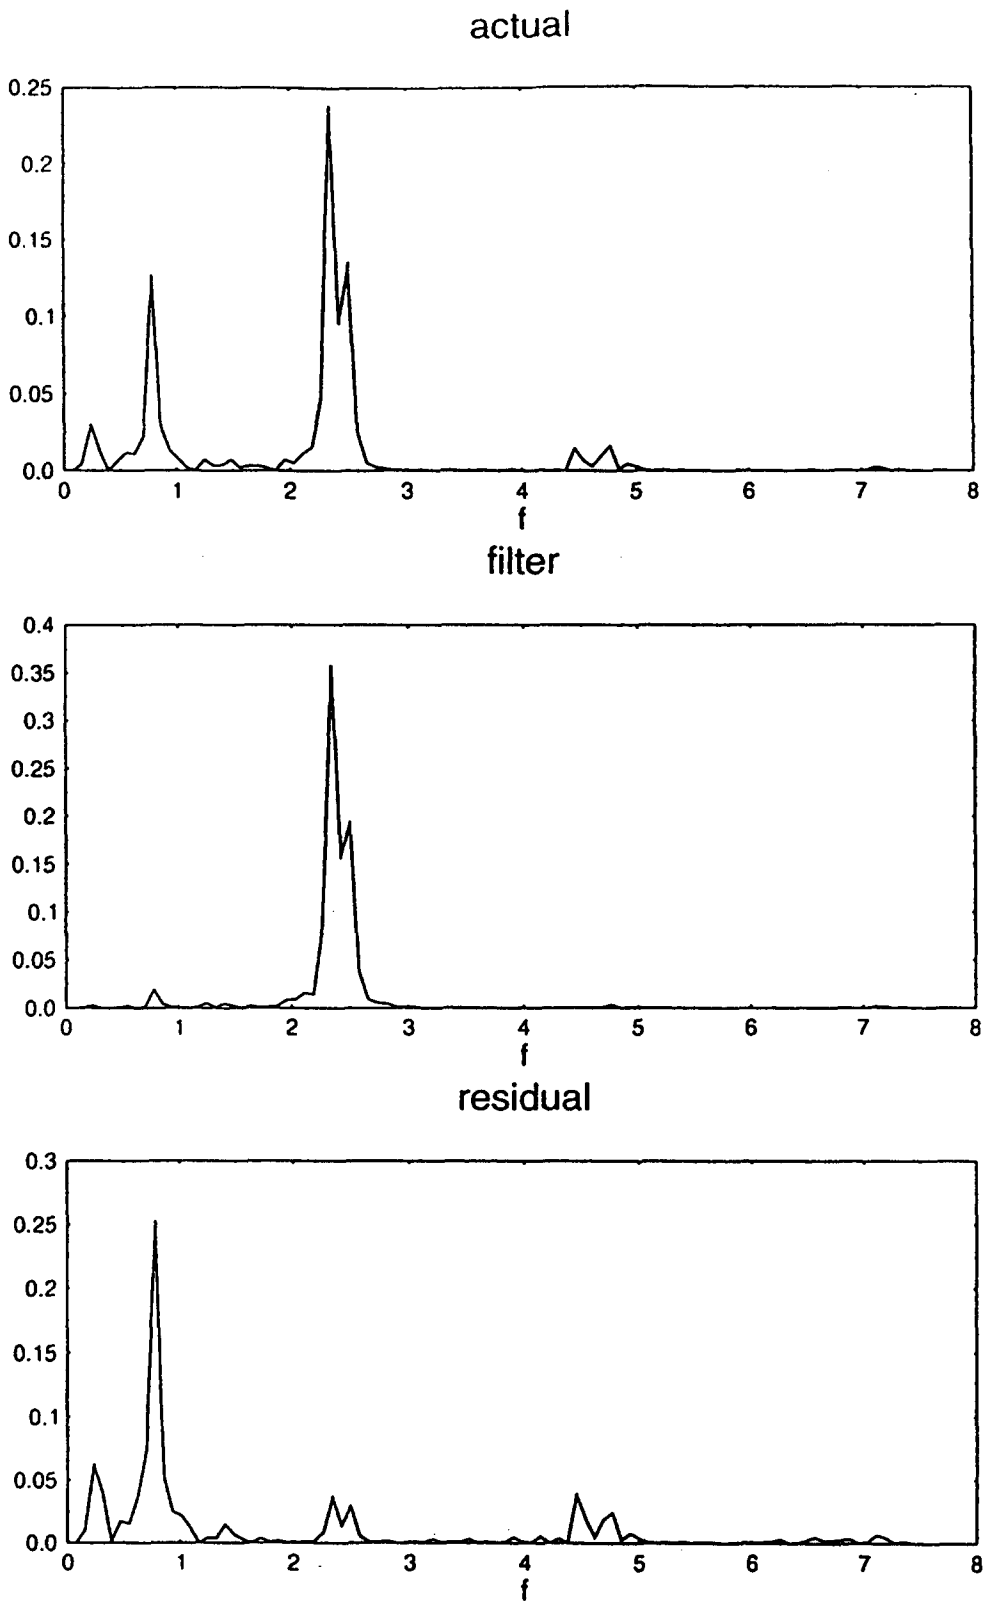

Figure 3. Power spectral density $P(f)$ for y, y_F, y_R in Figure 1 as a function of frequency f, computed with a square window (no data windowing) and normalised so that $\sum_0^{N/2} P(f_k) = \frac{1}{N}\sum_0^{N/2} |y_k|^2 \approx \frac{1}{T}\int_0^T |y(t)|^2\,dt$ for each series, where N is the number of data points (here 256), T is the corresponding time interval, $T \approx N\delta$, where δ is the sampling frequency = 0.05 seconds.

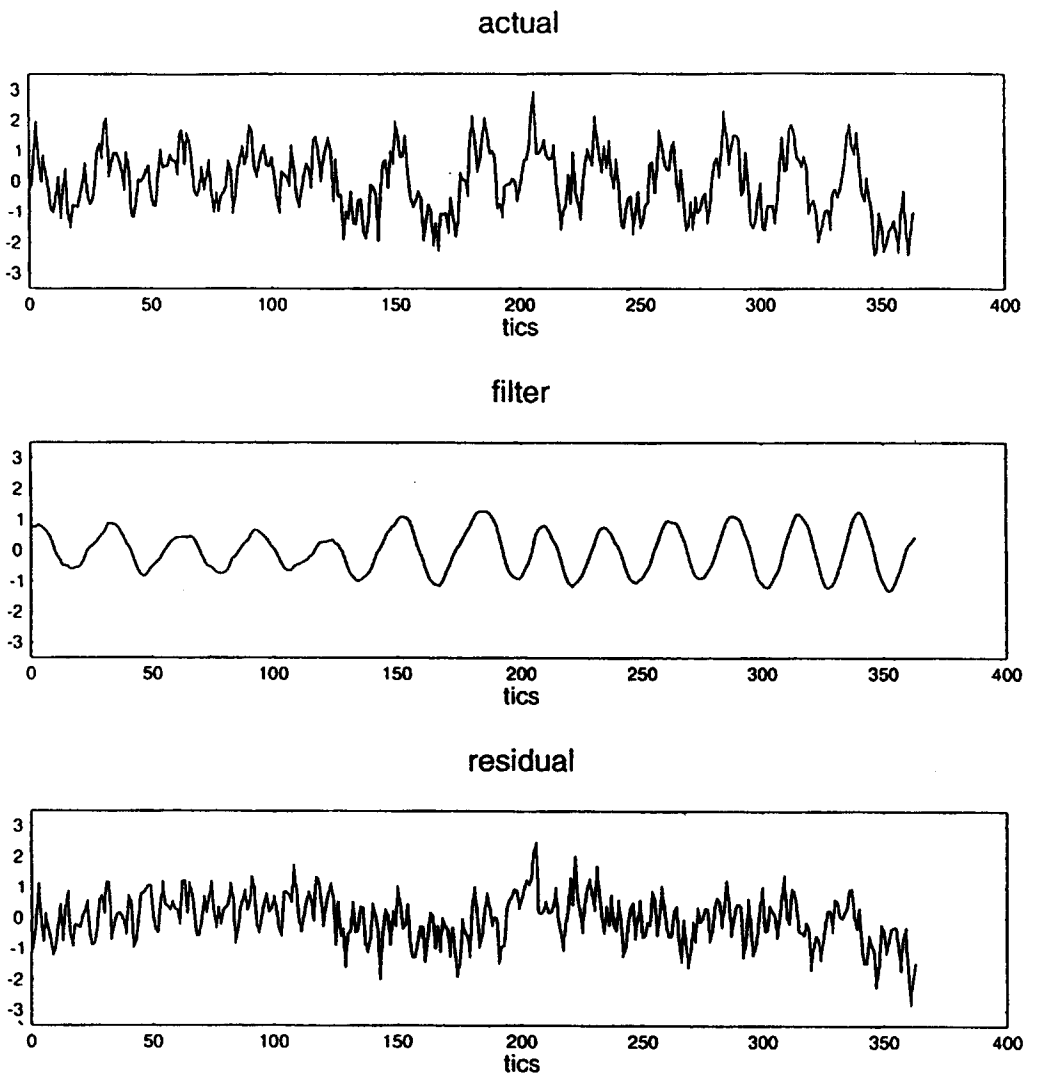

Figure 4. Actual residual oscillation y_R, filtered secondary oscillation y_{RF} and secondary residual y_{R^2} ($d_E = 20$, $\Delta = 1$).

APPLICATIONS

We illustrate our methodology by application to two time series from a pulse plethysmograph. Figure 2 shows the segment of pulse data $y(t)$ from the plethysmograph wave form of a pulse oximeter on an infant, together with the filtered signal y_F and the residual $y_R = y - y_F$. We have used other segments with similar results. Notice that only 400 data points (tics) are used. The time series is first normalised to have zero mean and unit variance. The power spectra in Figure 3 show two main oscillations of nearly pure frequencies, together with some low frequency noise. The higher frequency oscillation is the basic pulse pattern.

We choose $d_E = 5$ (which would be appropriate for a purely two-dimensional oscillator). A calculation of the SVF $f_{SV}(1)$ (see [12]) as a function of time lag Δ (measured in tics) shows a minimum at $\Delta = 2$, which we therefore choose as the embedding lag. With this choice, we anticipate that the primary oscillation is largely spanned by the first two singular vectors w_1 and w_2. With this in mind, the filtered signal y_F in Figure 2 corresponds to the projection of the embedded trajectory on to w_1 and w_2, while the residual series y_R is that obtained by projecting on to w_3, w_4, w_5. In spite of the crude resolution of the data set, we see that the $w_{1,2}$ projection picks up the primary pulse pattern, while a secondary oscillation is clearly visible in the residual.

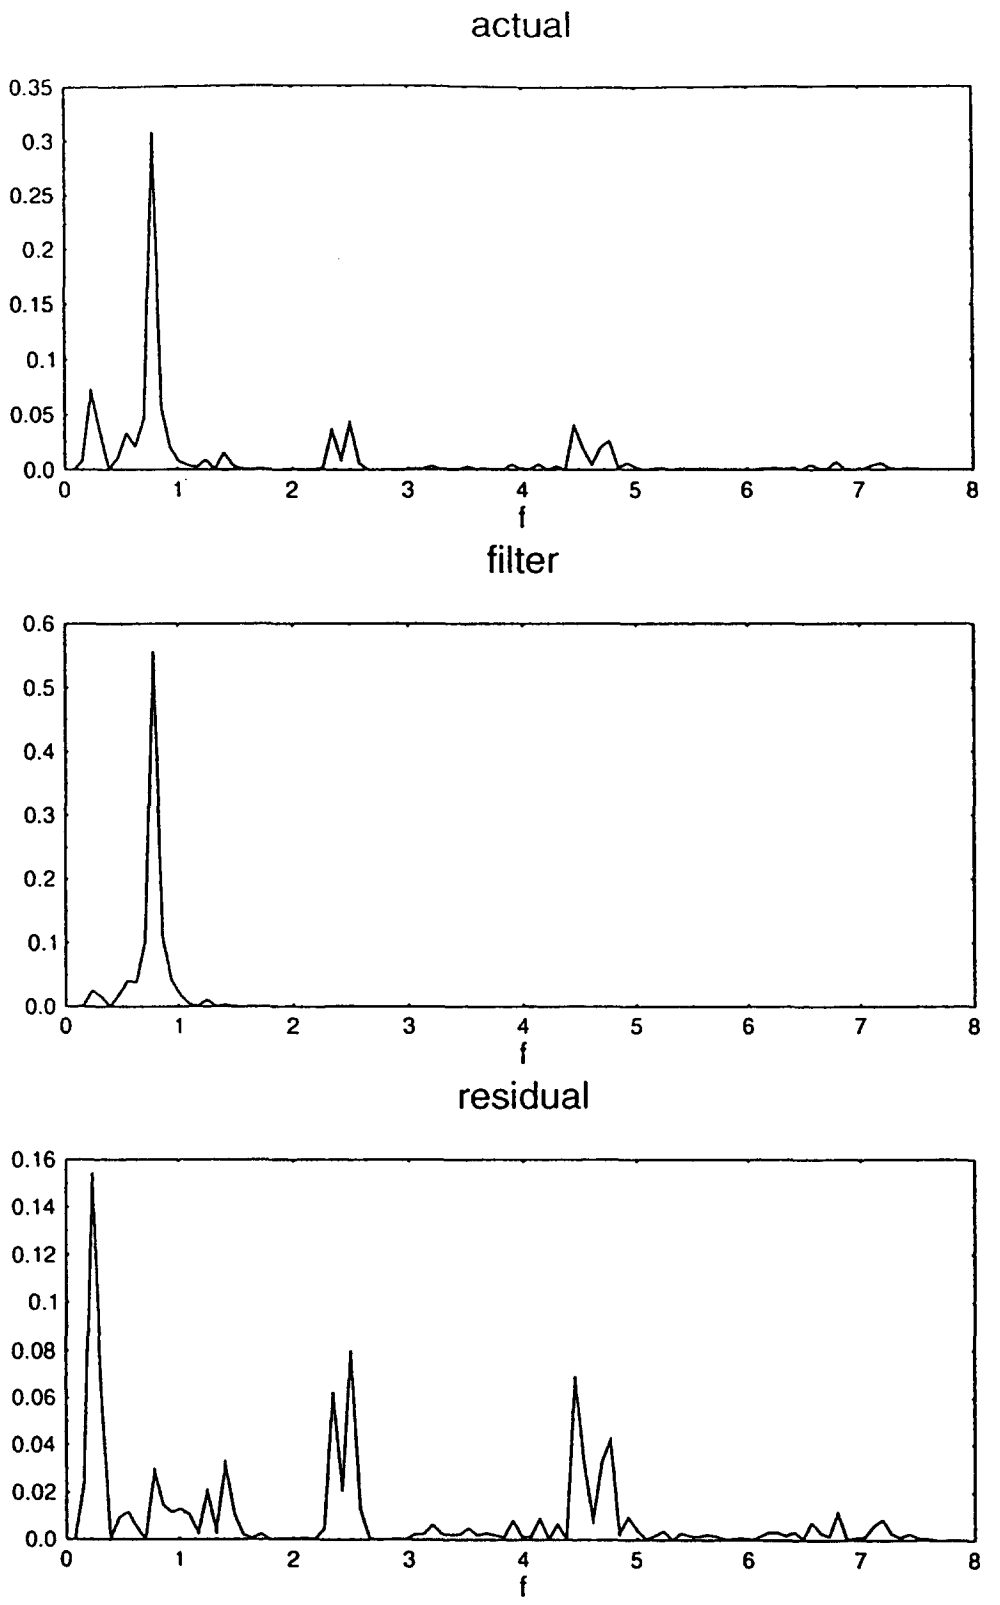

Figure 5. Power spectra of y_R, y_{RF} and y_{R^2} in Figure 4.

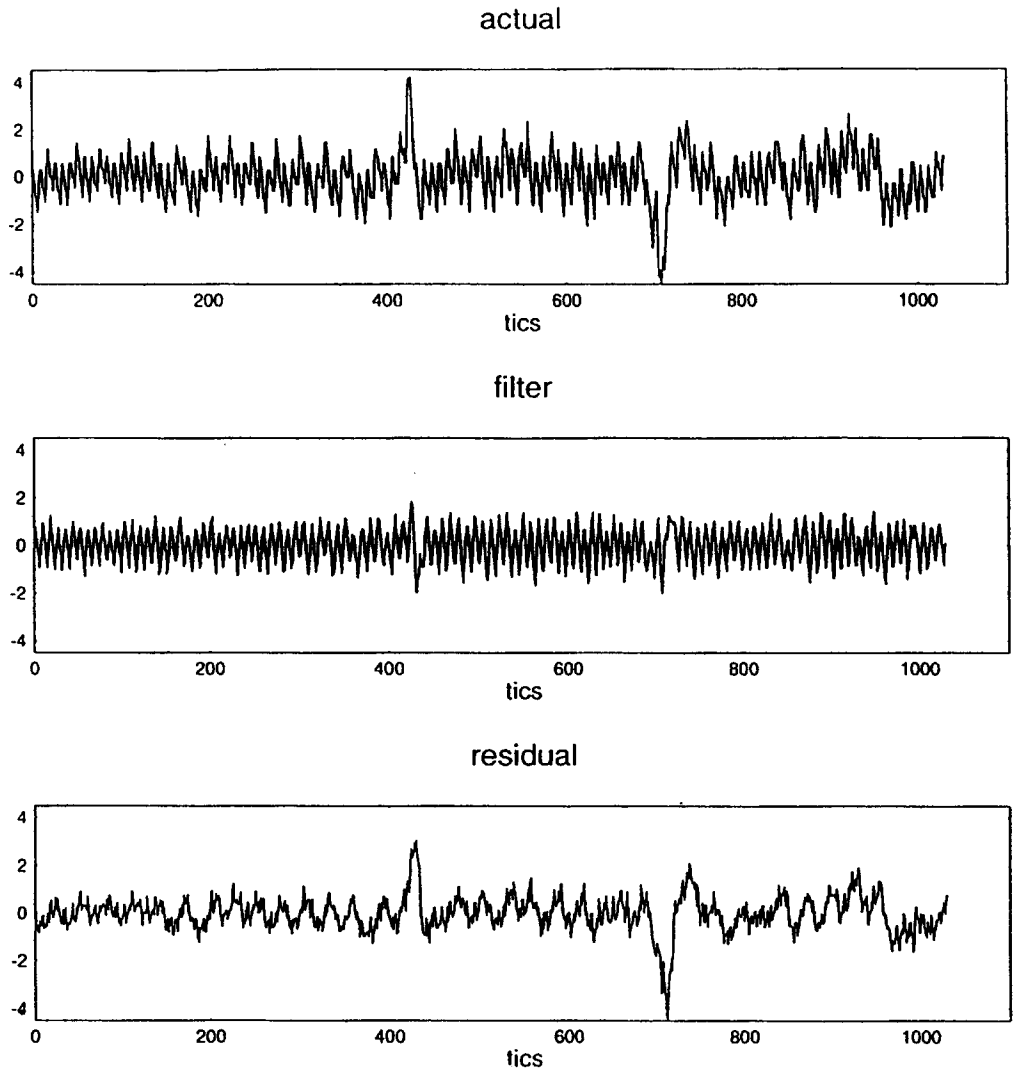

Figure 6. A second signal y from the plethysmograph record, incorporating two breathing movements, together with filtered pulse signal y_F and residual y_R.

Comparing the spectra in Figure 3, we see that that of y_F has picked up the basic frequency. The spectrum of y_R has power in the secondary frequency, but also contributions from the basic pattern frequency and harmonics, together with a significant low frequency component.

The key iterative step of the procedure is that we now apply the technique to the residual series, to extract its basic oscillatory characteristic. Now the main power of y_R is at a frequency of ~ 0.7, whereas there is significant power at frequencies of 2.3, 4.5 and indeed some at ~ 7. To pick up the high frequency component, we can only lower the time lag Δ to its minimum value, $\Delta = 1$. To cater for the low frequency component, the embedding window $\Delta(d_E - 1)$ must be of the order of half of the recurrence time. Eyeballing Figure 2 suggests this is ~ 17, whence $d_E \sim 18$. In Figure 4, we choose $d_E = 20$, $\Delta = 1$. The residual time series of Figure 2 is shown as the top series (note that the origin is shifted, since the first $d_E - 1$ points are not used in the embedding), together with the projection onto the newly computed singular vectors w_1 and w_2, and the residual of *this* projection. We see a very clear secondary oscillation associated with breathing, together with a noisy residual. Figure 5 shows the respective spectra.

Another example, using a different segment of the same data set, is shown in Figure 6. Note the two 'events' in the signal, associated with limb movements. With the same embedding as before (but projecting on to w_2 and w_3), we obtain the filtered signal y_F and the residual y_R

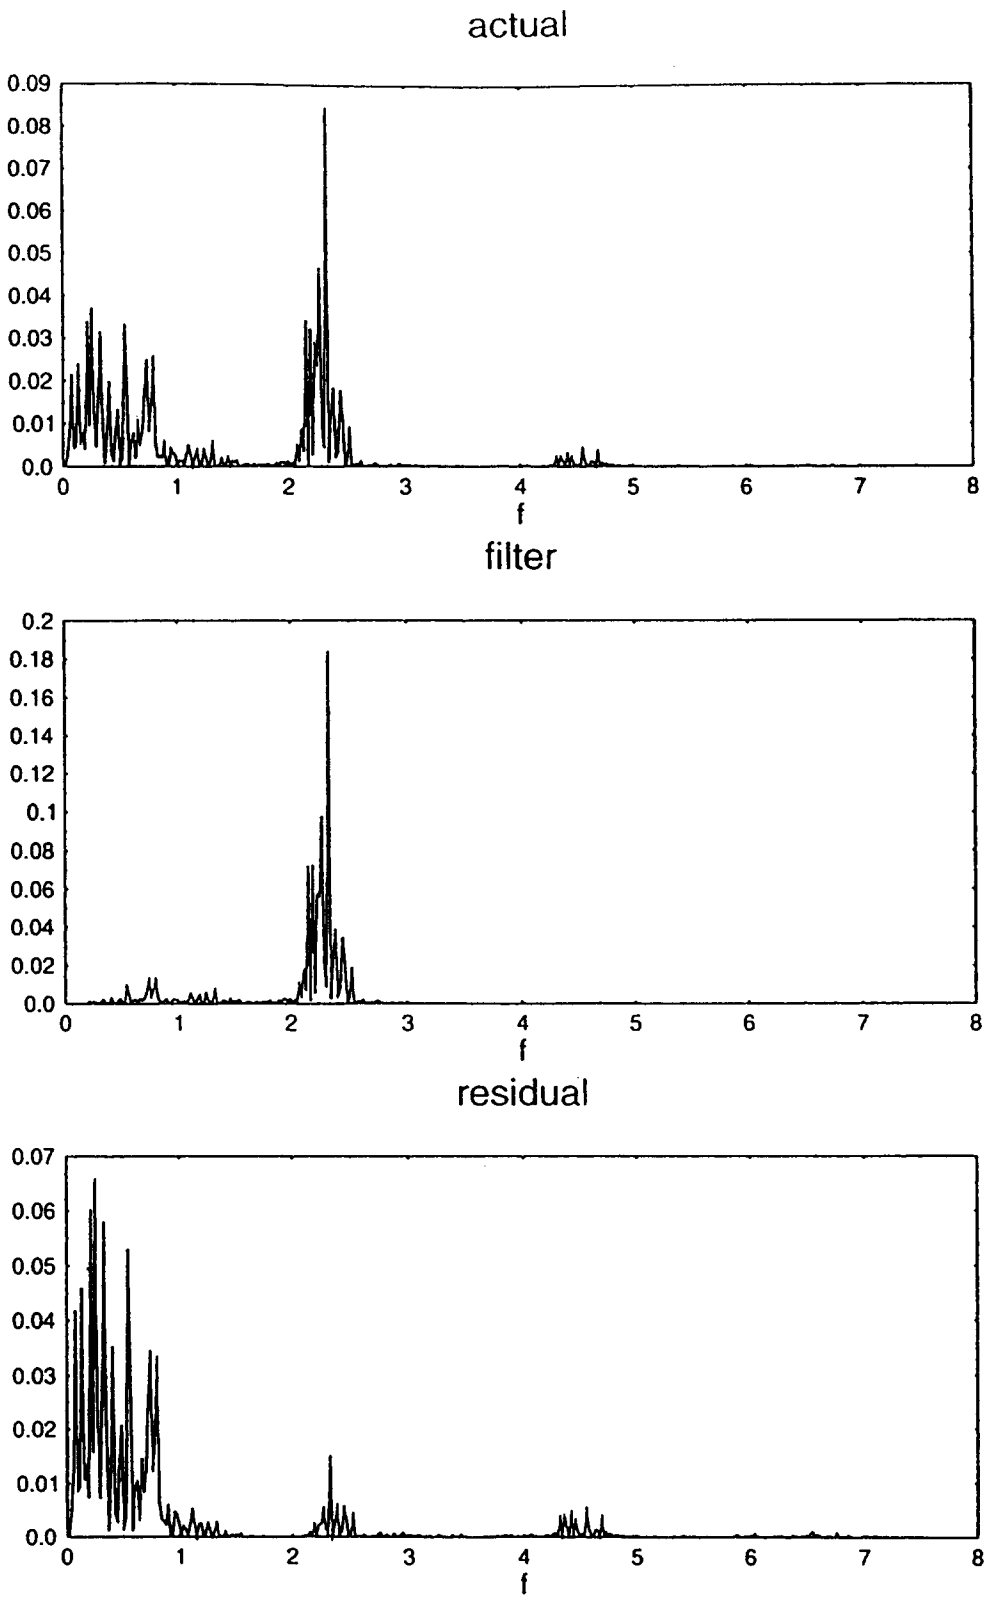

Figure 7. Power spectra of series in Figure 6.

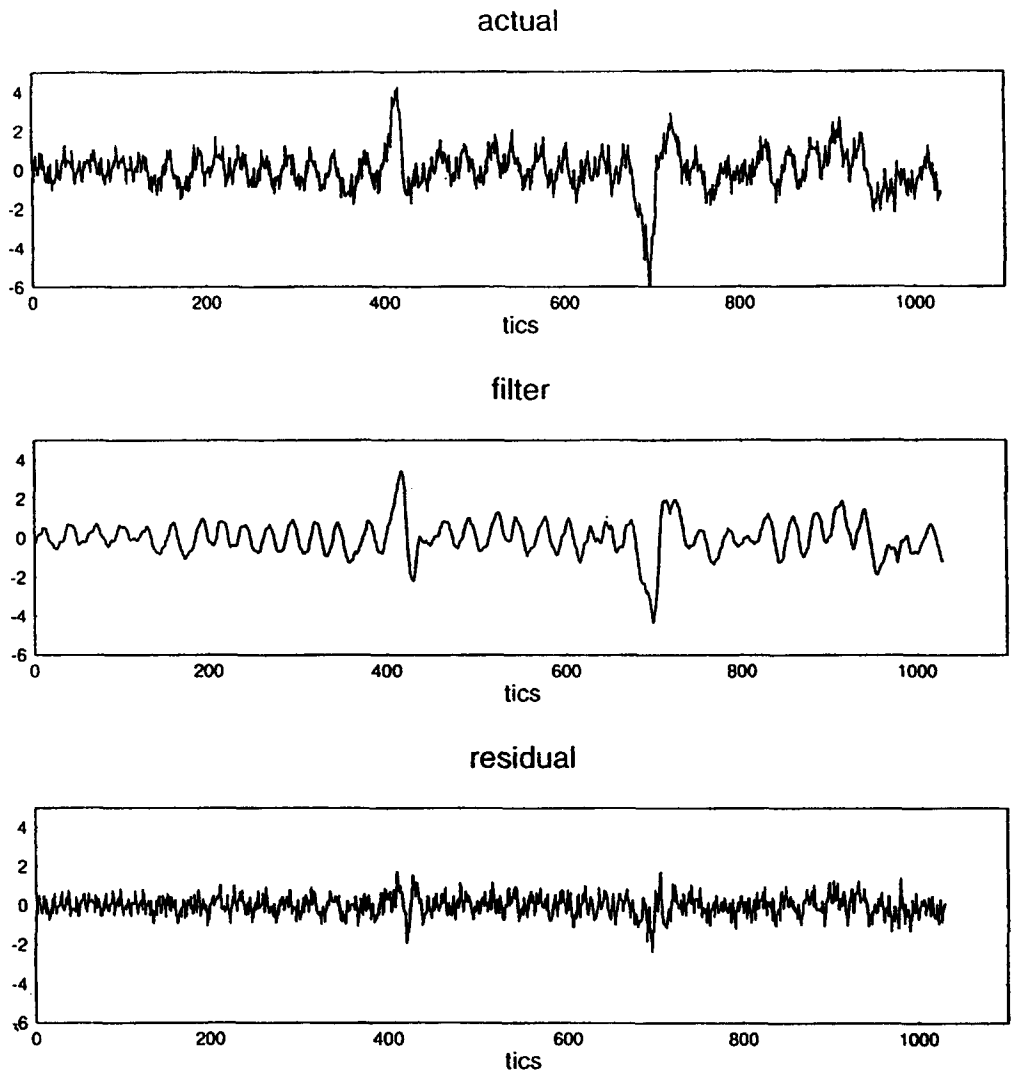

Figure 8. Residual signal y_R, filtered residual y_{RF}, and second residual y_{R^2}.

as shown. The respective power spectra are shown in Figure 7. Now we filter the residual to obtain y_{RF} and the second residual y_{RR} (or y_{R^2}), see Figure 8. The power spectra are shown in Figure 9. Notice also that the respiratory frequency is not easy to determine in Figure 9, although a spectral low pass filter at $f = 1.5$ is still able to separate the two signals adequately.

The potential advantages of the filter presented here are (i), its local character, which is able to resolve local cycle to cycle vacillations, and (ii), the iterative character of the filter, which (since the projection eigenfunctions ρ_k depend on the series $X(t)$) allows for a convergence of the filters which is unavailable in a linear method. Our method is similar to that of Vautard and coworkers [5,11], but differs in its recursive nature. In particular, the filter $\Phi(f) = P_{X_F}(f)/P_X(f)$ is akin to Wiener filtering [11], but, unlike that, repeated application of this nonlinear filter can be applied without destroying the signal.

In future work, we intend to address the thorny issue of filtering signal from signal, for example, to ask whether it is possible to separate two oscillators which interact nonlinearly. A simple example of the problems which arise can be seen in Figure 3. The power at $f \sim 4.6$ is removed in the filtered signal, although it is a second harmonic of the pulse frequency. Is this because the filter $\Phi(f)$ acts as a narrow band pass, or because the second harmonic is noise, or is it a forced response in the respiratory signal? We cannot answer that question as yet, but we know that in this example, at least, the filter gets it right. The plethysmograph signal in Figure 2 has clear plateaus, where $X(t_i) = X(t_{i-1})$; moreover, these are at identifiable discrete values of X.

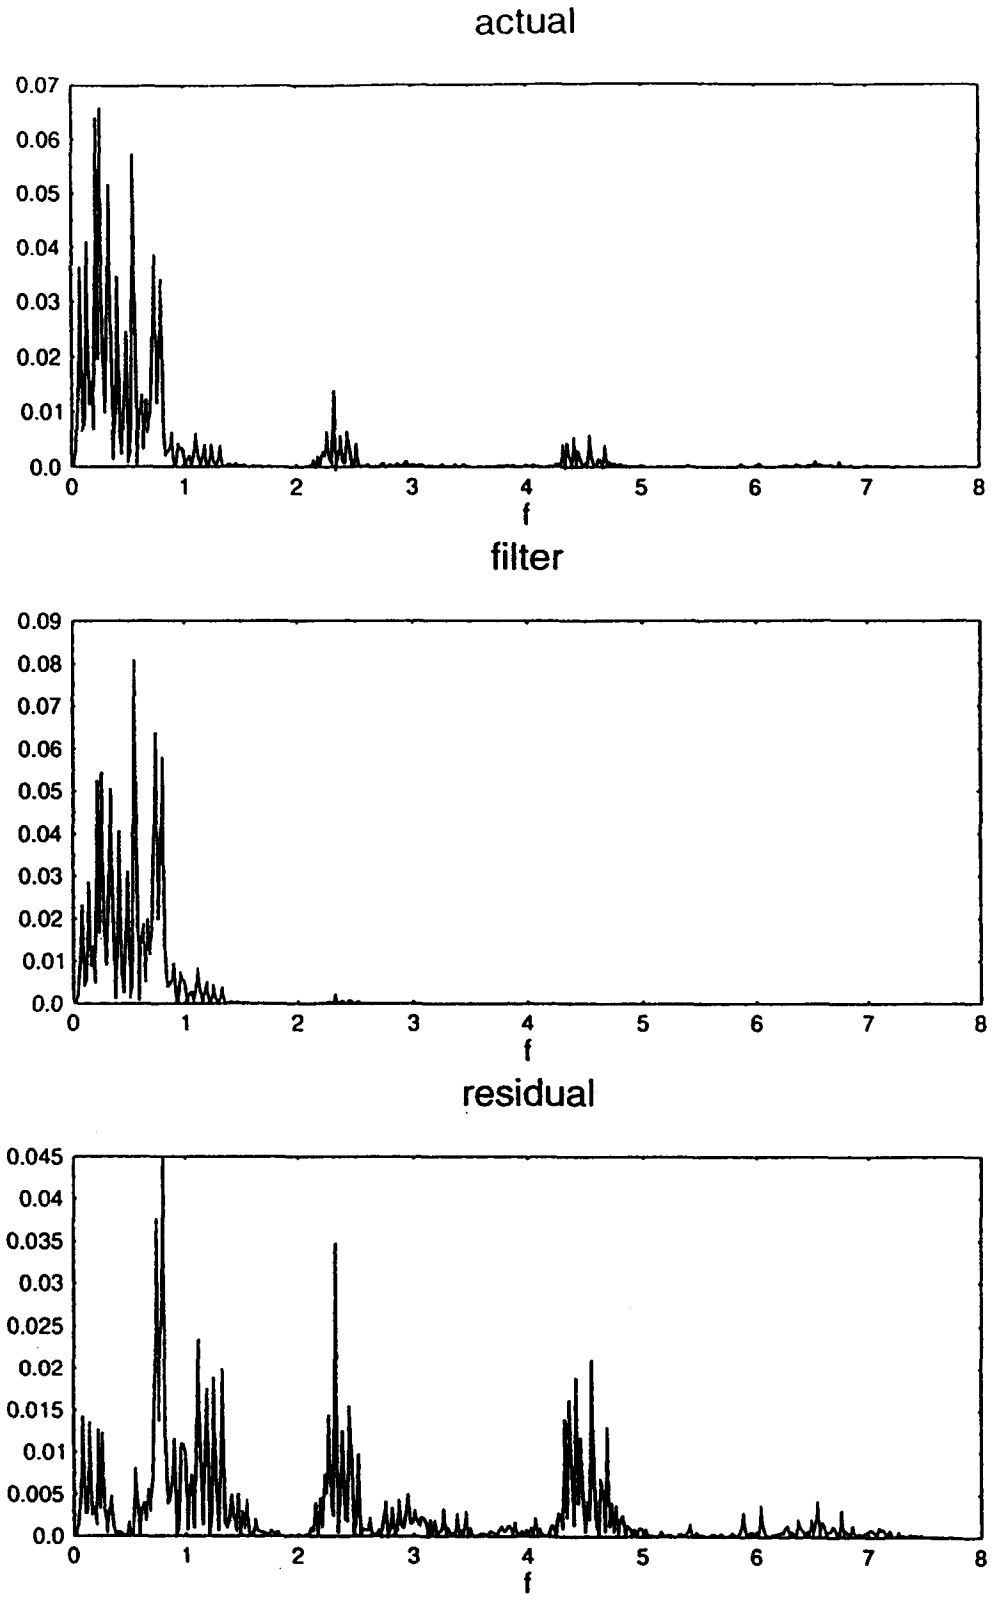

Figure 9. Power spectra of series in Figure 8.

It is then obvious that plateaus at fixed values X^* in a periodic signal will generate power in harmonics, but we can be sure that the plateaus arise as an artifact of the measuring process.

In summary, we have presented a technique for filtering separate oscillators in at least some types of multi-oscillator system, which has the beneficial effects of spectral analysis or linear filters, without the drawbacks to which these methods can be prone, when the underlying time series is chaotic and/or nonstationary.

REFERENCES

1. Th. Buzug, T. Reimers and G. Pfister, Optimal reconstruction of strange attractors from purely geometrical arguments, *Europhys. Letts.* **13**, 605–610 (1990).
2. J.D. Hays, J. Imbrie and N.J. Shackleton, Variations in the earth's orbit: Pacemaker of the ice ages, *Science* **194**, 1121–1132 (1976).
3. J.P. Gollub and S.V. Benson, Many routes to turbulent convection, *J. Fluid Mech.* **100**, 449–470 (1980).
4. D.S. Broomhead, and G.P. King, Extracting qualitative dynamics from experimental data, *Physica* **20 D**, 217–236 (1986).
5. R. Vautard and M. Ghil, Singular spectrum analysis in nonlinear dynamics, with applications to paleoclimatic time series, *Physica D* **35**, 395–424 (1989).
6. P.J. Fleming, A.L. Goncalves, M.R. Levine and S. Woollard, The development of stability of respiration in human infants: Changes in ventilatory responses to spontaneous sighs, *J. Physiol.* **347**, 1–16 (1984).
7. P. Johnson and D.C. Andrews, Hypoxia, temperature control and periodic breathing in postnatal life, In *Hypoxia: The Adaptations,* (Edited by J.R. Sutton, G. Coates and J.E. Remmers), pp. 84–87, Dekker, New York, (1990).
8. F. Takens, Detecting strange attractors in fluid turbulence, In *Dynamical Systems and Turbulence*, (Edited by D. Rand and L.-S. Young), pp. 366–381, Springer-Verlag, Berlin, (1981).
9. L.A. Smith, Identification and prediction of low dimensional dynamics, *Physica* **58 D**, 50–76 (1992).
10. L.A. Smith, K. Godfrey, P. Fox and K. Warwick, A new technique for fault detection in multi-sensor probes, In *Control 91, IEE Conf. Pub.*, 322, Vol. 2, pp. 1062–1067, (1991).
11. R. Vautard, P. Yiou and M. Ghil, Singular-spectrum analysis: A toolkit for short, noisy chaotic signals, *Physica D* **58**, 95–126 (1992).
12. G. Kember and A.C. Fowler, A correlation function for choosing time delays in phase portrait reconstructions, *Phys. Letts., A* **179**, 72–80 (1993).
13. J.F. Gibson, J.D. Farmer, M. Casdagli and S. Eubank, An analytic approach to practical state space reconstruction, *Physica D* **57**, 1–30 (1992).