J. Nonlinear Sci. Vol. 7: pp. 411–425 (1997)

Journal of
Nonlinear
Science

Local Nonlinear Filtering

G. Kember,[1] A. C. Fowler,[2] and H. B. Evans[3]

[1] Department of Applied Mathematics, Technical University of Nova Scotia, Halifax, NS, B3J 2X4, Canada

[2] Mathematical Institute, Oxford University, 24-29 St Giles', Oxford OX1 3LB, England

[3] Department of Physics and Astronomy, University of Wales College of Cardiff, 5 The Parade, Cardiff CF2 3YB, Wales

Received on September 1, 1994; revised manuscript accepted for publication on December 13, 1996

Communicated by Jaroslav Stark

Summary. Classical methods of filtering time series use Fourier power spectral analysis to separate signals from noise. Improved methods of signal separation can be developed by using projection techniques based on concepts of nonlinear dynamics. However, such methods are limited in their ability to distinguish between dynamically independent signals. Here we show how it is possible to combine Fourier projection with local nonlinear prediction to provide a methodology which can, in principle, recognise independent dynamical signals. We apply the methodology to a variety of chaotic signals with superimposed sine waves, and show how the sine wave frequency can be recognised dynamically (but not spectrally).

AMS/MOS classification numbers. 34A34, 34C35, 93E11.

1. Introduction

Classical noise suppression in signal processing is concerned with the removal of low amplitude noise from a time series. There is a whole literature on different methods of filtering, but a typical technique is to use the Fourier power spectrum. The noise spectrum may be assumed (for example) to be white, and its amplitude estimated from the spectral power at high frequencies. A low-frequency cut-off filter might then be applied to produce a cleaned version of the signal.

It is typical of such filters that they take no account of the dynamical nature of the signal, and they are thus inappropriate for filtering signals which are intrinsically chaotic, or for separating signals which are dynamically interactive. A typical instance occurs in a number of geophysical and medical applications where the signal may consist of a number of different oscillators. For example, certain oscillatory frequencies in the earth's

gravity field are thought to be due to the motions of the earth's core (Dehant et al. 1993), and it is of interest to separate them from the remainder of the signal. Respiratory signals typically consist of a number of different oscillatory signals such as the respiratory rhythm, the pulse rate, the vagal oscillator, and so on, and it is of interest to delineate the presence of particular oscillators in an apparently chaotic signal (Fleming et al. 1988).

Another problem in clinical monitoring is that of determining the baroreceptor reflex response. Ideally, one wishes to determine, given two time series of blood pressure and heart rate, each of them rather irregular, to what extent the first series forces the second. Current methods (Robbe et al. 1987) effectively use spectral methods although the dynamical coupling is almost certainly nonlinear.

Recently, the IPCC 130 year record of the earth's global mean temperature has been examined to identify oscillatory trends in the data (Ghil and Vautard 1991), using the methods of singular systems analysis (Broomhead and King 1986), (Vautard and Ghil 1989), but the interpretation of the results is controversial (Allen and Smith 1994).

In this paper we take as our motivating goal the identification of dynamically distinct signals. We first describe singular systems analysis and show how it fails to succeed in certain cases. Then we describe a method which allies spectral analysis to local nonlinear dynamics by incorporating the use of a nearest neighbour method predictor to distinguish between different oscillators, and we illustrate its applicability (although it is not in principle so limited) to the simple case of a sine wave superimposed on a chaotic deterministic time series. Our method is similar in concept to ideas developed by Kostelich and Yorke (1990), further elaborated in Kostelich and Schreiber (1993), and Broomhead (1994).

2. Singular Systems Analysis

Distinct, uncoupled linear oscillators can be identified via the power spectrum. Nonlinear oscillators produce harmonic peaks, and if they are coupled, will also produce peaks at resonant frequency mixes. However, coupled nonlinear oscillators will often produce a chaotic signal, with a broad-band spectral signature, and in this case it is difficult to interpret the spectrum (Fenstermacher et al. 1979). In particular, chaotic oscillators have cycle-to-cycle amplitude variation, and any band-pass filter tends to obscure this.

A more sophisticated idea is to analyse local windows of width $\tau_w = 2\tau$ centred at time t. That is, given a time series $X(t)$, we consider for each t the function $X(t+s)$ defined for $s \in [-\tau, \tau]$. A moving window spectral analysis then computes the Fourier spectrum for each window. If the spectral power at a frequency f is $P_{X(t)}(f)$, then the time series $P_{X(t)}(f)$ for fixed f gives a measure of the amplitude of the oscillation at this frequency as a function of time. For example, this is done by Dehant et al. (1993) analysing gravimeter data.

In practice, for discretely sampled data and finite window widths, one uses a discrete Fourier transform, and one may question whether this is the best spectral projection to use. The method of singular systems analysis (Vautard and Ghil 1989), (Vautard et al. 1992) is a mildly nonlinear generalisation of discrete Fourier analysis which uses the time series itself to determine the optimising sequence of orthogonal functions (for continuous $X(t)$; vectors for discrete $X(t)$) on which to project $X(t+s)$. For discrete

data, the algebraic method used is singular value decomposition (SVD). In the limit of large τ, SVD recovers Fourier analysis, but for small τ, the orthogonal basis functions tend to Legendre polynomials (Gibson et al. 1992), (Fowler and Kember 1996).

In the discrete formulation, Takens's theorem implies that, if $X(t)$ is determined by a finite-dimensional deterministic system, or at least lies on a finite-dimensional attractor, then $X(t+s)$ is diffeomorphic to this attractor, and this associates some properties of the dynamics of $X(t)$ with the trajectory kernel $X(t+s)$. In particular, this implies that projection onto the basis determined by singular systems analysis (SSA) can be used to filter dynamical signals, for example white noise from deterministic chaos (Broomhead and King 1986).

However, SSA is limited in its usefulness. It can certainly filter small-amplitude white noise from a deterministic chaotic signal. Depending on the system, however, a simple low-pass filter could do the same thing. It is not particularly successful at filtering higher amplitude noise, nor is it good at filtering anything, even red noise, which has some dynamic content. For example, Fowler and Kember (1996) show that filtering using SSA is in fact equivalent to spectral filtering, even if the precise nature of the filter is chosen by the time series. Moreover, the nature of the filter is very dependent on the choice of window width τ_w and on the embedding dimension d_E (Broomhead and King 1986).

The inability of SSA to distinguish dynamical signals is because dynamics are manifested at a local level. In Takens's concept, the trajectory kernel $X(t+s)$ (discretely, the trajectory matrix) represents a trajectory embedded in a phase space, and dynamics is manifested through the fact that nearby points *in the phase space* behave in a similar way: In fact, this is the basis of nonlinear prediction methods (Mack and Rosenblatt 1979), (Yakowitz and Karlsson 1987), (Farmer and Sidorovich 1987), (Casdagli 1989), (Sugihara and May 1990). SSA uses only the *global* geometric properties of the embedded trajectory, and since these global properties may be shared by dynamically different time series, SSA—like Fourier analysis—is unable to distinguish dynamics. We illustrate this point with an example. The forced Duffing equation,

$$\ddot{x} + \delta\dot{x} - x + x^3 = \gamma \sin ft, \tag{2.1}$$

is one of the 'standard' simple chaotic oscillators on which tests are often carried out. We take $\delta = 0.4$, $\gamma = 0.4$, $f = 1$, and follow methods we have described previously (Fowler and Kember 1995) to implement SSA. In Figure 1 we show the output, the filter which is the projection of the output onto the first two singular vectors (the embedding dimension $d_E = 10$), the residual, and the corresponding spectra. It seems interesting and perhaps encouraging that the projection filters out the forcing frequency, since this seems to indicate an ability to separate the forcing from the internal dynamics of the Duffing oscillator. However, now we repeat the procedure, but first adding $0.5\sin(0.75t)$ to the signal. In Figure 2 we again show the time series, the filter, and the residual together with the corresponding spectra. If we compare Figure 1 with Figure 2 we see that the filter is virtually the same and the added frequency is lumped with the forcing frequency. There is no particular recognition of the dynamically independent nature of the added sine wave, and there is nothing special in the spectrum to mark it out. The global SSA-based method filters depend very much on the choice of the projection dimension d_p, the embedding

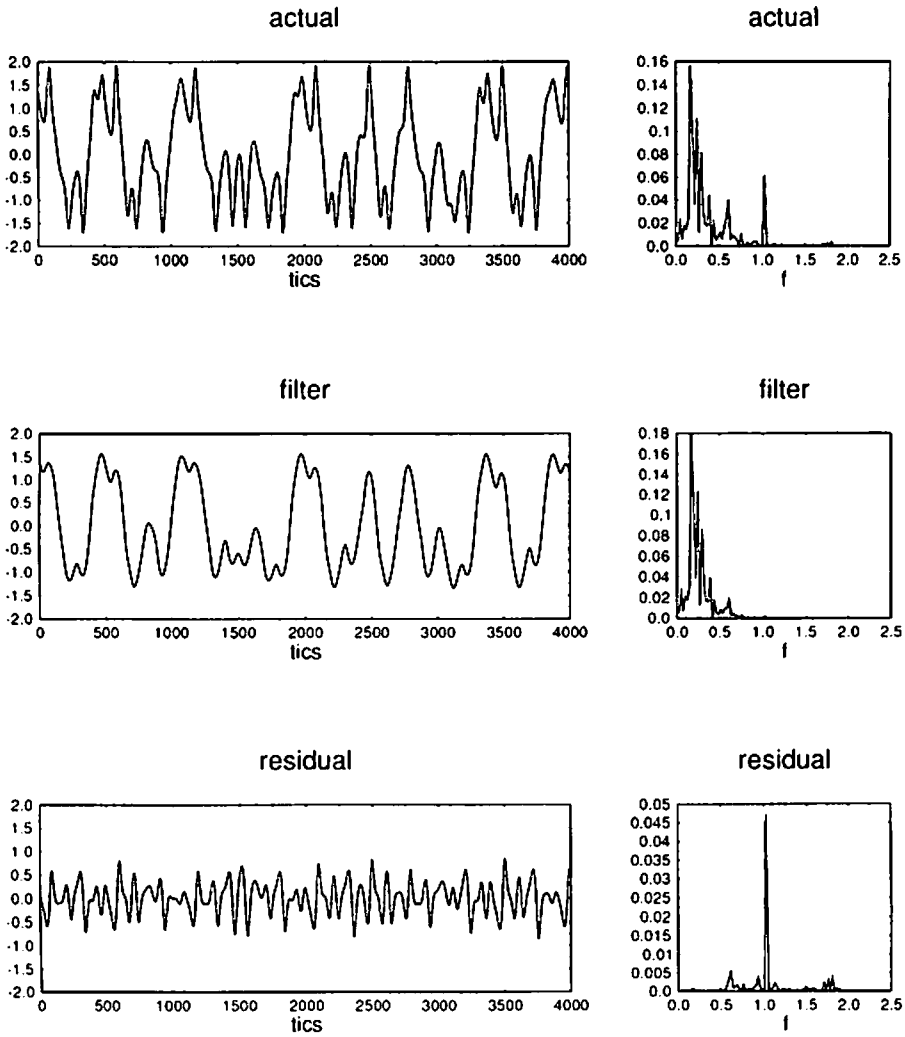

Fig. 1. Time series, filter, residual, and corresponding power spectra for the forced Duffing oscillator (2.1) using SSA (Fowler and Kember 1995). Embedding dimension $d_E = 10$, projection space dimension $d_p = 2$.

dimension d_E, and lag time Δ without regard to the signal dynamics. In particular, if one does not know what the series is, one *may* arrive at erroneous conclusions.

3. Nonlinear Prediction and Filtering

Kostelich and Yorke (1990) use the concept of phase-space embedding to filter a dynamical time series. They apply their method to an experimental time series and are able to remove small amplitude noise. The idea is to use local linear approximations to the

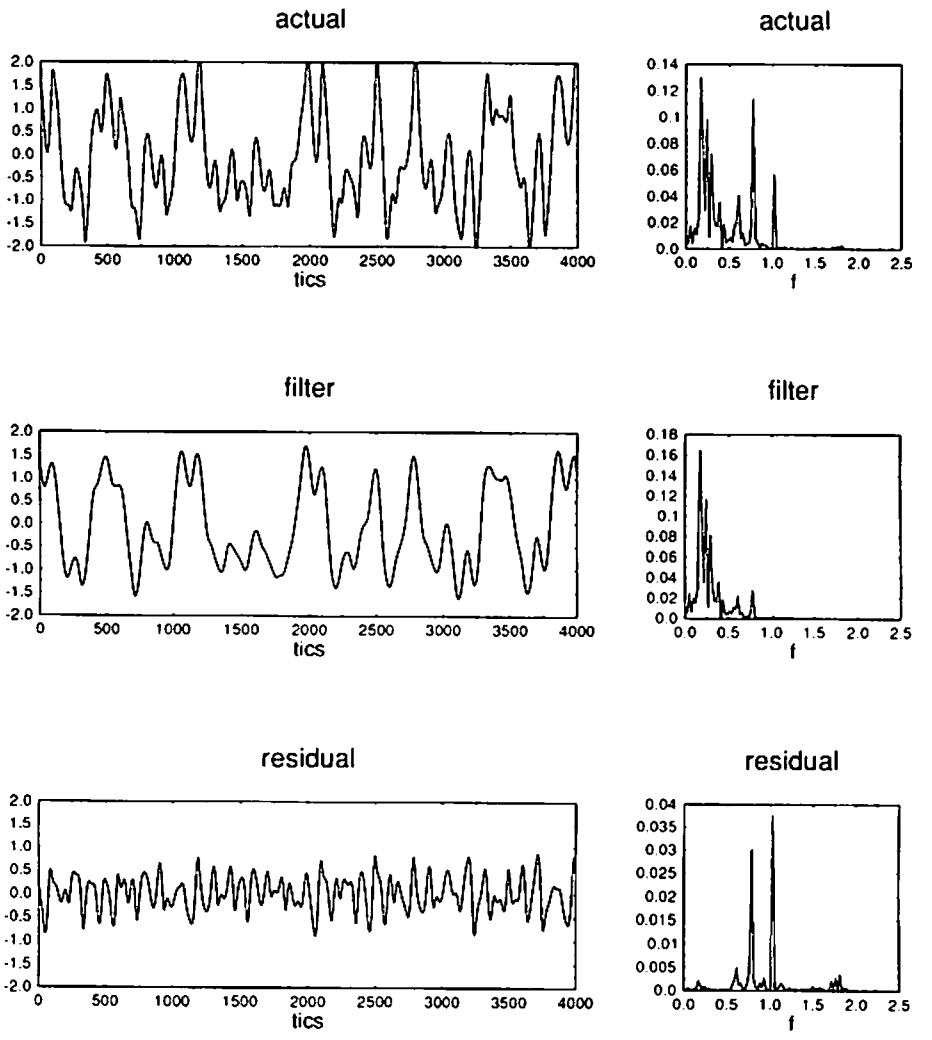

Fig. 2. Time series, filter, residual, and corresponding power spectra for the forced Duffing oscillator, but with $0.5\sin(0.75t)$ added. The SVD based filter provides essentially the same results as for Figure 1.

dynamics on the trajectory to improve its dynamic consistency through small trajectory adjustments. In Broomhead's (1994) method if one wants to separate a pure sine wave from a chaotic background, one first removes the sine wave by a bandstop linear filter and then recovers the chaotic background from its filtered version by using a reconstruction of the dynamics (Broomhead uses radial basis functions). The difficulty is that one needs to know the frequency of the offending cycle. Our aim is to dynamically distinguish the offending cycle, and we do this by minimizing local linear prediction error where the power spectrum is used to control the minimizing sequence.

Dynamics are manifested at a local level through the similar behaviour of nearby points, and this provides the basis for the algorithm we use, *local linear prediction*: the L-step-ahead forecast of the future evolution of a predictee point is taken to be the average evolution of nearby points in phase space L-steps into the future. The evolution of clouds of points in phase space for a linear signal is fundamentally different from that of a nonlinear signal. For a linear signal, nearby points are more or less mapped to other clouds of points of similar dimension. The same is not true for nonlinear signals where clouds of points may be mapped to other clouds with profoundly changed dimensions. The distortion of such a mapping caused by the addition of a dynamically independent signal as simple as a sine wave can cause a large relative increase in the prediction error *regardless of the embedding definition*. We demonstrate this for the forced Duffing equation described above, and compute the percent relative change in the mean absolute prediction error (MAPE) for the time series generated by (2.1), and for the same series with the added signal $0.5\sin(0.75t)$. In Figure 3 is depicted the percent relative increase to the MAPE for various embedding definitions and numbers of nearest neighbours. It is clear that the addition of the simple sine wave has increased the percent relative error in all cases on the order of at least one hundred percent. Most importantly the number of neighbours used in the local linear predictor has little effect, and increasing the embedding definition significantly increases the relative change in the MAPE for windows exceeding the optimal window. Although more sophisticated local predictors based on a richer class of maps could certainly reduce the effect seen in Figure 3, this would not serve our purpose since our objective is not to form better predictions but to use the change in predictive ability to distinguish dynamic signals. We comment further on this in the discussion. We now detail a method for the detection of dynamics which couples local linear prediction to classical spectral analysis.

4. A Diagnostic for Dynamic Detection

For a signal $X(t)$ which is the sum of a nonlinear signal $X_1(t)$ and a sine wave $X_2(t)$, distinguishing $X_1(t)$ and $X_2(t)$ corresponds here to minimizing prediction error of the filtered nonlinear signal $\hat{X}_1(t)$. The Fourier power spectrum can be used to control the optimizing sequence to minimize prediction error in the following way. The prediction error is computed for various bandstop filters of $X(t)$, i.e., $\hat{X}_1(t)$, constructed by zeroing small bands of selected frequencies in $X(t)$. If the amplitudes deleted are mostly associated with the nonlinear signal $X_1(t)$, then the prediction error of the band-passed signal $\hat{X}_1(t)$ will increase since the deleted band is part of the dynamic of $X_1(t)$. However, if the deleted band corresponds mostly to the added sine wave then the prediction error of $\hat{X}_1(t)$ will decrease since the sine wave is not part of the dynamic of $X_1(t)$.

The mean absolute prediction error (MAPE) $\varepsilon(f)$ is not only a function of the band deleted (f is the band-centre frequency) but also of the power P in that band. A difficulty arises if the power in the deleted band is relatively small. In this case the prediction error will be relatively small and may obscure the minimum of ε associated with the deletion of the offending cycle but whose power may be much larger. This difficulty is circumvented by defining a diagnostic $\hat{\varepsilon}$ as the product of the prediction error ε and power P in the

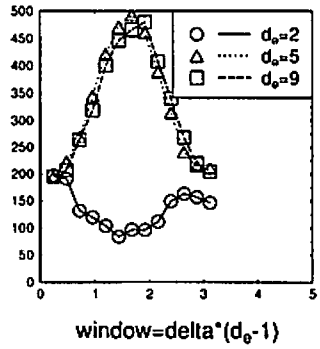

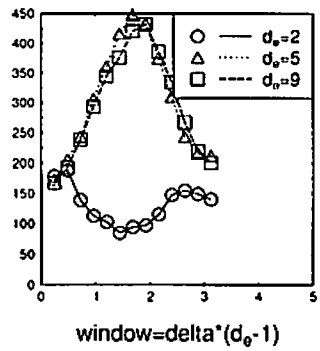

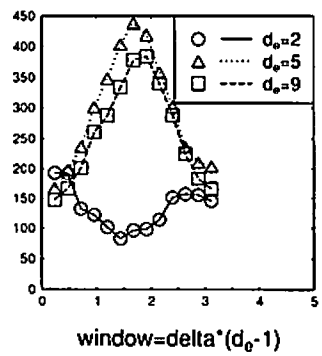

Fig. 3. Percent relative change in the mean absolute prediction error (MAPE) between the forced Duffing equation and the same with $0.5\sin(0.75t)$ added. The top figure is for two neighbours, the center figure for five neighbours, and the bottom for $d_E + 1$ neighbours. The horizontal axis is the embedding window where delta is the lag time in the Takens's embedding definition. The MAPE is much larger for the contaminated signal.

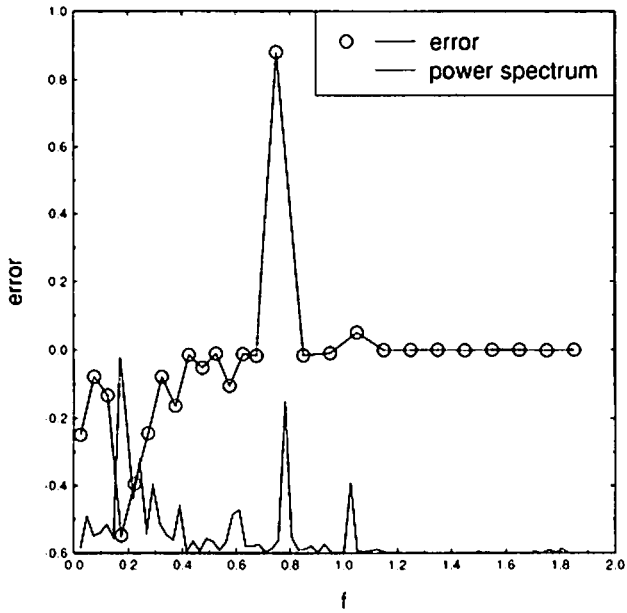

Fig. 4. Normalised error $\hat{\varepsilon}$ as a function of frequency f with superimposed power spectrum, for the contaminated Duffing series. The vertical axis is for the normalised error. The peak at $f = 0.75$ is interpreted as representing a spurious (nondynamic) signal.

deleted band, i.e.,

$$\hat{\varepsilon}(f) = \left(\frac{\varepsilon(f) - \varepsilon_\infty}{\varepsilon_{\min} - \varepsilon_\infty}\right)\left(\frac{P(f) - P_{\min}}{P_{\max} - P_{\min}}\right), \tag{4.1}$$

where $\varepsilon_{\min} = \min_f \varepsilon(f)$, and we have normalised the prediction error and the power so that $-\infty \leq \hat{\varepsilon} \leq 1$, and the prediction error ε_∞ is the error obtained when no bands are removed ($f \to \infty$). Importantly, we assume $\varepsilon_\infty \geq \varepsilon_{\min}$ (i.e., the full data set is less predictable than with the optimal band removed); otherwise our method fails. Ideally, if $\hat{\varepsilon} < 0$ (i.e., $\varepsilon > \varepsilon_\infty$) then the corresponding spectral component is *dynamic* since deletion of the band at f increases the prediction error; if $\hat{\varepsilon} > 0$, the component is interpreted as nondynamic. While there is no unique recipe for a 'good' diagnostic, we have found that $\hat{\varepsilon}$ works well and is robust to changes in embedding definitions. We now demonstrate the use of $\hat{\varepsilon}$ on a number of signals.

5. Results

We first illustrate the use of the diagnostic $\hat{\varepsilon}$ for the forced Duffing equation with an added sine wave $0.5\sin(0.75t)$ as in Figure 1 and Figure 2. In Figure 4, $\hat{\varepsilon}(f)$ is plotted with the power spectrum. It is not obvious that the added frequency at $f = 0.75$ is dynamically independent of the remainder of the power spectrum, but $\hat{\varepsilon}(f)$ reveals an

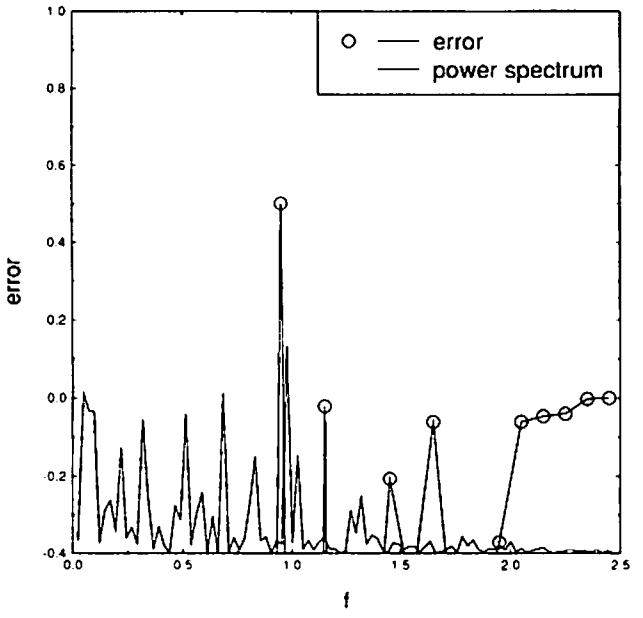

Fig. 5. As in Figure 4, but for the contaminated Lorenz series. The added frequency is at $f = 0.95$ with amplitude 2. The peak at $f = 0.95$ identifies the spurious (nondynamic) signal.

enormous peak at the spurious frequency. Interestingly, there is also a small peak at the forcing frequency $f = 1$. This peak becomes a minimum if d_E is increased from 2 to 5, which we interpret as being due to the natural dimension of the attractor being between 2 and 3.

We consider the Lorenz system (as in Broomhead and King 1986), a more complicated system, and add a spurious cycle of frequency at $f = 0.95$ and amplitude 2. In Figure 5 we again show $\hat{\varepsilon}(f)$ plotted with the power spectrum; it is not obvious that the added frequency at $f = 0.95$ is dynamically independent of the remainder of the power spectrum, but $\hat{\varepsilon}(f)$ reveals a clear peak at the spurious frequency. In this case, though, if the amplitude is dropped to one the method fails. Actually one then finds that $\varepsilon(f) > \varepsilon_\infty$ for all f, so that the assumption $\varepsilon_{\min} < \varepsilon_\infty$ is violated and we have no result. This is probably due to the limited length of the data set (3300 points were used). We have deliberately used this technique with relatively small data sets, as this is usually the case in practice.

One of the features of the diagnostic is that the embedding dimension used for the nearest neighbour predictor can be very low: $d_E = 2$ is often sufficient (see Figure 3). The point is that it is not the magnitude of the error which is important, but relative increases to it. We illustrate this by adding a spurious frequency to the output of the Mackey-Glass equation (Mackey and Glass 1977), which for large values of the delay has a high-dimensional attractor (Farmer 1982), (Ershov 1992). In Fowler and Kember (1993) it was found that the Mackey-Glass prediction error is significantly reduced for

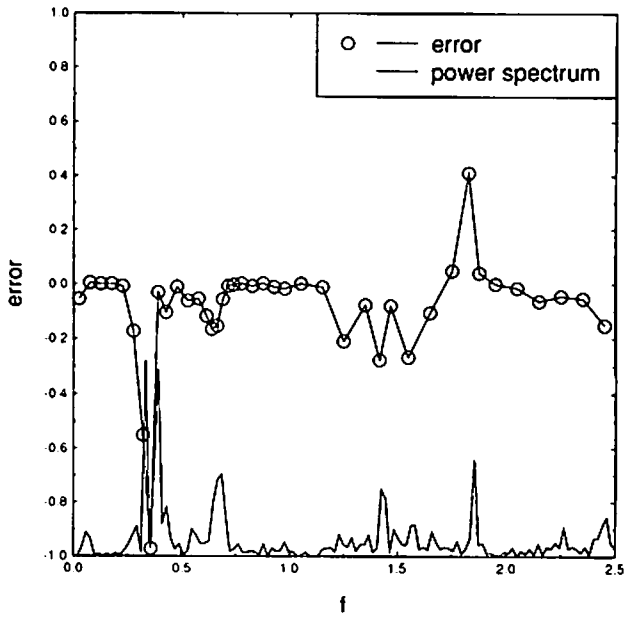

Fig. 6. As in Figure 4 but for the contaminated Mackey-Glass equation, in the form $\varepsilon \dot{x} = -x + f(x_1)$, $x_1 = x(t-1)$, $f(x) = 2x/(1+x^{10})$, $\varepsilon = 0.05$. The output is contaminated by addition of $0.1 \sin[2\pi(1.83)t]$ and corresponds to the peak at $f \sim 1.8$.

Takens embeddings that include a lag at the delay time of the Mackey-Glass equation. In Figure 6 we plot the power spectrum and $\hat{\varepsilon}(f)$ for a $d_E = 2$ embedding which includes a lag at the Mackey-Glass delay time, and there is a clear peak of $\hat{\varepsilon}$ exhibited at the spurious frequency. For other embeddings, though, the method must fail (all other phase-space methods would probably fail too) since the Mackey-Glass prediction error is that obtained for a white noise signal of the same mean and variance, so adding any spurious signal would have no effect on the prediction error.

Our final application is to a real data set, a pulse plethysmograph signal (Fowler and Kember 1995) which exhibits the output of two oscillators corresponding to heart rate ($f \sim 2.3$) and respiration ($f \sim 0.7$). The spectrum also reveals a harmonic frequency at $f \sim 4.6$ (see Figure 7). Part of our earlier interest in this series was to ask whether the harmonic was simply part of the cardiac oscillator, whether it was excited by the respiratory oscillator, or whether it is something else entirely. In fact, inspection of the time series or of its phase portrait (Figure 8) indicates that the second harmonic is an instrumental artifact associated with the plethysmograph 'sticking' at certain values. Therefore we in fact know that it is not dynamic. Figure 9 shows $\hat{\varepsilon}(f)$ for an embedding dimension of $d_E = 2$, with the power spectrum superimposed. We see peaks of $\hat{\varepsilon}(f)$ precisely at the spectral peaks. This suggests that the embedding space is unable to resolve any predictable structure. If we increase d_E to 5 we find that the respiratory oscillator is predictable, and in Figure 10 there is a clear jump in $\hat{\varepsilon}(f)$ at the second

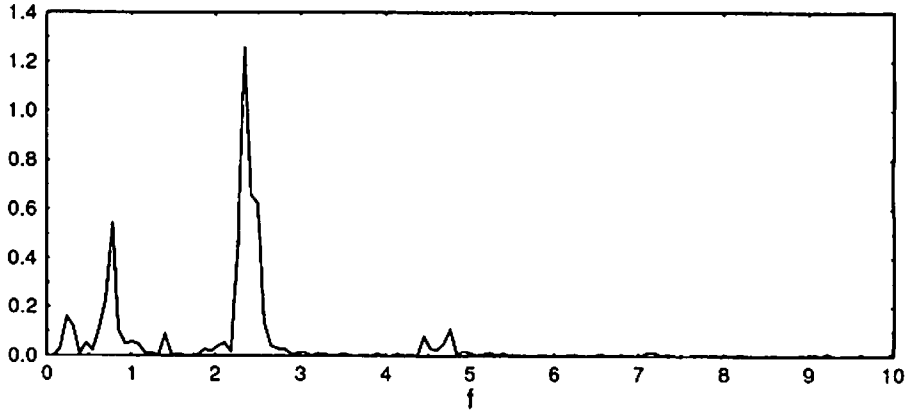

Fig. 7. Power spectrum of plethysmograph signal (Fowler and Kember 1995).

harmonic. This indicates that the second harmonic at $f \sim 4.6$ is nondynamic. The lower diagram of Figure 8 shows the phase portrait of the signal with a spectral band centred at $f = 4.4$ removed, which corresponds to the removal of the instrument artifact.

6. Discussion

Use of nonlinear methodologies to filter signal from noise, or signal from signal, has in the past been mostly limited to small-amplitude white noise and low-dimensional signals. One of the challenges facing nonlinear dynamics is to deal with signals where noise or contamination may be of large amplitude, and where two or more signals may contaminate each other. We have presented a method which seems able to distinguish spurious signals in a variety of systems with little restriction on their size, or on the size of the contaminating signal. We have only used bandstop filters here since our primary interest was to distinguish dynamics. An extension of the method to include variable amplitude reduction of the offending cycle would allow not only dynamic detection, but precision filtering of the offending cycle.

As with most algorithmic methods based on nonlinear dynamics, we cannot provide a guarantee that the method presented here will always work, even in the limited circumstances that we have described. However, to the extent that we have tested it, we have in fact found that the method performs robustly. It is important to emphasise that, although we use nonlinear prediction as a tool, we do *not* aim to provide the best predictor in so doing. More specifically, our method is based on the following two (related) tenets. First, that the added sine wave causes the embedded trajectory to be more 'complex'. Second, we suppose that the nonlinear predictor we use is essentially more accurate for the pure signal than when the sine wave is added. Thus it is essential that the predictor that we use is *not* optimised on the given time series (signal plus sine wave) since in that case, inevitably, the pure signal would be less predictable. What is important is the relative predictability of the time series compared to its bandstop filtered counterparts.

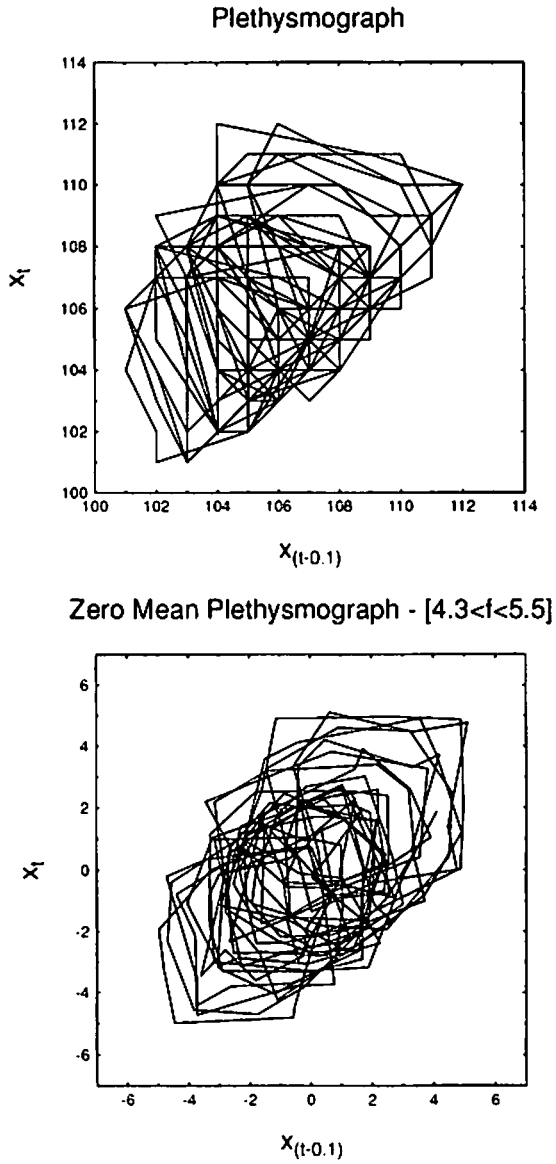

Fig. 8. Plethysmograph phase portrait (upper) and filtered portrait (lower) where a band at the spurious frequency $f \sim 4.6$ has been removed. The removed band corresponds to the instrument artifact.

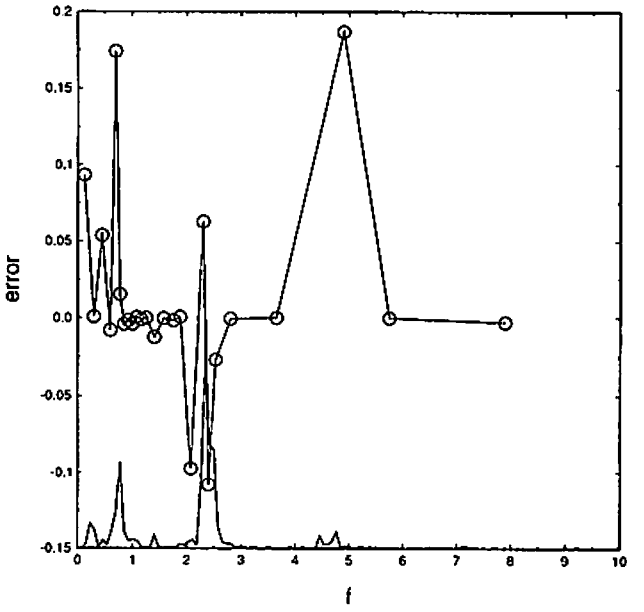

Fig. 9. $\hat{\varepsilon}$ versus f, plethysmograph signal. The peak between 4.3 and 6 represents nondynamic behaviour. Notice with $d_E = 2$ (here), we also see peaks at the pulse and respiratory frequencies, $f \sim 2.3$ and $f \sim 0.7$.

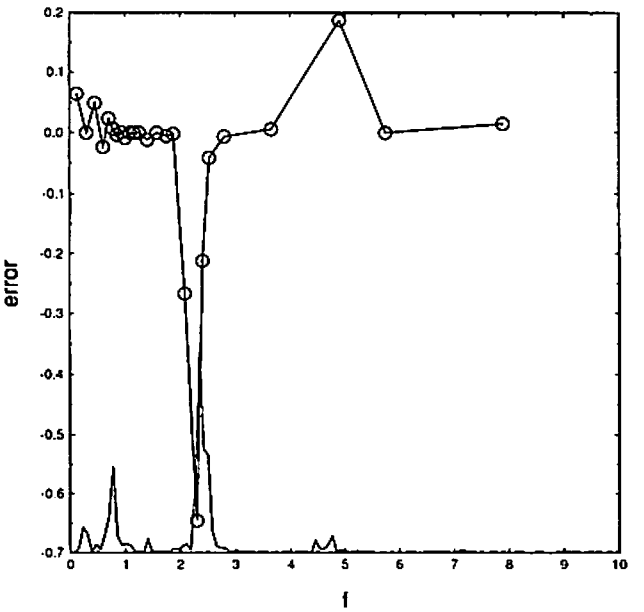

Fig. 10. As for Figure 9, but with $d_E = 5$. The peak due to the nondynamic signal at $f = 4.6$ is clearly apparent.

To put it another way, we can consider the prediction error as being a function of excised bandstop frequency f and prediction method parameter(s) p. Normally, one optimises prediction by seeking $\min_p \varepsilon(\infty, p) = \varepsilon^*$. On the other hand, here we fix p and study $\min_f \varepsilon(f, p) = \bar{\varepsilon}$. In general, we can expect $\varepsilon^* < \bar{\varepsilon}$. What we suggested here is that if prediction methods are chosen without specific regard to the system, then the diagnostic parameter $\bar{\varepsilon}$ can be of use in distinguishing different dynamical components, at least in simple situations. Whether the details of the diagnostic presented here will be appropriate in a wide variety of circumstances is unclear. However, we have shown how the combination of classical spectral methods with those of nonlinear dynamics can yield novel methods of discriminating between dynamically distinct signals, following the notions of Broomhead (1994).

References

Allen, M. R., and L. A. Smith 1994 Investigating the origins and significance of low-frequency nodes of climate variability. *Geophys. Res. Lett.* **21**, 883–886.

Broomhead, D. S. 1994 Nonlinear signal processing. In *Mathematics of Signal Processing III*, ed. J. G. McWhirter, IMA Conference Proceedings No. 49, 407–420. OUP, Oxford.

Broomhead, D. S. and G. P. King 1986 Extracting qualitative dynamics from experimental data. *Physica D* **20**, 217–236.

Casdagli, M. 1989 Nonlinear prediction of chaotic time series. *Physica D* **35**, 335–356.

Dehant, V., B. Ducarme, and P. D. Defraigne 1993 New analysis of the superconducting gravimeter data of Brussels. In *Dynamics of Earth's Deep Interior and Earth Rotation*, eds. J.-L. Le Mouël, D. E. Smylie, and T. Herring, Geophys. Monog. 72, 35–44. AGU, Washington, D.C.

Ershov, S. V. 1992 Asymptotic theory of multidimensional chaos. *J. Stat. Phys.* **69**, 781–812.

Farmer, J. D. 1982 Chaotic attractors of infinite-dimensional systems. *Physica D* **4**, 366–392.

Farmer, J. D. and J. J. Sidorovich 1987 Predicting chaotic time series. *Phys. Rev. Lett.* **59**, 845–848.

Fenstermacher, P. R., H. L. Swinney, and J. P. Gollub 1979 Dynamical instabilities and the transition to chaotic Taylor vortex flow. *J. Fluid Mech.* **94**, 103–128.

Fleming, P. J., M. R. Levine, A. M. Long, and J. P. Cleave 1988 Post neonatal development of respiratory oscillations. *Ann. N.Y. Acad. Sci.* **533**, 305–313.

Fowler, A. C. and G. Kember 1993 Delay recognition in chaotic time series. *Phys. Lett. A* **175**, 402–408.

Fowler, A. C. and G. Kember 1995 A nonlinear filtering technique for multi-oscillator systems. *Comput. Math. Appl.* **29**, 55–67.

Fowler, A. C. and G. Kember 1997 Moving window spectral analysis. *Eur. J. Appl. Math.*, submitted.

Ghil, M. and R. Vautard 1991 Interdecadal oscillations and the warming trend in global temperature time series. *Nature* **350**, 324–327.

Gibson, J. F., J. D. Farmer, M. Casdagli, and S. Eubank 1992 An analytic approach to practical state space reconstruction. *Physica D* **57**, 1–30.

Kostelich, E. J. and T. Schreiber 1993 Noise reduction in chaotic time-series data: a survey of common methods. *Phys. Rev. E* **48**, 1752–1763.

Kostelich, E. J. and J. A. Yorke 1990 Noise reduction: Finding the simplest dynamical system consistent with the data. *Physica D* **41**, 183–196.

Mack, Y. P. and Rosenblatt, M. 1979 Multivariate k-nearest neighbour density estimates. *Multivariate Anal.* **9**, 1–15.

Mackey, M. C. and L. Glass 1977 Oscillations and chaos in physiological control systems. *Science* **197**, 287–289.

Robbe, H. W. J., L. J. M. Mulder, H. Rüddel, W. A. Langewitz, J. B. P. Veldman, and G. Mulder 1987

Assessment of baroreceptor reflex sensitivity by means of spectral analysis. *Hypertension* **10**, 538–543.

Sugihara, G. and R. May 1990 Nonlinear forecasting as a way of distinguishing chaos from measurement error in time series. *Nature* **344**, 734–741.

Vautard, R. and M. Ghil 1989 Singular spectrum analysis in nonlinear dynamics, with applications to paleoclimatic time series. *Physica D* **35**, 315–424.

Vautard, R., P. Yiou, and M. Ghil 1992 Singular spectrum analysis: A toolkit for short, noisy chaotic signals. *Physica D* **58**, 95–126.

Yakowitz, S. and Karlsson, M. 1987 Nearest neighbour methods for time series with application to rainfall-runoff prediction. *Stochastic Hydrology*, 149–160.