Time-Frequency and Time-Scale Methods in the Detection and Classification of Non-Stationarities in Human Physiological Data

N. Darvish

Electrical Engineering Department
Imperial College
London SW7 2BX

R. I. Kitney

Electrical Engineering Department
Imperial College
London SW7 2BX

Abstract

In this paper, a variety of time-frequency and time-scale methods has been applied to the detection of transients in the heart rate and respiration of neonates during quiet sleep. The methods used include the well-known Short-Term Fourier Transform, the Wigner-Ville Distribution and the Wavelet Transform as well as the lesser-known Instantaneous Power Spectrum and the Choi-Williams Distribution. As far as the authors were aware, the latter two methods have not yet been applied to physiological signals. This study was carried out as part of a project to study the dynamic interactions in the cardiorespiratory system during transient phenomena.

1: Introduction

In the past two decades, a variety of signal processing techniques have been applied to human physiological data (in particular heart rate and respiratory data) in an effort to identify the mechanisms by which the various systems in the body interact and how the nature of these interactions change with age and disease [1]. Areas of interest include how the various control systems follow eachother ("entrain") and the extent and spectral distribution of the variability of the heart rate. In most cases, however, the studies were carried out on "stationary" data.

The reasons for this concentration on stationary techniques were twofold. One was that this approach yielded valuable information by studying the "steady-state" functioning of the systems and, by assuming that in this case the data would be stationary, provided valuable insights into the relations between various systems. Another reason was the lack of available signal processing tools that did not assume stationarity of the data series. The main tools used were the Fourier Transform and autoregressive methods which, at best, assume that the data is quasi-stationary within the window in which the analysis is being carried out.

The techniques used in this study consist of the time-frequency representations: the Short-Term Fourier Transform (STFT), the Wigner-Ville Distribution (WVD), the Instantaneous Power Spectrum (IPS) and the Choi-Williams Distribution (CWD) and the time-scale representation: the wavelet transform (WT).

It was found that using the data from a blind study of neonates in quiet sleep, the techniques were able to trace nonstationarities in the heart rate and respiration.

One reason for using such complicated and CPU-intensive methods would be to enable the results of any classification to be integrated into a monitoring device or model without requiring any further modelling or frequency detection. The results can be used to study the interactions between two variables (e.g. respiration and heart rate) during a transient phenomenon in one or both (e.g. during a sigh and its aftermath or during periodic breathing). Ultimately, the same methods could be used for the analysis of "steady-state" data so that the assumption of quasi-stationarity need no longer be used and the transition to non-stationary behaviour can be monitored to see if any further insight into the dynamic interactions of the various systems in the human body can be found.

2: Time-frequency methods

Cohen's class of time-frequency representations (TFR's) is defined as:

$$F(t,\omega) = \frac{1}{4\pi^2}\iiint_{\infty} e^{-j\theta t - j\tau\omega + j\theta u}\varphi(\theta,\tau).s^*\left(u-\frac{1}{2}\tau\right)s\left(u+\frac{1}{2}\tau\right)d\theta\, d\tau\, du$$

$$= \frac{1}{4\pi^2}\iiint_{\infty} e^{-j\theta t - j\omega\tau + j\tau u}\varphi(\theta,\tau).S^*\left(u+\frac{1}{2}\theta\right)S\left(u-\frac{1}{2}\theta\right)d\theta\, d\tau\, du$$

where $\varphi(\theta,\tau)$ is called the 'kernel' of the representation and S^* is the Fourier transform of the signal.

1058-6393/95 $4.00 © 1995 IEEE

Desirable qualities for TFR's include being real-valued (because of the difficulties inherent in interpreting a complex function as an energy surface), yielding the time and frequency marginals as the first conditional moments in time and frequency respectively, preserving the time and frequency support of the signal i.e. if the signal is limited to a section of the time axis (or the frequency axis) and is zero outside that section, then the TFR should also be zero outside that section.

It can be shown [4] that for each signal composed of n components, the quadratic TFR contains $n(n-1)/2$ cross-terms complicating the visual analysis of multi-component signals. Generally, it has been found that a trade-off exists between good T-F resolution and small interference terms.

For use as a method of tracking transients, the TFR must be computationally efficient and be robust with respect to small changes in the underlying system (e.g. slow changes in breathing frequency) while at the same time being quite sensitive to sudden changes in the structure of the signal and be able to accurately pinpoint the start and end of the transient in time.

2.1: Short-Term Fourier Transform

The Fourier Transform (FT) has long been used to obtain a spectral representation of a time-domain signal.

$$X(f) = \int x(t)\exp(-2\pi jft)dt$$

The technique does, however, assume that the data is stationary, an assumption that is seldom (if ever) true of physiological data. Hence the FT provides no time-localisation of the characteristics of the signal which would be vital if the start and end of the transient are to be located.

The STFT overcomes, to a certain extent this limitation of the FT by observing the signal $x(t)$ through a finite-length time window $w(t)$ (for the duration of which the signal is assumed to be stationary).

$$STFT_x(t,f) = \int x(\tau)w^*(\tau - t)\exp(-j2\pi f\tau)d\tau$$

The STFT is a linear TFR and, as such, has no cross-terms. The uncertainty principle, however, implies that the joint time-frequency resolution of the STFT is limited such that there is an inherent trade-off between the time-resolution (requiring the use of short windows) and frequency resolution (using long windows) [4]

The spectrogram (the square modulus of the STFT) is a member of Cohen's class of TFR's (with kernel $\varphi(\theta,\tau) = w(t+\tau/2)w^*(t-\tau/2)$) and may be loosely interpreted as the signal energy although it does not satisfy the necessary time and frequency marginal properties. The spectrogram does not have very good T-F resolution. It is, however, always non-negative-definite which makes it desirable for interpretation as an energy distribution. The STFT is a quadratic transform and satisfies the quadratic superposition principle, but the cross-terms of the spectrogram are greatly attenuated by the kernel and are localised to the neighbourhood of the auto-terms. This means that the results are easy to interpret visually[6].

2.2: Wigner-Ville Distribution

The WVD is a member of Cohen's class of time-frequency representations (with kernel $\varphi(\theta,\tau) = 1$) which satisfies a great number of the properties desirable in TFR's [4]. These include the fact that it is real-valued; preserves time and frequency shifts; satisfies the time and frequency marginals; its first moments in time and frequency give the group delay and the instantaneous frequency and it provides finite time and frequency support. The WVD is given by:

$$W_{x,x}(t,f) = \int x(t+\tau/2)x^*(t-\tau/2)\exp(-2\pi jf\tau)d\tau$$

One of the main advantages of the WVD is that the time and frequency resolution can be arbitrarily good. One of its main disadvantages is the presence of the cross-terms which will be nonzero irrespective of the position of the auto-terms in the T-F plane. The cross-terms appear at the arithmetic mean of the frequencies of the two auto-terms, oscillate in time at the frequency difference of the two tones and (if the auto-terms were of the same magnitude) could have up to twice their magnitude [6].This makes the visual interpretation of the WVD of multi-component or nonlinearly modulated signals difficult.

For the study of real signals, the WVD must be performed on the analytic signal which may be obtained using the Hilbert transform of the real signal. This is done in order to avoid the need to sample at twice the Nyquist rate and to avoid aliasing due to the negative frequencies.

2.3: Instantaneous Power Spectrum

The Instantaneous Power Spectrum (IPS) is a member of Cohen's class of time-frequency distributions with kernel:

$$\phi(\theta,\tau) = w(0)w(\tau)\cos(\pi\theta\tau)$$

It may be thought of as a smoothed version of the Margenau-Hill Distribution and thence the smoothed real part of the Rihaczek Distribution.

$$P_T(\omega) = \mathrm{Re}[s^*(t)\{S(\omega)\}e^{-j2\pi ft}]$$

The real part of the Rihaczek distribution is guaranteed to be zero anywhere in the T-F plane that the signal or its spectrum is zero [6], but the cross-terms at a particular frequency are not limited to the times when that frequency appears in the signal and appear unattenuated at all signal frequencies at all times that the signal is non-zero. In order to limit the cross-terms between two frequencies to the local T-F support of each signal component, Hippenstiel and De Oliviera [3] suggested using filters with decaying impulse responses. This is accomplished by replacing $e^{-j2\pi ft}$ in the above equation with $w(t)e^{-j2\pi ft}$ to obtain:

$$IPS(t,f)=\frac{1}{2}\int_{-\infty}^{\infty}\left[s(t)s^*(t-\tau)+s^*(t)s(t+\tau)\right]$$
$$.w(0)w(\tau)e^{-j2\pi ft}d\tau$$

The IPS has many properties that are useful in time-frequency analysis. It is time- and frequency-shift invariant, satisfies the marginals in time and frequency (the former only true if $w(0)=1$), is real, and yields the instantaneous frequency and group delay as the first moments in frequency and time respectively. The IPS does not require the use of an analytic signal , a sampling rate higher than the Nyquist rate or single spectral components at any time (to avoid cross-terms). It is also fairly insensitive to window parameters and retains good time-frequency localisation even at the analysis endpoints. The transitions between different states in non-stationary signals are also faster in the IPS, which yields better time resolution of the transition times. The IPS also has better SNR resolution than the WVD although at the expense of degraded frequency resolution. The correlation estimate of the IPS has half the effective lag range of that of the WVD so the effective spectral resolution of the IPS is also less than that of the WVD by the same factor. Pei and Tsai [7], have pointed out that the IPS suffers from oscillating cross-terms that are located at about the same location as the auto-terms They concluded that in applications where the modulation of the auto-terms is more acceptable than the introduction of cross-terms, the IPS is more useful than the WVD.

2.4: Choi-Williams Distribution

The Choi-Williams Distribution (CWD) [2] is a member of Cohen's class of TFR's with the kernel

$$\varphi(\theta,\tau)=(\sigma/4\pi\tau^2)^{1/2}\exp\left(\frac{-\sigma(t/\tau)^2}{4}\right)$$

It attenuates the cross-terms in the representation with the parameter σ. There is a trade-off between good auto-term resolution and good interference attenuation. Small values of σ attenuate the cross-terms to a large extent with accompanying loss of auto-term resolution while large values attenuate the cross-terms to a lesser extent but maintain better resolution. With $\sigma=\infty$, the WVD is obtained.

The CWD is real-valued, preserves time and frequency shifts and marginals, time and frequency moments and time frequency scalings as well as giving instantaneous frequency and group delay as the first mean conditional moments in frequency and time respectively. It does not, however, maintain finite time and frequency support.

2.5: Wavelet Transform

The Fourier transform resolves the signal onto sinusoidal bases of infinite extent. While this approach is useful for stationary signals, it creates problems in analysing non-stationary ones. One way of solving the problem is the use of TFR's as defined above. Another would be to resolve onto basis functions of finite extent (prototype wavelets g(t)) and this may be thought of as the basis of wavelet analysis [8].

$$W_x(a,b)=1/\sqrt{|a|}\int x(t)g^*((t-b)/a)dt$$

In the discrete-time, discrete-frequency plane:

$$W_g x(m,n) = a_0^{-m/2}\int_{-\infty}^{\infty} x(t)\,g(a_0^{-m}t-nb_0)dt$$

Increasing a dilates the wavelet while decreasing a causes it to contract. The WT thus uses short-duration windows at high frequencies and long windows at low frequencies, providing arbitrarily good frequency resolution at low frequencies and arbitrarily good time resolution at high frequencies. The WT is also governed by the uncertainty principle with equality given in the case of the gaussian wavelet of which the Morlet wavelet below is a modification.

$$g(t)=\exp(-t^2/2+j\omega_0 t)$$

where ω_0 is the central frequency of the wavelet.

3: Materials and Methods

The data was obtained from 7 neonates in a blind study. It was sampled at 16 Hz. One channel is heart rate as determined by a Hewlett Packard neonatal ECG monitor (the machine has a voltage output which is based approximately on the 'instantaneous' heart rate). Another channel is 'respitrace abdomen' data - a strain gauge attached to an abdominal belt which records the abdominal circumference and the data low-pass filtered at 6 Hz. The data was recorded onto a Racal VHS tape analogue recorder and then transferred to PC disk.

"Real-time" implementations of the WVD and IPS were achieved by adapting the algorithm presented in [1] for the real-time implementation of the WVD for the IPS. The WT was implemented using the toolbox developed by Carl Taswell (1990). All the algorithms were implemented in the Matlab 4.2 environment.

In the case of the TFR's, to provide comparability, 400 data point segments and a 64 point Hamming window was used in all cases. In the case of the CWD, the value of $\sigma=1$ was used as recommended in [2].

The wavelet transform algorithms were implemented using the Morlet wavelet with 1024 point data segments.

A peak detection and thresholding algorithm was used to differentiate between stationary and non-stationary data using the TFR's

4: Results and Discussion

The results of one set of analysis results are presented in fig. 2. It can be seen that the spectrogram provides a very "smooth" representation that traces changes in the phase of respiration at ~8s (to lower and higher frequency values to account for the missed beat and resultant two spikes). It can, however, be seen that, as explained, the time-frequency resolution is not very good and it does not localise the spikes.

The WVD on the other hand does localise the spikes in time but, because of the oscillating cross-terms in the frequency direction, the visual interpretation on the exact position of the spikes is difficult It also localised the second spike at approximately 18s which, in the spectrogram, was very diffuse.

The IPS provided good time and frequency resolution such that the two spikes at approx. 8s are well localised as is the spike at 18s.

The CWD again had no noticeable cross-terms but with lower frequency resolution than the WVD. It did, however, with its excellent interference attenuation properties, localise the spikes in the respiration. It also clearly shows the point at about 16s at which the respiratory frequency falls slightly. In some ways, by attenuating the cross-terms instead of localising them to the position of the auto-terms (and hence modulating the auto-terms), it deals with them in a way very suited to complex biological systems.

The WT provided a fast and efficient method of dealing with the large amount of data. The WT of the respiration signal, localised the spikes in the time-scale plane (fig. 1).

The thresholding algorithm provided a clearer way of studying the output of TFR's and detecting the non-stationarities was possible by taking moments (fig. 3).

5: Conclusions

The time-frequency and time-scale methods were shown to provide a new way of looking at the essentially non-stationary cardiorespiratory data in order to study the interaction between the HR and respiration.

The IPS, WVD and CWD algorithms were shown to be effective in localising events in the time series to the time-frequency plane. In studying the relationship between the heart rate and respiration, the TFR's of each show correlations. The cross-Wigner-Ville distribution can be used to perform time-frequency cross-correlation with good results.

Acknowledgements

The data was collected by Dr. Simon Bignall at St. Mary's Hospital.

N. Darvish is supported by an EPSRC research grant.

The CWD algorithm used was courtesy of R. van der Heiden.

The authors would also thank Dr. A. Bharath for valuable discussions.

References

[1] Bergmann, N., New Formulation of Discrete Wigner-Ville Distribution, *Electronic Letters*, Vol. 27, No. 2, pp. 111-112.

[2] Choi, H-I. and Williams, W.J., Improved Time-Frequency Representation of Multicomponent Signals Using Exponential Kernels, *IEEE Transactions on Acoustics, Speech and Signal Processing*, Vol. 37, No. 6, pp. 862-871,1989.

[3] Hippenstiel, R.D. and De Oliveira, P.M., Time-Varying Spectral Estimation Using the Instantaneous Power Spectrum (IPS), *IEEE Transactions on Acoustics, Speech, and Signal Processing*, Vol. 38, No. 10 pp 1752-1759, 1990

[4] Hlawatsch, F. and Boudreaux-Bartel, G.F. Linear and Quadratic Time-Frequency Representations, *IEEE Signal Processing Magazine*, , pp. 21-60, April 1992.

[5] Kitney , R. I. and Darvish, N., Techniques for Studying Short- Term Changes in Cardio-Respiratory Data II, in *Computer Analysis of Cardiovascular Signals*, IOS Press, 1994

[6] Loughlin, P. J., Pitton, J. W., Atlas, L. E., Bilinear Time-Frequency Representations: New Insights and Properties, *IEEE Transactions on Signal Processing*, Vol. 41, No. 2, pp. 750750-767, 1993

[7] Pei, S , and Tsai, E., Cross-Terms analysis in the Modified Instantaneous Power Spectrum, *IEEE Transactions on Signal Processing*, Vol. 41, No. 1, pp. 477-480,1993.

[8] Rioul, O., and Vetterli, M., Wavelets and Signal Processing, *IEEE Signal Processing Magazine* pp. 14-38 October 1991.

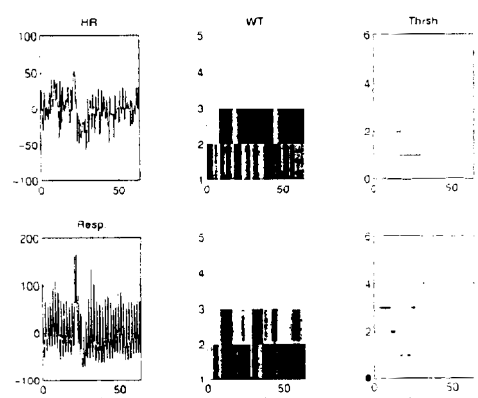

Fig1. The Wavelet Transform

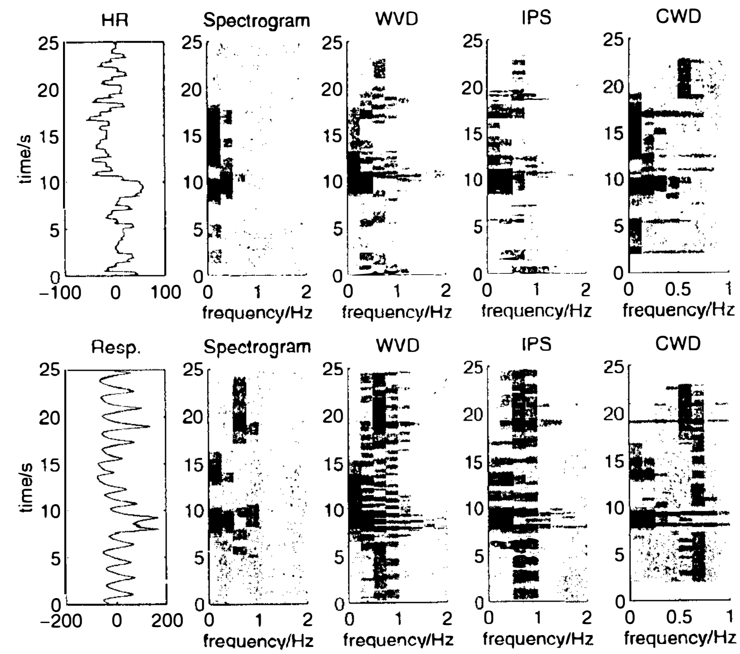

Fig2. The Time-Frequency Representations

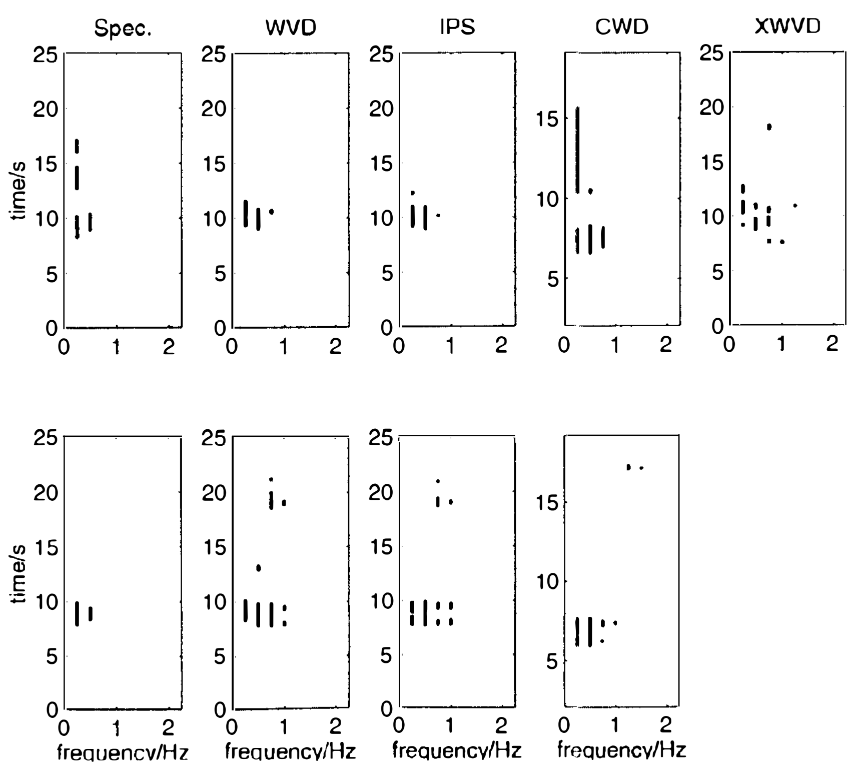

Fig.3 Thresholding applied to the algorithms in fig. 2