

Heart Sound Analysis comparing Wavelet and Autoregressive Techniques

J N Torry

University of Sussex, Brighton, UK

Abstract

Three different types of wavelets (Morlet, Bessel filter and Chirp) and four different types of autoregressive modelling techniques (Burg, Forward, Backward and Modified Covariance) are considered. A straightforward algorithm is suggested for all wavelet calculations. A modification is made to the standard autoregressive technique to reduce the dominance of certain peaks. The heart sounds of a set of normal patients and on a set of patients diagnosed as having aortic stenosis are investigated. After separating the sound into individual heart cycle and further separating into the four phases (S1, systolic, S2 and diastolic), the relative energy levels in different phases over a number of cycles are compared. Typical results are selected for individual cycles and show the benefit of the modified autoregressive approach.

1. Input data

The aim is to seek a better visualisation of heart sound using modern digital signal processing techniques. By recording the sound information on to a laptop computer both sound replay and graphical display can add to the understanding of various heart sound conditions. The value of such an approach [1] and the graphical display of results [2] have been seen as good aids to diagnosis.

Data has been collected over a number of years from patients at the Brighton and Sussex University Hospital (85 patients in 1995/6, 25 in 2000/01 and a recent batch of 15 in 2003) using a lab-built digital stethoscope and standard interface to a laptop. Analogue measurement used the HP21050A sensor for the early recordings and a good electret microphone for the recent set. The bandwidth is limited to 10Hz to 700Hz using third and fourth order Chebyshev filters respectively. The signal is digitised with 12-bit resolution at a relatively high sampling rate of 4096Hz. Simultaneous recording were taken of the patients' ECG. A bulky 486 PC and large interface box was used for the 1995/6 and 2000/1 recordings but has now been updated to a manageable laptop, making the 5 mile trip to the local hospital easier. All patients were seen

by a resident hospital cardiologist prior to undergoing probing from university postgraduates. In most cases there is supporting written documentation from the cardiologist and in some cases ultrasound scan information is available.

Once recorded, the off-line process of analysis can begin. Matlab™ provides, at ones fingertips, a range of signal processing algorithms which can be adapted for this heart sound application. Analysis begins by segmenting the input waveform into individual heart sound cycles. The R wave from the ECG signal is used as the cycle divider and as an indication of the approximate start of S1. Wavelet and AR techniques are best applied on a cycle-by-cycle basis. Further analysis requires the subdivision into the S1, systolic, S2 and diastolic phases. After which an energy calculation can be made to assess the relative energy in each phase and in particular give a value to the level of severity of any systolic murmur.

2. Wavelet analysis techniques

The Morlet wavelet has been used in recent years by a number of researchers [3]. This is not an orthogonal wavelet as required by those who do compression and decompression, but is merely a means of looking at some of the time-frequency information of a signal. Selecting a wavelet then becomes a process of seeing which wavelet suits the style of the signal to be investigated rather than being constrained by the compression process.

The 'wavelet manager' in Matlab™ is somewhat cumbersome but for a specific wavelet analysis it can be reduced to a few lines of code:

```
alf = 5; ncyc = 8; ind = 1; fs = 4096;
for af = fr
    ub =ncyc*pi/alf; lb =-ub; np =round(fs*ncyc/af);
    tw = linspace(lb,ub,np);
    wval = exp(-(tw.^2)/2).*(cos(alf*tw) + ...
        i*sin(alf*tw));
    rxy = conv(y,fliplr(wval));
    d = floor((length(rxy)-ly)/2); % offset dist
    x = rxy(d+1:d+ly); % centralise result
    coef(ind,:) = abs(x)*af;
    ind = ind + 1;
end
```

Table1 – Matlab™ source code for wavelet analysis

The Morlet wavelet is a complex wavelet which can be achieved by either using the $\cos + i \sin$ form or taking the Hilbert transform of either the wavelet or signal. 4 cycles either side of the centre give sufficient points for the wavelet to decay so that the convolution multiplication is not too long. The algorithm convolves the data set y (length ly) with scaled versions of the wavelet to give a set of coefficients from which a contour plot can be drawn. The frequency range, fr , is usually logarithmic but can be linear. Strictly speaking the vertical axis on a wavelet time-frequency plot is 'scale' and not 'frequency'. Since the Morlet wavelet has only a single constant frequency given by $\cos(af*tw)$ the direct equivalent to frequency can be used rather than the inverted logarithmic axis given by 'scale'.

An alternative wavelet can be devised based on the impulse response of a band pass filter. In general the impulse response is a set of m terms of the form:

$$\exp(al*t(m))'*(abs(res).*\cos(be*t(m)-angle(res)))$$

Despite the number of different frequency terms involved, the impression given is that the impulse transient response moves from a signal resonating at the upper cut-off frequency to one at the lower cut-off frequency. The set of Bessel filters give a relatively smooth transition between the limits of the cut-off frequencies. (Poles are approximately equi-distance from the $j\omega$ axis in a straight line between the cut-off frequency limits). Thus, for example, a 6th-order band pass filter, derived from a 3rd-order Bessel low-pass prototype, centred on 60Hz with a Q of 2.1 has a 3dB bandwidth between 76 and 47 Hz. The transient response shows a gradual move from a waveform at about 78Hz reducing after 3 cycles to 46Hz. If in another case one chooses a Q of 1.13 (therefore bandwidth between 92 and 39Hz), the output transient goes between 114 and 26Hz in about 2 cycles. Thus setting the Q factor establishes the range of 'instantaneous' frequency. As with all filters, the narrower the bandwidth, the higher the Q , the more ringing there is in the transient response. Thus the second example takes longer to move from its 78Hz to 46Hz. The end result is a waveform which changes from high frequency to low with an approximately linear frequency sweep modified by an envelope based on the exponential terms

A chirp wavelet, with a set frequency sweep, can be defined such that it is approximately equivalent to the Bessel wavelet. Defining a chirp wavelet in frequency terms is not difficult. Taking a sine wave at 80Hz and stretching it down, linearly, to 40Hz over a period of time forms the basis of the frequency content of the mother wavelet. The amplitude, or envelope shaping, which allows rising from time zero to a gradual decaying after some arbitrary time, can more or less be of any shape. A straightforward envelope can be generated from a damped second-order low-pass filter response with an exponential

rise followed by an exponential fall, $\exp(-25t) - \exp(-20t)$. The following code can be inserted suitably into the Morlet calculation:

```
ifr = fmin + linspace(0, (fmax-fmin), np)/2;
chp = sin(2*pi*tw.*ifr);
env = exp(-1.25*af*tw) - exp(-1.55*af*tw);
wval = env.*chp;
```

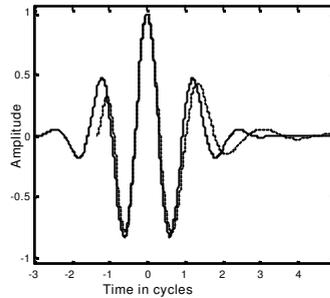


Figure 1. Morlet and Bessel

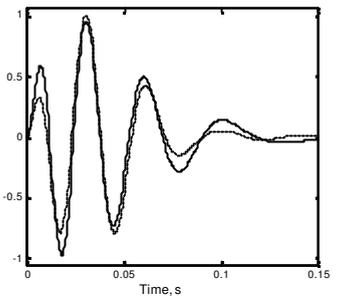


Figure 2. Bessel and chirp

The Bessel (dotted plot) and Chirp wavelets are not symmetrical and to compare with the Morlet they have been shifted in time to get the best match.

3. Autoregressive analysis techniques

Autoregressive modelling analysis techniques can usually be relied upon to give a better power spectral estimation than traditional FFT analysis, especially when the data set is fairly short. The theoretical foundation for the techniques can be tested by setting up coefficients for a low-pass digital filter, calculating its impulse response, adding noise and then finding the best autoregressive approach to estimate the filter coefficients. Hayes [5] gives Matlab™ algorithms for Forward Covariance page310, Burg p319, Modified Covariance p324. For assessing alternative techniques, his test process enables one to examine (i) spectral resolution, (ii) effects of added noise and (iii) effects of changing the order of the model filter. Results show that both the Burg and Modified Covariance techniques easily outshine the Forward and Reverse Covariance approaches. However both Burg and Modified Covariance although accurate in frequency resolution show a tendency for high peaks at the filter pole frequencies. This peakiness occurs when either there is low noise or when a higher than needed order is selected. The problem is that high peaks can saturate the system and prevent the identification of other lower-amplitude resonances. Since one does not know the noise level of any particular heart sound recording or anything about the nature of the resonances, it is difficult to select the appropriate order for the AR model. The approach suggested here is that these peaks should be smoothed with, for example, a low pass filter. This means taking the unusual step of filtering the frequency spectrum. The amount of filtering has been selected by trial and error, with the main test based on looking at the final time-

frequency contour plot and making sure that no peaks are over flattened.

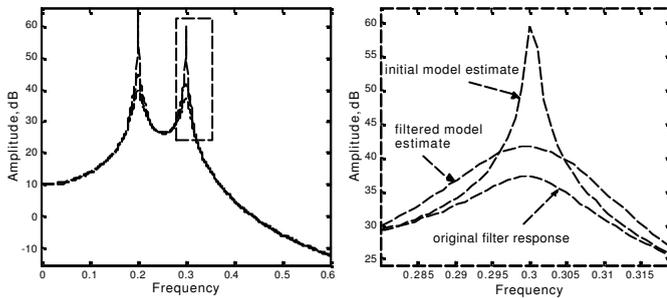


Figure 3. Full/detailed spectrum before and after filtering

Having started the idea of filtering the frequency spectrum, it can be continued if it is desired to emphasis a particular range of frequencies. In our recordings, allowing signals down to 10Hz gives extra information below the normal audible range. However heart sounds can get swamped by this low frequency content and if one wishes to emphasise murmurs then a gradual reduction of low frequency content is desirable. Here an arctan envelope is applied to achieve further filtering.

$$fo = 20; Q = 0.7;$$

$$envt = atan2(fo/Q*fr, (fo*fo - fr.*fr))/pi;$$

Time resolution is achieved by applying AR analysis to short sections of the waveform, analogous to the Short-Time Fourier Transform. The number of time windows taken within the heart cycle period affects the overall look of the combined contour plot. Individual time windows can overlap each other so as to achieve a sliding window across the data set. For this study, between 60 and 80 time windows were chosen with an overlap factor of between 0.8 and 1.3. An overlap of 1 indicates that the second time slot begins at the halfway point of the first slot, and thus the third begins at the end of the first. Within the time window, with its relatively few number of data points, the model order can be low. Model orders of 6 or 8 are quite sufficient to bring out the relevant features.

4. The severity of a systolic murmur

Applying the above techniques is done in one of two ways. Firstly as means of comparing the relative strengths of the S1, S2 and murmur signals and secondly by looking at individual cycles to see what information the various techniques can yield.

Initially the heart cycle must be segmented into the S1, systolic, S2 and diastolic phases. This is not a trivial task to do in general. In normal patients there is a clear distinction between the lub and dub of the heart sound. In the presence of a systolic murmur, such as generated by aortic stenosis, the boundaries between the end of S1 and the beginning of the murmur, and the end of the murmur

and the beginning of S2 can become very obscure. Frequency differences can assist this segmentation task with S1 and S2 being in the 40 to 80Hz range and murmur being higher, 80 to 150Hz, but is not always the case. Although the fundamental resonances of S1 and S2 are at relatively low frequencies in the audible range, the lub-dub effect can be still hear in the 100 to 200Hz range. An algorithm has been devised that looks for peaks of signal around the probable region where S1 and S2 occur i.e. 75ms and 0.4s after the R wave. The algorithm looks for minimums either side of the peaks within certain time limits.

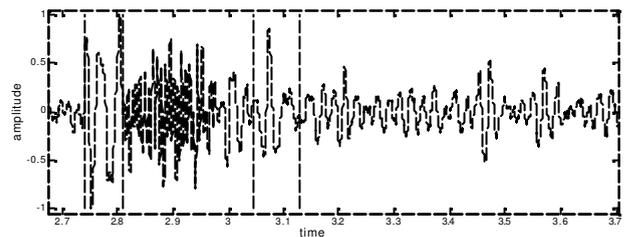


Figure 4. Single cycle segmented into S1 and S2

To estimate the severity of a murmur, a comparison is made of the energy level in the systolic period relative to the energy in S1 and S2. If there is no diastolic murmur then the signal level in the diastolic period is used to estimate the general noise level of the recording and subtracted from the other levels to eliminate the lower noise threshold.

5. Results for individual cycles

The first example is taken from cycle number 3 from patient AS12, chosen because there is a significant systolic murmur.

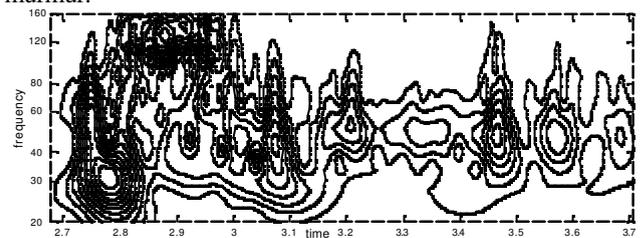


Figure 5. Morlet wavelet analysis

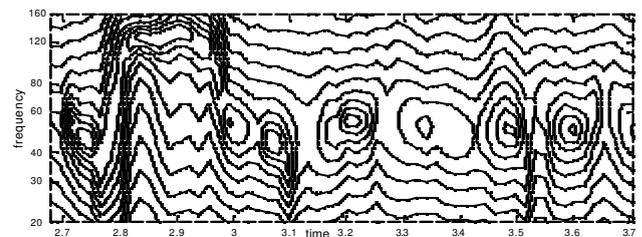


Figure 6. Modified Covariance analysis

Even in a black and white reproduction, the detail seems clearer using the autoregressive approach. In a second example, taken from cycle 9 of patient AS7, the problem is that in the wavelet representation the dominant features of S1 and S2 hide the lesser features of the murmurs.

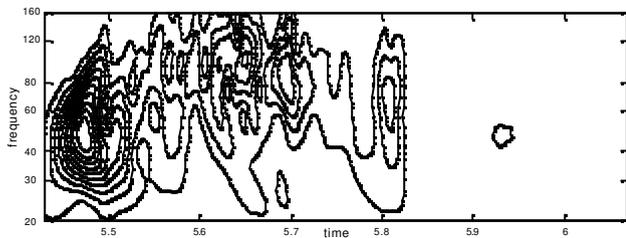


Figure 7. Bessel wavelet analysis

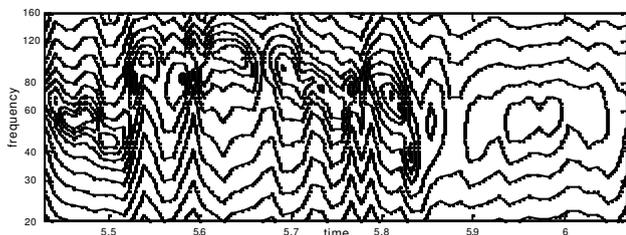


Figure 8. Burg analysis

Although this is a different patient, the Bessel wavelet gives almost identical results to the Morlet wavelet approach and, again, a dominant S1 swamps the detail in the rest of the cycle. The Burg analysis is very similar results to the Modified Covariance approach.

A third example, from cycle 9 of patient AS89, gives a chirp wavelet analysis and corresponding Modified Covariance analysis

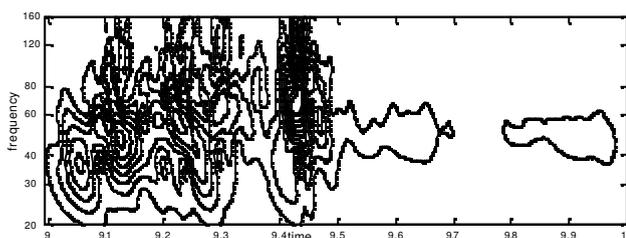


Figure 9. Chirp wavelet analysis

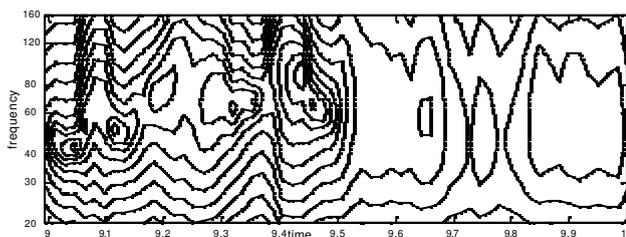


Figure 10. Modified Covariance analysis

Again the Chirp wavelet is very similar to other wavelet analysis. Even Forward and Reverse Covariance techniques emphasise some of the additional features but show sporadic peaks compared with Burg and Modified Covariance.

In these examples the individual cycles have been chosen so that the individual severity matches the average over some 25 to 35 cycles.

6. Conclusions

A conclusion drawn from the above suggests that there is little difference from a variety of different wavelets when applied to heart sounds, especially when the wavelets visually look closely matched. There is also not much difference in the various autoregressive modelling techniques, with a slight preference towards the Modified Covariance and Burg approaches. By filtering the frequency spectrum the traditional AR approach has been improved so that it yields greater detail than the well-respected wavelet approach.

Acknowledgements

The author wishes to thank past and present postgraduates, Blanca Tovar-Corona and Hosam Mgdob, and Professor Richard Vincent for assisting in the collection of the patient data.

References

- [1] Tilkian AG, Conover MB, Understanding Heart Sounds and Murmurs with an introduction to Lung Sounds, 1984, Philadelphia: Saunders.
- [2] Lukarinen S, Nojonen AL, Sikio K, Angerla A, A New Phonocardiographic Recording System, Computers in Cardiology 1997, Lund, IEEE Computer Society Press, vol 24:117-120.
- [3] Tovar-Corona B, Hind MD, Torry JN, Vincent R, Effects of Respiration on Heart Sounds Using Time-Frequency Analysis, Computers in Cardiology 2001, Rotterdam, IEEE Computer Society Press, 2001: 457
- [4] Bentley PM, Grant PM, McDonnell JTE, Time-Frequency and Time-Scale techniques for the classification of native and prosthetic heart valve sounds, IEEE Trans. Biom. Eng. 1998 vol 45(1):125-128.
- [5] Haynes MH, Statistical digital signal processing and modelling, John Wiley & Sons, 1996, ISBN 0-471-59431-8:101-104.

Address for correspondence.

Dr John N Torry
 Department of Engineering and Design
 University of Sussex
 Falmer, Brighton
 BN1 9QT, UK
 E-mail address: j.n.torry@sussex.ac.uk