

A Comparison of Signal Processing Techniques for the Extraction of Breathing Rate from the Photoplethysmogram

Susannah G. Fleming and Lionel Tarassenko

Abstract—The photoplethysmogram (PPG) is the pulsatile waveform produced by the pulse oximeter, which is widely used for monitoring arterial oxygen saturation in patients. Various methods for extracting the breathing rate from the PPG waveform have been compared using a consistent data set, and a novel technique using autoregressive modelling is presented. This novel technique is shown to outperform the existing techniques, with a mean error in breathing rate of 0.04 breaths per minute.

Keywords—Autoregressive modelling, breathing rate, photoplethysmogram, pulse oximetry.

I. INTRODUCTION

PULSE oximetry is frequently used in clinical situations for non-invasive measurement of heart rate and arterial oxygen saturation. It has been suggested that signal processing techniques can be used to extract the breathing rate from the photoplethysmogram (PPG), which is the pulsatile waveform produced by a pulse oximeter at one of its two wavelengths (red and infra-red). If this is possible, it would allow non-invasive measurement of breathing rate using a device (the pulse oximeter) that is already used in many clinical situations, and is known to cause a minimum of distress or inconvenience to the patient.

A number of methods for deriving the breathing rate from the PPG have been suggested in the literature. Results from the assessment of these methods using a consistent data set are presented in this paper, allowing comparisons between the methods to be made. A new method using autoregressive (AR) modelling has also been developed, and is shown to perform better than the existing techniques.

II. MATERIALS AND METHODS

Seven records from the MIMIC database in the Physiobank archive [1] were identified for use in assessing the accuracy of the algorithms for the extraction of breathing rate. These records all contain both the PPG waveform, and a synchronous respiratory waveform (believed to be obtained by nasal thermistry) for use as a reference. Two five-minute sections from each of the records were identified for use in the tests, resulting in a dataset consisting of fourteen five-minute sections from seven individuals. Sections were identified by looking for the




Fig. 1. A thirty second period of the PPG and reference respiratory waveform from Patient A. The positions of individual breaths detected in the reference respiratory waveform are show by circles.

first two non-overlapping five-minute sections which did not have missing data in either the PPG or the reference respiratory waveform channels.

An extrema detection algorithm was implemented to obtain timestamps for individual breaths in the reference respiratory waveform, with the results checked visually to ensure that all of the breaths in the respiratory waveform were identified, and that all of the detected breaths corresponded to actual breaths in the waveform. Fig. 1 shows the positions of these timestamps for a 30-second period within one of the test sections. As the reference respiratory waveform occasionally contains artefactual spikes, it was filtered prior to detection using an FIR band-pass filter. The filter uses a Kaiser windowing function, with a pass-band extending from 0.1–40Hz (i.e. from 6 breaths per minute to just below mains frequency), 30dB attenuation, and a 5% pass-band ripple.

Where visual inspection showed that the algorithm had identified artefactual breaths due to excessive levels of noise on the reference respiratory waveform, the artefactual breaths were removed from the timestamp data. In some cases, the reference respiratory waveform was so noisy that visual identification of breaths was not possible; in these cases, the section was replaced by another five-minute section from the same record.

The existing signal processing algorithms for obtaining the breathing rate from the PPG were implemented in Matlab, following the descriptions in the literature. Two methods using

Manuscript received September 13, 2007.

S. G. Fleming (corresponding author, phone: +44-1865-273096; email: susannah@robots.ox.ac.uk) and L. Tarassenko (email: lionel@robots.ox.ac.uk) are with the Department of Engineering Science, University of Oxford, Parks Road, Oxford OX1 3PJ, UK.




Fig. 2. PPG derived respiratory waveforms from the two digital filtering methods and reference respiratory waveform for Patient A, with detected breaths marked by circles.

digital filtering, as described by Nilsson *et. al.* and Nakajima *et. al.* were investigated, as well as the wavelet-based method described by Addison *et. al.*. Where the algorithms resulted in a breathing-synchronous waveform, the breath timings were obtained using the same extrema detection algorithm as was used for the reference respiratory waveform, omitting the initial band-pass filter.

The ensemble breathing rate in breaths per minute over the five-minute section was calculated and the absolute difference from the reference rate (derived from the reference respiratory waveform) was then used as the primary measure of performance for the algorithms for extracting the breathing rate from the PPG waveform.

A. Digital Filtering

Nilsson *et al.* suggest the use of a 3rd order Butterworth band-pass filter with a pass-band from 0.1–0.3Hz (6 to 18 breaths per minute) [2], [3]. This was found to perform well, with an average error of less than 0.5 breaths per minute when compared to the reference rate.

Filtering is also used in the method suggested by Nakajima *et al.* [4]. In this, the PPG is initially filtered with analogue low- and high-pass filters with cut-off frequencies at 0.1 and 5Hz. The breathing signal is then obtained using three different low-pass filters, with cut-off frequencies at 0.3, 0.4 and 0.55Hz. The choice of filter is determined by the heart rate (or pulse rate), so that a higher cut-off frequency is used at a higher heart rate, and hysteresis is employed to reduce the rate at which filters are switched. The filters were implemented using the characteristics described in [4], using a Kaiser windowing function. Despite the increased complexity of this algorithm, it does not perform as well as the one proposed by Nilsson *et al.*, and has an average error of over 3 breaths per minute when compared to the reference rate.

Fig. 2 shows the results of using the two different filtering methods on the relatively clean PPG signal from Patient A, together with the reference respiratory waveform.




Fig. 3. PPG derived respiratory waveform from the wavelet method (obtained from the ridge of the breathing band) and reference respiratory waveform from Patient A, with detected breaths marked by circles.

B. Wavelet decomposition

Addison *et al.* use wavelet transforms to extract a breathing waveform from the PPG [5], [6], [7], [8], [9], [10], [11], [12]. The PPG signal is decomposed by a continuous wavelet transform using the complex Morlet wavelet to produce a scalogram in which two bands can be identified: one at the frequency corresponding to the breathing rate, and one at the frequency corresponding to the pulse rate. The crest of the ridge corresponding to the pulse band is followed and projected as either an amplitude-time or frequency-time signal. These signals are referred to as the ridge amplitude perturbation (RAP) and ridge frequency perturbation (RFP) respectively.

The RAP and RFP are then subjected to further wavelet transformation, and the resulting scalograms are interrogated for the presence of breathing ridges. Therefore, up to three breathing signals may be obtained from a single PPG: from the original signal, the RAP, and the RFP. The ‘best’ source for each five-minute section was chosen as the one with the lowest absolute error when compared to the reference. This could not be used in practice, when a reference rate will not be available, but was used here to ensure that the best possible results are presented.

Despite the complexity and computational difficulty of this method, it was found to have an average error of 1 breath per minute when compared to the reference respiratory rate. These results are worse than those reported in [9], [10], [11], [12], although the reason for this is not clear. The published results are reported to have been obtained using a raw PPG signal, and it is likely that the waveforms in the MIMIC database have undergone some pre-processing, which may explain at least some of the discrepancy.

C. Novel AR method

We have developed a novel method of measuring the breathing rate from the PPG signal using autoregressive modelling. This technique has been applied to a number of other physiological signals, including the EEG [13] and the intrapartum

cardiotocogram [14], but has not yet been applied to the problem of extracting breathing rate information from the PPG waveform.

AR modelling can be formulated as a linear prediction problem, where the current value $x(n)$ can be modelled as a linearly weighted sum of the preceding p values. The parameter p is the model order, which is usually much smaller than the length of the sequence N .

$$x(n) = -\sum_{k=1}^p a_k x(n-k) + e(n) \quad (1)$$

The value of the output $x(n)$ is therefore a linear regression on itself, with an error $e(n)$, which is assumed to be normally distributed with zero mean and a variance σ^2 . The problem can also be visualised in terms of a system with input $e(n)$, and output $x(n)$, in which case the transfer function H can be formulated as shown below:

$$H(z) = \frac{1}{\sum_{k=1}^p a_k z^{-k}} = \frac{z^p}{(z-z_1)(z-z_2)\dots(z-z_p)} \quad (2)$$

As shown in (2), the denominator of $H(z)$ can be factorised into p terms. Each of these terms defines a root z_i of the denominator of $H(z)$, corresponding to a pole of $H(z)$. Since $H(z)$ has no finite zeros, the AR model is an all-pole model. The poles occur in complex-conjugate pairs, and define spectral peaks in the power spectrum of the signal, with higher magnitude poles corresponding to higher magnitude peaks. The resonant frequency of each spectral peak is given by the phase angle of the corresponding pole, such that the phase angle θ of a pole at frequency f is defined by (3), which shows that it is also dependent on the sampling interval Δt .

$$\theta = 2\pi f \Delta t \quad (3)$$

The PPG signal is typically sampled at rates between 100 and 250 Hz to ensure that the shape and heart rate information are preserved. At such high sample rates, the phase angles corresponding to breathing frequencies are very small, which is likely to lead to inaccuracy in identifying the frequency of the breathing pole, or possibly even the absence of a breathing pole in the AR model (since it would be subsumed into the real-axis, or d.c., poles). It is therefore necessary to downsample the signal to increase the angular resolution of the low frequency information. This also ensures that the cardiac-synchronous pulsatile component of the PPG is no longer dominant, as otherwise many or all of the poles will be used to model this signal, rather than the wanted breathing signal. To improve the stability of the AR model, it is also necessary to remove any DC offset from the downsampled PPG.

The PPG signal is therefore downsampled and detrended prior to AR modelling. A decimation algorithm, which filters the signal prior to resampling, is used to reduce the effect of aliasing in the downsampled signal.

A range of angles is defined by the expected breathing frequencies for a normal subject, and the poles with phase angles within this range are identified as possible breathing poles. For models with $p > 3$, multiple poles may be identified,




Fig. 4. Downsampling to 1 Hz and detrending of the PPG of Patient A prior to applying the AR model.




Fig. 5. Poles of AR model (11th order model, PPG downsampled to 1Hz) for Patient A, with breathing pole circled, and the limits of the sector of interest indicated by the dashed lines. The phase angle of the breathing pole corresponds to 0.27 Hz (16 breaths per minute)

as the range of possible breathing frequencies is quite large, and most other signals will have been removed by the filtering and downsampling steps. The choice of pole is made using the magnitude of the poles, as poles corresponding to breathing should have a high magnitude. The pole with the highest magnitude in the sector of interest is identified, and a threshold of 95% of the magnitude of this pole is used to find candidate poles in the sector of interest.

Of this candidate group, the pole with the smallest angle (corresponding to the lowest frequency) is identified as the breathing pole. This requirement is introduced because poles at a multiple of the breathing frequency occasionally occur with a slightly higher magnitude than that of the true breathing pole.

For adult subjects, the sector of interest is defined as that covering the angles corresponding to frequencies between 0.08 and 0.7 Hz (4.8–42 breaths per minute), although it should be noted that a different upper limit may be imposed by the Nyquist limit following downsampling (i.e. downsampling to 1 Hz would impose an upper limit of 0.5 Hz, or 30 breaths

TABLE I
ERRORS (IN BREATHS PER MINUTE) BETWEEN PREDICTED BREATHING RATE FROM THE PPG AND THE REFERENCE RESPIRATORY RATE

Method	Mean error	Range
Filtering (Nilsson <i>et al.</i>)	0.47	0.006–2.52
Filtering (Nakajima <i>et al.</i>)	3.02	0.004–8.2
Wavelets (Addison <i>et al.</i>)	1.01	0.06–2.73
Novel AR method	0.04	0.0008–0.17

per minute).

The optimal combination of AR model order and downsampled frequency was determined experimentally using six of the fourteen sections in the dataset. The lowest error was found with an 11th order model and a downsampled frequency of 1 Hz. This resulted in an average error of just 0.04 breaths per minute when compared to the reference respiratory rate.

Fig. 4 shows the pre-processing of a PPG signal before the AR model is applied, and Fig. 5 shows the poles of the corresponding AR model. The accuracy of the breathing rate estimation is underlined by the high magnitude of the breathing pole circled in Fig. 5, which lies very close to the unit circle. As the PPG is downsampled to 1 Hz, the upper limit of the sector of interest is reduced to 30 breaths per minute by the Nyquist criterion, as shown in Fig. 5 by the line along the negative real axis.

III. RESULTS AND DISCUSSION

Table I shows that the novel AR method presented in this paper performs better than both the digital filtering and wavelet decomposition methods for extracting the breathing rate from a PPG waveform on most signals. Further processing to reduce the likelihood of choosing a harmonic or subharmonic of the breathing rate should be possible and would improve the accuracy to that shown in the final row of Table I.

The model order and downsampled frequency of the AR model have been optimised in this section for 5-minute sections of the PPG waveform from adult subjects, and may need to be altered for use with sections of different lengths or paediatric patients, whose higher breathing rate may necessitate a higher downsampled frequency. As with most methods, the accuracy of the results is likely to be correlated with the length of the signal being analysed, as longer signals will contain more breathing cycles. However, the longer the sections, the greater the delay before a breathing rate estimate is produced and the lower the ability to track changes in breathing rates.

IV. REAL-TIME BREATHING RATE EXTRACTION USING AR

For most clinical applications, the breathing rate needs to be tracked over the period of monitoring. Further investigations were carried out using the novel AR method described above to assess its utility for real-time tracking of breathing rate.

To ensure that changes in breathing rate can be tracked accurately with minimal delay, it is necessary to reduce the length of the sections used to calculate the breathing rate. To simulate real-time measurement, the five-minute sections of data were windowed into 30-second sections, with a 25-second overlap, so that the start times of consecutive windows

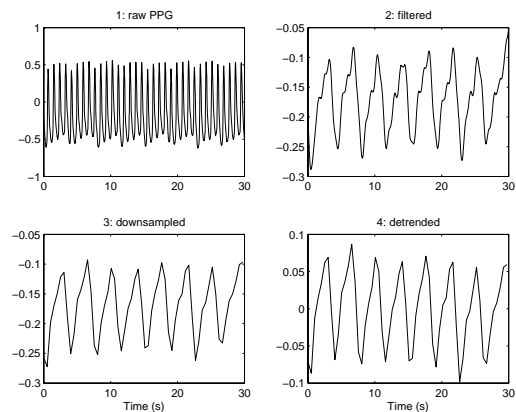


Fig. 6. Pre-processing of 30-second section of PPG from Patient A prior to AR modelling.




Fig. 7. Poles of AR model (9th order model, PPG downsampled to 2 Hz) on a 30-second section of PPG from Patient A, with breathing pole circled, and the limits of the sector of interest indicated by the dashed lines. The phase angle of the breathing pole corresponds to 0.26 Hz (16 breaths per minute)

differ by 5 seconds. The shorter window length reduces the amount of data available for estimating the parameters of the AR model (by a factor of 10), and so the accuracy of the breathing rate is likely to be degraded. The amount of available data can be increased by reducing the downsampling ratio, but the heart rate frequency then becomes more prominent in the downsampled signal, reducing the accuracy of the placement of the breathing pole. Poles due to the subharmonics of the heart rate can also appear in the breathing sector of the AR model, and may be of higher magnitude than the true breathing pole, leading to incorrect pole choice.

To allow the use of a lower downsampling ratio, pre-filtering of the PPG can be carried out to minimise the effect of the heart rate information. An FIR low-pass filter using the Kaiser windowing function was designed for this purpose. The filter has a transition band extending from 0.4–0.8 Hz (i.e. from 24–48 breaths or beats per minute). The upper frequency is low enough to ensure that all heart rate frequencies should lie outside the pass-band, without introducing excessive attenuation at adult breathing rates. The pass-band ripple was designed




Fig. 8. Poles of AR model (9th order model, PPG downsampled to 2 Hz) on a 30-second section of PPG from Patient B, with breathing pole circled, and the limits of the sector of interest indicated by the dashed lines. The phase angle of the breathing pole corresponds to 0.13 Hz (7.8 breaths per minute)

to be 5%, and the stop-band attenuation was chosen to be 30 dB. The filter requires 490 coefficients, which corresponds to a delay of less than 4 seconds at the 125 Hz sampling frequency used in the MIMIC dataset. After pre-filtering, the PPG is decimated and downsampled as before.

As the window size had been changed, and extra pre-processing had been introduced, it was necessary to repeat the tests to find an optimal combination of downsampled frequency and AR model order. This was determined experimentally in the same way as for the five-minute sections, and was found to be a 9th order model with a downsampled frequency of 2 Hz.

Figs. 6 and 7 shows how this model performs on a 30-second section of one of the PPG waveforms from the MIMIC dataset. The breathing pole is correctly placed and has high magnitude. As the downsampling ratio has doubled compared to that in Fig. 5, the angle of the sector of interest has decreased, and is no longer limited by the Nyquist criterion.

Artefacts in the PPG waveform will have a greater effect on the estimation of breathing rate in shorter windows of the PPG waveform, and this can lead to inaccurate breathing pole placement for some 30-second sections. A tracking method, such as Kalman filtering, could be used to reduce the error introduced in this way.

V. CONCLUSIONS

Current methods for extracting breathing rate from the PPG vary in their accuracy, and this accuracy is not necessarily correlated to the complexity of the method. A novel method using autoregressive model has been shown to significantly outperform the current methods, and could be used for real-time tracking of breathing rate using the PPG waveform.

Pulse oximeters are already used in clinical practice for continuous non-invasive monitoring of arterial oxygen saturation and heart rate. The addition of a measure of breathing rate to the pulse oximeter display could allow earlier recognition of deterioration in patients whose condition does not merit the use of a dedicated sensor to measure breathing rate, as their

breathing rate would otherwise only be measured manually and infrequently. It would also be possible to use data fusion methods on the three signals from the pulse oximeter to provide early warning if more than one of the measurements deviate from the norm [15].

REFERENCES

- [1] A. L. Goldberger, L. A. N. Amaral, L. Glass, J. M. Hausdorff, P. C. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C.-K. Peng, and H. E. Stanley, "PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals," *Circulation*, vol. 101, no. 23, pp. e215–e220, 2000, circulation Electronic Pages: <http://circ.ahajournals.org/cgi/content/full/101/23/e215>.
- [2] L. Nilsson, A. Johansson, and S. Kalman, "Respiration can be monitored by photoplethysmography with high sensitivity and specificity regardless of anaesthesia and ventilatory mode," *Acta Anaesthesiol. Scand.*, vol. 49, no. 8, pp. 1157–62, Sep. 2005.
- [3] L. Nilsson, T. Goscinski, A. Johansson, L. G. Lindberg, and S. Kalman, "Age and gender do not influence the ability to detect respiration by photoplethysmography," *J. Clin. Monit. Comput.*, vol. 20, no. 6, pp. 431–6, Dec. 2006.
- [4] K. Nakajima, T. Tamura, and H. Miike, "Monitoring of heart and respiratory rates by photoplethysmography using a digital filtering technique," *Med. Eng. Phys.*, vol. 18, no. 5, pp. 365–72, Jul. 1996.
- [5] P. Addison and J. Watson, "Secondary wavelet feature decoupling (SWFD) and its use in detecting patient respiration from the photoplethysmogram," in *Engineering in Medicine and Biology Society. Proceedings of the 25th Annual International Conference of the IEEE*, Sep. 2003, pp. 2602–5.
- [6] P. Leonard, T. F. Beattie, P. S. Addison, and J. N. Watson, "Standard pulse oximeters can be used to monitor respiratory rate," *Emerg. Med. J.*, vol. 20, no. 6, pp. 524–5, Nov. 2003.
- [7] P. S. Addison and J. N. Watson, "Secondary transform decoupling of shifted nonstationary signal modulation components: application to photoplethysmography," *Int. J. Wavelets Multiresolut. Inf. Process.*, vol. 2, no. 1, pp. 43–57, Mar. 2004.
- [8] P. Leonard, T. Beattie, P. Addison, and J. Watson, "Wavelet analysis of pulse oximeter waveform permits identification of unwell children," *Emerg. Med. J.*, vol. 21, no. 1, pp. 59–60, Jan. 2004.
- [9] P. Leonard, N. R. Grubb, P. S. Addison, D. Clifton, and J. N. Watson, "An algorithm for the detection of individual breaths from the pulse oximeter waveform," *J. Clin. Monit. Comput.*, vol. 18, no. 5–6, pp. 309–12, Dec. 2004.
- [10] P. A. Leonard, J. G. Douglas, N. R. Grubb, D. Clifton, P. S. Addison, and J. N. Watson, "A fully automated algorithm for the determination of respiratory rate from the photoplethysmogram," *J. Clin. Monit. Comput.*, vol. 20, no. 1, pp. 33–6, Feb. 2006.
- [11] P. A. Leonard, D. Clifton, P. S. Addison, J. N. Watson, and T. Beattie, "An automated algorithm for determining respiratory rate by photoplethysmogram in children," *Acta Paediatr.*, vol. 95, no. 9, pp. 1124–8, Sep. 2006.
- [12] D. Clifton, J. G. Douglas, P. S. Addison, and J. N. Watson, "Measurement of respiratory rate from the photoplethysmogram in chest clinic patients," *J. Clin. Monit. Comput.*, vol. 21, no. 1, pp. 55–61, Feb. 2007.
- [13] J. Pardey, S. Roberts, and L. Tarassenko, "A review of parametric modelling techniques for EEG analysis," *Med. Eng. Phys.*, vol. 18, no. 1, pp. 2–11, Jan. 1996.
- [14] S. Cazares, M. Moulden, W. G. Redman, and L. Tarassenko, "Tracking poles with an autoregressive model: a confidence index for the analysis of the intrapartum cardiocogram," *Med. Eng. Phys.*, vol. 23, no. 9, pp. 603–14, Nov. 2001.
- [15] L. Tarassenko, A. Hann, A. Patterson, E. Braithwaite, K. Davidson, V. Barber, and D. Young, "Biosign™: multi-parameter monitoring for early warning of patient deterioration," in *Proc. 3rd IEE International Seminar on Medical Applications of Signal Processing*, 2005, pp. 71–6.