

Detecting Heart-Rate while Jogging: Blind Source Separation of Gait and Heartbeat

Tuoi T. N. Vo¹, Mark McGuinness², Alan Hegarty³, Stephen O'Brien³, Kevin O'Sullivan³, and Sandra Healy⁴

¹MACSI, Department of Mathematics and Statistics, University of Limerick, Ireland.

Email: Tuoi.Vo@ul.ie

²School of Mathematics, Statistics and Operations Research, Victoria University of Wellington, New Zealand. Email: Mark.McGuinness@vuw.ac.nz

³MACSI, Department of Mathematics and Statistics, University of Limerick, Ireland

⁴Analog Devices, Raheen Industrial Estate, Limerick, Ireland

Abstract

A blind source signal separation problem that was brought to a Study Group in Limerick in 2013 required a way to prevent the gait of a jogger from masking the heart-beat, when detected by a simple photodiode that measures light transmission through a jogger's wrist tissues. The group was successful in discovering a singular value decomposition (SVD) approach, which not only allows accurate detection of heart-rate, but also allows recovery of a good facsimile of the entire blood pressure time series from the mixed photodiode signal.

Keywords: blind source separation, SVD, heart-rate monitor.

1 Introduction

We consider a blind source signal separation problem that was brought to the European Study Group with Industry held in June 2013 at the University of Limerick. The problem arises when considering how to detect heart rate in a simple non-intrusive way while walking or jogging.

Specifically, a very simple heart-rate monitor was considered, based on a single

wrist device similar to that pictured in Figure 1, that does not require a separate transmitting chest band. The device is to contain a single channel photodiode mounted against the skin so that light reaching it has to travel through wrist tissue, which detects changes in the intensity of the light that reaches it, and a processor chip to analyse the single channel signal and display the heart rate detected. The light reaching the photodiode is affected mainly by two things, variation in the perfusion of the tissues through which it passes, and variation in the incident intensity on the skin surface surrounding the monitor. The first effect is due to varying blood pressure, which causes a pulsatile perfusion of body tissues, hence a pulsatile variation in the absorption of the light that travels from skin through tissue to the photodiode. The second effect, which masks the first, is due to the gait of a person, walking or running perhaps, moving the wrist and causing the light levels incident upon the skin to vary rhythmically. The periods of these two effects are usually very close. Walking is typically about one step per second, and walking heart rate is about 60 beats per minute. Running gait may be three times as fast, and running heart rate is usually also similar to gait at about 180 beats per minute.



Figure 1: A heart rate monitor.

Although no data was available for the Study Group to assess, we were advised that in prototypes the gait was masking detection of heart rate, so that the instrument was reporting running rate rather than the desired heart rate. The Group was asked to investigate and assess ways to process the time series and separate the heart rate out of the combined signal coming from the photodiode.

The problem then is to separate two time series signals with the same or similar periods but with different characteristic structures or wave-shapes in time, when they are presented as a single combination. The input thus comprises a single combined signal, and the required output is the recovery of the two initial signals. In the literature this is referred to as *blind source separation*. Some information about the sources is needed to do this. For example, if one voice is deeper than the

other, it is possible to model the footprint of each voice in the frequency domain by assuming a certain shape to the power spectrum, fit the shapes to the observed power spectrum, and subtract out each voice in turn [9]. Or if a fetal heart rate is higher than the mother’s heart rate, this can be leveraged to extract it from a mixed signal obtained from a passive sound sensor on the mother’s abdomen [1].

It is noted in Syed *et al.* [12] that blind source separation problems are considered tractable if they involve linear mixtures of the sources, and that it is necessary to know something about the structure of the sources or something about the mixing matrix that combines them. In particular, if the sources are statistically independent random signals, Independent Component Analysis (ICA) is effective, or if the sources are nonnegative, matrix factorisation is effective, or if the sources are sparse, a sparse component analysis can be effective.

In a review of blind source separation (BSS) methods, Jutten *et al.* [7] note that linear problems with plain random sources with no temporal structure are routinely solved using ICA. However where there is temporal structure (as here), or worse, nonlinear combinations of sources (as here), BSS problems require additional prior information. Indeed, they point out that BSS “is simply impossible for general nonlinear transformations by resorting to statistical independence only without additional constraints or prior information”. Furthermore, their development is for random signals, which is not a good description of our problem.

Gao *et al.* [6] consider separating a single signal into foetal and maternal ECG sources, and use a combination of approaches. They first compute a spectrogram of the combined signal, then use Singular Value Decomposition (SVD, described in more detail later) to smooth the spectrogram by retaining the first k singular values and left and right eigenvectors. Then they use ICA on the right singular eigenvectors, generate new left eigenvectors, and use ICA on those separately. They note that this involves the crucial assumption that the spectral components of foetal and maternal signals are independent random variables. Then they select the two most significant new components and test how well they represent the maternal and foetal blind sources. Since they are working in frequency space (evolving slowly with time), they are leveraging the different spectra of the two sources – foetal ECG is typically higher frequency than maternal. They note that straight SVD does not do as well at their combination approach, at detecting foetal heartbeats. Their method does not reproduce a foetal source component, and only broadly locates foetal heartbeats in time.

A similar approach is taken by Murata *et al.* [10], who seek to separate voices in a conversation recorded by one microphone. They start with the windowed Fourier

transform, then seek to diagonalise the correlation functions. These need to be weakly stationary.

The position of ICA as a new paradigm for BSS, and issues involving stability and convergence, and the number of signals having to be greater than or equal to the number of sources (which is not the case here), and ways around the need for linearity, are discussed in the special issue of *Neurocomputing* edited by Amari *et al.* [2].

After the work discussed in this paper was done, we were alerted to a little-known paper by Fowler *et al.* [4], which recommends a similar procedure to the one we find here to be so effective. However, Fowler *et al.* were investigating a different problem, to filter different oscillatory mechanisms in data from the respiration of children in quiet and REM sleep. The filtering is to assist in diagnosis of ailments and to warn of incipient sudden infant death syndrome. In their work, a careful choice of lag is made when embedding data in a phase space, then SVD is applied to smooth the data in preparation for diagnosis. While the approach is similar in style to ours, there is no attempt or need to separate blind source signals in [4].

In the remainder of this paper, we will indicate how we overcame the shortcoming of there being no data available by creating an ersatz set of data incorporating two sources (gait and blood pressure), we discuss embedding our ersatz data in a high dimensional phase space, and we show how the use of SVD can not only prevent the gait source from masking heart rate, but how we can recover a remarkably good facsimile of the original blood pressure source.

2 Ersatz Data

The group found some high-resolution arterial blood-pressure data that would serve as a proxy for perfusion of tissue (Figure 2) [5]. The Beer-Lambert law for light transmission through an absorbing medium was then used together with a simple assumed formula for the effect of gait on incident light intensity, to manufacture an ersatz data set that we could use to test our ideas on separating the heart signal from the gait signal.

2.1 The Beer-Lambert law

The Beer-Lambert law gives the transmitted intensity of light $I(t)$ when light of intensity I_0 passes through an absorbing medium. Here we modify it slightly to

relate transmitted light intensity to blood pressure, as

$$I(t) = I_0(t)e^{-\lambda P(t)}, \quad (1)$$

where $P(t)$ is blood pressure, and λ is to convert the blood pressure to effective absorption by the tissue.

Figure 2 shows a typical arterial blood pressure signal, $P(t)$, on the left and a signal generated from P using equation (1), in the particular form

$$I(t) = [1 + 0.05 \sin(7t)] e^{-\frac{0.05P(t)}{\max(P(t))}},$$

on the right. The period of the sine function has been chosen to approximately match the average frequency of the blood pressure data $P(t)$, and the amplitudes in the exponent and of the sine are motivated by anticipated signal to noise ratios of 5%. Our study is based on this ersatz data set. It has the advantage that we know the input signals, so we can assess how well our blind-source separate technique works.




Figure 2: A typical arterial blood pressure signal, $P(t)$, on the left [5] and an ersatz signal generated from P using the Beer-Lambert Law, $I(t) = [1 + 0.05 \sin(7t)] e^{-\frac{0.05P(t)}{\max(P(t))}}$, on the right.

2.2 Spectral Characteristics

The group was interested in whether a spectrogram might indicate characteristics of the combined signal in the frequency domain, that differed in some way from the

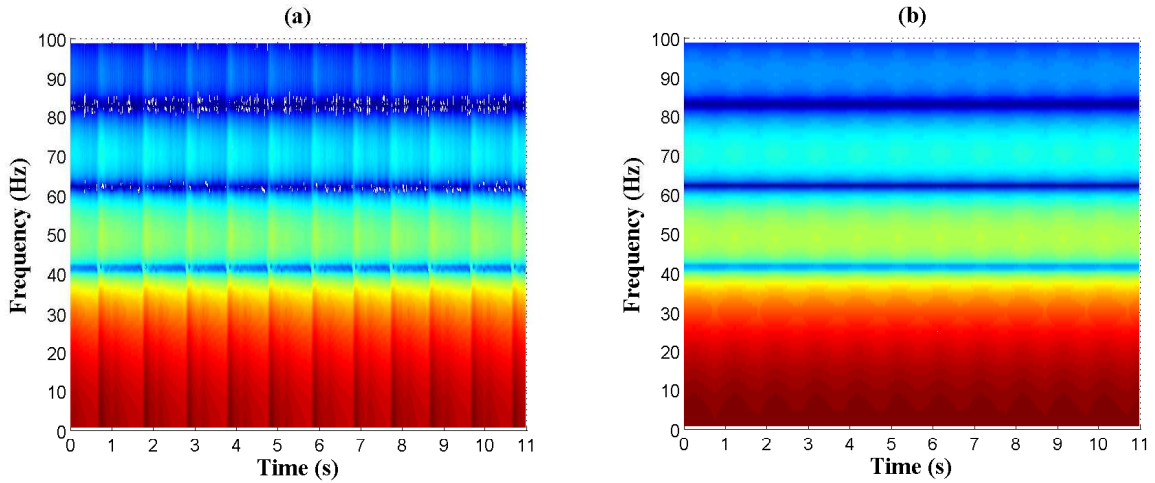


Figure 3: (a) A spectrogram of the blood pressure signal, $P(t)$, and (b) a spectrogram of the light intensity from the Beer-Lambert law $I(t)$.

pressure signal. Figure 3 (a) displays a spectrogram of light intensity corresponding to light intensity $I(t)$. Figure 3 (b) shows a spectrogram of the blood pressure signal, $P(t)$.

It is difficult to conclude anything from these spectrograms. The periods in both do not appear to vary much, although the blood pressure period does appear to vary from beat to beat in the time series plot. There is a more apparent jump in the colors of the blood pressure spectrogram at times that correspond to the rapid rise in pressure associated with systole, compared to the gentler sine-dominated oscillations of $I(t)$.

Spectral techniques are designed for linear systems, and rely on a frequency difference. So next we attempted a phase plane reconstruction [11], to see how we could leverage differences in temporal structure of the smooth gait and sharper blood pressure. Such techniques are motivated and based on the theory of dynamical systems and the time evolution is defined in an appropriate phase space.

2.3 Phase Space Reconstruction

Suppose that we have a sequence of measurements made on a system which is evolving on a finite dimensional attractor - does this time series carry information about the attractor dimension? In general it does and in fact it carries much more information. The established approach is to embed the time series onto a trajectory in a finite dimensional space. It has been demonstrated under quite

general circumstances that the reconstructed attractor is topologically equivalent to the attractor in the unknown space in which the original attractor is living. The particular method used here is the method of delays, based on the idea of a delay register [3, 11].

To do this we choose an embedding dimension and a time lag or delay. Embedding theorems guarantee faithful reproduction of the attractor if the embedding dimension is larger than twice the number of active degrees of freedom, regardless of how large the dimensionality of the true space is. The delay is not the subject of the embedding theorems since they consider data with infinite precision. In practice, the time delay must be found by experimentation: no rigorous way of establishing its optimal value has been determined [4, 8].

If the delay is small compared to the intrinsic time scales of the system, successive elements of the delay vectors are strongly correlated. However, if it is too small, there is almost no difference between the different elements of the delay vectors; if it is large enough, the different coordinates may be almost uncorrelated, or independent, providing a topologically correct view of dynamical behaviour. Delays are fed into the register and propagate sequentially until they are lost at the other end n clock cycles later. At any instant the register thus contains n consecutive data values $(v_i, v_{i-1}, \dots, v_{i-n})$.

Each sequence of data points can be thought of as an n -dimensional vector, usually written as a column vector. The sequence of n -vectors generated by clocking the data through the delay register can be thought of as a discrete trajectory in an n -dimensional euclidean space. If the signal arises from a finite-dimensional deterministic system, then provided the embedding dimension is sufficiently large, the trajectory gives a genuine image of the dynamics in its own phase space.

We then reconstructed a two-dimensional phase space view of our data in this way, using the ersatz values of $I(t)$ created using the Beer-Lambert law. The results appear in Figure 4, with a choice of lag of about 1/4 of the main period present in the signal, which also is close to the first zero of the autocorrelation of $I(t)$. The large D-shaped structure we associate with gait, while the small loops seen at various places on the straight line part of the D we associate with heart-beats.

The appearance of this phase-space reconstruction led us to think about ways to remove the larger D-shaped behaviour from the signal, then leaving behind mostly the effects of the blood pressure. The results of the remarkably effective SVD method are detailed in the following sections.




Figure 4: A two-dimensional phase-space reconstruction of $I(t)$, using a lag k of about $1/4$ of the main period present, and obtained by plotting I_{n+k} vs I_n for $k = 1, 2, \dots$

3 Singular Value Decomposition (SVD)

The set of all delay vectors forms an ellipsoidal cloud in the euclidean space and we wish to establish and then remove the two most important directions associated with the data, as these will correspond to the D shape seen in Figure 4. The procedure continues by subtracting off the mean value from each column vector creating a new set x_i and then diagonalising the covariance matrix

$$\frac{1}{N} \sum_{i=1}^N x_i x_i^T.$$

This is related to the singular value decomposition of a trajectory matrix M which has rows consisting of the row vectors x_i^T . The covariance matrix is actually the product $M^T M$ and the right singular vectors are the eigenvectors of the covariance matrix. The matrix is real and symmetric, hence its eigenvalues are real and its eigenvectors are orthogonal. Mathematically the singular values of M are the square roots of the eigenvalues of the covariance matrix.

Finding the singular vectors corresponds to finding the semi-major axes of the data cloud generated in the embedded phase space by the trajectory, and the singular values are then related to the lengths of these axes. Thus the singular vectors give a geometric description of where the trajectory lies, and the singular values

are a measure of the trajectory “surface” in the corresponding directions. The most relevant directions in space relate to the vectors corresponding to the largest eigenvalues; directions associated with small eigenvectors may be neglected.

The most relevant directions in space are given by the eigenvectors corresponding to the largest eigenvalues. The group thought that the SVD might then identify the two most important axes as those associated with the D-shape seen in two dimensions, and allow us to remove those dimensions from the data.

The first step in using SVD is to create a matrix M , each row of which is the vector $(I_{k+1}, I_{k+2}, \dots, I_{k+n})$, and the first row having $k = 0$, the second $k = n + 1$, the third $k = 2n + 1$, etc. We generated as many rows as the data allowed, say m rows. Then each row is averaged separately, and the average is subtracted from that row, to create a normalised M matrix. This moves the origin in phase space to the middle of the ellipsoid.

The SVD decomposes the normalised matrix M as

$$M_{mn} = U_{mm} S_{mn} V_{nn}^T,$$

where $U^T U = I, V^T V = I$; the columns of U are orthonormal eigenvectors of MM^T ; the columns of V are orthonormal eigenvectors of $M^T M$; and S is a diagonal matrix containing the singular values of M in descending order.

In order to remove the D-shape seen in two dimensions, we set the two biggest singular values in S to zero. We then compute a new M matrix using the modified S matrix, and the same matrices U and V as obtained from the original M matrix. This new M matrix then gives a time series that should represent or capture the notches in $I(t)$.

Figure 5 shows the blood pressure data, the gait data, the composite signal $I(t)$, the result (the residual) of removing just the most singular value from the time series by setting the first singular value in S to zero, and the result of also removing the second most singular value from $I(t)$. Note that the remaining time series has a pulse of activity located exactly at each place that blood pressures rise to systole. Figure 6 shows the result of repeating this analysis when the gait has a period that is about 50% longer than before. SVD seems to work just as well, in locating the places where blood pressure rises, that is, in finding heart rate.

Figure 7 shows more detail of just the blood pressure curve and the result of removing the two most singular values, otherwise the same case as in Figure 5. Figure 8 shows a close-up view of the blood pressure signal over one period, superposed with the combined $I(t)$ signal and the result of removing the two most singular values from that signal. Note the increased amplitude of the oscillations in this “after SVD” data, at the place where blood pressure rises suddenly to systole.




Figure 5: SVD numerical results using a gait of the same average period as blood pressure. The first curve is blood pressure, the second is the gait, the third is the recorded signal $I(t)$, the fourth is the result of removing the most singular value from $I(t)$, and the last is the result of also removing the next most singular value from $I(t)$.




Figure 6: SVD numerical results as in Figure 5, except that the period of the gait effect has been increased by about half.




Figure 7: Comparison of original blood pressure signal and the result of removing the two most singular values from the SVD of the combined signal $I(t)$. This is a different view of the same case as in Figure 5.

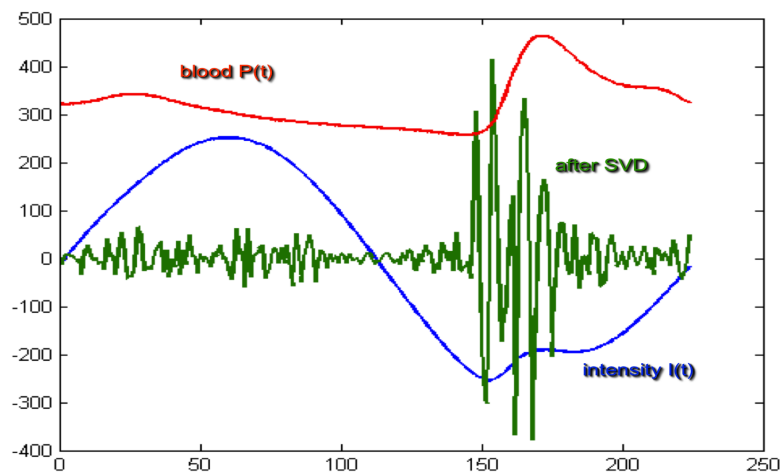


Figure 8: A close-up look at the blood pressure $P(t)$, the combined signal $I(t)$, and the SVD result after removing the two most singular values, over a period of one heartbeat. This is a different view of the same case as in Figure 5. The x-axis represents a time index for time $t = 0.1$ second, and the y-axis represents re-scaled values of three signals $P(t)$, $I(t)$ and the SVD result.

In fact, we discovered it is possible to recover curves that look just like the original blood pressure data, by first rectifying and then smoothing the residual signal obtained after removing the two most singular values. As illustrated in Figure 9, the resulting signal looks a lot like the original blood pressure data before modifying with the Beer-Lambert law. That is, buried in the apparently noisy SVD result is the original blood pressure data. Removing the two most singular values from the composite signal $I(t)$, then rectifying and smoothing, gives a filtered signal that looks very like the desired blood pressure curve, and will give correct heart-rate values. More extensive testing is required to be sure of the general efficacy of this approach, but it looks very promising and is visually compelling. We found that it is not sensitive to the length of vectors n chosen to make the matrix M from the time series.

4 Conclusions

A point of difference is needed, either in the time domain or in the frequency domain, to separate blind sources. For example, different frequency footprints can be used to reconstruct signals (male/female voices), or different appearance in the time domain, like notches, give some hope of distinguishing two mixed signals.

Our use of the SVD appears to be very effective in the time domain, for extracting a time series from the Beer-Lambert law data, that looks very similar to the blood pressure data used as one of the sources when generating the data for light intensity $I(t)$ as measured by a photodiode. Apparently removing the two most singular values in the SVD approach can recover the blood pressure signal, with no specific prior knowledge of shape required. This approach needs to be explored further, with other shapes used for gait, and with plethysmograph data for perfusion of finger tissue, or even better with actual data obtained from a photodiode worn by an active person.

This approach does not require linearity in the way that signals are combined, or stationarity, or that the signals be independent random processes, or that the periods of the two signals be very different, just that the nature of the blind source signals be temporally different. In principle it can be extended to separating more than two sources, provided that two singular values are used to remove each source in turn. This in itself is an assumption about the nature of the sources, that they be essentially two-dimensional in nature.

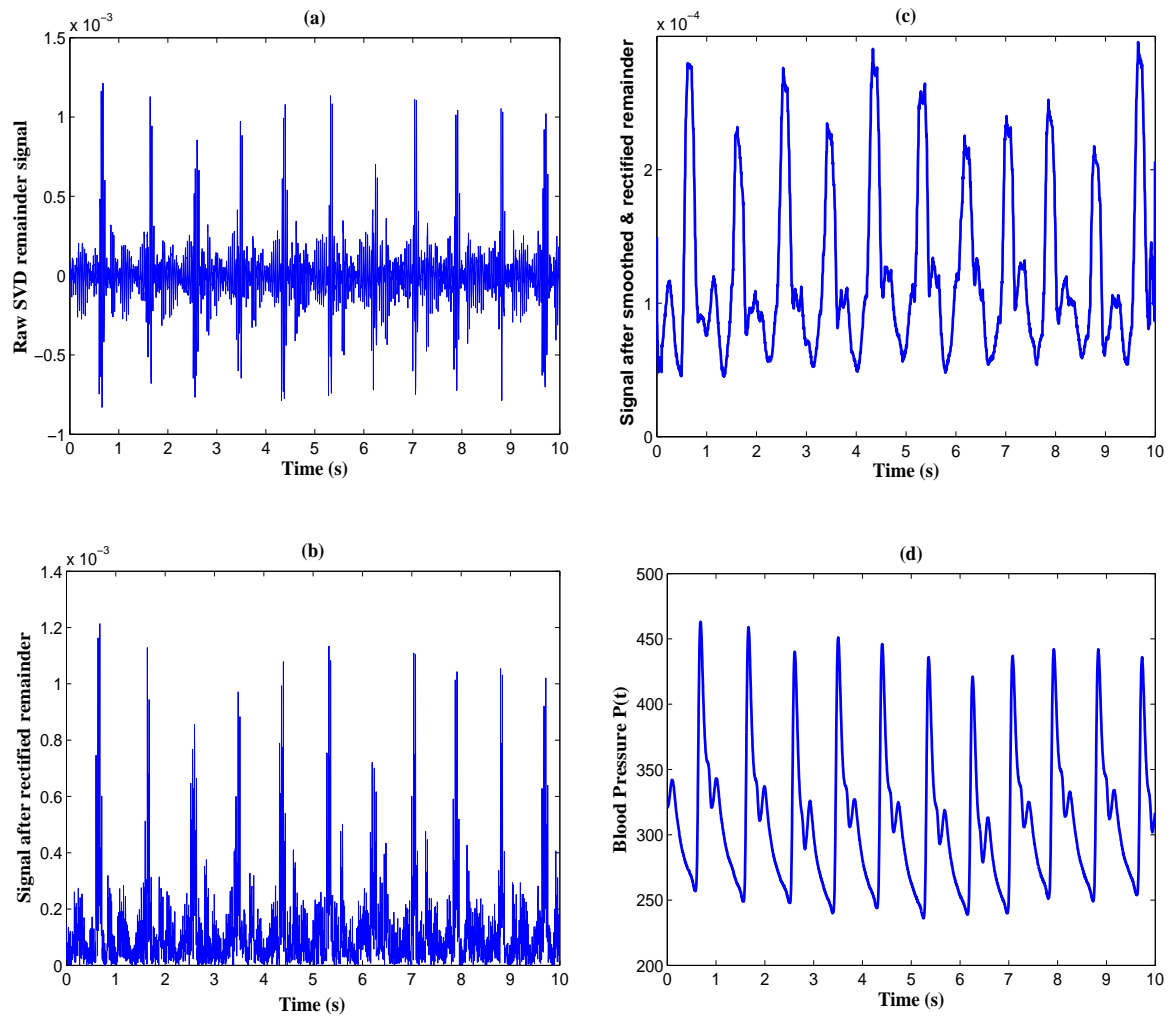


Figure 9: (a) The raw remainder signal after removing the two most singular values from $I(t)$, (b) The rectified signal (a), (c) The smoothed signal (b), and (d) The original blood pressure signal, $P(t)$. Note the similarity between the original blood pressure signal and the smoothed rectified SVD result.

Acknowledgements

We are grateful for the support of the Mathematics Applications Consortium for Science and Industry (www.macsi.ul.ie) funded by the Science Foundation Ireland (SFI) investigator award 12/IA/1683 and the SFI conference and workshop grant 13/CW/12578.

References

- [1] R. Acharyya, N. L. Scott, P. L. Teal, *Non-invasive foetal heartbeat rate extraction from an underdetermined single signal*, *Health* **1**(2): 111-116, 2009.
- [2] S. -I. Amari, A. Hyvarinen, S. -Y. Lee, T. -W. Lee, and V. D. Sanchez A, *Blind signal separation and independent component analysis*, *Neurocomputing* , **49**: 1-5, 2002.
- [3] D. S. Broomhead and G. P. King, *Extracting qualitative dynamics from experimental data*, *Physica* **20D**: 217-236, 1986.
- [4] A. C. Fowler, G. Kember, P. Johnson, S. J. Walter, P. Fleming, and M. Clements, *A method for filtering respiratory Oscillations*, *Journal of Theoretical Biology*, **170**: 273-281, 1994.
- [5] A. L. Goldberger, L. A. N. Amaral, L. Glass, J. M. Hausdorff, P. Ch. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C. K. Peng, H. E. Stanley, *PhysioBank, PhysioToolkit, and PhysioNet: Components of a New Research Resource for Complex Physiologic Signals*, *Circulation* **101** (23): e215-e220, 2000. <http://physionet.org/physiobank/database/fantasia/>.
- [6] P. Gao, E. Chang, and L. Wyse, *Blind separation of fetal ECG from single mixture using SVD and ICA*, *ICICS-PCM*, 1418-1422, 2003.
- [7] C. Jutten and J. Karhunen, *Advances in blind source separation (BSS) and independent component analysis (ICA) for nonlinear mixtures*, *International Journal of Neural Systems*, **14** (5): 267-292, 2004.
- [8] H. Kantz and T. Schreiber, *Nonlinear Time Series Analysis*, Cambridge University Press, 1997.
- [9] T. Kristjansson, H. Attias, and J. Hershey, *Single Microphone Source Separation Using High Resolution Signal Reconstruction*, in *Proc. ICASSP II* 817-820, 2004.

- [10] N. Murata, S. Ikeda, and A. Ziehe, *An approach to blind source separation based on temporal structure of speech signals*, Neurocomputing, **41**: 1-24, 2001.
- [11] D. Ruelle and F. Takens, *On the nature of turbulence*, Communications in Mathematical Physics, **20** (3): 167-192, 1971.
- [12] M. N. Syed, P. G. Georgiev, and P. M. Pardalos, *A hierarchical approach for sparse source blind signal separation problem*, Computers & Operations Research **41**: 386-398, 2014.