Blind Separation of Maternal and Fetal ECG Recordings using Adaptive Sparse Representations

Thomas Blumensath, Mike E. Davies
The University of Edinburgh
IDCOM & Joint Research Institute for Signal and Image Processing
King's Buildings, Mayfield Road, Edinburgh EH9 3JL, UK
thomas.blumensath,mike.davies@ed.ac.uk

ABSTRACT

In this paper we present a method to blindly separate maternal and fetal ECG signals using recordings made from the maternal abdomen. The method is general and is able to separate the signals using any number of recording electrodes, including the difficult case of single channel blind source separation. The approach is based on the assumption that the signals admit a sparse and shift-invariant representation. Adaptation of the representation then leads to the emergence of the PQRST complexes of the maternal and the fetal heart for each electrode so that we can separate the individual sources by reconstructing the signal using only a single PQRST complex.

Keywords: Sparse coding, fetal ECG, Single Channel Blind Source Separation

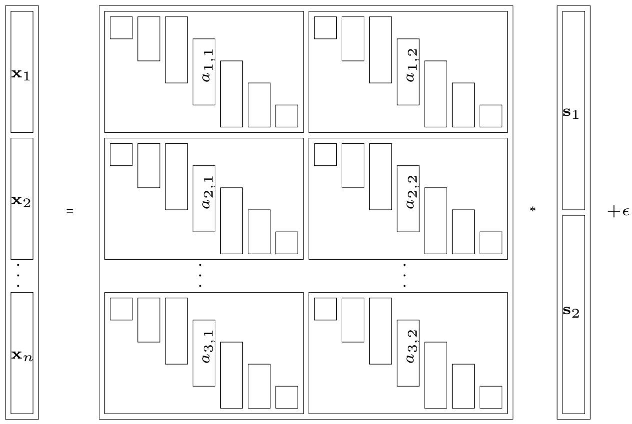

Figure 1: Graphical representation of the structures in the equation $\mathbf{x} = \mathbf{A}\mathbf{s} + \epsilon$.

1 INTRODUCTION

The electrocardiogram (ECG) is a valuable diagnostic tool to assess heart conditions. Details on general ECG diagnostics can be found in [4] and a description of pathological ECG data is given in [5] and [6]. It is often desired to record fetal ECG's which contain important diagnostic information for prenatal medical assessment [7]. Unfortunately, non-invasive techniques are currently not available to directly record fetal ECG's. Instead, recordings are taken from the maternal abdomen. However, the recorded signals are then contaminated by the signal from the maternal heart, which in general has a much stronger amplitude than the signal from the fetal heart. In addition, the recordings are often contaminated by substantial noise.

In this paper we present a method that can separate the signals from the fetal and maternal hearts and also removes the noise. The characteristic waveform associated with each individual heartbeat recorded by a particular ECG electrode is called a PQRST complex [4]. The shape of this waveform together with the information about the relative position between the electrode and the heart contain important diagnostic information. As the PQRST complex changes depending on the relative position between the recording electrode and the heart, it is important to separate the signals for each electrode location.

2 ALGORITHM

We assume that the timing and strength of maternal and fetal heartbeats can be represented as two sparse and positive time-series, which are convolved with the associated PQRST complexes. The PQRST complexes are different for each recording electrode so that we have twice the number of PQRST complexes than electrodes, i.e. we have different PQRST complexes for the maternal and the fetal heart for each of the recording electrodes. The observation at each of the recording electrodes is the summation of the two sparse time-series, each convolved with the PQRST complex for that electrode, and a noise term.

We denote the vector composed of the concatenation of the observation time-series as $\mathbf{x}$, the vector composed of the concatenation of the two sparse time-series as $\mathbf{s}$ and the error term as ϵ. The model can then be written as $\mathbf{x} = \mathbf{A}\mathbf{s} + \epsilon$, where the matrix $\mathbf{A}$ has a certain structure and contains the information of all PQRST complexes $\mathbf{a}_{i,j}$. This is shown graphically in figure 1. This additional structure is a result of the assumed convolutive model and the way in which the different time series have been concatenated.

We define the following probabilistic model:

$$p(s_n|u_n) = u_n\, p(s_n; \mu\sigma_R^2) + (1 - u_n)\delta_0(s).$$

Here s_n is used to denote any one sample from the time-series $\mathbf{s}$ and u_n is an indicator variable. We use $\delta_0(s_n)$ to denote the unit mass at zero, so that if u_n is zero, then so will be s_n. The non-zero coefficients s_n model the amplitudes of the heartbeats, which are assumed to follow the modified Rayleigh distribution

$$p(s_n; \mu, \sigma_R^2) = \frac{1}{Z_R} s e^{-(s_n - \mu)^2/2\sigma_R^2},$$

which was introduced in [2]. Note that the modified Rayleigh distribution is a conjugate prior for the Gaussian mean. The indicator variable u_n is assumed to follow the Bernoulli distribution:

$$P(u_n) = Z^{-1} e^{-0.5\lambda_u u_n}, u_n \in \{0, 1\}.$$

Finally, ϵ is assumed to be i.i.d. Gaussian noise with scale parameter λ_ϵ

To estimate the PQRST complexes we use a gradient method to estimate the maximum of the marginalised posterior $p(\mathbf{A}|\mathbf{x}) = \int p(\mathbf{A}, \mathbf{s}|\mathbf{x})\, d\mathbf{s}$ such that $\mathbf{A}$ has the necessary structure.

The gradient is then:

$$\Delta\mathbf{A} = \int \left[\frac{\partial \ln p(\mathbf{x}|\mathbf{s}, \mathbf{A})}{\partial \mathbf{A}} + \frac{\partial \ln p(\mathbf{A})}{\partial \mathbf{A}}\right] p(\mathbf{s}|\mathbf{A}, \mathbf{x})\, d\mathbf{s}, \tag{1}$$

where we use the notation $\frac{\partial}{\partial A}$ to denote the matrix of derivatives with respect to the individual values of the PQRST complexes $\mathbf{a}_{i,j}$. We assume that the prior $p(\mathbf{A})$ is relatively flat so that the second gradient is zero. However, evaluation of equation (1) involves non-standard integration and we approximate the gradient using a Gibbs sampling Monte Carlo estimate. Note that there is a scale ambiguity between the norm of $\mathbf{a}_{i,j}$ and the elements in $\mathbf{s}$. We therefore fix the variance of $\mathbf{s}$ a priori.

Sampling from $p(\mathbf{s}|\mathbf{A}, \mathbf{x})$ is done by sampling from $p(u_n, s_n|u_{\hat{n}\neq n}, s_{\hat{n}\neq n}, \mathbf{A}, \mathbf{x})$ iteratively for all n. We use u_n and s_n to refer to the n^{th} element of $\mathbf{u}$ and $\mathbf{s}$ respectively, while $u_{\hat{n}\neq n}$ and $s_{\hat{n}\neq n}$ are the sets of indicator variables and coefficients other than u_n and s_n.

To sample from $p(u_n, s_n|u_{\hat{n}\neq n}, s_{\hat{n}\neq n}, \mathbf{A}, \mathbf{x})$ we can sample alternately from $p(u_n|u_{\hat{n}\neq n}, s_{\hat{n}\neq n}, \mathbf{A}, \mathbf{x})$ and $p(s_n|\mathbf{u}, s_{\hat{n}\neq n}, \mathbf{A}, \mathbf{x})$. Note that the first distribution is the distribution of the n^{th} indicator variable conditional on all other variable apart from the associated coefficient s_n. The second distribution $p(s_n|\mathbf{u}, s_{\hat{n}\neq n}, \mathbf{A}, \mathbf{x})$, is the coefficient distribution conditioned on all other variables. This distribution is either a delta function at zero (if $u_n = 0$) or a modified Rayleigh distribution (if $u_n \neq 0$). Mathematical details and a method to draw samples from the modified Rayleigh distribution are presented in [2]. A thorough study of several Monte Carlo methods for a similar model can also be found in [1].

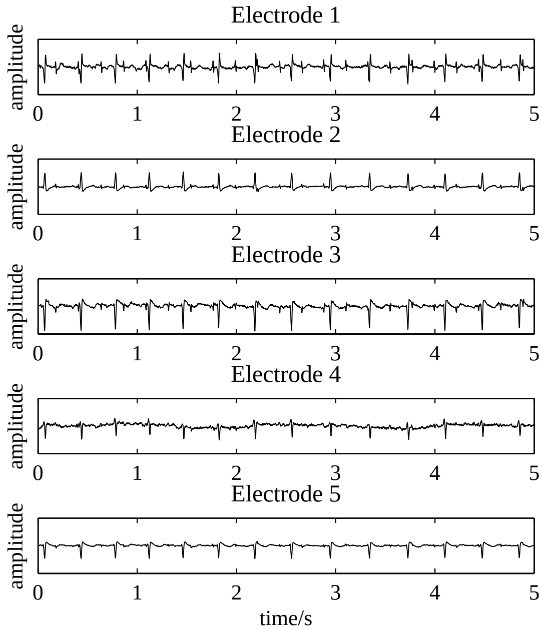

Figure 2: The five signals recorded from the maternal abdomen.

Table 1: Sketch of the learning algorithm

Input:
Signal $\mathbf{x}$, number and length of $\mathbf{a}_{i,j}$, number of sensors
Output:
$\mathbf{A}$, estimate of $\mathbf{s}$
Loop through next 3 steps till convergence of $\mathbf{A}$
Draw set of samples from $p(u_n, s_n|u_{\hat{n}\neq n}, s_{\hat{n}\neq n}, \mathbf{A}, \mathbf{x})$
Estimate gradient using Monte Carlo estimate
Update $\mathbf{A}$ with gradient step and normalise $\mathbf{A}$
Draw set of samples from $p(u_n, s_n|u_{\hat{n}\neq n}, s_{\hat{n}\neq n}, \mathbf{A}, \mathbf{x})$
Use samples drawn to estimate $\mathbf{s}$

After estimating $\mathbf{a}_{i,j}$, it is then possible to draw more samples from $p(\mathbf{s}|\mathbf{x}, \mathbf{A})$ from which mean or MAP estimates of $\mathbf{s}$ can be calculated.

A sketch of the algorithm can be found in table 1.

Note that this model can be seen as a shift-invariant version of the noisy overcomplete ICA model where the shift-invariant structure leads to the constraints on the mixing matrix. The source signals are here the sparse time-series and not the fetal and maternal ECG signals themselves. However, a reconstruction of the fetal and maternal ECG signals is possible for each electrode by convolving the estimated sparse time-series with each of the estimated PQRST complexes.

Here we have not addressed the problem of specifying or adapting other model parameters, such as λ_ϵ, μ, σ_R^2 and λ_u. This can also be done using gradient based optimisation, however, not all model parameters might be uniquely identifiable. In particular a trade off between sparsity, as defined by λ_u and noise variance λ_ϵ^{-1} has to be defined. In the experiments reported below we fixed the parameters and used, $\lambda_\epsilon = 100$, $\lambda_u = 10$, $\sigma_r^2 = 2$ and $\mu = 0$.

In early experiments we found that a direct implemen-

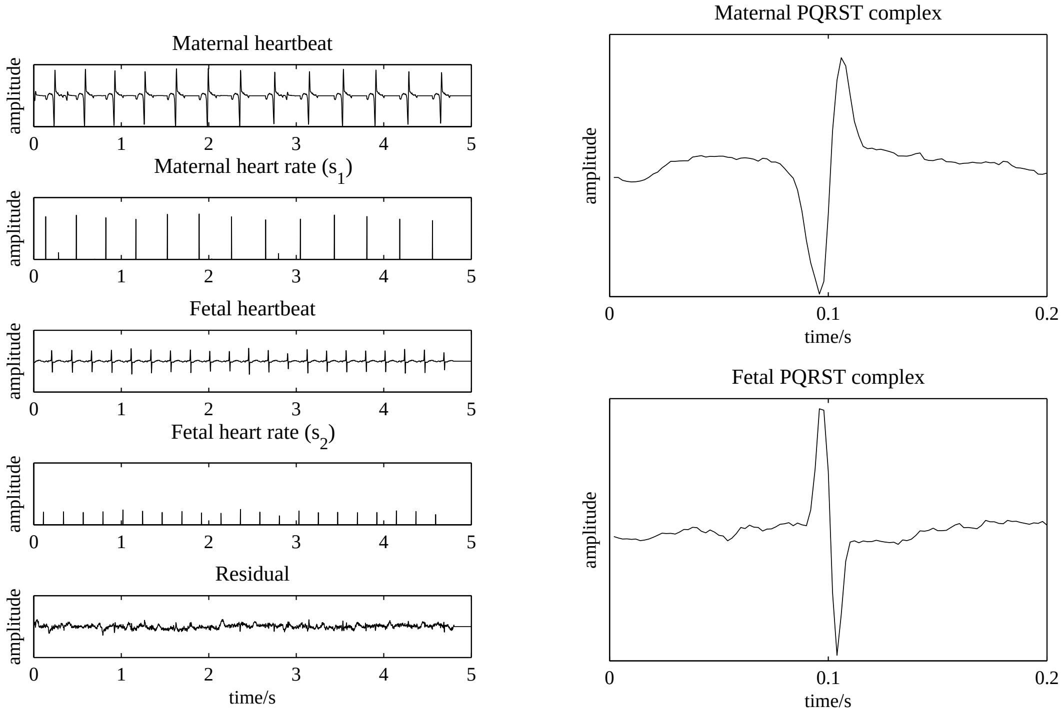

Figure 3: Separation from a single channel. Left from top to bottom: the separated maternal ECG sequence, the amplitude and position of heartbeats, the fetal ECG sequence, the amplitude and position of the fetal heartbeats and the residual. On the right are the maternal PQRST complex (top) and the fetal PQRST complex (bottom).

tation of the method detailed above was not able to find both, maternal and fetal, PQRST complexes. It was observed that this was due to the iterative and competitive learning strategy. The maternal PQRST complex is much stronger than the fetal one and the method is able to adapt and extract the maternal PQRST complex very fast. However, the maternal and fetal PQRST complexes can be very similar and once the maternal PQRST complex has been learned, the method often fails to learn the fetal PQRST complex. This problem is here avoided by adapting the learning rate for the maternal and fetal PQRST complex depending on the strength of the fetal and maternal contribution. This is done by weighting the learning rate by the inverse of the sum of the sparse time-series associated with the maternal and fetal heart respectively.

3 EXPERIMENTAL EVALUATION

To analyse the performance of the method we use a set of five signals recorded from the abdomen of a pregnant woman as shown in figure 2 (this is the data-set from [3]). The same data-set has previously been used in [8], so that our results are directly comparable to those reported in that paper. Note however that contrary to our method the method in [8] used eight channels. Three additional channels were recorded from the thoracic region and only contained maternal ECG contributions.

3.1 Single Channel Blind Source Separation

The main advantage of the newly proposed approach is its use with single channel recordings. To analyse the performance on a single channel we have run the method using individual channels from the data set introduced above.

We found that we could extract fetal PQRST complexes from all channels but channel (4), which has a low signal to noise ratio, a low contribution of the fetal signal and baseline wandering.

The best results were obtained from electrode (1). These results are shown in figure 3. Note that the maternal and fetal PQRST complex are quite dissimilar. However, similar results were obtained for the other electrodes.

3.2 Multi Channel Blind Source Separation

Using all five signals it is possible to separate all of the channels. The separated signals are shown in figure 4. As can be seen, separation using all five channels allowed us to also separate channel (4). However, the baseline wander in channel four introduces amplitude modulation in all fetal source reconstructions. This can be avoided by filtering channel (4) to remove the low frequency baseline wander. The results achieved by incorporating this preprocessing are shown in figure 5.

4 DISCUSSION AND CONCLUSION

In this paper we have presented an algorithm able to separate maternal and fetal ECG's from abdominal recordings. The method is able to use any number of electrodes and was shown to work even for the difficult single channel case. This is an important improvement to previous methods such as those in [8], which require at least 6 sensors.

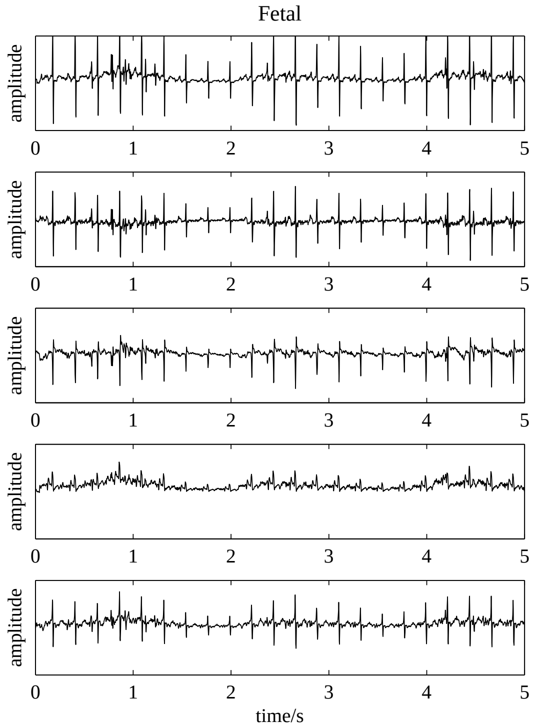

Figure 4: After separation we have clear fetal ECG signals, however, the baseline wander introduces modulation over all channels.

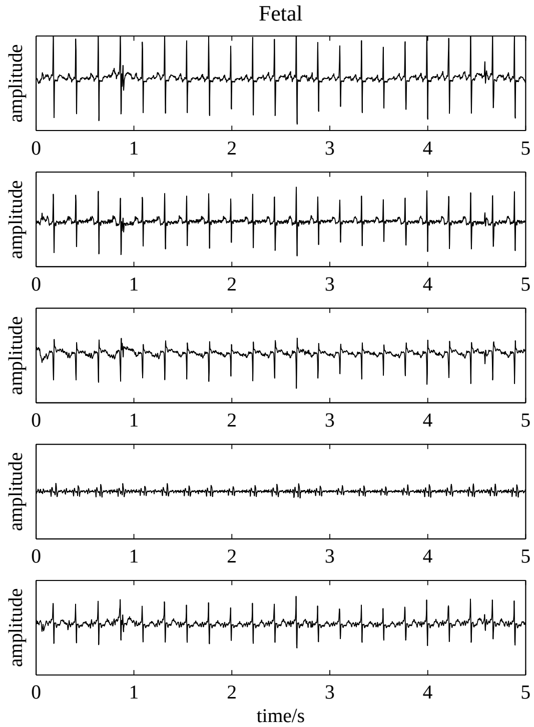

Figure 5: The separated fetal ECG signals after removal of baseline wander.

Otherwise, our results are comparable to previous results. The problem of baseline wander, which was addressed in [8] by explicitly modelling it as an extra source, further increasing the required number of sensors, was here addressed using standard pre-processing techniques such as filtering.

Our method not only separates the individual channels but also estimates the characteristic PQRST complexes for each heart and electrode. As these PQRST complexes contain much of the diagnostic information, such an estimation is of advantage. However, for some pathological heart conditions, the PQRST complexes change significantly between individual heartbeats. This behaviour is not modelled here. However, such deviations from the characteristic PQRST complexes will be present in the residual term. Also, for heart conditions for which individual PQRST complexes can show two different shapes, one normal and one pathological, the method presented here could be extended by including and learning an additional PQRST complex for the pathological heartbeats.

ACKNOWLEDGEMENTS

This research was supported by EPSRC grant D000246/1. MED acknowledges support of his position from the Scottish Funding Council and their support of the Joint Research Institute with the Heriot-Watt University as a component part of the Edinburgh Research Partnership.

References

[1] T. Blumensath and M. E. Davies. Monte Carlo methods for adaptive sparse approximations of time-series. *submitted for publication*.

[2] T. Blumensath and M. E. Davies. Enforcing sparsity, shift-invariance and positivity in a Bayesian model of polyphonic music. In *Proc. of the IEEE Workshop on Statistical Signal Processing*, July 2005.

[3] L. De Lathauwer, B. De Moor, and J. Vandewalle. Fetal electrocardiogram extraction by source subspace separation. In *Proc. IEEE SP / ATHOS Workshop on HOS*, pages 134–138, Girona, Spain, June 1995.

[4] D. J. Rowlands. *Understanding the Electrocardiogram: A New Approach: The Normal E.C.G Section 1*. Mosby, 1983.

[5] D. J. Rowlands. *Understanding the Electrocardiogram: Morphological Abnormalities Vol 2*. Mosby, 1982.

[6] D. J. Rowlands. *Understanding the Electrocardiogram: Rhythm Abnormalities Vol 3*. Mosby, 1987.

[7] C. Widmark, T. Jansson, K. Lindecrantz, and K. G. Rosen. ECG waveform, short term heart rate variability and plasma catecholamine concentrations in response to hypoxia in intrauterine growth related guinea-pig fetuses. *Journal of Developmental Physiology*, 15(3):161–168, 1991.

[8] V. Zarzoso and A. K. Nandi. Noninvasive fetal electrocardiogram extraction: Blind separation versus adaptive noise cancellation. *IEEE Transactions on Biomedical Engineering*, 48(1):12–18, Jan 2001.