Blind Separation of Maternal and Fetal ECG's using any Number of Channels

Thomas Blumensath, *Member, IEEE,* Mike E. Davies, *Member, IEEE*

***Abstract*— In this paper we report on the separation of maternal and fetal heartbeats from electrocardiogram (ECG) recordings based on a sparse generative signal model. The proposed algorithm uses Bayesian learning strategies to both learn the characteristic PQRST complexes and then infer their time location. Reconstruction of the signal based on each of the learned PQRST complexes leads to a separation of the maternal and fetal heartbeats. The method is flexible and can be used for single and multiple channel recordings. The extracted information is valuable for medical diagnostics, offering the fetal and maternal heart rate as well as the PQRST complexes from which diagnostic information such as pathological PQRST shapes become evident.**

***Index Terms*— Blind source separation, fetal ECG, sparse approximation, medical data analysis, Bayesian modelling.**

I. Introduction

In prenatal diagnostics of fetal heart conditions the electrocardiogram (ECG) signal of the fetal heart is of immense value [1]. Of particular interest are the fetal heart rate and the shape of the PQRST complex, e.g. the T/QRS ratio [1] and possible variations in the PQRST complex such as ectopic beats. The PQRST complex is the wave-shape of an ECG recording associated with a single heartbeat. Two PQRST wave-shapes, which have been learned with the method proposed here, can be seen on the right of figure 3. The first small excursion is the P wave, whilst the first large positive excursion is the R wave. If the R wave is preceded by a negative excursion, the negative excursion is called the Q wave and a possible negative excursion after the R wave is called the S wave. Finally, a small excursion follows, the T wave. Not all of these excursions are visible in all recordings of the ECG and the particular wave-shape depends crucially on the relative position between the heart and the recording electrodes. For more detail the reader should consult standard textbooks on ECG diagnostics such as [2] and for a description of pathological ECG data [3] and [4].

Unfortunately, non-invasive techniques are currently not available that are able to directly record the fetal ECG signal without substantial interference. The standard method to record fetal ECG signals is to place the electrodes on the maternal abdomen. This practice leads to a contamination of the recorded signals with noise and interference from the maternal heartbeat. What is more, the amplitude of the maternal heartbeat is in general much stronger than the fetal one. An example of one such recording is shown in the top left plot of figure 3.

In [5] and [6] two early methods were proposed to blindly separate the maternal and fetal ECG signals. These early methods used Singular Value Decompositions and in effect relied only on second order statistics of the signals. More recently, very good results have been achieved using higher order statistical methods such as independent component analysis (ICA) and independent subspace analysis [7], [8], [9]. These previous approaches use multiple simultaneous recordings and exploit the space diversity between different electrodes. The observation at each electrode is assumed to be a different projection of a three-dimensional current source dipole vector. The problem is then to infer the maternal and fetal bioelectric vector processes from the observations. (See [5] for details).

The aim of this paper is to show that non-linear methods based on a sparse generative signal model can be used to blindly separate maternal and fetal ECG signals, even if only a single sensor is available. The main assumptions used are:

1) The fetal and maternal PQRST complexes can each be modelled with a single unknown time-domain waveform.
2) The PQRST complexes can occur at arbitrary time locations.
3) The locations of the PQRST complexes are sparsely distributed.
4) The strength with which each PQRST complex contributes to the observation is non-negative.

In this paper we use assumption two instead of the more restrictive assumption of a periodic heart rate, i.e. we do not explicitly assume heartbeats to be periodic. Such an assumption could be included into the model described below by the specification of prior probabilities to enforce this periodicity. However, early experiments with such a more complex formulation did not offer significant advantages. Furthermore, for diagnostic purposes, deviations from the periodicity are of interest and in order not to mask such effects, we do not use such constraints here.

The first assumption is the most restrictive assumption on the model performance. Slight variations in the PQRST complexes are common and a single feature cannot represent these variations. Some pathological heart conditions lead to variations between PQRST complexes. As these variations are in general large, they will be detectable in the residual term. In general, a more complex model for the PQRST complexes could be envisaged and incorporated into our formalism. As such complications distract from the main theme of this

The authors are with IDCOM & Joint Research Institute for Signal and Image Processing, Edinburgh University, King's Buildings, Mayfield Road, Edinburgh EH9 3JL, UK (Tel.: +44(0)131 6505659, Fax.: +44(0)131 6506554, e-mail: thomas.blumensath@ed.ac.uk, mike.davies@ed.ac.uk).

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.

paper they are not discussed further here. Nevertheless, it should be noted that slow variations in the PQRST complexes can be tracked with the proposed method if the method is implemented using an online implementation as used in [10].

The exact form of the signal model is described in more detail in the next section in which the main model components are introduced. In particular we specify a linear generative model for time-series in which features such as the PQRST complexes can occur at arbitrary locations. To facilitate the adaptation of the model to a particular set of observations we also specify a set of probability densities, which then allow us to develop Bayesian learning strategies. A practical algorithm is introduced in section III. This algorithm is based on Monte Carlo approximations of the developed learning rule. The successful separation of maternal and fetal heartbeats is then demonstrated in section IV

II. Signal Model

In this paper we use the notation $\mathbf{a}_{k,r}$ to refer to the unknown PQRST complexes. Here $k \in \{1,2\}$ labels the maternal and the fetal PQRST complexes respectively and r is the label for the recording electrode. We further introduce two time-series of impulse trains $s_k[t]$ (one for the maternal and one for the fetal heartbeat) to model the strength and time-location at which the heartbeats occur. These time-series are zero most of the time and each heartbeat is then encoded by a single positive value.

ECG recordings often contain contaminating noise. Furthermore, PQRST complexes show slight variations from heartbeat to heartbeat. To model these two effects we introduce the noise term $e_r[t]$.

With these assumptions we can write the observation sequence at electrode r as:

$$x_r[t] = \sum_{k\in\{1,2\}} \sum_{l} a_{k,r}[t-l]s_k[l] + e_r[t].$$

The left sum is adding the maternal and fetal ECG signals, which are in turn modelled as a convolution of the characteristic PQRST complexes $\mathbf{a}_{k,r}$ and the associated time-series $s_k[t]$. For notational convenience we will write the above model in matrix notation as $\mathbf{x} = \mathbf{A}\mathbf{s} + \mathbf{e}$, where $\mathbf{x}$ is a vector in which the individual observation time-series have been concatenated and $\mathbf{s}$ is a similar vector for the unobserved sparse impulse trains. The matrix $\mathbf{A}$ is then a concatenation of convolution matrices. The exact structure is best understood from the visual representation in figure 1.

In the above model, both $\mathbf{a}_{k,r}$ as well as the sequences $s_k[t]$ are unknown and have to be estimated. To solve this problem we exploit the fact that the time-series $s_k[t]$ are known to be sparse, which means that most of the values are exactly zero, i.e. heartbeats occur at isolated positions. This prior knowledge is incorporated into a probabilistic model formulation by assuming that the unknown time-series $s_k[t]$ are sequences of independent variables that are zero with high probability. In order to describe, whether the time-series $s_k[t]$ are zero or not, we introduce indicator variables u. The non-zero values are then assumed to be positive and are

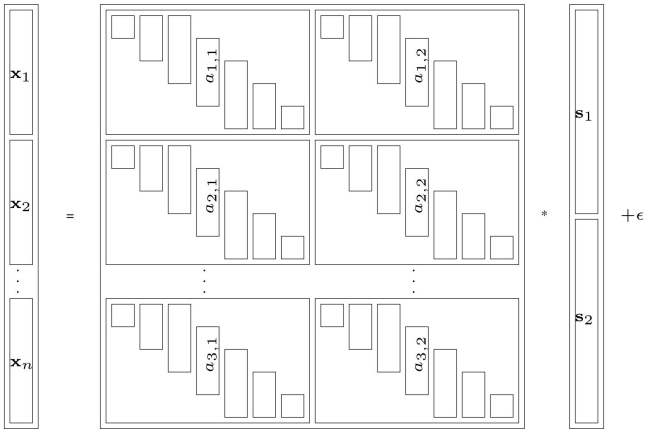

Fig. 1. Graphical representation of the mixture model in matrix notation.

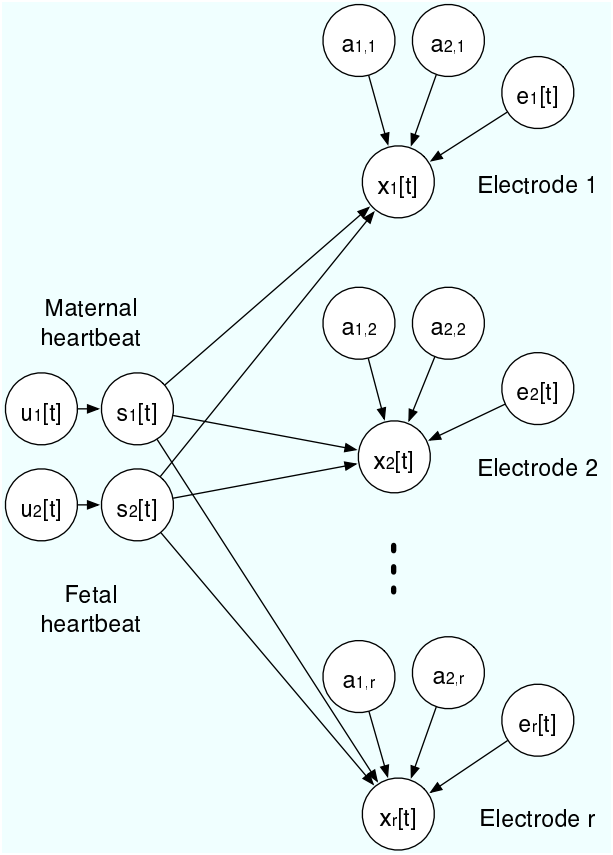

Fig. 2. Graphical model representation of the stochastic signal model. On the left are the stochastic impulse trains that model the maternal and fetal heartbeats. On the right are the models for each observation electrode, where each heartbeat is convolved with the associated PQRST complex and noise is added.

here modelled with the modified Rayleigh distribution [10]. The exact mathematical expressions for the model parameters introduced here are summarised in appendix I.

The probabilistic formulation is best described using a bayesian graphical model. This is done in figure 2. On the left are the impulse trains, which encode the timing and strength of the maternal and fetal heart. If at a particular time instance the indicator variable $u_k[t]$ is zero, then so will be $s_k[t]$ at

that time instance. If, however, $u_k[t]$ is one, then $s_k[t]$ will be a modified Rayleigh random variable. On the right we show how each of the observations is modelled. Each of the sparse time series is convolved with the associated PQRST complex for that observation electrode and the results added together. In addition, random noise contaminates the observations. The PQRST complexes themselves are assumed to be unknown random quantities.

All one is normally given are the individual observations $x_r[t]$ and the problem is then to infer the other random quantities in the model. A method to do this inference is derived in the next section. Once estimates have been found, it is possible to reconstruct the maternal and fetal ECG signals using any one of the PQRST complexes and the associated time-series:

$$x_{k,r}[t] = \sum_l a_{k,r}[t-l]s_k[l].$$

It is important to note that the model introduced here has two ambiguities. Firstly, multiplication of $\mathbf{a}_{k,r}$ by any scale factor for all r and division of $s_k[l]$ by the same scale factor leads to the same reconstructions. This problem can be solved by keeping the variance of the non-zero $s_k[t]$ fixed or by normalising $\mathbf{a}_{1,1}$ and $\mathbf{a}_{2,1}$. Secondly, there is an ordering ambiguity. Exchanging $\mathbf{a}_{1,r}$ and $\mathbf{a}_{2,r}$ for all r as well as $s_1[t]$ and $s_2[t]$ again does not change the reconstruction of the $x_r[t]$. Once estimates of $\mathbf{a}_{k,r}$ have been calculated, we therefore need to identify which one belongs to the fetal heart and which belongs to the maternal heart. This is relatively straight forward and has been done in this paper using the assumption that the fetal heart rate is faster than the maternal one and that the fetal PQRST complex is shorter than the maternal one.

III. Learning Algorithm

In order to estimate the PQRST complexes we treat the unknown $s_k[t]$ and $u_k[t]$ as nuisance parameters. We therefore calculate the maximum of the marginalised posterior

$$p(\mathbf{A}|\mathbf{x}) = \int p(\mathbf{A}, \mathbf{s}|\mathbf{x}) \; d\mathbf{s}.$$

This maximisation is done using a gradient descent approach that minimises $-\ln p(\mathbf{A}|\mathbf{x})$, the gradient of which can be written as:

$$\Delta \mathbf{a}_{k,r} = \int \frac{\partial \ln p(\mathbf{x}|\mathbf{s}, \mathbf{A})}{\partial \mathbf{a}_{k,r}} p(\mathbf{s}|\mathbf{A}, \mathbf{x}) \; d\mathbf{s}, \qquad (1)$$

where we assume that the density $p(\mathbf{A})$ is uninformative and flat.

The integral in equation (1) cannot be evaluated analytically. We therefore approximate the gradient using Monte Carlo integration. This is done by drawing J samples $\hat{\mathbf{s}}^j$ and $\hat{\mathbf{u}}^j$ from $p(\mathbf{s}, \mathbf{u}|\mathbf{A}, \mathbf{x})$ using a Gibbs sampler [11]. The approximation of the integral is then the average of the gradient of $\ln p(\mathbf{x}|\hat{\mathbf{s}}^j, \mathbf{A})$ evaluated at these sample points.

For the problem addressed here we found that a direct implementation of the above gradient method often did not find both, maternal and fetal, PQRST complexes. This was due to the iterative learning strategy and the much stronger maternal PQRST complex. The method quickly extracts the maternal PQRST complex. Once the maternal PQRST complex had been learned, the method was not able to also learn the fetal one. This problem was solved by weighting the gradient update by a term proportional to the sum of each of the time-series $\hat{s}_k^j[t]$. This ensured that both PQRST complexes were updated at a comparable rate.

The problem of estimating the time series $s_k[t]$ can also be solved based on the samples $\hat{\mathbf{s}}^j$ drawn from $p(\mathbf{s}, \mathbf{u}|\mathbf{A}, \mathbf{x})$ once the PQRST complexes have been found. We here estimated the time-series $s_k[t]$ using the sample mean of $\hat{\mathbf{s}}^j$. Further details of this Gibbs sampling based approach for a similar problem can be found in [10].

IV. Experimental evaluation

To evaluate the method we used the data-set from [7]. This data-set contains five signals recorded simultaneously with five electrodes from the abdomen of a pregnant woman. The three thoracic signals of the data-set were discarded for our experiments. [1]

A. Single Channel Blind Source Separation

In the first experiment we only used the data from a single channel. We fixed the length of the two PQRST complexes to 100 samples each and learned the PQRST complexes from the sequence with the algorithm outlined above.

The results obtained from channel (1) are shown in figure 3. In the top left corner we show the original signal used as a training sequence. It is clear that the maternal heartbeat is much stronger in this signal than the fetal one. In the second row on the left we show the reconstructed heartbeat of the maternal heart. The third row on the left shows the sequence $s_k[t]$ associated with the maternal heartbeat. From the occurrence of the non-zero values, the heart rate is easily calculated. In the next two rows we show the fetal heartbeat and the associated sequence $s_k[t]$. Finally, in the last row we show the residual signal.

On the right we show the two PQRST complexes learned from the data. The top panel shows the PQRST complex of the maternal heart and the bottom panel shows the fetal PQRST complex. For both PQRST complexes many of the important diagnostic features can be determined. Similar results could be obtained from the other channels with the exception of channel four. The recording from this sensor is challenging as the fetal ECG contribution is very weak. This can be seen in figure 5 where the fetal ECG signal cannot be detected by eye. For the other channels we found that occasionally a fetal heartbeat was not detected and that on other occasions a maternal heartbeat was also detected as a fetal one. (See for example the results in figure 6 below, where the maternal heartbeat at sample 2300 leads to a small fetal contribution.)

[1] Note that in [7] the sample frequency for these signals is specified as 500Hz, however, as noted in [9], this would lead to a very high maternal and fetal heart rate. We therefore label all graphs using sample number instead of time.

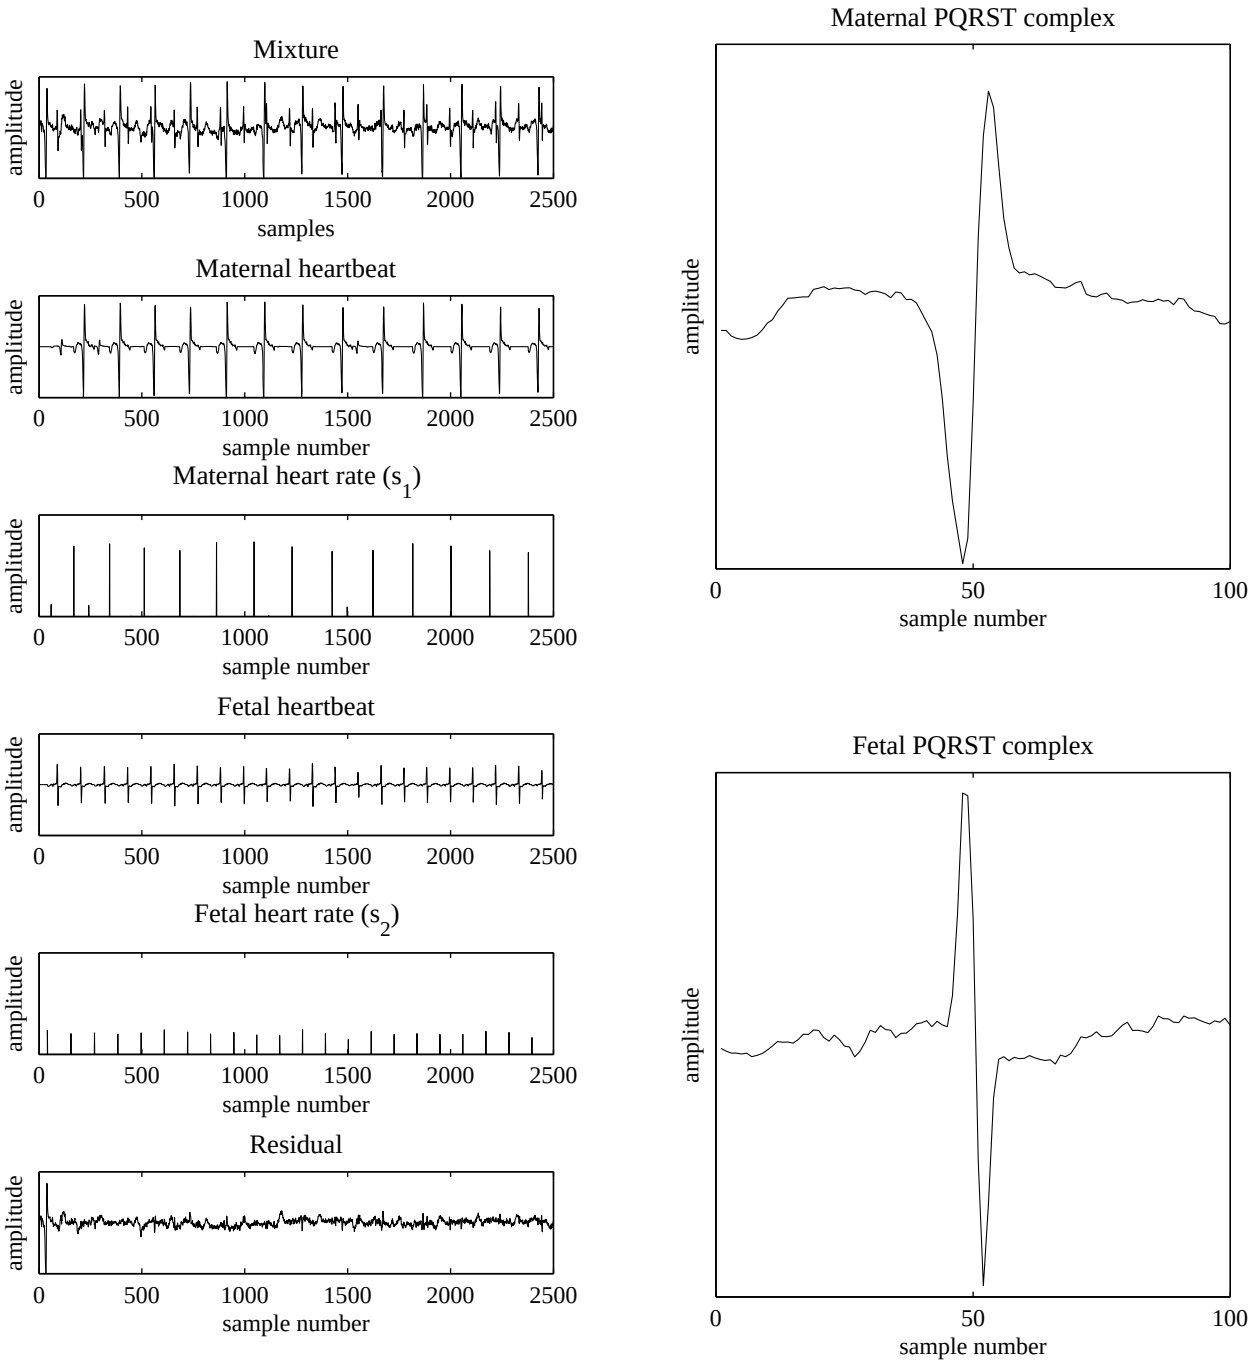

Fig. 3. Separation of fetal and maternal heartbeat. On the left we show the original single channel recording, in which the fetal heartbeat is much weaker than the maternal one and in which the SNR is low (top), the separated maternal heartbeat signal (second panel), the heart rate (third panel), the separated fetal heartbeat (fourth panel), the fetal heart rate (fifth panel) and the residual noise (last panel). On the right we show the maternal PQRST complex (top) and the fetal PQRST complex (bottom).

B. Using Multiple Channels

The signal recorded with electrode 4 has a low frequency contamination. A direct application of our method to multiple channels, which include channel four, led to a modulation of the strength of the fetal coefficient time-series $s_k[t]$. This problem is easily avoided by high-pass filtering of the recorded signals before analysis similar to the method in [12]. Alternatively, it might be possible to improve the noise model by dropping the independence assumption and using an auto regressive (AR) Gaussian noise model. However, this extension has not been investigated yet.

Using the high-pass filter pre-processing and all available data, the algorithm introduced here is able to separate all five channels. The ten PQRST complexes learned in this case are shown in figure 4. Figure 6 shows the separated fetal ECG signals for all five sensors. Even the fetal PQRST complex and ECG signal for sensor 4 could be found with this method.

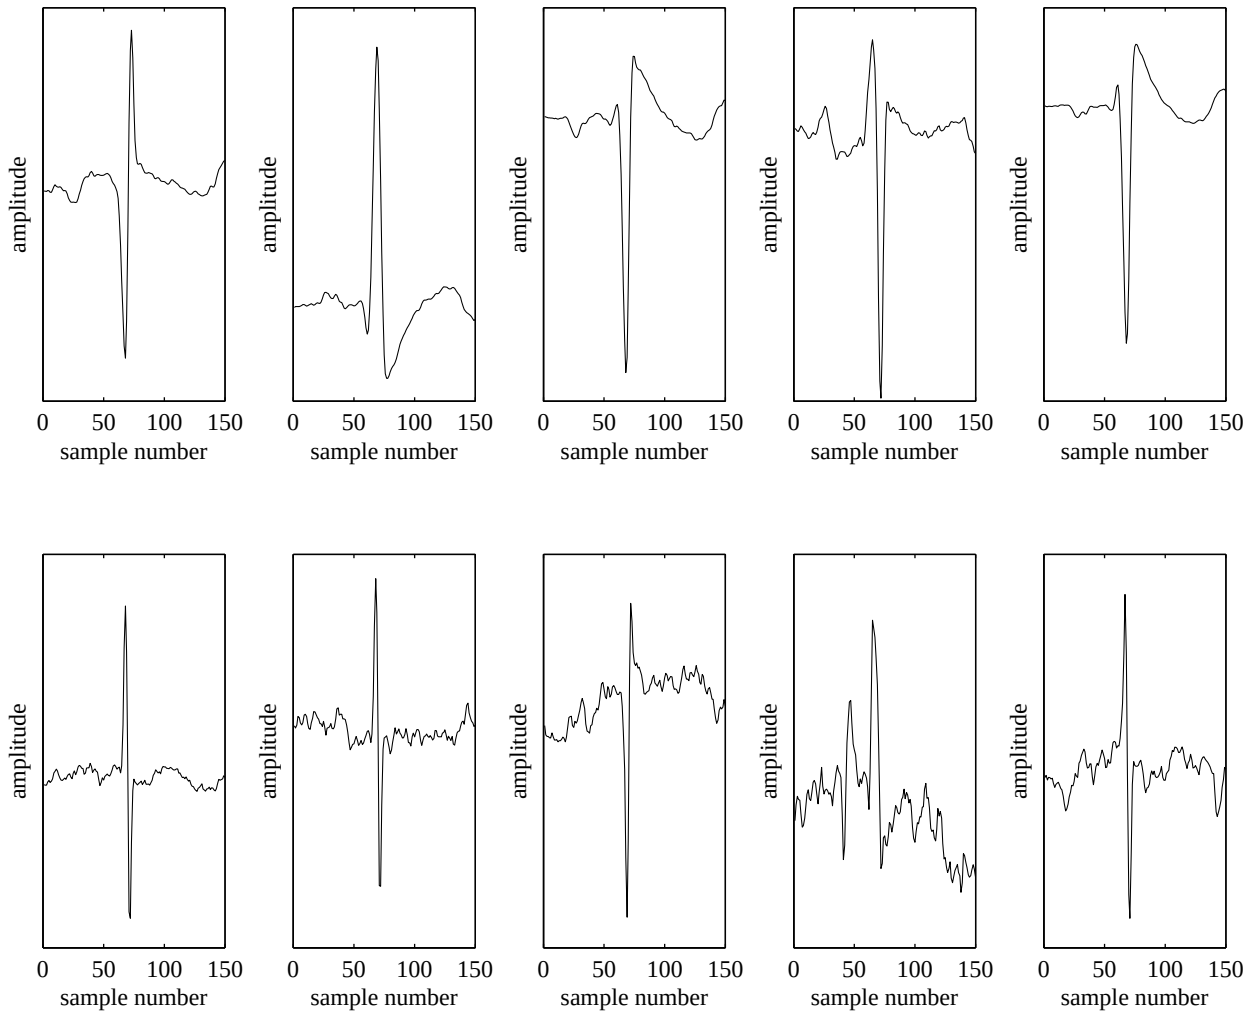

Fig. 4. The five maternal PQRST complexes (top) and the five fetal PQRST complexes (bottom), for each of the five sensors.

In figure 6 some weak spurious false detections can be seen in the fetal heartbeat. These are however much weaker than the actual fetal heartbeats. This problem was observed in several of our experiments. For the single channel separation in the previous section we also observed that occasionally some of the fetal heartbeats were not detected, this did not happen for the multi channel experiments.

For comparison, we separated the same five channels using an ICA method as described in [9] (here we used the fastICA algorithm [14]). As seen in figure 7, the results are comparable to our results. We have here plotted the separated fetal signal for both, the ICA method (upper plot) and our method (lower plot). The results computed with our approach are in general less noisy. This is due to the fact that our model explicitly includes a noise term. The main difference between the results are for channel four. It appears at a first glance, that the ICA method offers better results. However, this is deceptive as further analysis shows that the larger amplitude of the fetal ECG signal extracted from channel 4 with the ICA method is an artefact. This is clear when looking at the residual of channel four after subtracting the fetal signal. This residual has a stronger fetal contribution than the original signal, i.e. while in the original signal (figure 5) the fetal heartbeat is barely visible, after subtracting the fetal reconstruction, the residual has a noticeable fetal contribution. The sparse method did seem to extract only the actual fetal ECG contribution.

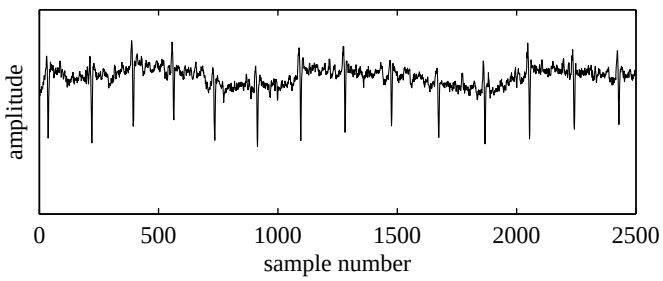

Fig. 5. The signal recorded with sensor four. The fetal heartbeat is buried in noise and low frequency baseline wander is evident.

V. Discussion and conclusion

Separation of maternal and fetal ECG signals from abdominal recordings are of immense diagnostic benefit. Previous methods such as [5], [6] and [9] have resorted to linear

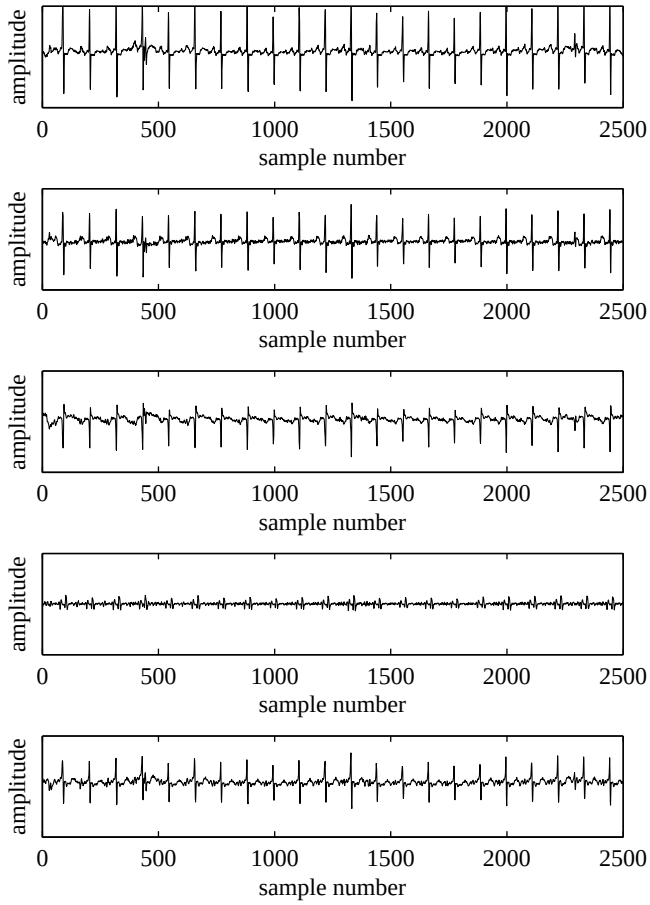

Fig. 6. The fetal ECG signals separated for each of the five sensors.

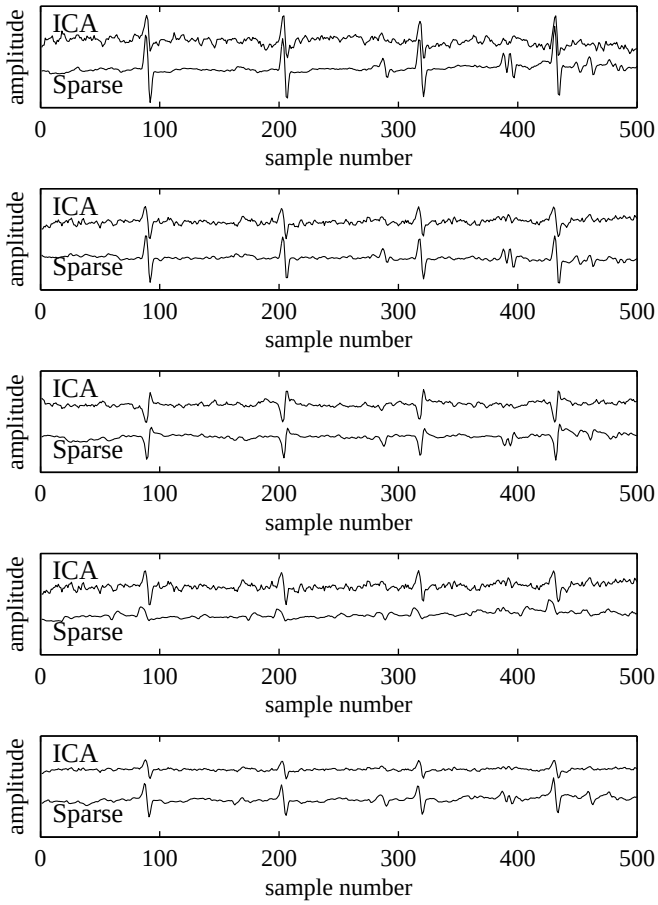

Fig. 7. Comparison between the results using an ICA method (upper graph in each panel) to the results with the method presented here (lower plot in each panel). Both methods used the same data with the exception that for our method, we used pre-processing to remove the baseline wander from channel four.

estimates and relied on multichannel recordings in order to separate the individual signals. These methods model the ECG signals as three-dimensional vector processes that are linearly projected onto one dimensional observations. A linear estimate of the two three-dimensional processes requires in general at least six observations, i.e. one would need at least six electrodes. As shown in [9] blind separation based on ICA type methods can then be applied successfully. It was also shown in [9], and confirmed in our experiments, that ICA is still able to recover some information from the fetal heart, using as few as three electrodes. However, as pointed out in [9], the results with three electrodes only identified a one dimensional subspace for the fetal heart and did therefore not recover all of the signal structure.

Here we have proposed a non-linear method able to extract PQRST complexes for the maternal and fetal heart. Based on this model a separation of even single channel recordings was possible. If more than one channel is available, this additional information can be used and we are able to separate signals with a very weak fetal contribution. Our model did not include an explicit model for the baseline wander, however, we found that standard ECG pre-processing techniques such as filtering could be used to remove this problem.

In the proposed mixture model two filters have to be learned from each of the observation sequences. In order to find possible and useful solutions to this problem we had to use strong prior assumptions. We here proposed the use of sparsity and non-negativity of the input sequences and did not use any other prior information about the expected structure and shape of the PQRST complexes. The found PQRST complexes and extracted heart rates are not biased towards the normal shape of these features and can therefore be used with more confidence in diagnostic settings.

The proposed estimation of $s_k[t]$ can be seen as a Bayesian approach to QRS detection ([12] [13]). The problem addressed by Pan and Tompkins in [12] and by Christov in [13] is the estimation of the occurrences of the QRS part of the PQRST complexes in ECG recordings (however, not in mixtures of ECG signals as discussed in this paper). Our method encompasses this estimation of the occurrence of the PQRST complexes and could therefore also be used to estimate the QRS locations in standard ECG recordings. An online implementation of our method could track changes in the PQRST complexes and would constantly adapt the QRS detection to the signal.

The main drawback of the proposed method is that the extracted PQRST complexes are averages over the PQRST complexes found in the signal. They do therefore not show the variability between different PQRST complexes of the

same heart. For the experiments conducted here, this was not found to be a problem. However, it is easy to replace the simple model for the PQRST complex with more involved models, which could capture such variability. Furthermore, strong anomalies in single PQRST complexes will lead to a large reconstruction error so that these anomalies would be visible in the residual. Nevertheless, many important diagnostic properties of the heartbeat [3] [4] could be diagnosed from the PQRST complexes extracted here.

The aim of this paper was a proof of concept, showing that a sparse generative signal model can be used successfully to extract fetal ECG signals from recordings made from the maternal abdomen. The challenge is now to refine the proposed approach and to develop a practical algorithm usable in clinical applications. Currently, the main drawback of the method is its computational complexity. However, it might be possible to replace the slow Gibbs sampling approach with much faster methods, such as an integral approximation by a delta function at the MAP estimate of $p(\mathbf{s}|\mathbf{A}, \mathbf{x})$ as used in [15]. A fast method could then hopefully be developed to find this MAP estimate, possibly based on previous work on QRS detection [13]. Such an approach could possibly lead to a real time implementation in which the impulse responses could be constantly updated, so that slow changes in the PQRST complexes due to movement of the fetus or the electrode could be tracked. However, much more research in this direction is required.

APPENDIX I
PROBABILISTIC FORMULATION

The model uses the following probabilistic distributions:

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

and

$$p(s|u) = up(s; \mu, \sigma_R^2) + (1 - u)\delta_0(s),$$

where we use the notation s to denote any one sample from the two time-series $s_k[t]$ and u to denote an indicator variable that is either zero or one. We use $\delta_0(s)$ to denote the delta function. The non-zero samples follow the modified Rayleigh distribution:

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

and the error term is modelled as independent and identically distributed Gaussian:

$$p(e_r[t]) = \frac{1}{\sigma_{\mathbf{e}}\sqrt{2\pi}} e^{-\frac{1}{2\sigma_{\mathbf{e}}^2} e_r^2[t]}.$$

REFERENCES

[1] C. Widmark, T. Jansson, K. Lindecrantz, and K. 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*, vol. 15, no. 3, pp. 161–168, 1991.

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

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

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

[5] J. Vanderschoot, D. Callaerts, W. Sansen, J. Vandewalle, G. Vantrappen, and J. Janssens, "Two methods for optimal mecg elimination and fecg detection from skin electrode signals," *IEEE Transactions on Biomedical Engineering*, vol. 34, no. 3, pp. 233–243, 1987.

[6] D. Callaertes, B. De Moor, J. Vandewalle, W. Sansen, G. Vantrappen, and J. Janssens, "Comparision of SVD methods to extract the foetal electrocardiogram from cutaneous electrode signals," *Medical and Biological Engineering and Computing*, vol. 28, pp. 217–224, 1990.

[7] L. De Lathauwer, B. De Moor, and J. Vandewalle, "Fetal electrocardiogram extraction by source subspace separation," in *Proc. IEEE SP / ATHOS Workshop on HOS*, (Girona, Spain), pp. 134–138, June 1995.

[8] V. Vigneron, A. Paraschiv-Ionescu, A. Azancot, O. Sibony, and C. Jutten, "Fetal electrocardiogram extraction based on non-stationary ICA and wavelet denoising," in *Proceedings of the Seventh International Symposium on Signal Processing and Its Applications*, pp. 69–72, 2003.

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

[10] T. Blumensath and M. 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.

[11] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 6, pp. 721–741, 1984.

[12] J. Pan and W. J. Tompkins, "A real-time QRS detection algorithm," *IEEE Transactions on Biomedical Engineering*, vol. 32, no. 3, pp. 230–236, 1985.

[13] I. Christov, "Real time electrocardiogram QRS detection using combined adaptive threshold," *BioMedical Engineering OnLine*, vol. 28, no. 3, 2004.

[14] A. Hyvärinen, "Fast and robust fixed-point algorithms for independent component analysis.," *IEEE Transactions on Neural Networks*, vol. 10, no. 3, pp. 626–634, 1999.

[15] T. Blumensath and M. Davies, "Sparse and shift-invariant representations of music," *IEEE Transactions on Audio, Speech and Language Processing*, vol. 14, pp. 50–57, Jan 2006.

PLACE PHOTO HERE

Thomas Blumensath received the BSc hons. degree in music technology from Derby University, Derby, U.K., in 2002 and his Ph.D. degree in electronic engineering from Queen Mary, University of London, U.K., in 2006. He is currently a postdoctoral research fellow in the Institute for Digital Communication at the University of Edinburgh, working on sparse representations for signal processing and coding. His research interests include mathematical and Bayesian methods in signal processing, unsupervised learning, sparse signal representations and applications to music and audio analysis.

PLACE PHOTO HERE

Mike Davies received a B.A. hons. degree in engineering from Cambridge University, Cambridge, U.K., in 1989 and the Ph.D. degree in nonlinear dynamics and signal processing from University College, London, U.K., in 1993. He was then awarded a prestigious Royal Society Research Fellowship which he held at UCL and then Cambridge. He is currently a Professor in the Institute for Digital Communications, at the University of Edinburgh and an Associate Editor for IEEE Transactions in Speech, Language and Audio Processing. His research interests include: statistical signal processing, audio processing, array signal processing, and information theory.