NON-INVASIVE REAL-TIME FETAL ECG EXTRACTION
A Block-on-Line DSP Implementation based on the JADE Algorithm

Silvia Muceli, Danilo Pani and Luigi Raffo
DIEE - Dept. of Electrical and Electronic Engineering, University of Cagliari, Piazza d'Armi, 09123 Cagliari, Italy

Keywords: Fetal monitoring, antenatal diagnosis, real-time processing, DSP.

Abstract: The possibility to access the fetal ECG non-invasively during the early stages of the pregnancy is a paramount requirement for cardiologists aiming to treat fetuses with congenital hearth diseases. Several research works have been presented during the past years to address this issue. In this paper we present a block-on-line blind source separation technique that combines the powerfulness of the batch JADE algorithm to the requirements of a separation able to adapt to a time-varying mixing process. To avoid estimated sources permutation, a simple preconditioning technique in conjunction with a proper parameters tuning has been developed and tested. The whole algorithm has been implemented on a powerful floating-point Digital Signal Processor, and it is ready to be embedded in an acquisition device for a deeper experimentation.

1 INTRODUCTION

Fetal electrocardiography would allow the early diagnosis of some fetal heart congenital diseases, effectively treatable by transplacental drugs administration, that cannot be discovered by means of other more expensive diagnostic instrumentation such as ultrasound-based devices. Fetal electrocardiogram (FECG) extraction from multichannel transabdominal maternal potential recordings can be accomplished resorting to Blind Source Separation (BSS) techniques. Off-line BSS methods exploiting fourth-order statistics have been proven to be more effective than other methods but they are very computationally intensive and unable to deal with long-lasting recordings with possibly time-varying mixing processes.

In this paper we present a possible solution to such problems introducing a block-on-line version of the famous JADE algorithm (Cardoso and Souloumiac, 1993), with the aim to take from it the high separation quality, at the same time allowing the processing in presence of a potentially time-varying mixing process. Estimated sources permutation was avoided by means of a preconditioning technique used in conjunction with a proper parameter tuning. To assess the actual real-time capabilities of the algorithm, we present its porting and simulation on a floating point Digital Signal Processor (DSP) by Texas Instruments (TI). Due to such implementation, which still preserves a great portability, the system is ready to be embedded in an acquisition device for a deeper experimentation. Separation quality results have been assessed on a publicly available database.

In Section 2 a short review of some related works is presented, whereas the proposed algorithm is described in Section 3. Section 4 deals with the DSP implementation of the system. Experimental results are presented in Section 5. Section 6 concludes this work.

2 STATE OF THE ART

Non-invasive fetal electrocardiography is not yet used in clinical practice because of the difficulty in extracting the signal of interest, particularly in real-time. Cutaneous recordings from a pregnant woman comprise fetal and maternal heart signals, electromyographic and breathing noises, and external interferences. Since FECG and maternal ECG (MECG) have overlapping spectra, they cannot be separated through conventional frequency selective filtering. Many methods yield an estimate of the maternal signal from her chest, then removing this contribution from the composite abdominal one to obtain the FECG. Widrow (Widrow and Stearns, 1985) proposed a method of adaptive filtering and noise cancelling but its performance is very dependent on the electrodes placement. Camps (Camps et al., 2001) provided an extension of Widrow's scheme by including in it a finite impulse response neural network. The drawback is the complexity of the choice of free pa-

Muceli S., Pani D. and Raffo L. (2008).
NON-INVASIVE REAL-TIME FETAL ECG EXTRACTION - A Block-on-Line DSP Implementation based on the JADE Algorithm.
In *Proceedings of the First International Conference on Bio-inspired Systems and Signal Processing*, pages 458-463
DOI: 10.5220/0001066804580463

rameters, whose number increases geometrically with the number of inputs, though without eliminating the problem of high sensitivity to electrodes placement. Kanjilal (Kanjilal et al., 1997) used Singular Value Decomposition (SVD) to recover fetal and maternal components from a single channel. Since this method is based on nearly periodicity of ECG signals, it is unsuitable to detect unexpected patterns. Mochimaru (Mochimaru et al., 2002) suggested a wavelet-based method, but this approach suffers in presence of overlapping maternal and fetal QRS complexes.

BSS aims to recover the source signals from a set of mixtures without a priori knowledge of the mixing process (blindly). Under the hypothesis that the mixing process is linear, if we denote with $\mathbf{s}$ the n-dimensional vector of sources, with $\mathbf{x}$ the m-dimensional vector of observed mixtures and with $\mathbf{A} \in R^{m\times n}$ the mixing matrix, the problem can be mathematically formulated as in (1):

$$\mathbf{x}(k) = \mathbf{A}\mathbf{s}(k) \tag{1}$$

where k represents a discrete time index. The solution consists in finding an unmixing matrix $\mathbf{B} \in R^{n\times m}$ such that $\hat{\mathbf{s}}(k) = \mathbf{B}\mathbf{x}(k)$ is a good estimation of $\mathbf{s}$ up to a permutation and a multiplicative constant (the two BSS ambiguities).

Source separation can be accomplished by Principal Component Analysis (PCA) and Independent Component Analysis (ICA). PCA looks for components which are uncorrelated, whereas ICA looks for components which are statistically independent, i.e. whose higher order cumulants are all diagonal tensors. Since uncorrelated sources are not necessarily independent, unless they have a Gaussian distribution, and ECG signals are known to be super-Gaussian, ICA is more appropriate than PCA, producing better results (Bacharakis et al., 1996). Sources involved in this application descend from different bioelectric phenomena and can be fairly considered statistically independent. It is also well-accepted that MECG gives rise to 3 source signals and FECG can be represented by means of 2 source signals (Nandi and Zarzoso, 1997). Hence it suffices to have a number of observed mixtures not lower than the number of sources to ensure the identifiability of ICA model. In the following, we assume $m = n$, then all vectors will be n-dimensional and all matrices will be $n \times n$.

Since De Lathauwer (De Lathauwer, 1995) has first used ICA to separate FECG successfully, many researchers resorted to the same approach, proved to perform better than the Widrow's method (Zarzoso and Nandi, 2001), and rather robust with respect to electrodes placement. De Lathauwer (De Lathauwer et al., 2000) showed ICA capability to reveal ectopic beats in a regular ECG, differently from other approaches making assumptions on the signal characteristics. This property makes BSS by means of ICA suitable for medical applications.

3 THE PROPOSED ALGORITHM

All the algorithms cited in Section 2 process data in batch mode, so they can be used only if the mixing process does not change over time. Since both mother and fetus can move, the separation algorithm must be able to track such changes. Some on-line solutions have been conceived but, unlike batch ones, the quality of the estimated sources is quite poor. We identified a good on-line algorithm in terms of separation quality in Mermaid (Marossero et al., 2003) but the main drawback is the parameter tuning, since several parameters have to be chosen empirically to achieve a good-quality signal. Balancing pros and cons of both batch and on-line techniques, we chose to derive a block-on-line method from a batch one: JADE (Cardoso and Souloumiac, 1993).

3.1 Background: The Jade Algorithm

In the JADE algorithm the ICA problem is solved by means of a two-stage procedure consisting of a preliminary processing performed by employing Second-Order Statistics (SOS) and a second one performed by resorting to Higher-Order Statistics (HOS). The first stage, multiplying by a whitening matrix $\mathbf{W}$ the observed mixtures, decorrelates and normalizes them; the second stage aims to obtain higher-order independence by multiplying decorrelated mixtures by an orthogonal rotation matrix $\mathbf{G}$ which minimizes the sum of squared fourth order cross cumulants of whitened observations:

$$\Psi = \sum_{ijkl \neq iikl} {K_{ijkl}}^2 \tag{2}$$

The $\mathbf{G}$ matrix is found by joint diagonalization of the cumulant matrices, efficiently performed by Jacobi method. This technique works by repeated sweeps of plane rotations, each one applied to a pair of rows of the cumulant matrices. During a sweep, for each pair (i, j) with $1 \leq i \leq n$, the Givens angle θ_{ij} which minimizes the contrast is calculated and if $\theta_{ij} > \theta_{th}$ (θ_{th} is a threshold angle whose value determines the optimization accuracy) the pair is rotated. The convergence is obtained when no pairs have been rotated over a sweep. This optimization procedure does not suffer from problems of convergence, as opposed to the gradient descent algorithm, and does not show the difficulty of parameters tuning.

The blind identifiability requires the diagonalization of the whole fourth-order cumulant set (n^2 cumulant matrices). The original JADE algorithm deals with complex signals and finds the $\mathbf{G}$ matrix by approximate joint diagonalization of the n most significant cumulant matrices, to reduce the computational load. Since we are dealing with real signals, we can exploit the symmetries of the cumulants to reduce to $n(n+1)/2$ the number of matrices to be diagonalized, without statistical loss (Cardoso, 1999).

3.2 On-Line Jade Implementation

For a batch algorithm, the permutation ambiguity is simply the obvious outcome of the mathematical formulation of the problem: the order of the sources is unpredictable and usually insignificant. It comes to be a critical problem in block-on-line algorithms, where the estimated sources can be differently ordered in different blocks, so that channels can swap producing meaningless signals. We solved this problem by combining a sliding window strategy (window length: L samples, overlap: $(L-T)$ samples) with a two-step HOS stage. First of all, the algorithm performs the preprocessing stage (SOS, i.e. centering and whitening) on a sample-by-sample basis. The data stream is subdivided into data blocks $\mathbf{X}$ of length T, with $\mathbf{X} = \{\mathbf{x}(k), 0 \le k < T\}$. Recorded data of a block are centered subtracting their mean value by means of a running average given by:

$$\overline{\mathbf{x}}(k) = (1-\gamma)\overline{\mathbf{x}}(k-1) + \gamma\mathbf{x}(k), \qquad (3)$$

where γ is a forgetting factor. We chose $\gamma = 1/L$ so that the memory depth is equal to the sliding window length. Centered data are whitened as follows:

$$\mathbf{z}(k) = \mathbf{W}(k)(\mathbf{x}(k) - \overline{\mathbf{x}}(k)) \qquad (4)$$

and whitening matrix is updated using an approach adapted from (Cardoso and Laheld, 1996):

$$\begin{aligned} &\mathbf{W}(k+1) = \\ &= \mathbf{W}(k) - \frac{\lambda_0}{1+\lambda_0\mathbf{z}(k)^T\mathbf{z}(k)}(\mathbf{z}(k)\mathbf{z}(k)^T - \mathbf{I})\mathbf{W}(k) \end{aligned} \qquad (5)$$

where λ_0 is a constant. The smaller λ_0 the slower the convergence, but with high values it is easier to incur in instability. The newest T preprocessed data are inserted in a L-wide window $\mathbf{Z} = \{\mathbf{z}(k), 0 \le k < L\}$ with the last $L-T$ preprocessed samples.

After that, in the HOS stage, we precondition the rotation process by a coarse separation of the actual block of whitened mixtures with the rotation matrix of the previous block. This procedure considers that fetal movements can change the physical configuration and consequently the mixing matrix but, if T is chosen so that two consecutive blocks are close enough from a temporal point of view with respect to the dynamics of the mixing process, the coarse separation rotates whitened observations so that the basis vectors (the columns of the mixing matrix) of the current block are close to the basis vectors of the previous block. Coarse-separated mixtures are then used for the computation of the cumulant matrices and their joint diagonalization. At last, the fine separation matrix $\mathbf{G}$ is provided in output, the optimization process being carried out in the direction started from the coarse separation. The algorithm can be described by the following steps:

1. Acquire a block $\mathbf{X}$ of T new samples $\mathbf{X} = \{\mathbf{x}(k), 0 \le k < T\}$;
2. For all k so that $0 \le k < T$:
   (a) Center $\mathbf{x}(k)$ by subtracting $\overline{\mathbf{x}}(k)$ calculated as in (3), and keep $\overline{\mathbf{x}}(T-1)$ as the $\overline{\mathbf{x}}(-1)$ for the next block;
   (b) Whiten $(\mathbf{x}(k) - \overline{\mathbf{x}}(k))$ according to (4) to obtain $\mathbf{z}(k)$;
   (c) Update the whitening matrix following (5) and keep $\mathbf{W}(T-1)$ as the $\mathbf{W}(-1)$ for the next block;
   (d) Insert $\mathbf{z}(k)$ in $\mathbf{Z}(L-T+k)$;
3. Perform the coarse separation of $\mathbf{Z}$: $\mathbf{Y}' = \mathbf{Q}_{prev}\mathbf{Z}$;
4. Apply JADE algorithm to $\mathbf{Y}'$ to obtain the matrix $\mathbf{G}$;
5. Perform the fine separation $\mathbf{Y} = \mathbf{G}\mathbf{Y}'$;
6. Update the rotation matrix $\mathbf{Q}_{cur} = \mathbf{G}\mathbf{Q}_{prev}$, $\mathbf{Q}_{prev} = \mathbf{Q}_{cur}$;
7. Go to 1 for another block.

3.3 Parameters Setting

The proposed algorithm has been applied to a real dataset consisting of $n = 8$ potential recordings, the first 5 abdominal and the last 3 thoracic. The dataset (BIOMED, 2005) is recorded at a sampling rate of 250 Hz, and it is composed of 2500 samples (Figure 1). It represents the benchmark used in most of research works about fetal ECG extraction. We set λ_0 to 0.001. The choice of L and T aims at achieving a good trade-off among permutation rejection, separation quality and computational efficiency. L and T are set to be 1024 and 256 respectively. A higher T gives rise to ECG channels swapping, a higher L increases the needed floating point operations (FLOPS) without improving the separation quality.

Figure 1: The real signal mixtures (BIOMED, 2005).

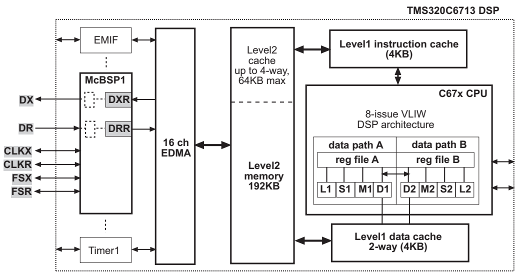

Figure 2: A block diagram of the simulated system.

4 THE DSP SYSTEM DESIGN

Modern electrocardiographs are based on DSPs, sometimes joined to a microcontroller for the user interface. The main problems of such implementations are the limited amount of memory available both on chip and on board and the low operating frequency, compared to PC processors. Conversely, DSPs are significantly less power hungry and highly optimized for signal processing operations. Almost all the devices on the market can be programmed in C or C++, so enhancing code portability. However, some performance improvements can be achieved only by exploiting at the most the architectural characteristics of the processor and then also assembly-written routines. All these aspects have been carefully taken into account in the proposed real-time system design.

The algorithm has been coded in C to work by file I/O on a PC, and then ported on a TMS320C6713 DSP, a TI 1800MFLOPS floating point VLIW processor up to 300MHz. Having to rely only on recorded signals, we avoided any hardware board implementation. To evaluate the real-time capabilities of the proposed DSP solution, it has been simulated under Code Composer Studio (CCS) 3.1, the Integrated Development Environment for the TI processors, with the Device Cycle Accurate Simulator as target. It is able to give an indication of the cycle count of the application with cycle accuracy even at peripherals level. The target DSP provides a rich set of on-chip advanced peripherals. Among them, we used one of the two Multi-channel Buffered Serial Ports (McBSP) available on chip, and the Enhanced Direct Memory Access Controller (EDMA).

Data enter the system through the McBSP1 in frames composed of 8 channels of the ECG, 16 bit/sample, with the proper sampling rate. The EDMA manages the acquisition performing a data ordering in an array placed in the internal L2 memory, so that at the end of a block acquisition such array is composed of 8 contiguous blocks of T samples, each one consisting only of the input samples for one channel. At the same time, the output samples are outputted through the same serial port with the inverse reordering procedure. The CCS Port Connect and Pin Connect features were used to provide the McBSP1 with the external clocks and frame syncs needed to correctly perform both the sample acquisition and the outputs collection, thus emulating a real hardware system. A simplified representation of the simulated system is depicted in Figure 2, where the externally driven pins (CLKX, CLKR, FSX, FSR) and ports (DXR and DRR) are shadowed. To guarantee stable input signals during the acquisition of a next block, EDMA performs automatic ping-pong buffering. The newest T samples $\mathbf{X}_i$ are acquired during the processing of the previous block $\mathbf{Z}_{i-1}$, so that the overall processing time must be less than $T_X = T/f_s$. This way the interrupt signal period is also equal to T_X. After sample-by-sample pre-processing, a Quick DMA (QDMA) call performs the sliding window mechanism needed for the HOS stage. We used fastRTS and DSPlib for highly optimized math functions and DSP array operations respectively. The access to the external memory dramatically increases the latency, whilst the wide data arrays substantially reduce the code space. With the chosen target, which comes with 256KB L2 internal RAM, and several optimizations, no external memory is required.

5 EXPERIMENTAL RESULTS

The sources estimated by means of the on-line algorithm are depicted in Figure 3. ECG channels do not suffer from the permutation problem, whereas there is a suspect permutation in the y_2 and y_4 noise channels. Such different robustness to permutations can be explained in the light of fact that cumulant based methods need a large number of samples to reach good separation results. The L value has been chosen so

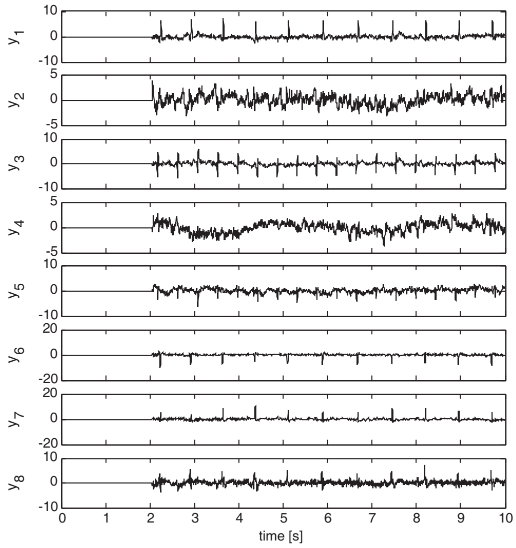

Figure 3: The sources estimated with the proposed on-line solution. The blanked parts of the signals at the beginning are due to the initialization of the on-line whitening procedure. Traces y_3 and y_5 clearly show the FECG, the hearth rate being almost twice the maternal one.

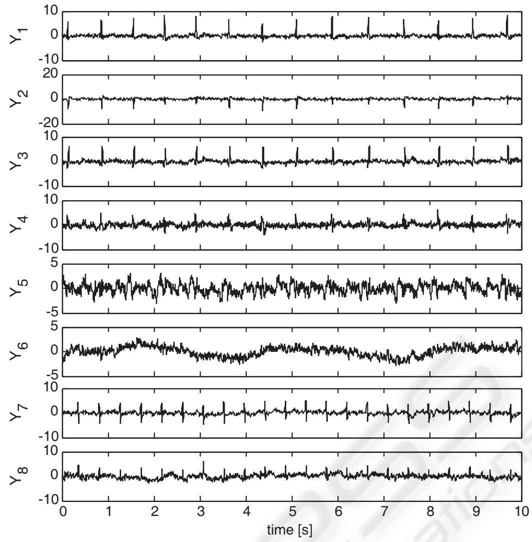

Figure 4: The sources estimated with the batch JADE algorithm. Traces Y_7 and Y_8 clearly show the FECG.

that a window contains some periods of ECG signals. Such noises can be interpreted as maternal respiration and electromyogram. Since respiration baseline wander is characterized by a pseudo-period higher than the one of ECG, we should work with a wider window to make up to a window containing some periods of fluctuation due to the maternal respiration, but this would increment the computational load. Furthermore some more unstructured noises could take no advantage of such resizing. Since our target are the ECG signals, we can disregard noise channels, maintaining $L = 1024$. With this choice, ECG waveforms are reconstructed with the same quality of the original JADE algorithm (Figure 4). To compare the performance of the two methods, we used the parameter proposed in (Bacharakis et al., 1996):

$$P_K = \frac{|K_{40}| + |K_{04}|}{\sum_{m+n=4} |K_{mn}|} \tag{6}$$

P_K represents the ratio between the sum of the modulus of the fourth-order auto-cumulants of two estimated sources and the sum of the modulus of all the fourth-order cumulants related to the same sources. If two components are really independent their cross-cumulants are close to 0, then the ratio tends to 1. In Table 1 are reported the cumulative results for all the signal couples, in terms of average value and standard deviation. It proves that the two methods allow to reach very similar source independence.

Table 1: Original batch JADE vs. block-on-line JADE. Average and standard deviation of the parameter P_K.

	avg	*STD*
Batch Algorithm	0.918	0.066
On-line Algorithm	0.912	0.109

5.1 Profiling Results

Performance have been also evaluated by means of cycle profiling to assess the real-time constraints satisfaction. From Table 2 we can see that the system largely respects the real-time requirements, since T_X is almost equal to 1 second and the DSP is running at 300 MHz. Hence, it will be possible either to add further pre/post processing to the actual implementation (such as signal enhancement, fetal QRS detection, fetal ECG delineation and measurement) and to safely reduce the clock rate. From Table 2 it is possible to notice that the time spent for the execution of the algorithm with the number of sweeps required (on the same block $\mathbf{Z}_i$), T_{tot}, represents about 8.3% of T_X, pre-processing requires less than 7% of T_{tot}, and the coarse separation plus $\mathbf{Q}_{prev}$ update less than 0.9% of T_{tot}. On the available records, we counted a maximum of 6 sweeps. In the HOS stage cycle count, there is a considerable offset due to the cumulant matrices computation that can be quantified in 21,399,260 clock cycles, 83.8% of T_{tot}. It can be estimated that every pairwise Givens rotation requires 13,534 clock cycles, 0.05% of T_{tot}. Hence, considering that 1 sweep, in the worst case, consists of maximum 28 Givens rotations, and that the coarse-separation dramatically reduces

Table 2: Profiling results for the DSP implementation: *(a).* for the whole algorithm on a single block allowing all the required sweeps, and *(b)* for only the 2^{nd} step of the HOS stage allowing a different (imposed) number of sweeps.

	Algorithm section	# ck ticks
(a)	Pre-processing	1,714,300
	Coarse-sep. + $\mathbf{Q}_{prev}$ update	220,258
	step of the HOS stage	23,599,929
(b)	1 sweep	21,557,940
	2 sweep	21,936,878
	6 sweeps (threshold control)	23,599,929

the need for many sweeps, the number of sweeps is not a performance limiter.

6 CONCLUSIONS

In this paper a block on-line version of the JADE algorithm and its real time implementation on a floating point DSP has been presented and evaluated in terms of both separation quality and profiling. Separation quality results are in line with those achievable with the batch JADE algorithm. This can be interpreted in the light of the considerations presented in Section 5 since the signal is extremely clear and stable so that it is possible to assume that mixing process does not changes a lot over time. At the same time, longer signals will benefit from the block-on-line implementation with preconditioning since the probability to have a stable mixing process over the whole recording is very low. The DSP implementation proved the possibility to use this algorithm in real-time on a portable device. The system is powerful enough to use only less than 10% of the overall available processing time to execute the algorithm, then leaving space for further processing actually required in a real implementation. The realization of a front-end analog circuitry for signals acquisition will enable the trial with custom databases.

ACKNOWLEDGEMENTS

The authors wish to thank Dr. R. Tumbarello (MD. Head, Division of Pediatric Cardiology) and Dr. S. Mura (Division of Gynecology and Obstetrics) of the Hospital “G. Brotzu”, Cagliari, Italy.

REFERENCES

Bacharakis, E., Nandi, A. K., and Zarzoso, V. (1996). Foetal ECG extraction using Blind Source Separation methods. In *Proc. EUSIPCO96*, pages 395–398.

BIOMED (2005). Katholieke Universiteit Leuven, Belgium, http://homes.esat.kuleuven.be/~biomed/biosource/data006/foetal_ecg.dat.

Camps, G., Martinez, M., and Soria, E. (2001). Fetal ECG Extraction using an FIR Neural Network. *Computers in Cardiology*, pages 23–26.

Cardoso, J. and Souloumiac, A. (1993). Blind beamforming for non-Gaussian signals. In *IEEE PROCEEDINGS F*, volume 140, pages 362–370.

Cardoso, J. F. (1999). High-order contrasts for independent component analysis. *Neural Computation*, 11(1):157–192.

Cardoso, J. F. and Laheld, B. (1996). Equivariant adaptive source separation. *IEEE Trans. on S.P.*, 44(12):3017–3030.

De Lathauwer, L. (1995). Fetal Electrocardiogram Extraction by Source Subspace Separation. In *Proc. IEEE Workshop on HOS*, pages 134–138, Girona (Spain).

De Lathauwer, L., Moor, B. D., and Vandewalle, J. (2000). Fetal Electrocardiogram Extraction by Blind Source Subspace Separation. *IEEE Trans. on Biomedical Engineering*, 47(5):567–572.

Kanjilal, P. P., Palit, S., and Saha, G. (1997). Fetal ECG extraction from single-channel maternal ECG using singular value decomposition. *IEEE Transactions on Biomedical Engineering*, 44(1):51–59.

Marossero, D., Erdogmus, D., Euliano, N., Principe, J., and Hild, K. E. (2003). Independent component analysis for fetal electrocardiogram extraction: a case for the data efficient MERMAID algorithm. pages 399–408.

Mochimaru, F., Fujimoto, Y., and Ishikawa, Y. (2002). Detecting the fetal electrocardiogram by wavelet theory-based methods. *Progress in Biomedical Research*, 7(3):185–193.

Nandi, A. and Zarzoso, V. (1997). Foetal ECG separation. In *IEE Colloquium on the Use of Model Based Digital Signal Processing Techniques in the Analysis of Biomedical Signals (No. 1997/009)*, pages 8/1–8/6.

Widrow, B. and Stearns, S. D. (1985). *Adaptive Signal Processing*. Prentice-Hall Signal Processing Series.

Zarzoso, V. and Nandi, A. K. (2001). Noninvasive fetal electrocardiogram extraction: Blind separation versus adaptive noise cancellation. *IEEE Transactions on biomedical engineering*, 48(1).