



HAL
open science

Extraction of Fetal Cardiac Signals from an Array of Maternal Abdominal Recordings

Reza Sameni

► **To cite this version:**

Reza Sameni. Extraction of Fetal Cardiac Signals from an Array of Maternal Abdominal Recordings. Signal and Image processing. Institut National Polytechnique de Grenoble - INPG; Sharif University of Technology (SUT), 2008. English. NNT: . tel-00373361

HAL Id: tel-00373361

<https://theses.hal.science/tel-00373361v1>

Submitted on 6 Apr 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Sharif University of Technology



Extraction of Fetal Cardiac Signals from an Array of Maternal Abdominal Recordings

Reza Sameni

Submitted in partial fulfillment of the requirements
for the joint degree of Doctor of Philosophy in
Signal Processing and Telecommunications

in the

Institut Polytechnique de Grenoble, Grenoble, France
Grenoble Images Signal Parole et Automatique Laboratory (GIPSA-lab)

and

Sharif University of Technology, Tehran, Iran
School of Electrical Engineering
Biomedical Signal and Image Processing Laboratory (BiSIPL)

under the joint supervision of

Prof. Christian Jutten and Dr. Mohammad B. Shamsollahi

July 2008

INSTITUT POLYTECHNIQUE DE GRENOBLE

N° attribué par la bibliothèque:

--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--

THÈSE EN COTUTELLE INTERNATIONALE

pour obtenir le grade de

DOCTEUR de L'Institut Polytechnique de Grenoble

et

de L'Université Technologique Sharif

Spécialité: Signal, Image, Parole, Télécoms

préparée aux laboratoires

Grenoble Images Signal Parole et Automatique (GIPSA-lab)

dans le cadre de l'École Doctorale

Électronique, Électrotechnique, Automatique et Traitement du Signal

et au laboratoire

Biomedical Signal and Image Processing Laboratory (BiSIPL)

présentée et soutenue publiquement par

Reza SAMENI

le 7 juillet 2008


**Extraction of Fetal Cardiac Signals from an
Array of Maternal Abdominal Recordings**

Directeurs de thèse:

Christian JUTTEN et Mohammad B. SHAMSOLLAHI

JURY

M. Jeanny HÉRAULT, Président
M. Adrian van OOSTEROM, Rapporteur
M. Jean-François CARDOSO, Rapporteur
M. Christian JUTTEN, Co-directeur de thèse
M. Mohammad B. SHAMSOLLAHI, Co-directeur de thèse
M. Massoud BABAIE-ZADEH, Examineur
M. Vicente ZARZOSO, Examineur



تقدیم به

شہید دکتر مصطفیٰ چمران

Dedicated to a Great Martyr,

Dr. Mostafa Chamran

Acknowledgments

I praise God for all His mercy and grace, the beauties of the world and the chance that he has given us to discover its laws.

I would like to thank Dr. Mohammad Bagher Shamsollahi and Prof. Christian Jutten, the supervisors of this research, for their exceptional characters, the scientific and moral guidance of this work, and their availability. This research would not have been accomplished without their constant support. Christian and Mohammad have been far beyond supervisors for me and I hope to carry the lessons that I learned from them in my personal life.

I sincerely thank my committee members: the president of the jury Prof. Jeanny Hérault, the reviewers Prof. Adrian van Oosterom and Prof. Jean-François Cardoso for the evaluation of this work and their invaluable comments, and the examiners Dr. Massoud Babaie-Zadeh and Dr. Vicente Zarzoso for their insightful feedbacks.

I would like to thank my beloved wife Maryam for all her love and patience during this very first year of our marriage. This thesis would have been dedicated to her, if Mostafa Chamran had not been a beau ideal for both of us.

I would like to thank my parents, my sister and my brother, for all their support and encouragement throughout my life.

I owe much of my signal processing insights to Dr. Mohammad Bagher Shamsollahi and Dr. Ali Asghar Eftekhar, who gave me the courage to work with signal processing theories in real world applications.

I would like to give my special thanks to Dr. Mohammad Ali Massoumnia, to whom I owe all my knowledge of estimation theory. Many of the ideas of this thesis would not have been developed without the wonderful and insightful course that I took with him in Sharif University.

I would like to thank Dr. Vincent Vigneron, Dr. Gari D. Clifford, Dr. Dirk Hoyer, and Dr. Raphael Schneider, for the fruitful discussions that we had and for providing some of the fetal cardiac signals used in this research.

I also acknowledge Dr. Jeroen G. Stinstra and Dr. Charles D. Lawrence for permitting us to use some of the figures in Chapter 3 for illustration.

I would like to thank all the teachers and professors that I have had throughout my life, especially my first teacher, Mrs. Oskourchi, who taught me the Persian alphabet.

I should also thank my dear friends and colleagues in Sharif University and GIPSA-lab, with whom I shared many of the ideas of this research and who gave me invaluable feedbacks.

I would finally like to thank Shiraz University, the French embassy in Tehran, Sharif University of Technology, the French ministry of foreign affairs, the Égide organization, the Rhône-Alpes region, and Iran Telecommunication Research Center who provided us with scholarships and research grants during this work.

چکیده

نارسایی‌های قلبی جزء شایع‌ترین بیماری‌های مادرزادی می‌باشند. این نارسایی‌ها انواع مختلفی داشته و همه انواع آن در بدو تولد و یا تا سالها پس از تولد بروز ظاهری از خود نشان نمی‌دهند و ممکن است تنها بر رشد کودک تأثیر بگذارند. از این رو مانیتور کردن قلب جنین در تشخیص زود هنگام نارسایی‌های قلبی بسیار مؤثر است. در حال حاضر متداولترین روش‌های موجود برای مانیتور کردن قلب جنین سیستم‌های اولتراسوند می‌باشند که تنها ضربان قلب جنین و دسته معدودی از نارسایی‌های قلبی جنین را می‌توان بوسیله آنها تشخیص داد. این در حالیست که منشأ بسیاری از نارسایی‌های قلبی، در سیستم عصبی قلب و آریتمی‌های موجود در نحوه انقباض عضله قلبی می‌باشند. به همین علت مانیتور کردن سیگنال الکتروکاردیوگرام (ECG) که حاوی نحوه فعالیت الکتریکی قلب است، بسیار مفید بوده و به لحاظ کلینیکی می‌تواند جایگزین و یا مکملی برای سیستم‌های اولتراسوند فعلی باشد.

از سوی دیگر بعلت ضعیف بودن سیگنال ECG جنین، در مقابل سایر تداخل‌های محیطی و بویژه سیگنال ECG مادر، مانیتور کردن دقیق شکل موج ECG جنین تاکنون با روش‌های تهاجمی و از روی مجموعه جنین امکان‌پذیر بوده است که این روش‌ها هم منحصر به زمان زایمان می‌باشند.

هدف از این تحقیق بهبود جنبه‌های پردازشی مسئله استخراج سیگنال‌های قلبی جنین از آرایه‌ای از کانال‌های ثبت شده از روی شکم مادر و همچنین تحلیل بهتر این سیگنال‌ها می‌باشد. در مقایسه با کارهای گذشته که عموماً مبتنی بر تکنیک‌های کلاسیک پردازش سیگنال می‌باشند، نوآوری روش‌های پیشنهادی در استفاده مناسب از اطلاعات پیشین (*a priori*) موجود نسبت به سیگنال‌های قلبی، همچون ساختار شبه‌پریودیک این سیگنال‌ها می‌باشد. در این تحقیق با استفاده از این اطلاعات سعی شده است تا علاوه بر بهبود کیفیت روش‌های موجود، تکنیک‌های پردازشی مختص سیگنال‌های قلبی طراحی گردد.

از آنجا که سیگنال‌های قلبی جنین در حوزه‌های مختلف، همچون زمان، مکان، فرکانس و ویژگی‌ها با سایر سیگنال‌ها و نویزها تداخل دارند، روش‌هایی که تنها بر اطلاعات موجود در یکی از این حوزه‌ها استوار می‌باشند، قادر به تفکیک کامل سیگنال‌های ECG نمی‌باشند. لذا در روش‌های ارائه شده سعی شده است که از نقاط قوت حوزه‌های مختلف توأم استفاده شود. به لحاظ تئوریک، روش‌های پیشنهادی ترکیبی از مدل‌های مرفولوژیک ECG، فیلترهای Bayesian مطرح در تئوری تخمین، و انواع خاصی از فیلترهای مکانی مطرح در تئوری تفکیک کور و شبه‌کور منابع می‌باشند.

متمدهای ارائه شده مبتنی بر ساختار مرفولوژیک سیگنال‌های قلبی می‌باشند و با توجه به جامع بودن این روش‌ها می‌توان از آنها در پردازش سیگنال‌های قلبی بزرگسالان و همچنین در طراحی سیستم‌های مانیتورینگ بلادرنگ نیز بهره جست. همچنین با توجه به شباهت مرفولوژیک سیگنال‌های الکتریکی و مغناطیسی قلب (MCG)، کلیه روش‌های پیشنهادی بر روی سیگنال‌های مغناطیسی جنین و بزرگسالان نیز قابل اعمال می‌باشند. بطور خاص، نمونه‌ای از کاربرد این روش‌ها در تفکیک سیگنال‌های MCG جنین‌های دوقلو ارائه می‌گردد.

در کنار روش‌های فوق، تکنیک پردازشی بازگشتی ارائه خواهد شد که قادر به تفکیک سیگنال‌های زیرفضاهای مطلوب یک سیگنال از ترکیب‌های تکین سیگنال و نویز می‌باشند. این تکنیک دارای کاربردهای بسیار متنوعی در مباحث پردازش سیگنال می‌باشد.

Résumé

Les malformations cardiaques congénitales sont parmi les malformations les plus communes à la naissance et la première cause de décès des nouveau-nés. La plupart des anomalies cardiaques sont visibles dans la morphologie des signaux électriques cardiaques, qui sont enregistrés par l'électrocardiographie qui semble contenir plus d'informations par rapport aux méthodes conventionnelles sonographiques. Par conséquent, l'étude non invasive des signaux cardiaques du fœtus peut fournir un moyen efficace pour contrôler le bon fonctionnement du cœur du fœtus et peut être utilisé pour la détection précoce des anomalies cardiaques.

Dans les précédentes études, diverses méthodes ont été mises au point pour le traitement et l'extraction d'électrocardiogramme (ECG) du fœtus, à partir des signaux enregistrés de la surface du corps de la mère. Toutefois, en raison du faible rapport signal/bruit de ces signaux, l'application d'électrocardiographie fœtale a été limitée à l'analyse des battements cardiaques et à des enregistrements ECG invasifs pendant l'accouchement.

Dans cette recherche, l'objectif est d'améliorer les méthodes de traitement du signal utilisées en cardiographie du fœtus et d'apporter de nouvelles solutions à ce problème, en développant de nouvelles techniques de modélisation et de filtrage des signaux d'ECG du fœtus enregistrés par un réseau d'électrodes placées sur le ventre maternel. L'idée de base derrière les méthodes développées, consiste à utiliser les informations *a priori* des signaux cardiaques, tels que leur pseudo-périodicité, afin d'améliorer les performances des méthodes existantes et de concevoir de nouvelles techniques de filtrage qui sont spécifiques aux signaux cardiaques. En raison du recouvrement des signaux du fœtus avec les interférences/bruits dans différents domaines, les méthodes qui utilisent l'information dans un seul de ces domaines, ne réussissent pas à extraire les ECG fœtaux. Par conséquent, nous proposons des méthodes de traitement qui utilisent les informations provenant de différents domaines, afin d'améliorer la qualité des signaux extraits.

Théoriquement, les méthodes proposées sont des combinaisons de modèles morphologiques de l'ECG, de techniques de filtrage bayésienne *ad hoc* basées sur la théorie de l'estimation et de classes spéciales de filtres spatiaux issus du contexte de la séparation aveugle et semi-aveugle de sources. Il est montré que, en raison de la généralité des méthodes proposées, les mêmes procédures sont également applicables aux signaux ECG multicapteurs chez l'adulte, et peuvent être utilisées en temps réel dans les systèmes de surveillance cardiaque.

En outre, les méthodes développées sont fondées sur la morphologie du signal cardiaque, sans prendre en compte les particularités de la théorie du volume conducteur et la propagation électromagnétique dans les milieux physiologiques. Par conséquent, les mêmes méthodes sont applicables à d'autres modalités de surveillance cardiaque, comme le magnétocardiogramme (MCG), qui sont morphologiquement similaire à l'ECG. En particulier, nous présentons une étude de cas sur l'extraction des signaux MCG de jumeaux.

Nous présentons également une technique originale de déflation, qui vise à séparer les sous-espaces formés par les signaux d'intérêt dans des mélanges sous-déterminés. Cette idée s'avère très performante et débouche sur des applications diverses dans d'autres contextes.

Abstract

Congenital heart defects are among the most common birth defects and the leading cause of birth defect-related deaths. Most cardiac defects have some manifestation in the morphology of cardiac electrical signals, which are recorded by electrocardiography and are believed to contain much more information as compared with conventional sonographic methods. Therefore, the noninvasive study of fetal cardiac signals can provide an effective means of monitoring the well-being of the fetal heart and may be used for the early detection of cardiac abnormalities.

In previous studies, various methods have been developed for the processing and extraction of fetal electrocardiogram (ECG) signals recorded from the maternal body surface. However, due to the low signal-to-noise ratio of these signals, the application of fetal electrocardiography has been limited to heartbeat analysis and invasive ECG recordings during labor.

In this research, the objective is to improve the signal processing aspects of fetal cardiography and to provide better insights of this problem, by developing new techniques for the modeling and filtering of fetal ECG signals recorded from an array of electrodes placed on the maternal abdomen. The basic idea behind the developed methods is to use *a priori* information about cardiac signals, such as their pseudo-periodic structure, to improve the performance of the currently existing techniques and to design novel filtering techniques that are customized for cardiac signals. Due to the overlap of the fetal signals and interferences/noises in different domains, the methods that use the information in only one of these domains do not usually succeed in extracting the fetal ECG. Therefore, we design methods that use the information from various domains, in order to improve the quality of the extracted signals.

Theoretically, the proposed methods are combinations of morphological models of the ECG, *ad hoc* Bayesian filtering techniques based on estimation theory, and special classes of spatial filters adapted from the blind and semi-blind source separation context. It is shown that due to the generality of the proposed methods, the same procedures are also applicable to multichannel adult ECG recordings and can be used in real-time cardiac monitoring systems.

Moreover, the developed methods are based on the cardiac signal morphology without going into the details of volume conduction theory and the conductivities of the propagation media. Hence, the same methods are applicable to other cardiac monitoring modalities such as the magnetocardiogram (MCG), which are morphologically similar to the ECG. We specifically present a case study on the extraction of twin fetal MCG signals.

We also present an advanced deflation technique, which is able to separate subspaces of desired signals from degenerate mixtures of signal and noise. This idea has found various applications in other contexts.

Contents

1	Introduction	1
1.1	Overview of the Thesis and Contributions	2
2	State of the Art	7
2.1	Introduction	7
2.2	Historical Review of the Early Works	7
2.3	Objectives	9
2.4	Methodologies	9
2.4.1	Data Collection	9
2.4.2	Data Analysis	9
2.5	Forward Modeling vs. Inverse Solutions	10
2.6	Alternative Measurement Techniques	11
2.7	Current Challenges and Problem Definition	11
2.8	Summary and Conclusions	12
3	Electrophysiology of the Fetal Heart	13
3.1	Introduction	13
3.2	Fetal Cardiac Development	13
3.2.1	Fetomaternal Compartments	13
3.2.2	Fetal Presentations	14
3.3	Physiology of the Fetal Heart	16
3.3.1	Electrical Activity of the Fetal Heart	16
3.3.2	The Electrocardiogram	16
3.4	The Electromagnetic Basis of the Electrocardiogram	17
3.4.1	Multipole Expansion of Body Surface Potentials	19
3.4.2	Monopole and Dipole Approximations	20
3.5	Summary and Conclusions	20
4	Synthetic ECG Generation	21
4.1	Introduction	21
4.2	The Cardiac Dipole vs. the Electrocardiogram	22
4.3	A Synthetic ECG Generator	23
4.4	Cardiac Dipole Vector and ECG Modeling	24
4.4.1	Multichannel ECG modeling	25
4.4.2	Modeling maternal abdominal recordings	26
4.4.3	Fitting the model parameters to real recordings	27
4.5	ECG Noise Modeling	28
4.6	Illustrations	30
4.6.1	The model accuracy	30
4.6.2	Fetal ECG extraction	31
4.7	Time-varying volume conductor models	35
4.8	Summary and Conclusions	35

5	A Bayesian ECG Filtering Framework	37
5.1	Introduction	37
5.2	Review of the Bayesian Filtering Theory	37
5.2.1	The Extended Kalman Filter	38
5.2.2	The Extended Kalman Smoother	39
5.2.3	The Unscented Kalman Filter	39
5.3	Methods	39
5.3.1	Linearization of the Nonlinear Dynamic ECG Model	40
5.3.2	Observation Equations	40
5.3.3	Estimation of the Model Parameters	41
5.3.4	Stability and Convergence Issues	43
5.3.5	Practical Filtering Scheme	44
5.4	Evaluation on Simulated Noisy Mixtures	44
5.4.1	The Dataset	44
5.4.2	Noise Generation	45
5.4.3	Implementation	45
5.4.4	Benchmark Methods	46
5.4.5	Results	46
5.5	Evaluation on Real Maternal/Fetal Mixtures	50
5.6	Multichannel Extension	52
5.7	Summary and Conclusions	53
6	Linear Multichannel Analysis of Cardiac Signals	55
6.1	Introduction	55
6.2	Electrophysiology of the Heart: A Signal Processing Perspective	55
6.3	A Study of the Components Extracted from Multichannel Cardiac Recordings	57
6.3.1	ECG Dimensionality	57
6.3.2	Interpretation of Independent Components Extracted from Cardiac Recordings	60
6.3.3	Necessity of Preprocessing	63
6.4	Curse of Dimensionality and Over-fitting	65
6.5	Optimal Sensor Number and Positioning for Fetal ECG Extraction	68
6.6	Summary and Conclusions	68
7	Multichannel Electrocardiogram Decomposition using Periodic Component Analysis	69
7.1	Introduction	69
7.2	Background	70
7.2.1	Generalized Eigenvalue Decomposition	70
7.2.2	ICA versus Generalized Eigenvalue Decomposition	70
7.2.3	Periodic Component Analysis	71
7.3	Modifications for Time-varying Periods	72
7.4	Results	73
7.5	Multichannel Fetal MCG in Multiple Pregnancy: A Case Study	76
7.5.1	Baseline Wander Removal and Preprocessing	77
7.5.2	Maternal MCG Removal	77
7.5.3	Fetal MCG Extraction - First Run	78
7.5.4	Fetal Peak Detection and HRV Calculation - First Run	80
7.5.5	Fetal MCG Extraction - Second Run	82
7.5.6	Fetal Movement Tracking	83
7.6	Summary and Conclusions	85
8	A Deflation Procedure for Subspace Decomposition	89
8.1	Introduction	89
8.2	Data Model	89
8.3	Method	90
8.3.1	Single Channel Denoising	90
8.3.2	Linear Decomposition using GEVD	90

8.3.3	Iterative Subspace Decomposition	92
8.4	Application in Fetal ECG Extraction	93
8.4.1	Simulated Data	94
8.4.2	Real Data	97
8.5	Summary and Conclusion	100
9	Model-Based Bayesian Filtering of Cardiac Contaminants from Biomedical Recordings	103
9.1	Introduction	103
9.2	Cardiac Artifact Removal	103
9.3	Bayesian Filtering versus Conventional Techniques	105
9.4	Applications	106
9.4.1	Simulated Data	106
9.4.2	Single-Channel Fetal ECG Extraction and HRV Analysis	108
9.4.3	EEG Denoising	109
9.4.4	EMG Signal Denoising	112
9.5	Summary and Conclusions	113
10	Conclusions and Future Works	115
10.1	Future Works	116
A	Angles Between Subspaces	121
B	On the Spikes and Bumps Phenomenon	123
C	Multichannel Decomposition of Distributed Sources	127
D	The Open Source ECG Toolbox (OSET)	129
D.1	Baseline Wander Removal	129
D.2	R-Peak Detectors	130
D.3	Power-Line Noise Notch Filter	131

List of Figures

1.1	The amplitude and frequency range of biosignals that can interfere with fetal cardiac signals [218, 48, 187]. The labels in this figure stand for the maternal electrocardiogram (mECG), electroencephalogram (mEEG), electrohystrogram (mEHG), electrooculogram (mEOG), electromyogram (mEMG), electrohystrogram (mEHG), and the fetal ECG (fECG). Note that the amplitude of these signals also depends on the site from which the data is recorded.	2
1.2	An illustration of the overlap of the fetal ECG signals with the maternal ECG and other interferences and noise. We can see that the fetal signals overlap with the undesired signals in time, space, frequency, and feature domains; therefore, methods that only work in one specific domain are not capable to fully separate the fetal cardiac signals.	2
1.3	The conceptual link between the chapters.	3
2.1	Number of publications in the field of fetal electro- and magneto-cardiography, registered in PubMed [204].	8
2.2	Number of publications in the field of electro- and magneto-cardiography, registered in PubMed [204].	8
2.3	The percentage of publications in the field of electro- and magneto-cardiography normalized by the number of publications in all subjects, registered in PubMed [204].	8
3.1	The fetus and its heart in the early stages of development; adopted with permission from [118].	13
3.2	Developments of the fetal heart during gestation; adopted with permission from [118].	14
3.3	The major fetomaternal compartments that influence the fetal cardiac surface potentials; adopted with permission from [190]. The <i>vernix caseosa</i> is formed over the fetal skin.	15
3.4	Different fetal presentations and the percentage of incidence at the end of gestation; adopted with permission from [190].	15
3.5	Different fetal vertex positions and their incidence at the end of gestation; adopted with permission from [190]. In the right side abbreviations, L stands for left, R for right, O for occiput, A for anterior, T for transverse and P for posterior. The LOA position is seen in the left graph.	15
3.6	The anatomy of the fetal heart; adopted with permission from [190].	16
3.7	The activation cycle of the fetal heart; adopted with permission from [190].	17
3.8	The increasing complexity in fetal HRV time-series in different gestational ages; adopted from [214].	18
3.9	Illustration of the vectorcardiogram loop.	20
4.1	The three body axes, adapted from [127].	22
4.2	Synthetic ECG signals of the Frank lead electrodes	25
4.3	A typical synthetic VCG loop. Arrows indicate the direction of rotation. Each clinical lead is produced by mapping this trajectory onto a 1-dimensional vector in this three-dimensional space.	25
4.4	Illustration of the fetal and maternal VCGs vs. their body coordinates	27
4.5	Typical segment of ECG BW Noise (a) Original (b) Synthetic	29
4.6	Frequency response magnitudes of 32 segments of the time-varying AR filters for the Baseline Wander noises of the NSTDB. This figure illustrates how the AR filter responses are evolving in time.	29

4.7	Original vs. synthetic VCGs and ECGs using 5 and 9 Gaussian functions. For comparison, the ECG reconstructed from the Dower transformation is also depicted in (d)-(f) over the original ECGs. The synthetic VCGs and ECGs have been vertically shifted 0.2mV for better comparison. Refer to text for details.	31
4.8	Synthetic maternal abdomen signals with $\theta_{mf} = 10^\circ$, SIR=-30dB, $\text{SNR}_\eta=-20\text{dB}$, and $\text{SNR}_n=5\text{dB}$	33
4.9	Synthetic maternal abdomen signals with $\theta_{mf} = 60^\circ$, SIR=-30dB, $\text{SNR}_\eta=-20\text{dB}$, and $\text{SNR}_n=5\text{dB}$	33
4.10	Independent Components extracted from the synthetic multichannel recordings of Fig. 4.8. Strong maternal components can be seen in the first three components. Components four to six, more or less correspond with the baseline wander, muscle artifacts, and electrode movement noises. Noisy fetal R-peaks are seen in the seventh component and the last component is principally noise, with some minor traces of the fetal peaks.	34
4.11	Independent Components extracted from the synthetic multichannel recordings of Fig. 4.9. Strong maternal components can be seen in the first three components. Components four to six, more or less correspond with the baseline wander, muscle artifacts, and electrode movement noises. Fetal R-peaks are seen in the last two components.	34
5.1	An illustration of the phase assignment approach	41
5.2	Several cycles of the ECG phase-wrapped in the state space	41
5.3	An average and standard deviation-bar of 30 ECG cycles of a noisy ECG	42
5.4	Typical filtering results for an input signal of 6dB.	47
5.5	The mean (top) and Standard Deviation (bottom) of the filter output SNR improvements versus different input SNRs. In these curves EKF, UKF, and EKS correspond to the results without $\sigma_{v_k}^2$ adaptation, and EKF2, UKF2, and EKS2 refer to the ones with $\sigma_{v_k}^2$ adaptation. Refer to text for further details.	48
5.6	The mean (top) and Standard Deviation (bottom) of the filter output SNR improvements over the ST-segment of the ECG, versus different input SNRs. In these curves EKF, UKF, and EKS correspond to the results without $\sigma_{v_k}^2$ adaptation, and EKF2, UKF2, and EKS2 refer to the ones with $\sigma_{v_k}^2$ adaptation. Refer to text for further details.	49
5.7	Estimated signal SNR versus true SNR, from the isoelectric line segment of the ECG between the T-waves and the P-waves (solid), and from the whole ECG cycle (dashed).	50
5.8	Effect of noise color on the EKS results. (top) without $\sigma_{v_k}^2$ adaptation, (bottom) with $\sigma_{v_k}^2$ adaptation. Both curves merge for input SNRs below 10dB.	50
5.9	The first channel of the DaISy dataset [46], recorded from a maternal abdominal lead before and after the EKS procedure.	51
5.10	Several fetal ECG beats adopted from Fig. 5.9, before and after the post-processing EKS, together with the $\pm\sigma$ and $\pm 3\sigma$ confidence envelopes.	52
6.1	A typical segment of the standard 12-lead and the 3 Frank lead ECGs, adopted from the PTB diagnostic database [158, 22, 112].	58
6.2	Principal components extracted from the ECG channels of Fig. 6.1 in descending order of their corresponding eigenvalues.	59
6.3	Independent components extracted from the ECG channels of Fig. 6.1 using the JADE algorithm in descending order of their energy contribution (explained variance).	59
6.4	Three-dimensional scatter plot of the signals of Fig. 6.1.	60
6.5	The correlation dimensions of (a) a typical ECG, and (b) white Gaussian noise	61
6.6	Sensitivity of different ECG components to dimension reduction by applying PCA to eleven-dimensional data	61
6.7	Sensitivity of different ECG components to dimension reduction by applying JADE to eleven-dimensional data	62
6.8	Different views of the VCG representation of a typical ECG signal recorded from the Frank lead electrodes, together with column vectors of the mixing matrix estimated by JADE	62
6.9	A segment of independent components extracted from a typical ECG dataset.	63
6.10	A segment of independent components extracted from a typical ECG signal containing baseline wander. The baseline wander can be seen in most of the extracted components.	64

6.11	Different views of the VCG representation of a typical ECG signal having baseline wander.	64
6.12	The effect of PDF asymmetry on the extracted ICs for an (a) asymmetric PDF, and the (b) PDF made symmetric using equation (6.5). The vectors represent the mixing vectors estimated by JADE (in positive and negative directions), centered at the mean of the distribution.	65
6.13	Typical signals extracted from 72 maternal abdominal channels using FastICA. The last three components are typical ‘bumpy’ signals without any physical origin and due to the ‘spikes and bumps’ phenomenon [90, 179, 178].	66
6.14	Typical bumpy signals extracted from 40 channels of brown noise using FastICA.	67
7.1	General scheme of ICA algorithms with spatial whitening	70
7.2	Illustration of the phase assignment procedure used for calculating τ_t in each ECG beat.	72
7.3	Polar representation of a noisy ECG using the ECG phase $\phi(t)$.	72
7.4	Periodic components extracted from the ECG channels of Fig. 6.1 in descending order of periodicity.	74
7.5	Sensitivity of different ECG components to dimension reduction by applying π CA to eleven-dimensional data of Fig. 6.1	74
7.6	The DaISy dataset consisting of five maternal abdominal and three thoracic channels [46].	74
7.7	Independent components extracted from the dataset of Fig. 7.6, using JADE. Notice that components 1, 2, 3, and 5 correspond to the maternal subspace and components 4 and 8 to the fetal subspace.	75
7.8	Periodic components extracted from the dataset with maternal ECG beat synchronization. The maternal ECG contribution reduces from top to bottom.	75
7.9	Periodic components extracted from the dataset with fetal ECG beat synchronization. Notice that the fetal ECG contribution reduces from top to bottom.	76
7.10	Periodic components extracted from the dataset with maternal & fetal ECG beat synchronization. The maternal (fetal) ECG contribution reduces (increases) from top to bottom.	76
7.11	A typical segment of five channels after the preprocessing step	78
7.12	The first ten periodic components extracted from a typical segment of the data corresponding to the maternal MCG	79
7.13	The logarithm of the eigenvalues normalized by the first eigenvalue (cf. [120])	80
7.14	Independent components extracted from a typical segment of the data after preprocessing and maternal MCG cancellation using π CA	81
7.15	Overlaid and zoomed plot of components 4, 9, and 10 from Fig. 7.14. We can see that the green and red plots correspond to one of the fetuses while the blue plot corresponds to the other fetus	81
7.16	A primary estimate of the fetal HRVs	82
7.17	A zoomed view of the second fetal HRV plot of Fig. 7.16	82
7.18	The first ten periodic components extracted by using a mixed measure of maternal/fetal MCGs, proposed in Section 7.5.5. Note that these components are not as clean as the maternal components in Fig. 7.12, which were based on the maternal periodicity alone.	84
7.19	The last ten periodic components extracted by using a mixed measure of maternal/fetal MCGs, proposed in Section 7.5.5	85
7.20	Independent components extracted from a typical segment of the data in the post-processing step (Section 7.5.5)	86
7.21	Rotation angles (in degrees) calculated according to the procedure presented in Section 7.5.6	86
8.1	The general iterative subspace decomposition scheme	92
8.2	Maternal phase calculation procedure	94
8.3	The overall iterative procedure for maternal ECG cancellation	94
8.4	The L -channel Kalman filter for maternal ECG removal	94
8.5	The mean and standard deviation bar of SINR improvements achieved in different input SINRs for (a) $\theta_{mf} < 10^\circ$ (b) $\theta_{mf} > 60^\circ$	96
8.6	The results of the proposed method on the DaISy fetal dataset in five iterations	98
8.7	Comparison of the fetal ensemble average extracted by (a) ICA and (b) the deflation method followed by π CA from the DaISy database	99

8.8	Residual fetal components extracted by ICA (in red) and the deflation procedure (in blue) from the DaISy database	99
8.9	The results of the proposed method on the second fetal dataset in five iterations	101
9.1	The overall denoising scheme. As shown in this figure the R-peaks of the contaminating signals (CC) may either be detected from an arbitrary reference ECG or from the noisy biosignal after baseline wander (BW) removal.	104
9.2	Results of the EKS on a mixture of simulated signal plus ECG artifacts. The (a) original, (b) noisy, and (c) denoised signals can be seen in this figure. The SNR of the noisy signal was improved from 7.6dB to 15.1dB using the EKS.	106
9.3	SNR improvement results achieved on simulated data in different input SNRs and different spectral colors.	107
9.4	Periodicity measures (PM) achieved on simulated data in different input SNRs and different spectral colors, (a) before and (b) after denoising. The decrease of PM is used as a necessary, but not sufficient, measure of the filtering performance.	108
9.5	A typical segment of channel 1 of the dataset, including the fetal and maternal ECG . . .	108
9.6	The signal of Fig. 9.5, after maternal ECG removal using the EKS	109
9.7	(a) A segment of the original data, (b) the fetal signals after mECG removal, and (c) the estimated fetal HRV. We can see that the fetal R-peaks have been misdetected at some points and have caused errors in the HRV sequence.	109
9.8	The average and SD-bars of the (a) ECG, (b) EEG before denoising, and (c) EEG after denoising, by synchronous averaging of the data using the ECG R-peaks.	111
9.9	Results of the proposed method on a segment of EEG signals recorded during an fMRI experiment. (a) Reference ECG channel, (b) Noisy AF ₈ EEG channel, (c) AF ₈ channel after ECG removal.	111
9.10	The spectral estimate of the original and denoised EEG signals, from the fMRI experiment data. The difference of the two spectra is plotted at the bottom.	112
9.11	Results of the proposed method on a segment of EEG signals from the MIT-BIH Polysomnographic Database. (a) Reference ECG channel, (b) Noisy C ₄ -A ₁ EEG channel, (c) C ₄ -A ₁ channel after ECG removal.	112
9.12	The spectra of the original and denoised EEG signals, from the MIT-BIH Polysomnographic Database. The difference of the two spectra is plotted at the bottom.	113
9.13	Results of the proposed method on EMG recordings highly contaminated with ECG. . . .	114
C.1	Distributed source model and the electrode positions	128
D.1	The mean and standard deviation bars of SNR improvements achieved in different input SNRs for BW removal based on (MD2) two-stage MWM filter, (MD1) single-stage MWM filter, (MN2) two-stage MA filter, (LP) zero-lag lowpass filter, and (BP) bandpass DFT filter	130
D.2	A segment of diluted, denoised, and original ECG signal (from top to bottom)	132
D.3	Snapshot of the main interface and workspace of the OSET GUI	133
D.4	Snapshot of some of the functions of the OSET GUI	133
D.5	Multichannel analysis panel containing some of the implemented procedures	134
D.6	Signal plot panel	134
D.7	Kalman filter parameter optimization panel	135

List of Tables

4.1	Parameters of the synthetic ECG model in (4.4)	23
4.2	Parameters of the synthetic model presented in (4.7) for the ECGs and VCG plotted in Figs. 4.2 and 4.3	26
4.3	The normalized MSE (%) in the synthetic VCG channels using five and nine Gaussian functions	30
4.4	The percentage of MSE (%) in the ECGs reconstructed by Dower transformation from the original VCG and from the synthetic VCG using five and nine Gaussian functions	30
4.5	Parameters of the synthetic fetal dipole used in Section 4.6.2	32
5.1	Parameters of the process and observation noises	45
5.2	WD parameter combinations tested over the database	46
7.1	Description of the recorded channels	77
8.1	fetal ECG data PM in different iterations for the DaISy database	98
8.2	fetal ECG data PM in different iterations for the second database	100
9.1	Rules chart for the manipulation of the filter parameters	105
10.1	Comparison of the different methods and their applications	118

Abbreviations

AF	Adaptive Filter
BCG	Ballistocardiogram
BSS	Blind Source Separation
BW	Baseline Wander
CC	Cardiac Contaminants
EA	Ensemble Average
ECG	Electrocardiogram
EEG	Electroencephalogram
EKF	Extended Kalman Filter
EKS	Extended Kalman Smoother
fECG	Fetal Electrocardiogram
fMRI	Functional Magnetic Resonance Imaging
HRV	Heart Rate Variability
ICA	Independent Component Analysis
JADE	Joint Approximate Diagonalization of Eigen-matrices
KF	Kalman Filter
mECG	Maternal Electrocardiogram
MCG	Magnetocardiogram
π CA	Periodic Component Analysis
PCA	Principal Component Analysis
PF	Particle Filter
PL	Power Line
PVC	Premature Ventricular Contraction
SINR	Signal-to-Interference plus Noise Ratio
SIR	Signal-to-Interference Ratio
SNR	Signal-to-Noise Ratio
SVD	Singular Value Decomposition
UKF	Unscented Kalman Filter
VCG	Vectorcardiogram
WSS	Wide-Sense Stationary

Chapter 1

Introduction

Heart defects are among the most common birth defects and the leading cause of birth defect-related deaths [129, 136]. Every year, about one out of 125 babies are born with some form of congenital heart defects¹ [6]. The defect may be so slight that the baby appears healthy for many years after birth, or so severe that its life is in immediate danger. Congenital heart defects originate in early stages of pregnancy when the heart is forming and they can affect any of the parts or functions of the heart. Cardiac anomalies may occur due to a genetic syndrome, inherited disorder, or environmental factors such as infections or drug misuse [150, 129]. They can also occur due to specific fetal positioning that chokes the umbilical cord [230]. In any case, the regular monitoring of the fetal heart and the early detection of cardiac abnormalities can help obstetrics and pediatric cardiologist² to prescribe proper medications in time, or to consider the necessary precautions during delivery or after birth.

Most cardiac defects have some manifestation in the morphology of cardiac electrical signals, which are recorded by electrocardiography and are believed to contain much more information as compared with conventional sonographic methods [155]. However, due to the low SNR of fetal electrocardiogram (ECG) recorded from the maternal body surface, the application of fetal electrocardiography has been limited to heart-beat analysis and invasive ECG recordings during labor.

In Fig. 1.1 the amplitude and frequency range of fetal ECG have been compared with other noises and artifacts. Accordingly, the fetal ECG is much weaker than the other interfering biosignals. Moreover, from the signal processing perspective, there is no specific domain (time, space, frequency, or feature) in which the fetal ECG can be totally separated from the interfering signals (Fig. 1.2).

Although in previous works, a large body of research has been devoted to the filtering of fetal cardiac signals, due to the complexity of the problem there are still many open issues that require improved signal processing techniques.

In this research, the objective is to improve the signal processing aspects of fetal cardiography and to provide better insights of this problem, by developing new techniques for the modeling and filtering of fetal ECG signals recorded from an array of electrodes placed on the maternal abdomen. The basic idea behind the developed methods is to use *a priori* information about cardiac signals, such as their pseudo-periodic structure, to improve the performance of the currently existing techniques and to design novel filtering techniques specific to cardiac signals. Due to the overlap of the fetal signals and interferences/noises in different domains (Fig. 1.2), the methods that use the information in only one of these domains do not usually succeed in extracting the fetal ECG. Therefore, the objective is to design methods that use the information from various domains (including time, frequency and space), in order to improve the quality of the extracted signals. It is shown that due to the generality of the proposed methods, the same procedures are also applicable to multichannel adult ECG recordings and can be used in real-time cardiac monitoring systems.

The hereby developed methods are based on the cardiac signal morphology and we do not go into the details of volume conduction theory and the conductivities of the propagation media. We will show that the same methods are applicable to other cardiac monitoring modalities such as the magnetocardiogram (MCG), which are morphologically similar to the ECG. Therefore, throughout the manuscript, unless specifically noted, all the methods developed for the ECG are also applicable to MCG recordings.

¹ A condition is called ‘congenital’ when it is present at birth [129].

² Pediatric cardiologist: children’s heart disease specialist




Figure 1.1: The amplitude and frequency range of biosignals that can interfere with fetal cardiac signals [218, 48, 187]. The labels in this figure stand for the maternal electrocardiogram (mECG), electroencephalogram (mEEG), electrohystrogram (mEHG), electrooculogram (mEOG), electromyogram (mEMG), electrohystrogram (mEHG), and the fetal ECG (fECG). Note that the amplitude of these signals also depends on the site from which the data is recorded.

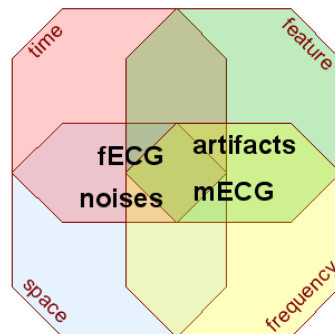


Figure 1.2: An illustration of the overlap of the fetal ECG signals with the maternal ECG and other interferences and noise. We can see that the fetal signals overlap with the undesired signals in time, space, frequency, and feature domains; therefore, methods that only work in one specific domain are not capable to fully separate the fetal cardiac signals.

1.1 Overview of the Thesis and Contributions

In Fig. 1.3, the conceptual links between the different chapters of the manuscript are depicted. While the next two chapters provide the necessary background, the main contributions of this research are presented in Chapter 4 to Chapter 9, plus some minor contributions in the formulations at the end of Chapter 3. The discrete classical methods used throughout the different chapters for benchmarking the proposed techniques are otherwise highlighted with citation to the original works. In what follows, we will have a brief overview of each chapter.

Chapter 2

In this chapter, the history of fetal electrocardiography and a selection of the related literature are reviewed. At the end of this chapter we will define the problem of interest and the objectives of this

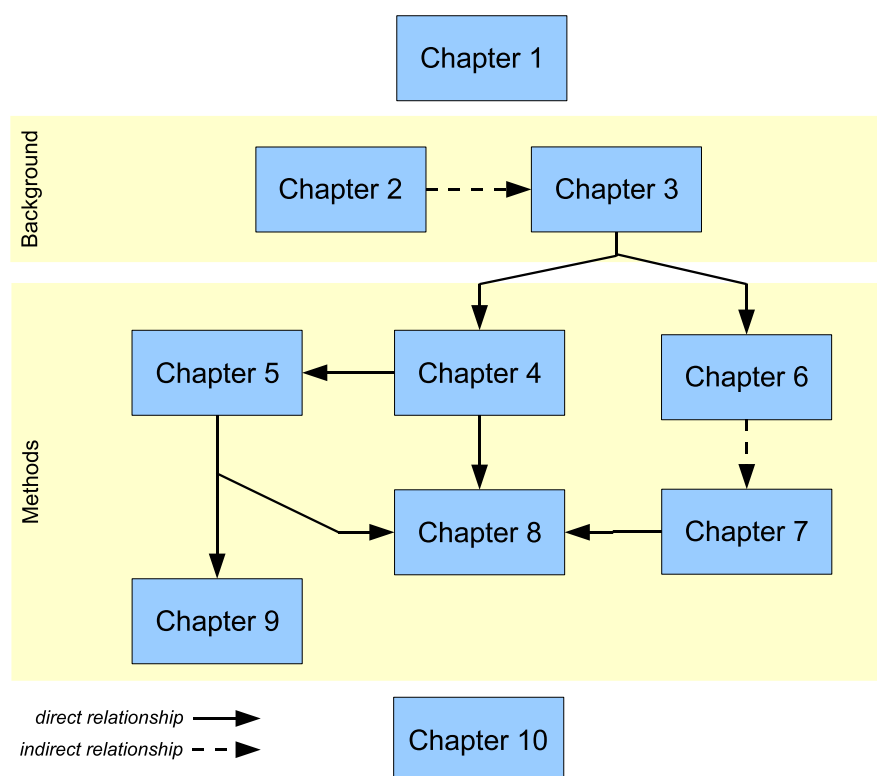


Figure 1.3: The conceptual link between the chapters.

work.

Chapter 3

In this chapter, the electrophysiology of the fetal heart and fetal cardiac development stages are presented. We also review the dipole model of the heart, which may be considered as the theoretical basis of electrocardiography. These issues are required in later chapters for the development of the proposed techniques and the interpretation of the extracted fetal signals. Although the theory of electrocardiography is rather developed and based on well-known electromagnetic principles, some of the physical properties presented in Section 3.4.1, have not yet been considered in the signal processing community.

Chapter 4

In order to evaluate single and multi-channel filtering techniques, suitable models are required for generating synthetic mixtures of fetal cardiac signals plus interferences and noise. In this chapter, based on the *dipole theory* of the heart and a previously existing dynamic model for single-channel adult ECG generation, a morphologic model is presented that can be used to generate an arbitrary number of synthetic ECG channels in single and multiple pregnancies, for different fetal positions, and with variable amounts of maternal ECG interference and noise. This model is used in later chapters for the evaluation and comparison of the proposed filtering techniques.

A shorter version of this chapter was published in [169]. More recently the model was also extended to abnormal beat generation, using Hidden Markov Models [40].

Chapter 5

In previous works many methods have been proposed for ECG denoising. These methods are based on the separation of the signal and noise in time, frequency, scale, or feature domain. In this chapter, based on the pseudo-periodic structure of ECG signals, a Bayesian filtering framework is presented for ECG denoising. The idea is based on a dynamic model originally proposed for generating synthetic adult ECG

[133]. It is shown that the presented framework is a logical extension of some of the existing methods and that it outperforms the most common filtering schemes. We will also show that this framework can be effectively used for denoising fetal ECG recordings and finding confidence intervals for the estimated components.

Preliminary versions of this Bayesian framework were reported in [172, 174] and a full version of the idea with quantitative comparisons with classical methods was presented in [175]. A three-dimensional extension of this model was also presented in [176].

Chapter 6

Multichannel linear transforms are commonly used for the decomposition of mixtures of fetal ECG and interferences. While in previous studies much attention has been paid to the denoising aspects of these approaches, the interpretation of linear decomposition of mixtures of ECG signals has not been well-studied. In this chapter, after a short overview of the multipole theory and electrophysiological properties of ECG signals, through several examples and case studies, we will present some general explanations and interpretations for linearly extracted ECG signals. Additional discussions concerning the dimensionality of multichannel ECG, necessity of preprocessing, practical problems concerning high-dimensional data and electrode positioning are also presented. The illustrations in this chapter are rather subjective and intended to give a better insight for designing ECG processing schemes and interpreting their results. Nevertheless, the presented ideas are based on theoretical facts and we have experimentally validated them over many existing datasets. These ideas are used in later chapters for designing specific filtering schemes that are specific to ECG signals.

Preliminary results of the finding of this chapter were presented in [177, 170, 100].

Chapter 7

The existing blind and semi-blind decomposition techniques are commonly based on the maximization of measures of component separation such as uncorrelatedness or independence. While these criteria are rather effective for blind decomposition, for semi-blind decomposition problems, i.e., where we have considerable *a priori* information about the desired signals, they are not necessarily the optimal approach. In these cases, it is more appropriate to design *ad hoc* source separation techniques that are customized to the prior information. In this chapter, we will present a linear multichannel source separation technique that is specifically designed for pseudo-periodic signals such as the ECG, with possible inter-beat variations of the RR-interval. It is shown that by this method, we can find ECG components ranked in order of relevance to the ECG periodicity. This procedure is of especial interest for the automatic removal of maternal ECG contaminations from fetal recordings. At the end of this chapter a case study is presented, in which a combination of techniques are used for the extraction and tracking of twin MCG signals in long magnetic recordings. It is shown that the movements of the fetuses can also be tracked from the components linearly extracted from multichannel MCGs. Through this example we show how different methods can be integrated together for extracting and denoising fetal cardiac signals.

The idea developed in the first part of this chapter was presented in [171], and was later used in a similar context for improving the performance of sequential source extraction methods [203].

Chapter 8

Following the ideas of Chapter 6, adult and fetal ECG signals should not be considered as signals having a finite and fixed number of dimensions. Instead, they should be considered as infinite dimensional signals that can only be ‘approximated’ with lower-dimensional subspaces. Therefore, by applying linear decomposition methods to mixtures of maternal and fetal ECG signals (with few number of electrodes), we can always have problems of rank deficiency, which may result in spurious residues of the maternal components within the fetal signals, or vice versa. This is a special case of the more general problem of extracting desired signals from degenerate mixtures of signal plus interference and noise. In this chapter, a deflation procedure is presented by which a desired signal subspace can be extracted from a (possibly) degenerate mixture of signals and full-rank noise. The proposed method consists of decomposition and denoising steps that are repeatedly applied to the noisy signal. While the idea is rather general and may be adapted to different applications, the presented results are based on the decomposition method presented in Chapter 7 and the Bayesian approach of Chapter 5. The method is specifically evaluated on real and simulated mixtures of fetal ECG signals.

Based on this method, we have recently applied for two patents in the domain of “fetal ECG/MCG denoising” and the “removal of cardiac interferences from other biosignal recordings”. Other applications of this framework were also reported in [75, 7].

Chapter 9

In this chapter, the Bayesian ECG filtering framework developed in Chapter 5, is extended to the removal of cardiac artifacts, such as ECG, MCG, Ballistocardiogram, etc., from other biomedical recordings, using as few as one single cardiac reference channel. This framework may be considered as a byproduct of the methods that we developed. The applications of this method are studied using simulated and real signals.

This framework together with various examples were recently published in [173].

Chapter 10

In the last chapter, we summarize the findings of the current work and their points of strength and weakness, as compared with previous methods. We will also present some of the possible directions of research that are left as open problems for future studies.

Appendix A

In this appendix, the theory of principal angles between subspaces is reviewed.

Appendix B

An over-fitting phenomenon, known as the problem of ‘spikes and bumps’, is introduced in Section 6.4. This phenomenon is very common in independent decomposition of high-dimensional noisy signals. In this appendix, some of our preliminary findings for explaining this phenomenon are presented.

Appendix C

In this appendix, the relationship between the physical properties of distributed sources and the covariance matrix of the corresponding observation signals are found in a simplified case. The derived equations, illustrate some of the ideas developed in Chapter 6.

Appendix D

The proposed methods also had some interesting byproducts that are not limited to fetal ECG extraction and denoising. In this appendix we present general explanations about some of these byproducts which had methodological novelty and may be considered as additional contributions of this work.

Matlab[®] implementations of many of the hereby presented methods are provided in our Open Source ECG Toolbox (OSET), online available at [167].

Chapter 2

State of the Art

2.1 Introduction

In this chapter, the state of the art in fetal cardiac signal extraction and analysis, before and during the current work are reviewed. Due to the rather old history of the problem and the rich literature in this field, it is not possible to cover all the existing methods in their details. Also due to the complexity of the problem, many of the existing methods have used a combination of approaches, some of which have been borrowed from other contexts. Therefore, we will review a selection of the available literature with special focus on the most significant ones, which have been specifically developed for the problem of interest. The more detailed literature is presented separately in each chapter.

2.2 Historical Review of the Early Works

The fetal electrocardiogram was first observed by M. Cremer in 1906 [42]. The early works in this area were done by using galvanometric apparatus of that time, which were limited by the very low amplitude of the fetal signals. As measurement and amplification techniques improved, fetal electrocardiography became more feasible and popular [124]. The limiting factor was then the low fetal SNR, especially in presence of the strong maternal cardiac interference; a problem which exists up to now. A few decades later, with the developments in computer science and signal processing techniques, automatic signal processing and adaptive filtering techniques were used for fetal R-wave detection [57], and maternal cardiac interference cancellation [220, 207]. The issue has since been considered as a challenging problem for both biomedical and signal processing communities.

In order to give an idea of the previous and current research interest in this field, the number of publications in the field of fetal electro and magneto-cardiography, which have been registered in PubMed [204], can be seen in Fig. 2.1. We can see that after a sharp peak in the 1960's, the trend seems to have been decreasing until the year 2000 (of course in terms of the number of publications). But in the last decade, the interest has again increased, especially for fetal magnetocardiography. This should be partially seen as a result of novel low-noise and low-price measurement and digitizing systems, and partially due to the developments in array signal processing and adaptive filtering techniques. Of course, comparing this with the total number of publications in the field of electro and magneto-cardiography (both for adults and fetuses) in Fig. 2.2, we notice that fetal cardiography is still in its preliminary stages and there is still a long way to go, for making fetal cardiography a clinically trustable means of fetal cardiac monitoring. It should also be noted that as illustrated in Fig. 2.3, despite the increase in the number of ECG/MCG researches in Fig. 2.2, when normalizing the number of these works by the total number of publications registered in PubMed over the same period, we notice that the percentage of researches in ECG studies has decreased since the 1980's, while MCG research has gained more interest.

In the following, we categorize some of the existing methods in this area from different viewpoints and study their cons and pros.

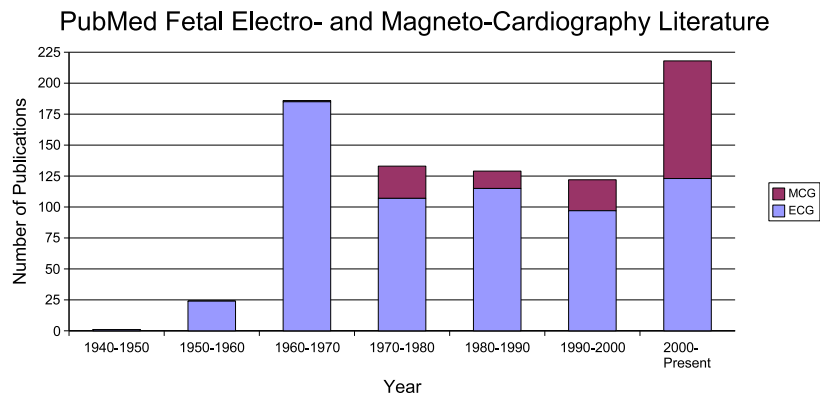


Figure 2.1: Number of publications in the field of **fetal** electro- and magneto-cardiography, registered in PubMed [204].




Figure 2.2: Number of publications in the field of electro- and magneto-cardiography, registered in PubMed [204].

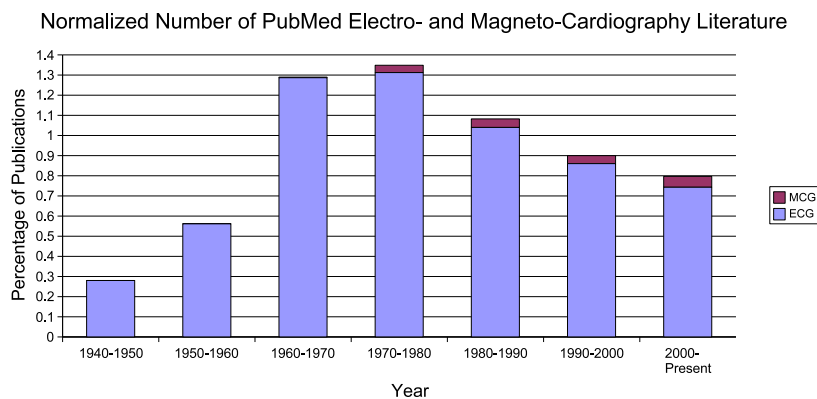


Figure 2.3: The percentage of publications in the field of electro- and magneto-cardiography normalized by the number of publications in all subjects, registered in PubMed [204].

2.3 Objectives

Previous studies have followed one of these objectives:

- Fetal heart-rate analysis
- Fetal ECG morphology analysis

The fetal ECG morphology contains much more clinical information as compared to the heart-rate alone. However, due to the low SNR of the fetal signals, it is a more challenging task. Therefore, most of the previous works have only achieved in extracting the fetal RR-intervals using the R-peaks or ensemble averages of the fetal ECG waveforms. The complete morphologic study of the fetal ECG, on a beat-to-beat basis, is therefore left as a challenging problem.

2.4 Methodologies

2.4.1 Data Collection

Fetal ECG data collection is either *invasive* or *noninvasive*. In invasive methods, the recording electrodes have direct contact with the fetal skin or scalp, which may only be achieved by using an *intrauterine* electrode during labor [149, 69, 114]. The signals recorded by invasive methods have better quality as compared with noninvasive methods; but the procedure is rather inconvenient and is limited to during labor. On the other hand, noninvasive methods use the signals recorded from the maternal abdomen, they can be done in any stage of pregnancy using dozens of electrodes. However, the low fetal ECG SNR and the other interferences are the limiting factors of this method. Nevertheless, due to the numerous advantages of the noninvasive method, a large body of research has been conducted towards the development of signal processing techniques for retrieving the fetal ECG from noninvasive recordings.

2.4.2 Data Analysis

We can categorize the existing fetal data analysis literature by their methodologies. The existing methods in this area include:

Direct Fetal ECG Analysis

In early works, fetal ECG detection was done over raw data without any processing. For instance in [117], some special cases were reported in which due to the *vertex* presentation of the fetus, the fetal R-peaks appeared as positive peaks while the maternal peaks had negative peaks. In such cases, the detection of the fetal RR-intervals is rather simple and may be achieved by simple peak detection, even without the removal of the maternal ECG. However, these methods are not always applicable and highly depend on the fetal presentation and gestational age.

Adaptive Filtering

Different variants of adaptive filters have been used for maternal ECG cancellation and fetal ECG extraction. These methods consist of training an adaptive or matched filter for either removing the maternal ECG using one or several maternal reference channels [220, 149], or directly training the filter for extracting the fetal QRS waves [57, 153]. *Ad hoc*, adaptive filters such as ‘partition-based weighted sum filters’ [185], and least square error fittings [131], have also been used for this purpose.

Note that the existing adaptive filtering methods for maternal ECG artifact removal, either require a reference maternal ECG channel that is morphologically similar to the contaminating waveform, or require several linearly independent channels to roughly reconstruct any morphologic shape from the three references [220]. Both of these approaches are practically inconvenient and with limiting performance, since the morphology of the maternal ECG contaminants highly depends on the electrode locations and it is not always possible to reconstruct the complete maternal ECG morphology from a linear combination of the reference electrodes¹. Therefore, a maternal ECG cancellation method that would not require any

¹See Chapter 3 for the limitations of the single dipole model of the heart.

excess reference electrodes or at most a single reference without the morphologic similarity constraint is of great interest.

Linear Decomposition

Decomposition of single or multi-channel recordings is another common approach. In this method, the signals are decomposed into different components by using suitable basis functions. The basis functions can be selected from classes that are somehow in coherence with the time, frequency, or scale characteristics of the fetal components. Wavelet decomposition [122, 108], and matching pursuits [3], are among these methods.

Spatial filtering techniques such as singular value decomposition (SVD) [43, 207, 210, 102], blind and semi-blind source separation [227], can be considered as ‘data-driven’ decomposition methods, which construct the required basis functions from the data itself, by maximizing some statistical measure of signal separation. In [226, 228], it has been shown that for fetal ECG extraction blind source separation methods outperform adaptive filters such as the one proposed in [220]. One of the advantages of spatial filtering over conventional adaptive filters is that they can also separate the maternal and fetal complexes with temporal overlap.

Different variants of blind and semi-blind source separation methods have been used for fetal ECG extraction [28, 45]. These methods are commonly based on the assumption of independent components (or more generally independent subspaces) for the maternal and fetal signals, or the existence of some temporal structure for the desired signals [14, 229, 123]. Blind source separation methods have also been combined with wavelet decomposition for extracting and denoising the fetal ECG signals [212, 95].

Decomposition methods are currently the most common and effective way of fetal ECG extraction and denoising. However, the existing methods are rather generic and have not been fully customized to the periodic structures of ECG. Therefore, a challenging issue is to design multichannel signal processing methods (linear or nonlinear) that are specific to ECG/MCG signals.

Nonlinear Decomposition

Linear decomposition methods using either fixed basis functions (e.g. wavelets), or data-driven basis functions (e.g. singular vectors) have limited performance for nonlinear or degenerate mixtures of signal and noise. In fact, fetal signals and other interferences and noises are not always ‘linearly separable’ (cf. Chapter 6). One solution for such cases is to use nonlinear transforms for separating the signal and noise parts of the observations.

Of course, nonlinear transforms are rather *ad hoc* and require some prior information about the desired and undesired parts of the signal. In [183, 182, 164, 103], using nonlinear projections, a series of maternal ECG cancellation and fetal ECG enhancement methods have been developed. These methods consist of using the noisy signal and its delayed versions for constructing a state-space representation of the signal, smoothing the state-space trajectory using conventional or PCA smoothers [110], and transferring the samples back to the time-domain representation. These methods are very attractive from the point that they are applicable to as few as one single maternal abdominal channel. However, the selection of the required time-lags is empirical and the important inter-beat variations of the cardiac signals can be wiped-out during the state-space smoothing. Moreover, they have higher computational complexity as compared with linear methods.

2.5 Forward Modeling vs. Inverse Solutions

An important aspect of noninvasive cardiac signal studies (either for adults or fetuses) is to find relationships between the cardiac potentials generated at the heart level and the potentials recorded on the body surface. This problem is known as the *forward problem* of electrocardiography, for which electromagnetic principles are used with electrophysiological models of cardiac potentials and volume conduction models, to predict the potentials that can appear on the body surface. Forward modeling also provides valuable insight for the more practical problem of estimating the cardiac potentials from body surface recordings, namely the *inverse solution*. Forward and inverse problems have long been studied for adult cardiac signals [71, 79]; but there are only few works of this sort for fetal cardiac signals.

Some pioneer studies on fetomaternal compartment conductivities and fetal ECG modeling were presented in [144, 147, 146, 145], where the fetal ECG and the volume conductor medium from the fetal heart up-to the maternal abdominal leads were modeled and simulated. They have also questioned some previous results on the conductivity of the maternal internal body tissues and note that the fetal ECG –as seen from the maternal abdomen– does not have the same dimensionality throughout pregnancy. This may partially be due to the physiological changes of the fetal heart itself, and partially due to the low conductivity of the *vernix caseosa* that surrounds the fetus and electrically shields it from the outside world in certain gestational ages. Some of these findings have also been verified with experimental results.

In a more recent study, the problem of reliability of the fetal magnetocardiogram (MCG) and ECG has been studied using forward modeling in normal and pathologic cases [190]. They use different models for different stages of gestation. Specifically, in the forward modeling of the last trimester of gestation, they considered a layer of *vernix caseosa* containing two holes and achieved fetal ECG maps that resembled the actual measured maps. The holes in the *vernix caseosa* were considered over the mouth of the fetus and the onset of the umbilical cord, corresponding to the ‘preferred’ current pathways. Other issues concerning the low-amplitude of the fetal T-waves are also presented in this work. However, for the processing of their real data, they used simple methods such as average waveform subtraction of the maternal MCG.

2.6 Alternative Measurement Techniques

In previous works, besides electrocardiography, the well-being of the fetal heart has also been monitored by other techniques [155], including:

- Echocardiography [224, 53]; also known as sonography of the heart which is based on standard ultrasound techniques.
- Phonocardiography [230, 111, 211]; which is a graphic registration of the heart sounds and murmurs produced by the contracting heart (including its valves and associated great vessels), picked up as vibrations and transformed by a piezoelectric crystal microphone into a varying electrical output according to the pressure imposed by the sound waves [33, 56].
- Cardiotocography [186]; which is the simultaneous measurement of the fetal heart rate with an ultrasound transducer, and the uterine contractions with a pressure-sensitive transducer (called a tocodynamometer), for measuring the strength and frequency of uterine contractions [221].
- Magnetocardiography [104, 190, 34, 198, 215]; which is a technique to measure the magnetic fields of cardiac signals using extremely sensitive devices such as the Superconducting Quantum Interference Device (SQUID) [223].

Among these methods, echocardiography is the most common –and commercially the most available– means of fetal cardiac monitoring. However, the ECG and the MCG contain more information, since most cardiac abnormalities have some manifestation on the ECG or MCG morphology or RR-interval timing [155]. The current work is therefore focused on the ECG and partially the MCG, which is the magnetic counterpart of the ECG. Note that due to the morphologic similarity of the MCG and ECG, MCG processing methods are similar to the ECG-based ones; although using the current SQUID technology for magnetic recordings, the SNR of the fetal MCG is usually higher than its ECG. However, ECG recording apparatus are simpler and currently more affordable as compared with MCG systems.

2.7 Current Challenges and Problem Definition

Reviewing previous works, we notice that despite of the rich literature, there are still several key elements that require further studies. Following our previous explanations, among the different data collection setups, we concentrate on ECG-based systems using multichannel noninvasive maternal abdominal measurements, and the objective is to retrieve the fetal ECG morphology with the highest possible fidelity, as required for morphological studies. In this context, the limiting factors and challenging signal processing issues include:

- weakness of fetal cardiac potentials and low-conductivity layers surrounding the fetus which lead to low amplitude fetal ECG at the maternal body surface;
- high interference of maternal ECG, uterus contraction, maternal respiratory, and motion artifact signals;
- possible movements of the fetus and the need for a sort of ‘canonical representation’ of the fetal cardiac signals with respect to the fetal body axis;
- development of automatic procedures that can be applied on long datasets with minimal interaction with an expert operator;
- providing confidence measures for the estimated cardiac signals and finding theoretical bounds for the amount of ‘retrievable information’ from body surface recordings in presence of noise.

Moreover, conventional ECG filtering methods are commonly based on some measure of time, frequency, or scale-separability of signals and noise, which is common with all denoising techniques. However, cardiac signals have additional pseudo-periodic structure, which we believe, have not been well-utilized in ECG denoising schemes.

In previous studies multichannel decomposition methods have been commonly applied to the observed signals rather ‘blindly’ and there is generally no guarantee that the fetal components be extracted as separate components. Therefore, an important issue is to increase the probability of extracting the fetal components and also to improve the quality of the extracted components, through appropriate preprocessing and using *a priori* information about the signal/noise mixtures. This is a necessary step for developing robust fetal ECG/MCG extraction algorithms.

Linear decomposition methods are very common, not only due to the validity of linear model itself²; but also for the simplicity of these models. However, as mentioned before, there are cases in which the desired signals are not linearly separable and nonlinear decomposition is inevitable³. Therefore, an interesting field of study is to combine linear and nonlinear techniques to benefit from the simplicity of linear transforms and the power of nonlinear methods at the same time.

Another related issue is to find physiological interpretations for components extracted by multichannel source separation techniques. While these methods are commonly based on the maximization of rather abstract statistical criteria, such as statistical independence, it is not very evident what the resultant components physically correspond to, when applied to real data. For cardiac signals, this issue is more crucial, when we consider that the heart is a distributed source and not a point (punctual) source.

Morphologic modeling of the fetal ECG is another issue of interest. While previous fetal ECG/MCG models concentrate on forward models based on electromagnetic and volume conduction theories [144, 190], for evaluating signal processing techniques based on body surface potentials, more abstract models are required. In fact, in order to evaluate and compare single or multi-channel processing techniques, we need models that enable us to manipulate the signal processing aspects of the simulated signals such as their morphology, RR-interval timing, fetal position, dimensionality, and SNR, without going into the details of signal propagation and volume conduction theories. For adult ECG, an example of such models was developed in [133], where single-channel adult ECG was modeled with a dynamic model. However, the existing models have not considered the multi-dimensional nature of the ECG and are not appropriate for the evaluation of multichannel techniques that use the ‘mutual information’ of the different channels.

2.8 Summary and Conclusions

In this chapter, we briefly reviewed the fetal cardiac signal extraction literature and its challenging issues. It was noted that in the current study we are interested in improving the signal processing aspects of this problem, in order to facilitate the extraction of fetal cardiac signals. In later chapters several methods are proposed for achieving these objectives, together with additional byproducts of this work that are not limited to fetal cardiac signals.

²A discussion about the validity of linear models for ECG decomposition is presented in Chapters 3 and 6.

³See Chapter 6 for related discussions.

Chapter 3

Electrophysiology of the Fetal Heart

3.1 Introduction

In this chapter the prerequisite physiological and electrophysiological aspects of fetal cardiac development and monitoring are presented. Some of these issues are used in later chapters for explaining the results achieved from multichannel analysis of maternal abdominal recordings.

3.2 Fetal Cardiac Development

The heart is among the first organs developed in the fetus and undergoes a considerable amount of growth in the very early stages of pregnancy [96] (Fig. 3.1). The most critical period of this development is between three and seven weeks after fertilization, when a simple heart tube assumes the shape of a four-chambered heart (Fig. 3.2).

The heart is believed to begin beating by the 22nd day of life and can be externally monitored by ultrasound imaging in the 7th to 9th week [96]; although only vague images are recordable at this step. The cardiac waveforms and beat-to-beat variability of the heart rate are not measurable in ultrasound imaging. Therefore, the fetal ECG and MCG that contain morphological information of the cardiac activity have received much interest. These signals can be recorded from the maternal abdomen as early as the eighteenth to twentieth week after conception [155, 206].

3.2.1 Fetomaternal Compartments

A simplified anatomy of the fetomaternal compartments can be seen in Fig. 3.3. Accordingly, the fetus is surrounded by several different anatomical layers with different electrical conductivities [144, 147]. The




Figure 3.1: The fetus and its heart in the early stages of development; adopted with permission from [118].




Figure 3.2: Developments of the fetal heart during gestation; adopted with permission from [118].

highest and lowest conductivity are found in the amniotic fluid and the *vernix caseosa*, respectively. Both of these layers surround the fetus completely.

In the maternal abdomen compartments, the skin and the subcutaneous fat also have a poor conductivity – about ten times smaller than the muscle tissue [147, 190]. Therefore these two layers, which are the interface of the surface electrodes and the internal tissues, have considerable influence on the recorded fetal ECG. All of these different tissues and layers form a so called, *volume conductor*, in which the fetal cardiac signals propagates up-to the maternal body surface. This volume conductor is not a steady conductor and its electric conductivity and geometric shape constantly change throughout gestation. Specifically, in the second half of gestation (20th week onward), in which the ECG and MCG are recordable from the surface electrodes, it is known that the amniotic fluid [23, 126], the placenta, and the fetus itself [12] are all increasing in volume. The very low conductivity *vernix caseosa* layer is formed between the 28th and 32nd weeks of gestation [144, 147]. It almost electrically shields the fetus and makes the recording of fetal surface ECG very difficult. However, for normal pregnancies (non-premature deliveries), the layer slowly dissolves in the 37th to 38th weeks of pregnancy [190]. In previous studies different interpretations have been presented for the fetal ECG signals recorded during the third trimester of gestation (after the forming of the *vernix caseosa*), based on preferred current pathways such as the umbilical cord, the oronasal cavity, or even random holes in the *vernix caseosa* [144, 147]. The preferred current pathway hypothesis, besides the intrinsic weakness of the fetal ECG and its gradual development, also explains some of the differences between adult and fetal ECG and vectorcardiogram shapes.

3.2.2 Fetal Presentations

During the first two trimesters of pregnancy the fetus does not have a specific presentation and moves about a lot. By the beginning of the third trimester it commonly settles in a head-down position known as the *vertex* presentation, which is more appropriate for birth [148]. However, as shown in Figs. 3.4 and 3.5 the fetus may also settle in other, but less-probable, presentations. The presentation of the fetus influences the fetal cardiac signals recorded from the maternal body surface over different leads [166].




Figure 3.3: The major fetomaternal compartments that influence the fetal cardiac surface potentials; adopted with permission from [190]. The *vernix caseosa* is formed over the fetal skin.

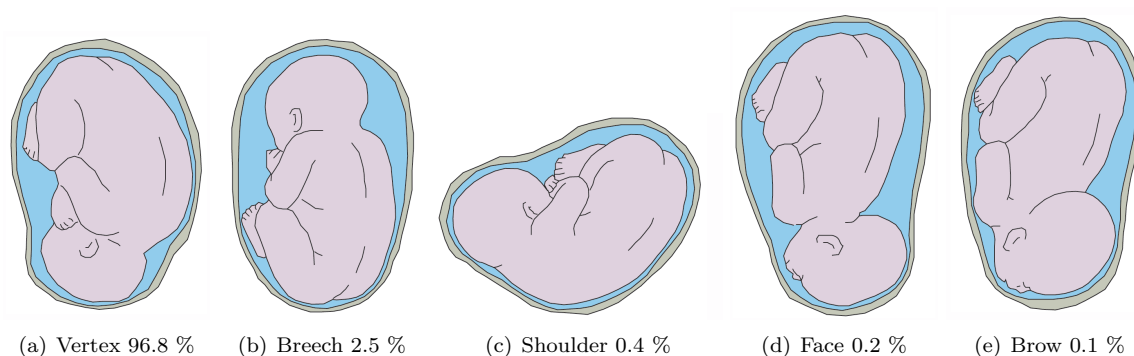


Figure 3.4: Different fetal presentations and the percentage of incidence at the end of gestation; adopted with permission from [190].




Figure 3.5: Different fetal vertex positions and their incidence at the end of gestation; adopted with permission from [190]. In the right side abbreviations, L stands for left, R for right, O for occiput, A for anterior, T for transverse and P for posterior. The LOA position is seen in the left graph.




Figure 3.6: The anatomy of the fetal heart; adopted with permission from [190].

3.3 Physiology of the Fetal Heart

The anatomy of a fully developed fetal heart is depicted in Fig. 3.6. There are some functional differences between the fetal and adult hearts. It is known that after birth, the left ventricle pumps blood to the body and the right ventricle pumps the blood to the lungs for acquiring oxygen. However, for the fetus the fetal oxygen is supplied by the *placenta*; therefore the blood is no longer pumped to the lungs for this purpose. Instead both ventricles pump the blood throughout the body (including the lungs) [190]. For this purpose there are two shunts, namely the *foramen ovale* and the *ductus arteriosus* that link the outgoing vessels of both ventricles. This allows blood to enter the right atrium and to bypass the pulmonary circulation. A similar adaptation in the fetus is the *ductus venosus*, which is a vessel that allows blood to bypass the liver. It carries blood with oxygen and nutrients from the umbilical cord straight to the right side of the fetal heart [184].

After birth, the *foramen ovale* closes with the first breaths and the *ductus arteriosus* partially closes in 10 to 15 hours after birth and takes up-to three weeks for complete closure. The *ductus venosus* also closes shortly after birth, when the umbilical cord is cut and blood flow between the mother and fetus stops [184]. There are also other minor changes in the physiology of the baby's heart and its circulatory system that take place within the first year after birth.

3.3.1 Electrical Activity of the Fetal Heart

While the mechanical function of the fetal heart differs from an adult heart, its beat-to-beat electrical activity is rather similar. The wave-like pumping action of the heart is controlled by a network of neural fibers that are distributed throughout the myocardium and coordinate its regular contraction and relaxation [26]. The myocardial stimulation starts from the *sinoatrial node* (SA-node), which serves as the natural pacemaker for the heart. The SA-node is a cluster of cells located in the upper-right posterior wall of the right atrium, which sends the electrical impulse that triggers each heartbeat. This impulse further stimulates the second cluster of cells, namely the *atrioventricular node* (AV-node) that is situated in the lower posterior wall of the right atrium. After the AV-node the depolarization front enters the bundle of His, the left and right bundles, and ends in the Purkinje fibers, depolarizing the ventricular muscles in its way.

The procedure of myocardium contraction is known as the *depolarization* (or *systole*) cycle that is followed by the *repolarization* (or *diastole*) cycle, in which the myocardium relaxes and becomes ready for the next activation. A complete cardiac cycle is depicted in Fig. 3.7.

3.3.2 The Electrocardiogram

The ECG measured on the body surface is a result of the stage-wise activation of the myocardium and results in the so-called PQRST-complex depicted in Fig. 3.7. This letter representation was first coined by




Figure 3.7: The activation cycle of the fetal heart; adopted with permission from [190].

Einthoven in 1895 [55, 88]. In this notation, the P-wave accounts for the spreading of the depolarization front through the atria. During the next 50ms, only very weak signals are recordable, as it takes some time for the depolarization front to travel through the AV-node and since only a small number of myocardial cells take part in the atrioventricular conduction the signals are very small¹ [91]. Next, the ventricles are depolarized resulting in the QRS-complex. At the same time the atria are repolarized; however this repolarization is obscured by the depolarization of the ventricles. Finally we have the T-wave, which corresponds to the repolarization of the ventricles. In some measurements the T-wave is followed by a small wave known as the U-wave which, for normal ECG, is believed to be due to repolarization of the His-Purkinje system [88]. Other hypothesis for the origin of the U-wave have been suggested in [13]. The segment between the end of the U-wave up-to the beginning of the next P-wave is known as the *isoelectric segment* of the ECG, in which the myocardium does not have a measurable activity.

It should be noted that despite the similarities in the electrical properties of fetal and adult cardiac systems, there are some differences in their RR-interval and morphology. The fetal heartbeat is almost twice as fast as an adult heartbeat with considerable changes in different stages of fetal cardiac development [86]. The *heart-rate variability* (HRV) of the fetus is also known to be simpler (less dynamic) than an adult. However, as the fetal autonomic nervous systems evolves the HRV patterns also become more and more complex [205, 214]. A typical example of this complexity is depicted in Fig. 3.8.

Morphologically, adult and fetal ECGs have rather similar patterns; but the relative amplitudes of the fetal complexes undergo considerable changes throughout gestation and even after birth. The most considerable change concerns the T-waves, which are rather weak for fetuses and newborns [206].

3.4 The Electromagnetic Basis of the Electrocardiogram

The electrical currents and potentials generated in the heart are the result of opening and closure of ionic channels at a cellular level. The coherent activation of numerous cellular reactions of this sort results in electrical fields that propagate in the so called body *volume conductor*, resulting in measurable potentials at the body surface. Neglecting the underlying chemical reactions, from an electromagnetic point of view, the biological volume conduction problem is somewhat unusual in the sense that the current sources are within the conductor rather than being due to induction [70]. Through the study of electrical and magnetic properties of body tissues, it is found that the cardiac volume conductor problem has the following important features [71]:

1. The electric displacement current is negligible. Therefore, the electromagnetic problem is *quasi-static*, which means that the electric and magnetic fields are decoupled, the electric field is proportional to the gradient of the electric scalar potential, and the divergence of the current density is zero.
2. The problem is linear, meaning that superposition holds for the potentials due to several sources.

¹Note that although frequently addressed as being *isoelectric*, the PQ segment is not isoelectric [91].




Figure 3.8: The increasing complexity in fetal HRV time-series in different gestational ages; adopted from [214].

3. For frequencies below several kilohertz, the capacitive component of the electrical impedance of body tissues is negligible, which means that the tissues are with a very good approximation, resistive.

From these properties, it has been shown that the electrical potential ϕ , recorded at the body surface satisfies the well-known *Poisson* equation:

$$\nabla^2 \phi = \frac{1}{\sigma} \nabla \cdot \mathbf{J} \quad (3.1)$$

where \mathbf{J} is the cardiac impressed current dipole moment density and σ is the conductivity at the measurement point of the potential ϕ [70].

To solve this differential equation, we require a model for the conduction media. For our problem of interest, we consider the volume conductor, or the propagation media, to be formed of M (a finite number of) homogeneous regions separated from one another by closed surfaces s_j having the conductivities σ_j' and σ_j'' on its interior and exterior. These regions correspond to the different anatomical compartments in Fig. 3.3. With this simplification and by using the Green's theorem, (3.1) may be transformed into a more useful representation for calculating the body surface potentials [70, 71]:

$$4\pi\sigma\phi = \int_v \mathbf{J} \cdot \nabla\left(\frac{1}{r}\right) dv + \sum_{j=1}^M (\sigma_j'' - \sigma_j') \int_{s_j} \phi_j \nabla\left(\frac{1}{r_j}\right) \cdot d\mathbf{s}_j \quad (3.2)$$

On the left-hand-side of (3.2), ϕ is the potential difference between a fixed point outside the charge distribution and the reference of potentials. It corresponds to the potentials recorded on the body surface. The volume v is an arbitrary volume containing the current dipole distribution (i.e. the heart). r is the radial distance of the element of volume or area from the point from which the potential ϕ is recorded. The potentials ϕ_j , inside the integral on the right-hand-side of (3.2) are the potentials of surface elements of the integral. $\nabla \cdot \mathbf{J}$ is the divergence of \mathbf{J} , $\nabla(1/r)$ is the gradient of $1/r$, and the dot sign (\cdot) represents vectorial inner-product. Finally, dv is the volume element and $d\mathbf{s}_j$ is a vectorial surface element, normal to the j -th surface element. An equivalent equation that is more appropriate for the later presented models is as follows [70]:

$$4\pi\sigma\phi = - \int_v \frac{1}{r} \nabla \cdot \mathbf{J} dv - \sum_{j=1}^M \frac{\sigma_j'' - \sigma_j'}{\sigma_j'' + \sigma_j'} \int_{s_j} \frac{2\mathbf{E}_j}{r_j} \cdot d\mathbf{s}_j \quad (3.3)$$

where $\mathbf{E}_j \doteq (\mathbf{E}' + \mathbf{E}'')/2$, and \mathbf{E}' and \mathbf{E}'' are the electric fields on either sides of the surfaces with conductivity discontinuity.

From (3.2) and (3.3), we can see that the body surface potentials consist of two parts: a term corresponding to the divergence of the cardiac impressed current dipole moments, and a term corresponding to the ‘reflections’ of the cardiac sources onto the surfaces of conductivity mismatch.

3.4.1 Multipole Expansion of Body Surface Potentials

Equations (3.2) and (3.3) are rather complicated to envisage and we can use *multipole expansions* to derive alternative representation for them. To do so, we need to expand the term $1/r$, in (3.3) using some basis function [193, 93]. While this expansion can be done using various basis functions, the *spherical harmonic* expansion is the most common expansion, as it is in accordance with the rather localized nature of our sources of interest. If we consider \mathbf{x} as the location vector of the point from which the potential ϕ is measured, and \mathbf{x}' as the location vector of the volume or surface elements of the integrands of the right-hand-side of (3.3), $1/r$ may be represented in spherical coordinates as follows [93]:

$$\frac{1}{r} = \frac{1}{|\mathbf{x} - \mathbf{x}'|} = 4\pi \sum_{l=0}^{\infty} \sum_{m=-l}^l \frac{1}{2l+1} \frac{r_{<}^l}{r_{>}^{l+1}} Y_{lm}^*(\theta', \phi') Y_{lm}(\theta, \phi) \quad (3.4)$$

where $r_{<} \doteq \min(|\mathbf{x}|, |\mathbf{x}'|)$, $r_{>} \doteq \max(|\mathbf{x}|, |\mathbf{x}'|)$, and $Y_{lm}(\theta, \phi)$ and $Y_{lm}^*(\theta', \phi')$ are the spherical harmonic functions that are related to the well-known *Legendre polynomials* as follows:

$$Y_{l,m}(\theta, \phi) = \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} P_l^m(\cos \theta) e^{jm\phi} \quad (3.5)$$

$$Y_{l,-m}(\theta, \phi) = (-1)^m Y_{l,m}^*(\theta, \phi) \quad (3.6)$$

where $P_l^m(\cos \theta)$ are Legendre polynomials.

Although the spherical harmonic expansion might seem complicated in the first view, its benefit is that the terms corresponding to the source locations (\mathbf{x}') are completely factorized from the terms corresponding to the observation point (\mathbf{x}). Next, by noting that for the points on the body surface $r_{<} = |\mathbf{x}'|$ and $r_{>} = |\mathbf{x}|$, and using (3.4) in (3.3), the potential ϕ recorded from the body surface (outside the volume conductor v), may be represented in the following compact form:

$$\phi = \lim_{L \rightarrow \infty} \sum_{l=0}^L \sum_{m=-l}^l a_{lm} s_{lm} \quad (3.7)$$

where s_{lm} and a_{lm} are defined as follows:

$$\begin{aligned} s_{lm} &\doteq q_{lm} + \sum_{j=1}^M w_{lm}^j \\ q_{lm} &= \int_v |\mathbf{x}'|^l Y_{lm}^*(\theta', \phi') \nabla \cdot \mathbf{J} dv \\ w_{lm}^j &= \frac{2(\sigma_j'' - \sigma_j')}{(\sigma_j'' + \sigma_j')} \int_{s_j} |\mathbf{x}'|^l Y_{lm}^*(\theta', \phi') \mathbf{E}_j \cdot d\mathbf{s}_j \\ a_{lm} &= \frac{-Y_{lm}(\theta, \phi)}{\sigma(2l+1)|\mathbf{x}|^{l+1}} \end{aligned} \quad (3.8)$$

Equation (3.7) is a key representation that requires further explanations. The terms a_{lm} are called the *multipole coefficients* and as noticed from (3.8), they only depend on the location of the observation point of the potential ϕ and do not rely on the cardiac source distributions. On the other hand, the s_{lm} are called the *multipole moments*, consisted of two parts: q_{lm} and w_{lm}^j , which are due to the impressed current densities and the electric fields caused by discontinuities in the surface conductivities, respectively. They are therefore independent of the observation point and solely depend on the distributions of the impressed current sources.

Note that although we focus on cardiac potentials in this study, the same equations hold for any biosignal (such as the EEG), which satisfy the same underlying assumptions. Moreover, due to the duality of the electric and magnetic fields, similar equations can be derived for the magnetocardiogram.




Figure 3.9: Illustration of the vectorcardiogram loop.

3.4.2 Monopole and Dipole Approximations

Two special cases of (3.7) are achieved for $L = 0$ and $L = 1$, namely the *monopole* and *dipole* expansions [127]. For $L = 0$, (3.7) reduces to:

$$\phi_0 = a_{0,0}s_{0,0} \quad (3.9)$$

which states that for observation points far from the charge distribution (far enough for a monopole model to hold), the potentials are proportional to the overall monopole moment. However, in bioelectricity we can never have a single isolated monopole current source due to the need of charge conservation [127], meaning that the monopole term is zero. Now neglecting the monopole term, for $L = 1$ we have:

$$\phi_1 = a_{1,-1}s_{1,-1} + a_{1,0}s_{1,0} + a_{1,1}s_{1,1} \doteq \mathbf{a}^T \mathbf{d} \quad (3.10)$$

where $\mathbf{a} = [a_{1,-1} \ a_{1,0} \ a_{1,1}]^T$ and $\mathbf{d} = [s_{1,-1} \ s_{1,0} \ s_{1,1}]^T$. Note that for ECG signals, \mathbf{d} is a pseudo-periodic function of time, representing the cardiac dipole vector. This last representation is the basis of the vectorcardiogram (VCG). Accordingly, we can model the electrical cardiac activity with a single rotating dipole located at the center of the heart with its end sweeping a pseudo-periodic loop in the space. The VCG loop is illustrated in Fig. 3.9. The single dipole model is believed to explain 80%–90% of the representation power of the body surface potentials [208].

3.5 Summary and Conclusions

In this chapter the physiology and electrophysiology of the fetus was reviewed. We also presented the theoretical basis for electrocardiogram signals recorded from the body surface. The theory of electrocardiography, presented in Section 3.4, has been readily developed in previous works [70, 71], and is based on well-known electromagnetic principles. However, to our knowledge, the compact representation in (3.7) and the reflection multipole coefficients in (3.8) have not been considered in previous works that have studied the problem of electrocardiography from the signal processing perspective. These properties will be used in later chapters for interpreting the components extracted by multichannel analysis of an array of maternal abdominal recordings.

Chapter 4

Synthetic ECG Generation

4.1 Introduction

The electrical activity of the cardiac muscle and its relationship with the body surface potentials has been studied with different approaches ranging from *single dipole models* to *activation maps* [51]. The goal of these models is to represent the cardiac activity in the simplest and most informative way for specific applications. Depending on the application of interest, any of the proposed models have some level of abstraction, which makes them a compromise between simplicity, accuracy, and interpretability for cardiologists. Specifically, as explained in Chapter 3, the single dipole model and its variants are equivalent source descriptions or multipole expansions of the true cardiac potentials. This means that they can only be used as far-field approximations of the cardiac activity and do not have evident interpretations in terms of the underlying electrophysiology [208]. However, despite these intrinsic limitations, the single dipole model remains a popular model, as it provides a three-dimensional representation of cardiac signals that accounts for 80% to 90% of the power of the body surface potentials [208, 127].

In recent years research has been conducted towards the generation of synthetic ECG signals to facilitate the validation of signal processing algorithms. Specifically, in [133, 132], a dynamic model has been developed, which reproduces the morphology of the PQRST complex and its relationship to the beat-to-beat (RR-interval) timing in a single nonlinear dynamic model. Considering the simplicity and flexibility of this model it is reasonable to assume that it can be easily adapted to a broad class of normal and abnormal ECGs. However previous works are restricted to single channel ECG modeling, meaning that the parameters of the model should be re-calculated for each of the recording channels. Moreover, for the maternal and fetal mixtures recorded from the abdomen of pregnant women, there are very few works which have both considered the cardiac source and the propagation media [144, 19, 134].

Real ECG recordings are always contaminated with noise and artifacts; hence besides the modeling of cardiac sources and propagation media, it is also very important to have realistic models for noise sources. Since common ECG contaminants are non-stationary and temporally correlated, time-varying dynamic models are required for the generation of realistic noises.

In this chapter, a three dimensional canonical model of the single dipole vector of the heart is proposed. This model, which is inspired by the single-channel ECG dynamic model presented in [133], is later related to the body surface potentials through a linear model that accounts for the temporal movements and rotations of the cardiac dipole, together with a model for the generation of realistic ECG noise. The ECG model is then generalized to fetal ECG signals recorded from the maternal abdomen.

The hereby proposed model is a *morphological* representation of the ECG. Therefore, due to the duality of the electric and magnetic field potentials of the heart, the same model may be used for modeling adult or maternal/fetal MCG recordings. The presented model is therefore believed to be an effective means of simulating realistic multichannel maternal/fetal ECG or MCG mixtures, in single and multiple pregnancies. This model is used in later chapters for the quantitative evaluation of the proposed multichannel filtering methods.




Figure 4.1: The three body axes, adapted from [127].

4.2 The Cardiac Dipole vs. the Electrocardiogram

The single dipole model of the heart was introduced in Section 3.4. Accordingly, the myocardium's electrical activity may be represented by a time-varying rotating vector, the origin of which is assumed to be at the center of the heart as its end sweeps out a quasi-periodic path through the torso. This vector may be mathematically represented in Cartesian coordinates, as follows:

$$\mathbf{d}(t) = x(t)\hat{\mathbf{a}}_x + y(t)\hat{\mathbf{a}}_y + z(t)\hat{\mathbf{a}}_z, \quad (4.1)$$

where $\hat{\mathbf{a}}_x$, $\hat{\mathbf{a}}_y$, and $\hat{\mathbf{a}}_z$ are the unit vectors of the three body axes shown in Fig. 4.1. With this definition, and by assuming the body *volume conductor* as a passive resistive medium which only attenuates the source field [71, 24], any ECG signal recorded from the body surface would be a linear projection of the dipole vector $\mathbf{d}(t)$, onto the direction of the recording electrode axes $\mathbf{v} = a\hat{\mathbf{a}}_x + b\hat{\mathbf{a}}_y + c\hat{\mathbf{a}}_z$:

$$\text{ECG}(t) = \langle \mathbf{d}(t), \mathbf{v} \rangle = a \cdot x(t) + b \cdot y(t) + c \cdot z(t) \quad (4.2)$$

As a simplified example, consider the dipole source of $\mathbf{d}(t)$ inside a homogeneous infinite volume conductor. The potential generated by this dipole at a distance $|\mathbf{r}|$ is:

$$\phi(t) - \phi_0 = \frac{\mathbf{d}(t) \cdot \mathbf{r}}{4\pi\sigma|\mathbf{r}|^3} = \frac{1}{4\pi\sigma} \left[x(t) \frac{r_x}{|\mathbf{r}|^3} + y(t) \frac{r_y}{|\mathbf{r}|^3} + z(t) \frac{r_z}{|\mathbf{r}|^3} \right] \quad (4.3)$$

where ϕ_0 is the reference potential, $\mathbf{r} = r_x\hat{\mathbf{a}}_x + r_y\hat{\mathbf{a}}_y + r_z\hat{\mathbf{a}}_z$ is the vector which connects the center of the dipole to the observation point, and σ is the conductivity of the volume conductor [127, 71]. Now consider the fact that the ECG signals recorded from the body surface are the potential differences between two different points. Equation (4.3) therefore indicates how the coefficients a , b , and c in (4.2) can be related to the radial distance of the electrodes and the volume conductor conductivity. Of course, in reality the volume conductor is neither homogeneous nor infinite, leading to a much more complex relationship between the dipole source and the body surface potentials. However as discussed in Chapter 3, even with a complete volume conductor model, the body surface potentials are linear instantaneous mixtures of the transmembrane potentials of the cardiac myocytes [71].

A three-dimensional vector representation of the ECG, namely the *vectorcardiogram* (VCG), is also possible by using three of such ECG signals. Basically any set of three linearly independent ECG electrode leads can be used to construct the VCG. However, in order to achieve an orthonormal representation that best resembles the dipole vector $\mathbf{d}(t)$, a set of three orthogonal leads that correspond with the three body axes are selected. The normality of the representation is further achieved by attenuating the different leads with *a priori* knowledge of the body volume conductor, to compensate for the non-homogeneity of the body thorax [127]. The *Frank lead system* [62], or the *corrected Frank lead system* [60], which has better orthogonality and normalization, are conventional methods for recording the VCG.

Based on the single dipole model of the heart, Dower *et al* have developed a transformation for finding the standard 12-lead ECGs from the Frank electrodes [52]. The Dower transform is simply a 12×3 linear transformation between the standard 12-lead ECGs and the Frank leads, which can be found from the

Table 4.1: Parameters of the synthetic ECG model in (4.4)

Index(i)	P	Q	R	S	T
$\theta_i(\text{rad})$	$-\pi/3$	$-\pi/12$	0	$\pi/12$	$\pi/2$
a_i	1.2	-5.0	30	-7.5	0.75
b_i	0.25	0.1	0.1	0.1	0.4

minimum mean square error (MMSE) estimate of a transformation matrix between the two electrode sets. Apparently the transformation is influenced by the locations of the recording leads and the attenuations of the body volume conductor with respect to each electrode [81]. The Dower transform and its inverse [54], are evident results of the single dipole model of the heart with a linear volume conductor. However, since the single dipole model of the heart is not a perfect representation of the cardiac activity, cardiologists usually use between six to twelve electrodes in their studies [127].

4.3 A Synthetic ECG Generator

The dynamic models presented in this chapter and the next are modifications of the synthetic ECG generator proposed by McSharry *et al.* [133]. This model has a variable number of free parameters that make it adaptable to many normal and abnormal ECGs. The dynamic model consists of a set of nonlinear dynamic state equations in Cartesian coordinates:

$$\begin{cases} \dot{x} = \rho x - \omega y \\ \dot{y} = \rho y + \omega x \\ \dot{z} = -\sum_{i \in \{P, Q, R, S, T\}} a_i \Delta\theta_i \exp\left(-\frac{\Delta\theta_i^2}{2b_i^2}\right) - (z - z_0) \end{cases} \quad (4.4)$$

where x , y , and z are the state variables, $\rho = 1 - \sqrt{x^2 + y^2}$, $\Delta\theta_i = (\theta - \theta_i) \bmod (2\pi)$, $\theta = \text{atan2}(y, x)$ is the four quadrant arctangent of the elements of x and y , with $-\pi \leq \text{atan2}(y, x) \leq \pi$, and ω is the angular velocity of the trajectory as it moves around the limit cycle in the $x-y$ plane. As it is seen in (4.4), each of the P, Q, R, S, and T-waves of the ECG waveform are modeled with a Gaussian function located at specific angular positions θ_i . The a_i , b_i , and θ_i terms in (4.4) correspond to the amplitude, width, and center parameters of the Gaussian terms of this equation. Some typical values of these parameters taken from [133] are listed in Table 4.1. In this model, the *baseline wander* of the ECG is modeled with the parameter z_0 that is assumed to be a relatively low amplitude sinusoidal component coupled with the respiratory frequency.

The three-dimensional trajectory generated by (4.4), consists of a circular limit cycle in the $x-y$ plane that is pushed up and down as it approaches each of the θ_i . The z coordinate of this three-dimensional trajectory, when plotted versus time gives the synthetic ECG. In [38], some methods have been developed for estimating the values of the model parameters for realistic ECGs, based on nonlinear optimization of the parameters of (4.4) and for a given ECG dataset.

Modification of the Dynamic ECG Model

The dynamic equations proposed in [133] are in Cartesian coordinates. As a first simplification, these equations can be transferred into polar coordinates. Moreover, assuming the z state variable in (4.4) to be in millivolts, b_i 's and θ_i 's in radians, and time in seconds, it is clear that the a_i 's are in $\text{mV}/(\text{rad} \times \text{s})$. So in order to simplify the dimensions and later relate the model parameters with real ECG recordings, the a_i terms in (4.4) will be replaced with:

$$a_i = \frac{\alpha_i \omega}{b_i^2} \quad i \in \{P, Q, R, S, T\},$$

where the α_i are the peak amplitudes of the Gaussian functions used for modeling each of the ECG components, in millivolts. This definition may be verified from (4.4), by neglecting the baseline wander term $(z - z_0)$ and integrating the \dot{z} equation with respect to t . With these changes, the new form of the

dynamic equations in polar coordinates is as follows:

$$\begin{cases} \dot{r} = r(1 - r) \\ \dot{\theta} = \omega \\ \dot{z} = -\sum_{i \in \{P, Q, R, S, T\}} \frac{\alpha_i \omega}{b_i^2} \Delta \theta_i \exp\left(-\frac{\Delta \theta_i^2}{2b_i^2}\right) - (z - z_0) \end{cases} \quad (4.5)$$

where r and θ are respectively the radial and angular state variables in polar coordinates. These new set of equations have some benefits compared with the original equations proposed in [133]. First of all, the polar form is much simpler and its interpretation is straightforward. Accordingly, the first equation in (4.5) represents the radial behavior of the generated trajectory, and converges to the limit cycle of $r = 1$ for any initial value of $r \geq 1$. However, the second and third equations of (4.5) are independent of r , making the first differential equation redundant. Therefore, this first equation may be excluded as it does not affect the synthetic ECG (the z state variable). Another benefit of this representation is that the phase parameter θ , is an explicit state-variable that indicates the angular location of the P, Q, R, S and T-waves (Table 4.1). This point is further used in Chapter 5 for the implementation of Bayesian ECG filters. For the problem of interest, (4.5) may be further simplified by discarding the baseline wander term $(z - z_0)$. The simplified ECG model is as follows:

$$\begin{cases} \dot{\theta} = \omega \\ \dot{z} = -\sum_i \frac{\alpha_i \omega}{b_i^2} \Delta \theta_i \exp\left(-\frac{\Delta \theta_i^2}{2b_i^2}\right) \end{cases} \quad (4.6)$$

This model is very generic and, as we will show, it may be used to model a broad class of periodic waveforms.

4.4 Cardiac Dipole Vector and ECG Modeling

From the single dipole model of the heart, it is known that the different ECG leads can be assumed as projections of the heart's dipole vector onto the recording electrode axes. All leads are therefore time synchronized with each other and have a quasi-periodic shape. Therefore, we propose to use a three-dimensional extension of (4.6) to model the dipole vector $\mathbf{d}(t)$ defined in (4.1):

$$\begin{aligned} \dot{\theta} &= \omega \\ \dot{x} &= -\sum_i \frac{\alpha_i^x \omega}{(b_i^x)^2} \Delta \theta_i^x \exp\left[-\frac{(\Delta \theta_i^x)^2}{2(b_i^x)^2}\right] \\ \dot{y} &= -\sum_i \frac{\alpha_i^y \omega}{(b_i^y)^2} \Delta \theta_i^y \exp\left[-\frac{(\Delta \theta_i^y)^2}{2(b_i^y)^2}\right] \\ \dot{z} &= -\sum_i \frac{\alpha_i^z \omega}{(b_i^z)^2} \Delta \theta_i^z \exp\left[-\frac{(\Delta \theta_i^z)^2}{2(b_i^z)^2}\right] \end{aligned} \quad (4.7)$$

where as with (4.5)-(4.6), $\Delta \theta_i^x = (\theta - \theta_i^x) \bmod (2\pi)$, $\Delta \theta_i^y = (\theta - \theta_i^y) \bmod (2\pi)$, $\Delta \theta_i^z = (\theta - \theta_i^z) \bmod (2\pi)$, and $\omega = 2\pi f$, where f is the beat-to-beat heart rate. Accordingly, the first equation in (4.7) generates a circular trajectory rotating with the frequency of the heart rate. Each of the three coordinates of the dipole vector $\mathbf{d}(t)$ is modeled by a summation of Gaussian functions with amplitudes α_i^x , α_i^y , and α_i^z ; widths b_i^x , b_i^y , and b_i^z ; located at rotational angles θ_i^x , θ_i^y , and θ_i^z . As before, the intuition behind this set of equations is that the baseline of each of the dipole coordinates is pushed up and down as the trajectory approaches the centers of the Gaussian functions, generating a moving variable-length vector in the (x, y, z) space. Moreover, by adding some deviations to the parameters of (4.7), i.e. considering them as random variables rather than deterministic constants, it is possible to generate more realistic cardiac dipoles with inter-beat variations.

This model of the rotating dipole vector is rather general, since due to the *universal approximation* property of Gaussian mixtures, any continuous function (as the dipole vector is assumed to be so), can be modeled with a sufficient number of Gaussian functions up-to an arbitrarily close approximation [18].




Figure 4.2: Synthetic ECG signals of the Frank lead electrodes

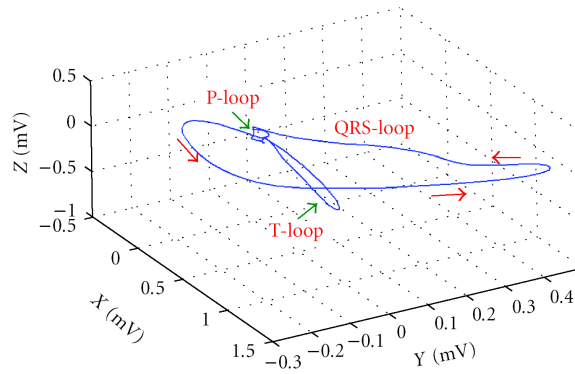


Figure 4.3: A typical synthetic VCG loop. Arrows indicate the direction of rotation. Each clinical lead is produced by mapping this trajectory onto a 1-dimensional vector in this three-dimensional space.

Note that the synthetic models in (4.6) and (4.7), could have also been simply presented as a sum of Gaussian functions (by simply integrating these equations with respect to time). However, this state-space representation can be used to study the evolution of the signal dynamics using state-space approaches [68]. This dynamic model is used in later chapters to represent the underlying dynamics of noisy ECG recordings within a Bayesian filtering framework.

Equation (4.7) can also be thought as a model for orthogonal lead VCG coordinates, using appropriate scaling factors for the attenuations of the volume conductor. This analogy between the orthogonal VCG and the dipole vector can be used to estimate the parameters of (4.7) from the three Frank lead VCG recordings. As an illustration typical signals recorded from the Frank leads and the dipole vector modeled by (4.7) are plotted in Figs. 4.2 and 4.3. The parameters of (4.7), used for the generation of these figures are presented in Table 4.2. These parameters have been estimated from the best MMSE fitting between N Gaussian functions and the Frank lead signals. As it can be seen in Table 4.2, the number of the Gaussian functions is not necessarily the same for the different channels and can be selected according to the shape of the desired channel.

4.4.1 Multichannel ECG modeling

The dynamic model in (4.7) is a representation of the dipole vector of the heart (or equivalently the orthogonal VCG recordings). In order to relate this model to realistic multichannel ECG signals recorded from the body surface, we need an additional model to project the dipole vector onto the body surface, which also accounts for the conductions of the body volume conductor, the possible rotations and scalings

Table 4.2: Parameters of the synthetic model presented in (4.7) for the ECGs and VCG plotted in Figs. 4.2 and 4.3

Index(i)	1	2	3	4	5	6	7	8	9	10	11
α_i^x (mV)	0.03	0.08	-0.13	0.85	1.11	0.75	0.06	0.10	0.17	0.39	0.03
b_i^x (rad)	0.09	0.11	0.05	0.04	0.03	0.03	0.04	0.60	0.30	0.18	0.50
θ_i (rad)	-1.09	-0.83	-0.19	-0.07	0.00	0.06	0.22	1.20	1.42	1.68	2.90
α_i^y (mV)	0.04	0.02	-0.02	0.32	0.51	-0.32	0.04	0.08	0.01		
b_i^y (rad)	0.07	0.07	0.04	0.06	0.04	0.06	0.45	0.30	0.50	-	-
θ_j (rad)	-1.10	-0.90	-0.76	-0.11	-0.01	0.07	0.80	1.58	2.90		
α_i^z (mV)	-0.03	-0.14	-0.04	0.05	-0.40	0.46	-0.12	-0.20	-0.35	-0.04	
b_i^z (rad)	0.03	0.12	0.04	0.40	0.05	0.05	0.80	0.40	0.20	0.40	-
θ_k (rad)	-1.10	-0.93	-0.70	-0.40	-0.15	0.10	1.05	1.25	1.55	2.80	

of the dipole, and the ECG measurement noises. Following the discussions of Section 4.2, a rather simplified linear model which accounts for these measures and is in accordance with (4.2) and (4.3) is suggested as follows:

$$\mathbf{ECG}(t) = H \cdot R \cdot \Lambda \cdot \mathbf{s}(t) + \boldsymbol{\eta}(t) + \mathbf{n}(t) \quad (4.8)$$

where $\mathbf{ECG}(t) \in \mathbb{R}^N$ is a vector of ECG signals in N leads, $\mathbf{s}(t) = [x(t), y(t), z(t)]^T \in \mathbb{R}^3$ contains the three components of the dipole vector $\mathbf{d}(t)$, $H \in \mathbb{R}^{N \times 3}$ corresponds to the body volume conductor model (as for the Dower transformation matrix), $\Lambda = \text{diag}(\lambda_x, \lambda_y, \lambda_z) \in \mathbb{R}^{3 \times 3}$ is a diagonal matrix corresponding to the scaling of the dipole in each of the x , y , and z directions, $R \in \mathbb{R}^{3 \times 3}$ is the rotation matrix for the dipole vector, $\boldsymbol{\eta}(t)$ is low-rank or *structured* noise, representing other biological sources that contaminate the ECG, and $\mathbf{n}(t)$ is full-rank observation noise that always exist in physiological measurements. In (4.8), we have discriminated between low-rank and full-rank noise to emphasize their different origins.

Note that H , R , and Λ are generally functions of time. Although the product of $H \cdot R \cdot \Lambda$ may be assumed to be a single $N \times 3$ matrix, the representation in (4.8) has the benefit that the rather stationary features of the body volume conductor that depend on the location of the ECG electrodes and the conductivity of the body tissues can be considered in H , while the temporal inter-beat movements of the heart can be considered in Λ and R , meaning that their average values are identity matrices in a long term study: $E_t\{R\} = I$, $E_t\{\Lambda\} = I$. In Section 4.7, by using the Givens rotation, a means of coupling these matrices with external sources such as the respiration and achieving non-stationary mixtures of the cardiac source will be presented.

4.4.2 Modeling maternal abdominal recordings

By utilizing a dynamic model like (4.7) for the dipole vector of the heart, the signals recorded from the abdomen of a pregnant woman, containing the fetal and maternal heart components can be modeled as follows:

$$\begin{aligned} \mathbf{x}(t) &= \mathbf{x}_m(t) + \mathbf{x}_f(t) + \boldsymbol{\eta}(t) + \mathbf{n}(t) \\ &= H_m \cdot R_m \cdot \Lambda_m \cdot \mathbf{s}_m(t) + H_f \cdot R_f \cdot \Lambda_f \cdot \mathbf{s}_f(t) + \boldsymbol{\eta}(t) + \mathbf{n}(t) \end{aligned} \quad (4.9)$$

where the noises $\mathbf{n}(t)$ and $\boldsymbol{\eta}(t)$ and the matrices H_m , H_f , R_m , R_f , Λ_m , and Λ_f have similar definitions as the ones in (4.8), with the subscripts m and f referring to the mother and fetus, respectively. Moreover, R_f has the additional interpretation that its mean value $E_t\{R_f\} = R_0$ can be assumed as the relative position of the fetus with respect to the axes of the maternal body (and not equal to identity, as assumed above). This is an interesting feature for modeling the fetus in the different typical positions such as *vertex* or *breech* positions previously shown in Figs. 3.4 and 3.5.

As illustrated in Fig. 4.4, $\mathbf{s}_f(t) = [x_f(t), y_f(t), z_f(t)]^T$ can be assumed as a canonical representation of the fetal dipole vector with respect to the fetal body axes. Therefore, to calculate this vector with respect to the maternal body axes, $\mathbf{s}_f(t)$ should be rotated by the three-dimensional rotation matrix R_0 :

$$R_0 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \theta_x & \sin \theta_x \\ 0 & -\sin \theta_x & \cos \theta_x \end{bmatrix} \begin{bmatrix} \cos \theta_y & 0 & \sin \theta_y \\ 0 & 1 & 0 \\ -\sin \theta_y & 0 & \cos \theta_y \end{bmatrix} \begin{bmatrix} \cos \theta_z & \sin \theta_z & 0 \\ -\sin \theta_z & \cos \theta_z & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (4.10)$$

where θ_x , θ_y , and θ_z are the angles of the fetal body planes with respect to the maternal body planes.




Figure 4.4: Illustration of the fetal and maternal VCGs vs. their body coordinates

The model presented in (4.9) may be simply extended to multiple pregnancies (twins, triplets, quadruplets, etc.), by considering additional terms for the other fetuses.

4.4.3 Fitting the model parameters to real recordings

As previously stated, due to the analogy between the dipole vector and the orthogonal lead VCG recordings, the number and shape of the Gaussian functions used in (4.7) can be estimated from typical VCG recordings. This estimation requires a set of orthogonal leads, such as the Frank leads, in order to calibrate the parameters. There are different possible approaches for the estimation of the Gaussian function parameters of each lead. *Nonlinear least square error* (NLSE) methods, as previously suggested in [39], have been proved as an effective approach. Otherwise, one can use the A^* optimization approach adopted in [18], or benefit from the algorithms developed for *radial basis functions* (RBF) from the *neural network* context [20].

It should be noted that (4.7) is some kind of canonical representation of the heart's dipole vector; meaning that the amplitudes of the Gaussian terms in (4.7) are not the same as the ones recorded from the body surface. In fact, using (4.7) and (4.8) to generate synthetic ECG signals, there is an intrinsic indeterminacy between the scales of the entries of $\mathbf{s}(t)$ and the mixing matrix H , since the cardiac dipole vector is only a modeling tool and not a physical source that can be recorded¹. To solve this ambiguity, and without loss of generality, we assume the power of the dipole vector to be normalized in each of its three coordinates.

As mentioned before, the H mixing matrix in (4.8) depends on the location of the recording electrodes. So in order to estimate this matrix, we first calculate the optimal parameters of (4.7) from the Frank leads of a given database. Next, the H matrix is estimated by using a MMSE estimate between the synthetic dipole vector and the recorded ECG channels of the database. In fact, by using the previously mentioned assumption that $E_t\{R\} = I$ and $E_t\{\Lambda\} = I$, the MMSE solution of the problem is:

$$\hat{H} = E\{\mathbf{ECG}(t) \cdot \mathbf{s}(t)^T\} [E\{\mathbf{s}(t) \cdot \mathbf{s}(t)^T\}]^{-1} \quad (4.11)$$

For the case of abdominal recordings the estimation of the H_m and H_f matrices in (4.9) are more difficult and require *a priori* information about the location of the electrodes and a model for the propagation of the maternal and fetal signals within the maternal thorax and abdomen [134]. A coarse estimation of H_m can be achieved for a given configuration of abdominal electrodes by using (4.11) between the abdominal ECG recordings and three orthogonal leads placed close to the mother's heart for recording her VCG. Yet, the accurate estimation of H_f requires more information about the maternal body and non-homogeneous models of the volume conductor (cf. Section 3.2).

The ω term introduced in (4.7) is in general a time-variant parameter that depends on physiological factors such as the speed of electrical wave propagation in the cardiac muscle and the HRV [133]. Furthermore, since the phase of the respiratory cycle can be derived from the ECG (or through other means such as amplifying the differential change in impedance in the thorax; *impedance pneumography*) and Λ

¹Unless theoretically, if we had the complete body conductivities and the surface distributions, which could be used in the Gabor-Nelson theorem to calculate the cardiac dipole vector [66].

is likely to vary with respiration, it is logical that an estimation of Λ over time can be made from such measurements.

The average (static) orientation of the fetal heart with respect to the maternal cardiac source is represented by R_0 which could be initially determined through a sonogram, and later inferred by referencing the signal to a large database of similar-term fetuses. Of course, both Λ and R_0 are functions of the respiration and heart rates and therefore tracking procedures such as *expectation maximization* (EM) [63], or *Kalman filter* (KF) may be required for online adaptation of these parameters [174, 175].

However, as we will show in Chapter 8, for multichannel statistical studies, where we need to analyze the performance of a given algorithm with *Monte Carlo* simulations, it is not necessary to derive the matrices defined in (4.9) from real observations. Instead, we may use random matrices with a desired angle between their column subspaces to represent the overall products $H_m \cdot R_m \cdot \Lambda_m$ and $H_f \cdot R_f \cdot \Lambda_f$. This idea is revisited in the case study presented in Section 4.6.2.

4.5 ECG Noise Modeling

An important issue that should be considered in the modeling of realistic ECG signals is to model realistic noise sources. Following [64], the most common high-amplitude ECG noises that cannot be removed by simple in-band filtering, are:

- baseline wander (BW)
- muscle artifact (MA)
- electrode movement (EM)

For the fetal ECG signals recorded from the maternal abdomen the following may also be added to this list:

- maternal ECG
- fetal movements
- maternal uterus contractions

These noises are typically very non-stationary in time and colored in spectrum (having long-term correlations). This means that white noise or stationary colored noise are generally insufficient to model ECG noise. In practice, researchers have preferred to use real ECG noises such as those found in the MIT-BIH non-stress test database (NSTDB) [139, 87], with varying SNRs.

As explained in the following, parametric models such as time-varying *autoregressive* (AR) models can be used to generate realistic ECG noises, which follow the non-stationarity and spectral shape of real noise. The parameters of this model can be trained by using real noises such as the NSTDB. Having trained the model, it can be driven by white noise to generate different instances of such noises, with almost identical temporal and spectral characteristics. There are different approaches for the estimation of time-varying AR parameters. An efficient approach that we adopt, is to reformulate the AR model estimation problem in the form of a standard *Kalman filter* (KF) [68].

For the time series y_n , a time-varying AR model of order p can be described as follows:

$$\begin{aligned} y_n &= -a_{n1}y_{n-1} - a_{n2}y_{n-2} - \dots - a_{np}y_{n-p} + v_n \\ &= -[y_{n-1}, y_{n-2}, \dots, y_{n-p}] \begin{bmatrix} a_{n1} \\ a_{n2} \\ \vdots \\ a_{np} \end{bmatrix} + v_n, \end{aligned} \quad (4.12)$$

where v_n is the input white noise and the a_{ni} ($i = 1, \dots, p$) coefficients are the p time-varying AR parameters at the time instant n . So by defining $\zeta_n \doteq [a_{n1}, a_{n2}, \dots, a_{np}]^T$ as a state vector, and $\mathbf{h}_n \doteq -[y_{n-1}, y_{n-2}, \dots, y_{n-p}]^T$, we can reformulate the problem of AR parameter estimation in the KF form as follows:

$$\begin{cases} \zeta_{n+1} = \zeta_n + \mathbf{w}_n \\ y_n = \mathbf{h}_n^T \zeta_n + v_n, \end{cases} \quad (4.13)$$




Figure 4.5: Typical segment of ECG BW Noise (a) Original (b) Synthetic




Figure 4.6: Frequency response magnitudes of 32 segments of the time-varying AR filters for the Baseline Wander noises of the NSTDB. This figure illustrates how the AR filter responses are evolving in time.

where we have assumed that the temporal evolution of the time-varying AR parameters follow a random walk model with a white Gaussian input noise vector \mathbf{w}_n . This approach is a conventional and practical assumption in the KF context when there is little information about the dynamics of a state vector [68].

To solve the standard KF equations [68], we also require the expected initial state vector $\bar{\zeta}_0 = E\{\zeta_0\}$, its covariance matrix $P_0 = E\{\bar{\zeta}_0 \bar{\zeta}_0^T\}$, the covariance matrices of the process noise $Q_n = E\{\mathbf{w}_n \mathbf{w}_n^T\}$, and the measurement noise variance $r_n = E\{v_n v_n^T\}$.

$\bar{\zeta}_0$ can be estimated from a global time-invariant AR model fitted over the whole samples of y_n , and its covariance matrix (P_0) can be selected ‘large enough’ to indicate the imprecision of the initial estimate. The effects of these initial states are of less importance and under some general convergence properties of KFs, they usually vanish in time.

By considering the AR parameters to be uncorrelated, the covariance matrix Q_n can be selected as a diagonal matrix. The selection of the entries of this matrix, depends on the extent of y_n ’s non-stationarity. For quasi-stationary noises, the diagonal entries of Q_n are rather small, while for highly non-stationary noises they are large. Generally the selection of this matrix is a compromise between convergence rate and stability. Finally, r_n is selected according to the desired variance of the output noise.

To complete the discussion, the AR model order should also be selected. It is known that for stationary AR models, there are information-based criteria such as the *Akaike information criterion* (AIC) for the selection of the optimal model order. However, for time-varying models the selection is not as straightforward, since the model is dynamically evolving in time. In general, we expect the model order to be less than the optimal order of a global time-invariant model.

Now having the time-varying AR model, it is possible to generate noises with different variances. As an illustration, in Fig. 4.5 a one minute long segment of BW with a sampling rate of 360Hz, from the NSTDB [139, 138], and the synthetic BW noise generated by the proposed method using a time-varying AR model of order twelve, are depicted. The frequency response magnitude of the time-varying AR filter designed for this BW noise is depicted in Fig. 4.6. As it can be seen, the time-varying AR model is acting as an adaptive filter that is adapting its frequency response to the contents of the non-stationary noise.

The proposed method generates single-channel noises. This method may be extended to multichannel noise in various ways. One approach is to use *multivariate autoregressive* models for simulating the noise sources [196]. A simpler approach that we adopt in the later presented results, is to generate M single-channel noises independently, and to project them to the ECG signal subspace with a projection matrix $A_n \in \mathbb{R}^{N \times M}$. By this way, we can model low-rank observation noises from different origins, such as

Table 4.3: The normalized MSE (%) in the synthetic VCG channels using five and nine Gaussian functions

VCG Channel	5 Gaussians	9 Gaussians
V_x	1.24	0.09
V_y	1.68	0.15
V_z	3.60	0.12

Table 4.4: The percentage of MSE (%) in the ECGs reconstructed by Dower transformation from the original VCG and from the synthetic VCG using five and nine Gaussian functions

ECG Channel	Original VCG	5 Gaussians	9 Gaussians
V_1	0.78	2.06	0.86
V_2	0.67	3.14	0.72
V_6	0.16	1.12	0.19

muscle contractions or respiration.

4.6 Illustrations

The proposed synthetic ECG model is believed to have interesting applications from both the theoretical and practical point of view. Here, we study the accuracy of the model and a special case study. Further applications of this model are presented in later chapters.

4.6.1 The model accuracy

As a first illustration, the accuracy of the model presented in (4.8) will be studied for a typical ECG signal of the Physikalisch-Technische Bundesanstalt diagnostic ECG database (PTBDB) [158, 22, 112]. The database consists of twelve standard ECG channels and three Frank lead VCGs. In order to have a clean template for extracting the model parameters, the signals are preprocessed by a bandpass filter to attenuate the baseline wander and high frequency noises. The ensemble average of the ECG is then extracted from each channel. Next, the parameters of the Gaussian functions of the synthetic model are extracted from the ensemble average of the Frank lead VCGs, using the nonlinear least squares procedure explained in Section 4.4.3. The Original VCGs and the synthetic ones generated from five and nine Gaussian functions are depicted in Figs. 4.7(a)–4.7(c) for comparison. The percentage of *mean square error* (MSE) of the two synthetic VCGs with respect to the true VCGs are listed in Table 4.3.

The H matrix defined in (4.8) may also be calculated by solving the MMSE transformation between the ECG and the three VCG channels (similar to (4.11)). As with the Dower transform, H can be used to find approximate ECGs from the three original VCGs or the synthetic VCGs. In Figs. 4.7(d)–4.7(f), the original ECGs of channels V_1 , V_2 , and V_6 , and the approximate ones calculated from the VCG are compared with the ECGs calculated from the synthetic VCG using five and nine Gaussian functions for one ECG cycle. As it can be seen in these results, the ECGs reconstructed from the synthetic VCG model have significantly improved as the number of Gaussian functions have been increased from five to nine, and the resultant signals very well resemble the ECGs reconstructed from the original VCG by using the Dower transform. The model improvement is especially notable around the asymmetric segments of the ECG such as the T-wave.

However, it should be noted that the ECG signals reconstructed from the Dower transform (either from the original VCG or the synthetic ones), do not perfectly match the true recorded ECGs, especially in the low amplitude segments such as the P-wave. This in fact shows the intrinsic limitation of the single dipole model in representing the low-amplitude components of the ECG which require more than three dimensions for their accurate representation (cf. Section 3.4). The percentage of MSE of the calculated ECGs of Figs. 4.7(d)–4.7(f) with respect to the true ECGs are listed in Table 4.4.

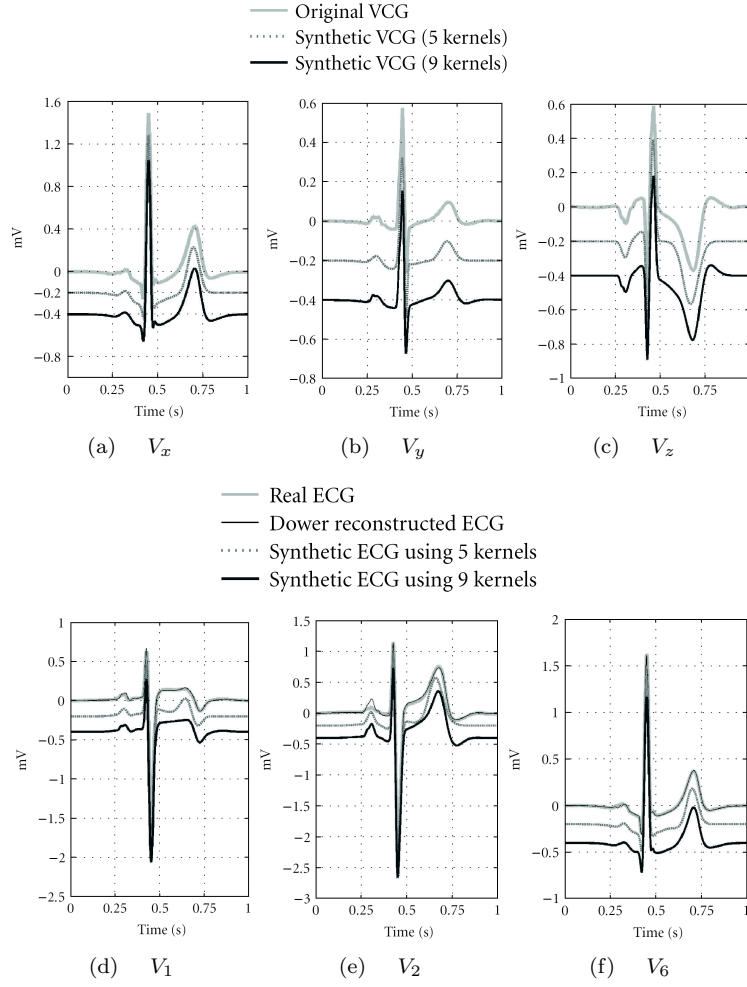


Figure 4.7: Original vs. synthetic VCGs and ECGs using 5 and 9 Gaussian functions. For comparison, the ECG reconstructed from the Dower transformation is also depicted in (d)-(f) over the original ECGs. The synthetic VCGs and ECGs have been vertically shifted 0.2mV for better comparison. Refer to text for details.

4.6.2 Fetal ECG extraction

In [169], a detailed example of the proposed method for modeling maternal abdominal recordings was presented. The thereby presented example was based on a specific electrode configuration and fetal positioning, and the matrices H_m and H_f were calculated by assuming the maternal thorax as an infinite homogeneous volume conductor. Here, we will not go into the details of that example; but will instead present an example that will be used in later chapters for evaluating multichannel source separation methods.

Following (4.9), each of the maternal and fetal parts of the simulated signals belongs to a three-dimensional subspace and the structured noise belongs to an M -dimensional one ($M < N$). We can therefore represent (4.9) in the following compact form:

$$\begin{aligned}
 \mathbf{x}(t) &= \mathbf{x}_m(t) + \mathbf{x}_f(t) + \boldsymbol{\eta}(t) + \mathbf{n}(t) \\
 &= [A_m \quad A_f \quad A_n] \cdot \begin{bmatrix} \mathbf{s}_m(t) \\ \mathbf{s}_f(t) \\ \mathbf{v}(t) \end{bmatrix} + \mathbf{n}(t) \\
 &\doteq A \cdot \mathbf{s}(t) + \mathbf{n}(t)
 \end{aligned} \tag{4.14}$$

where $A_m, A_f \in \mathbb{R}^{N \times 3}$, $A_n \in \mathbb{R}^{N \times M}$, $\mathbf{v}(t) \in \mathbb{R}^M$ represents M -dimensional noise, and $\boldsymbol{\eta}(t) = A_n \mathbf{v}(t)$ is the structured noise defined in (4.9). This form of the signal model is of special interest for the evaluation

Table 4.5: Parameters of the synthetic fetal dipole used in Section 4.6.2

Index(i)	1	2	3	4	5
α_i^x (mV)	0.07	-0.11	1.3	0.07	0.03
b_i^x (rad)	0.1	0.03	0.05	0.02	0.3
θ_i (rad)	-0.7	-0.17	0	0.18	1.4
α_i^y (mV)	0.04	0.30	0.45	-0.35	0.05
b_i^y (rad)	0.1	0.05	0.03	0.04	0.3
θ_j (rad)	-0.9	-0.08	0	0.05	1.3
α_i^z (mV)	-0.01	0.03	-0.40	0.46	-0.01
b_i^z (rad)	0.1	0.4	0.03	0.03	0.3
θ_k (rad)	-0.8	-0.3	-0.1	0.06	1.35

of blind and semi-blind source separation algorithms. In fact, according to (4.14), the source signal $\mathbf{s}(t)$ has $M + 6$ dimensions. Now depending on whether the number of synthetic channels N is equal, smaller, or greater than $M + 6$, the source separation problem becomes determined, under-determined, or over-determined, respectively.

Now, instead of calculating A_m and A_f from a volume conductor model, we choose them as random matrices having specific gains and a maximum *principal angle* of θ_{mf} in between. Principal angles are introduced in Appendix A; they are frequently used to define the angle between two subspaces [73]. Therefore, for our application of interest a small θ_{mf} indicates that the maternal and fetal subspaces are close to each other, which makes their separation more difficult. Now by considering the maternal ECG $\mathbf{x}_m(t)$ as interference and $\boldsymbol{\eta}(t)$ and $\mathbf{n}(t)$ as noise for the fetal signals $\mathbf{x}_f(t)$, the overall fetal *signal-to-interference ratio* (SIR) and *signal-to-noise ratio* (SNR) over all channels may be defined as follows:

$$\text{SIR} \doteq \frac{E_{t,i}\{x_{f_i}(t)^2\}}{E_{t,i}\{x_{m_i}(t)^2\}} \quad (4.15)$$

$$\text{SNR}_\eta \doteq \frac{E_{t,i}\{x_{f_i}(t)^2\}}{E_{t,i}\{\eta_i(t)^2\}} \quad (4.16)$$

$$\text{SNR}_n \doteq \frac{E_{t,i}\{x_{f_i}(t)^2\}}{E_{t,i}\{n_i(t)^2\}} \quad (4.17)$$

where $x_{f_i}(t)$, $x_{m_i}(t)$, $n_i(t)$, and $n_i(t)$ are respectively the entries of $\mathbf{x}_f(t)$, $\mathbf{x}_m(t)$, $\boldsymbol{\eta}(t)$, and $\mathbf{n}(t)$ in channel i , and $E_{t,i}\{\cdot\}$ represents summation over time and channels. Following these definitions, the gains of A_m , A_f , A_n , and the full-rank noise $\mathbf{n}(t)$ are defined such that the fetal SIR and SNR be equal to their desired values.

In this example, we use the parameters listed in Tables 4.2 and 4.5, to generate the maternal and fetal dipoles $\mathbf{s}_m(t)$ and $\mathbf{s}_f(t)$, respectively. The maternal and fetal heart rates are assumed to be $f_m = 0.9\text{Hz}$ and $f_f = 2.2\text{Hz}$, respectively.

We consider eight synthetic ECG channels, in two cases: once with a maximum principal angle of $\theta_{mf} = 10^\circ$ between A_m and A_f , and another time with a maximum principal angle of $\theta_{mf} = 60^\circ$.

For the structured noise, we consider three independent noise sources: real baseline wander, muscle artifacts, and electrode movement taken from the NSTDB [138, 139]. To project these noises to the sensor subspace, we use a random 8×3 matrix A_n . The full-rank noise $\mathbf{n}(t)$ is considered as white Gaussian noise. The gains of all the mixing matrices and the full-rank noise in (4.14), are adjusted such that we have $\text{SIR} = -30\text{dB}$, $\text{SNR}_\eta = -20\text{dB}$, and $\text{SNR}_n = 5\text{dB}$ for the fetus, which are in accordance with typical experimental recordings. A ten second segment of eight channels generated with this method can be seen in Figs. 4.8 and 4.9, for $\theta_{mf} = 10^\circ$ and $\theta_{mf} = 60^\circ$, respectively.

In this example, considering the number of sources and mixtures, we have an under-determined mixture of sources. Therefore, linear source separation algorithms that seek as many sources as sensors are not expected to extract all the signal components. For illustration, we apply the JADE algorithm [30], to these synthetic signals in order to extract eight independent components. The resultant components can be seen in Figs. 4.10 and 4.11, corresponding to the inter-subspace angles $\theta_{mf} = 10^\circ$ and $\theta_{mf} = 60^\circ$, respectively. From these results we can see that JADE has extracted the three maternal components in both cases. On the other hand, for the case in which the maternal and fetal subspaces were rather far

($\theta_{mf} = 60^\circ$) two fetal components have been extracted, while when the maternal and fetal subspaces are close ($\theta_{mf} = 10^\circ$), only one noisy fetal component has been extracted. A closer look at these figures shows that there are traces of the fetal components left in the maternal signal, especially for $\theta_{mf} = 10^\circ$. The other components in both figures correspond to the structured and full-rank noises. Further discussions concerning the interpretation of components extracted by multichannel source separation techniques are presented in Chapters 6 and 8.

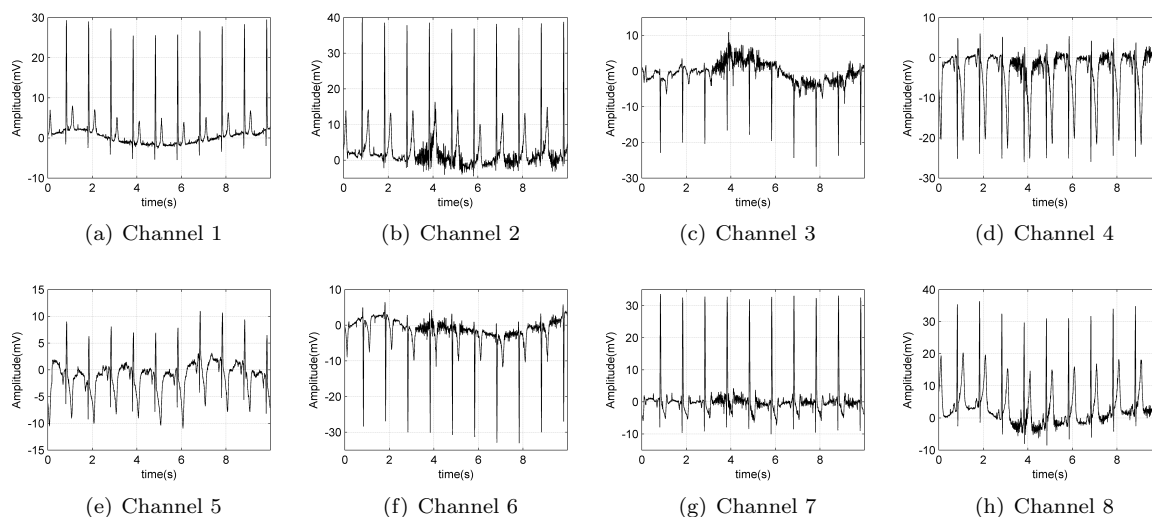


Figure 4.8: Synthetic maternal abdomen signals with $\theta_{mf} = 10^\circ$, $SIR = -30\text{dB}$, $SNR_\eta = -20\text{dB}$, and $SNR_n = 5\text{dB}$.

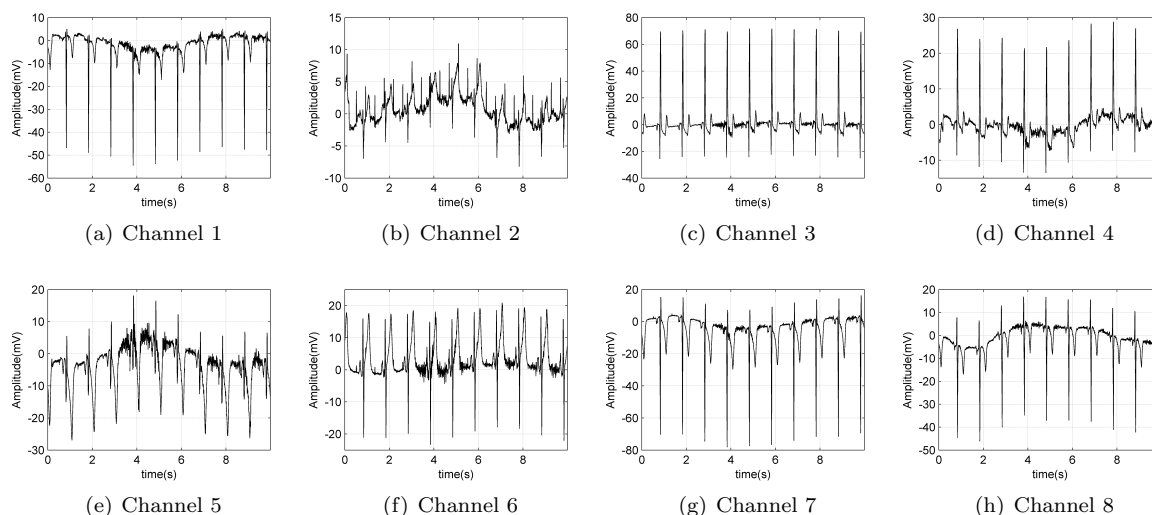


Figure 4.9: Synthetic maternal abdomen signals with $\theta_{mf} = 60^\circ$, $SIR = -30\text{dB}$, $SNR_\eta = -20\text{dB}$, and $SNR_n = 5\text{dB}$.

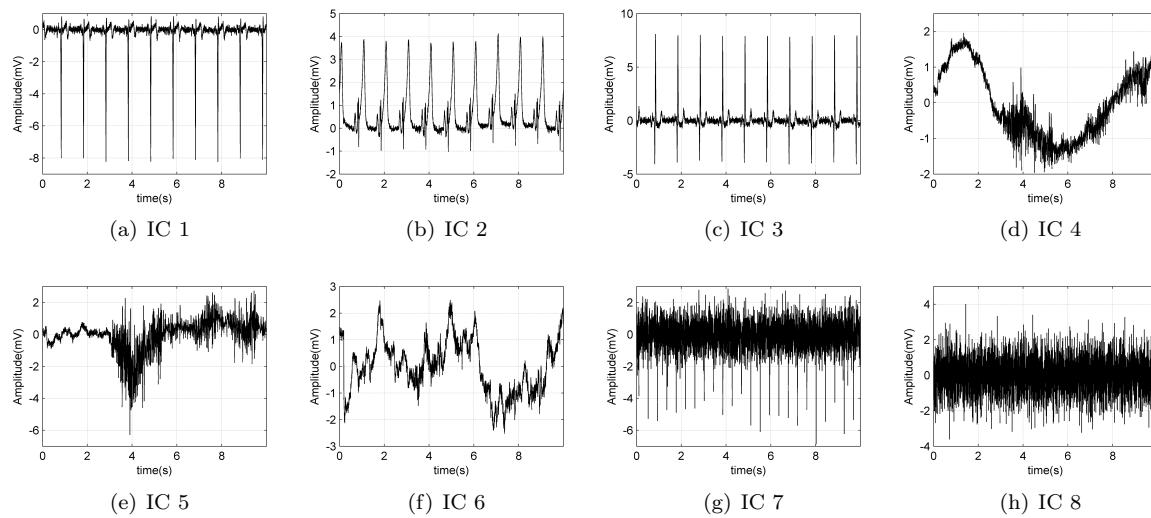


Figure 4.10: Independent Components extracted from the synthetic multichannel recordings of Fig. 4.8. Strong maternal components can be seen in the first three components. Components four to six, more or less correspond with the baseline wander, muscle artifacts, and electrode movement noises. Noisy fetal R-peaks are seen in the seventh component and the last component is principally noise, with some minor traces of the fetal peaks.

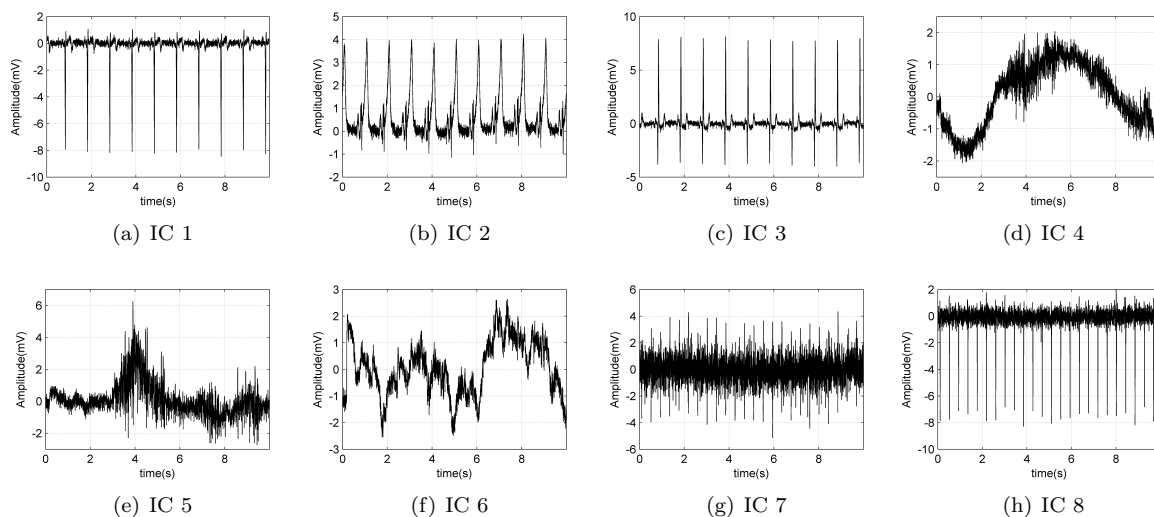


Figure 4.11: Independent Components extracted from the synthetic multichannel recordings of Fig. 4.9. Strong maternal components can be seen in the first three components. Components four to six, more or less correspond with the baseline wander, muscle artifacts, and electrode movement noises. Fetal R-peaks are seen in the last two components.

4.7 Time-varying volume conductor models

As mentioned in previous sections, the mixing matrices of cardiac signals are generally functions of time, having oscillations coupled with the respiration rate or the heart beat. Such oscillatory couplings may be modeled in our synthetic ECG model by using the Givens transform [73].

It is known that any rotation (orthonormal) matrix $R \in \mathbb{R}^{N \times N}$ can be decomposed into $L = N(N-1)/2$ single rotations corresponding to the number of possible rotation planes in this space. This decomposition, known as the Givens decomposition, is as follows:

$$R = \prod_{i=1 \dots N-1, j=i+1 \dots N} R_{ij}, \quad (4.18)$$

where R_{ij} is the Givens rotation matrix of the i - j plane, derived from an N -dimensional identity matrix with the four following changes in its entries:

$$\begin{aligned} R_{ij}(i, i) &= \cos(\theta_{ij}), & R_{ij}(i, j) &= \sin(\theta_{ij}) \\ R_{ij}(j, i) &= -\sin(\theta_{ij}), & R_{ij}(j, j) &= \cos(\theta_{ij}), \end{aligned} \quad (4.19)$$

where θ_{ij} is the rotation angle between the i and j axes, in the i - j plane. The R_0 matrix presented in (4.10) is a three-dimensional example of the general rotation in (4.18).

Now in order to achieve a time-varying rotation matrix that is coupled with an external source, such as the respiration rate or heart beat (either for the adult or the fetus), any of the θ_{ij} rotation angles can oscillate with the external source frequency, as follows:

$$\theta_{ij}(t) = \theta_{ij}^{\max} \sin(2\pi ft) \quad (4.20)$$

where θ_{ij}^{\max} is the maximum deviation of the θ_{ij} rotation angle, and f is the frequency of the external source. The axes that are coupled with the oscillatory source depend on the nature of the sources of interest and the geometry of the problem (i.e. the relative location and distance of the sources). Apparently depending on this geometry other means of coupling are also possible.

The hereby presented idea can be used to achieve time-varying mixtures in (4.8)-(4.9), by either using them as multiplicative factors for the mixing matrices, or for modifying the left or right eigenmatrices of these mixing matrices.

4.8 Summary and Conclusions

In this chapter a three-dimensional model of the dipole vector of the heart was presented. The model was used for the generation of synthetic multichannel signals recorded from the body surface of adults and pregnant women. A practical means of generating realistic ECG noises recorded in real conditions was also developed and the application of the proposed model, especially for fetal ECG studies, was illustrated through a simulated example. Due to the multidimensional nature of cardiac signals, the hereby presented models are of especial interest for multichannel analysis methods presented in later chapters.

Considering the simplicity and generality of the proposed model, there are many other issues which may be addressed in future works.

The proposed model is based on the single dipole model of the heart, which limits the cardiac subspace dimensions to three. In Chapter 6, it is shown that up-to five or six dimensions may be necessary for the better representation of the cardiac dipole. Therefore, the accuracy of the model can be increased by using more dimensions to represent the cardiac subspace.

In future works, the idea of extending the single dipole model to *moving dipoles* [208], can also be studied. For such an approach, the dynamic representation in (4.7) can be very useful. In fact the moving dipole would be simply achieved by adding oscillatory terms to the x , y , and z coordinates in (4.7), to represent the speed of the heart's dipole movement. Therefore, besides the modeling aspects of the proposed model, it can also be used as a means of verifying the performance of different heart models.

Following the discussions in Section 4.4, it is known that Gaussian mixtures can be used to model any ECG signal, even with asymmetric shapes such as the T-wave, which are rather common in real recordings. However, in these cases, two or more Gaussian terms or a *log-normal* function may be required to model each asymmetric shape. For such applications, it could be simpler to substitute the Gaussian functions

with naturally asymmetric functions, such as the *Gumbel function* which has a Gaussian shape that is skewed towards the right- or left-side of its peak [80].

The modeling of abnormal ECG signals (either for adults or fetuses) should also be considered in future works. The dynamic models presented in this chapter can effectively generate multichannel normal or abnormal beats with consistent morphology. However, cardiac abnormalities usually appear as occasional odd-beats between a set of normal ones. We believe that the Hidden-Markov Model (HMM) can be very effective for modeling such abnormalities. A preliminary idea of this extension was recently presented in [40].

Finally, we note that we should be cautious about how we interpret the components of dipole or multipole expansions of the cardiac potentials. In Chapter 6 we will discuss that multipole expansions of electromagnetic field potentials are not generally unique and a change of observation point (sensor location) can lead to different dipole expansions. Therefore, the dipolar sources that we modeled in (4.7) should not be envisaged as a truly existing dipolar source that is independent of sensor locations. This is of especial importance for interpreting the results of source separation methods applied to real and synthetic data.

Chapter 5

A Bayesian ECG Filtering Framework

5.1 Introduction

Up to now, many approaches have been developed for ECG denoising. Despite of the rich literature in this field, there are still many applications, such as the fetal ECG extraction problem, which lack reliable signal processing tools to extract the weak ECG components contaminated with background noise and permit the measurement of subtle features in the ECG. The numerous non-cardiac ECG contaminants overlap with the cardiac components in the frequency domain, particularly in the 0.01Hz to 100Hz range. Bandpass filtering is therefore inadequate to suppress such contaminants [37], [74].

Ensemble averaging (EA) is another common approach for the extraction of small cardiac components from the noise contaminated ECG. However, as EA requires the averaging of many beats, the subtle but important inter-beat variations in the cardiac cycle are lost in the averaging procedure [116]. As an improvement over EA, classical *adaptive filter* (AF) architectures have also been used for the noise cancellation of ECGs containing baseline wander, power line interference, EMG noise, and motion artifacts [220, 199, 113].

For stationary signals, the *Wiener filter* (WF) is the optimal linear filtering technique in the *minimum mean square error* (MMSE) sense, applied either in a causal sense in the time-domain, or as the non-causal WF applied in the frequency domain. However, the WF is not expected to (and does not) give good results for a noisy ECG, due to the non-stationary nature of the cardiac signal. In some related works, filtering approaches have been proposed based on time-frequency [115], [116], and time-scale [107], [141] WFs. The intuition behind the use of the time-frequency or wavelet transforms in these applications is to apply the WF in two domains, to facilitate the tracking of ECG nonstationarities.

Wavelet denoising (WD) is now a common practice for denoising of signals having multi-resolution characteristics such as the cardiac signal in the ECG. Donoho [49], proposed a soft thresholding method for the so-called *shrinkage* of the noise components in the wavelet domain. Their approach together with some *ad hoc* variants of it, have since been used for many applications, including high-resolution ECG denoising [107, 162, 2]. In these cases, the model of the ECG is essentially based on the frequency content of the ECG and to some degree, the localization of the ECG peaks in time.

In this chapter, we demonstrate that by using a realistic model to describe the quasi-periodic behavior of the ECG, the idea of model-based filtering may be further extended to a general Bayesian filtering framework for adult and fetal ECG denoising.

5.2 Review of the Bayesian Filtering Theory

A classical problem in estimation theory is the estimation of the hidden states that are observable through a set of measurements of a system with an underlying dynamic model. The well-known *Kalman filter* (KF) is one such method and under certain general constraints, it can be proved to be the optimal filter in the MMSE sense [106]. The conventional KF assumes a linear model for the system dynamics and observation equations. In practice however, most systems are nonlinear in nature and in order to extend

the idea of conventional KF to such systems, several variants of the original KF have been developed. In this section, the theoretical foundations of some of these extensions are briefly reviewed to facilitate the presentation of the proposed methods.

5.2.1 The Extended Kalman Filter

The *extended Kalman filter* (EKF) is an extension of the standard KF to nonlinear systems. Consider a discrete-time nonlinear system with the unobserved underlying state vector \mathbf{x}_k and observation vector \mathbf{y}_k at time instant k . A dynamic model of this system may be represented as follows:

$$\begin{cases} \mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{w}_k, k) \\ \mathbf{y}_k = g(\mathbf{x}_k, \mathbf{v}_k, k) \end{cases} \quad (5.1)$$

where $f(\cdot)$ is the state evolution function and $g(\cdot)$ represents the relationship between the state vector and the observations. The process and measurement noise vectors are \mathbf{w}_k and \mathbf{v}_k respectively, with associated covariance matrices $Q_k = E\{\mathbf{w}_k \mathbf{w}_k^T\}$ and $R_k = E\{\mathbf{v}_k \mathbf{v}_k^T\}$. The initial estimate of the state vector is also assumed to be known and is given by $\bar{\mathbf{x}}_0 = E\{\mathbf{x}_0\}$, with $P_0 = E\{(\mathbf{x}_0 - \bar{\mathbf{x}}_0)(\mathbf{x}_0 - \bar{\mathbf{x}}_0)^T\}$.

Our objective is to find the MMSE estimate of the vector \mathbf{x}_k . In order to use the KF formalism for this problem, it is necessary to derive a linear approximation of (5.1) near a desired reference point $(\hat{\mathbf{x}}_k, \hat{\mathbf{w}}_k, \hat{\mathbf{v}}_k)$ [84], [142]. This leads to the following linear approximate model:

$$\begin{cases} \mathbf{x}_{k+1} \approx f(\hat{\mathbf{x}}_k, \hat{\mathbf{w}}_k, k) + A_k(\mathbf{x}_k - \hat{\mathbf{x}}_k) + F_k(\mathbf{w}_k - \hat{\mathbf{w}}_k) \\ \mathbf{y}_k \approx g(\hat{\mathbf{x}}_k, \hat{\mathbf{v}}_k, k) + C_k(\mathbf{x}_k - \hat{\mathbf{x}}_k) + G_k(\mathbf{v}_k - \hat{\mathbf{v}}_k) \end{cases} \quad (5.2)$$

where

$$\begin{aligned} A_k &= \left. \frac{\partial f(\mathbf{x}, \hat{\mathbf{w}}_k, k)}{\partial \mathbf{x}} \right|_{\mathbf{x}=\hat{\mathbf{x}}_k} & F_k &= \left. \frac{\partial f(\hat{\mathbf{x}}_k, \mathbf{w}, k)}{\partial \mathbf{w}} \right|_{\mathbf{w}=\hat{\mathbf{w}}_k} \\ C_k &= \left. \frac{\partial g(\mathbf{x}, \hat{\mathbf{v}}_k, k)}{\partial \mathbf{x}} \right|_{\mathbf{x}=\hat{\mathbf{x}}_k} & G_k &= \left. \frac{\partial g(\hat{\mathbf{x}}_k, \mathbf{v}, k)}{\partial \mathbf{v}} \right|_{\mathbf{v}=\hat{\mathbf{v}}_k} \end{aligned} \quad (5.3)$$

Moreover, to simplify the matrix notations, the matrices F_k and G_k are usually absorbed into the noise covariance matrices as follows:

$$Q_k \leftarrow F_k Q_k F_k^T \quad , \quad R_k \leftarrow G_k R_k G_k^T$$

With these notations, the EKF algorithm may be summarized as follows:

$$\begin{aligned} \hat{\mathbf{x}}_{k+1}^- &= f(\hat{\mathbf{x}}_k^+, \mathbf{w}, k) \Big|_{\mathbf{w}=\bar{\mathbf{w}}_k} & \mathbf{r}_k &= \mathbf{y}_k - g(\hat{\mathbf{x}}_k^-, \mathbf{v}, k) \Big|_{\mathbf{v}=\bar{\mathbf{v}}_k} \\ K_k &= P_k^- C_k^T [C_k P_k^- C_k^T + R_k]^{-1} & \hat{\mathbf{x}}_k^+ &= \hat{\mathbf{x}}_k^- + K_k \mathbf{r}_k \\ P_{k+1}^- &= A_k P_k^- A_k^T + Q_k & P_k^+ &= P_k^- - K_k C_k P_k^- \end{aligned} \quad (5.4)$$

where by definition \mathbf{r}_k is the *innovation* signal, $\bar{\mathbf{w}}_k = E\{\mathbf{w}_k\}$, $\bar{\mathbf{v}}_k = E\{\mathbf{v}_k\}$, $\hat{\mathbf{x}}_k^- \doteq \hat{\mathbb{E}}\{\mathbf{x}_k | \mathbf{y}_{k-1}, \dots, \mathbf{y}_1\}$ is the *a priori* estimate of the state vector in the k -th stage using the observations \mathbf{y}_1 to \mathbf{y}_{k-1} , and $\hat{\mathbf{x}}_k^+ \doteq \hat{\mathbb{E}}\{\mathbf{x}_k | \mathbf{y}_k, \dots, \mathbf{y}_1\}$ is the *a posteriori* estimate of this state vector after using the k -th observation \mathbf{y}_k . P_k^- and P_k^+ are defined in the same manner to be *a priori* and *a posteriori* estimates of the state vector covariance matrices before and after using the k -th observation, respectively, i.e.

$$P_k^- \doteq E\{\mathbf{e}_k^- \mathbf{e}_k^{-T}\} \quad (5.5)$$

$$P_k^+ \doteq E\{\mathbf{e}_k^+ \mathbf{e}_k^{+T}\} \quad (5.6)$$

where $\mathbf{e}_k^- \doteq \mathbf{x}_k - \hat{\mathbf{x}}_k^-$ and $\mathbf{e}_k^+ \doteq \mathbf{x}_k - \hat{\mathbf{x}}_k^+$.

The matrix P_k^+ is an essential part of the Kalman filter and it is calculated and updated as the filter propagates in time. The eigenvalues of this matrix can be used to form an *error likelihood ellipsoid*, also known as *concentration ellipsoid* [209, p. 79], that represents the region of highest likelihood for the state vector \mathbf{x}_k . This likelihood ellipsoid provides a confidence region for the estimated signals.

5.2.2 The Extended Kalman Smoother

As with the *Kalman smoother*, the *extended Kalman smoother* (EKS) uses the information of *future* observations to give better estimates of the *current* state. Due to this non-causal nature, the EKS is expected to have a better performance compared with the EKF. The EKS algorithm basically consists of a forward EKF stage followed by a backward smoothing stage. Depending on the smoothing strategy, smoothing algorithms are usually classified into *fixed-lag* or *fixed-interval* smoothers [68]. In this work the fixed interval EKS is used, since the filtering procedure is carried out offline on the entirety of each ECG signal. For real-time applications of the proposed EKS methods, the fixed-lag smoother is usually more appropriate.

5.2.3 The Unscented Kalman Filter

For highly nonlinear systems, the linear estimate of the nonlinear model does not provide a good approximation of the model, and hence the EKF will not track the desired signal around sharp turning points (such as for the ECG). In recent years there has been great interest towards the extensions of the KF to highly nonlinear systems [35]. The *unscented Kalman filter* (UKF) is a filter based on the *unscented transform* (UT), a method for the estimation of the first and second order statistics of the outputs of highly nonlinear systems with Gaussian inputs [84]. In fact, for the UKF the linearization of the system model is no longer necessary since the prior estimate of the state covariance matrix, which is required for the Kalman gain calculations in (5.4), is directly estimated using the UT. The theory of the UKF and its implementation issues have already been discussed in the literature and the reader is referred to [84, ch. 7] for a detailed mathematical description. Note that the UKF is numerically sensitive and the covariance matrices estimated by the UT may become semi-definite; therefore much effort has been made to achieve numerically stable versions of this algorithm. The UKF algorithm used in this work is based on the *ReBEL* Matlab[®] library, previously developed for nonlinear Bayesian filtering [217] and is optimized to prevent the estimated covariance matrices from becoming semi-definite.

5.3 Methods

In Chapter 4, it was shown that realistic ECG signals can be produced by using a set of state-space equations to model the ECG. In the following, similar equations are used to model the temporal dynamics of ECG signals, for designing a Bayesian filter for ECG denoising. In order to do so, we use a variant of (4.6) in its discrete form with the assumption of a small sampling period δ :

$$\begin{cases} \theta_{k+1} = (\theta_k + \omega\delta) \bmod (2\pi) \\ z_{k+1} = -\sum_i \delta \frac{\alpha_i \omega}{b_i^2} \Delta\theta_i \exp\left(-\frac{\Delta\theta_i^2}{2b_i^2}\right) + z_k + \eta \end{cases} \quad (5.7)$$

where z_k represents the ECG at time instant k , $\Delta\theta_i = (\theta_k - \theta_i) \bmod (2\pi)$, η is a random additive noise that models the inaccuracies of the dynamic model (including the baseline wander), the summation i is taken over the number of Gaussian functions used for modeling the shape of the desired ECG channel, and the other parameters follow the definitions given in Section 4.4¹. As we showed in Chapter 4, due to the flexibility of Gaussian mixtures, by using a sufficient number of Gaussian functions they can be fitted to signals recorded from different ECG leads. However, in order to illustrate the general filtering framework, in this chapter we only use five Gaussians to model the ECG channels containing the P, Q, R, S, and T-waves.

Here forth, θ_k and z_k are assumed as the state variables, and ω , α_i , θ_i , b_i and η are assumed as i.i.d Gaussian random variables considered to be process noises². Following the notation in (5.1), the system state and process noise vectors are defined as follows:

$$\begin{aligned} \mathbf{x}_k &= [\theta_k, z_k]^T, \\ \mathbf{w}_k &= [\alpha_P, \dots, \alpha_T, b_P, \dots, b_T, \theta_P, \dots, \theta_T, \omega, \eta]^T, \end{aligned} \quad (5.8)$$

¹Note that in this notation, θ_k is a time-varying phase at time instant k , while θ_i is the phase representing the center of the i -th Gaussian function (with i varying over the indexes P, Q, R, Q, T, or etc.).

²In our case, the Gaussianity of the process noises is a ‘working assumption’ that helps us in using the Kalman filtering formulation in practice; but does not exactly hold in theory for some of the parameters such as ω .

and the process noise covariance matrix is given as $Q_k = E\{\mathbf{w}_k \mathbf{w}_k^T\}$.

5.3.1 Linearization of the Nonlinear Dynamic ECG Model

In order to set up an EKF model based on the nonlinear synthetic model (5.7), it is necessary to have a linearized version of the model. Consequently, the state-equation (5.7) requires linearization using (5.2) and (5.3). By defining:

$$\begin{cases} \theta_{k+1} = F_0(\theta_k, \omega, k) \\ z_{k+1} = F_1(\theta_k, z_k, \omega, \alpha_i, \theta_i, b_i, \eta, k), \end{cases} \quad (5.9)$$

the following equations represent the linearized model with respect to the state variables θ_k and z_k :

$$\begin{aligned} \frac{\partial F_0}{\partial z_k} &= 0 & \frac{\partial F_0}{\partial \theta_k} &= \frac{\partial F_1}{\partial z_k} = 1 \\ \frac{\partial F_1}{\partial \theta_k} &= -\sum_{i \in \{P, Q, R, S, T\}} \delta \frac{\alpha_i \omega}{b_i^2} \left[1 - \frac{\Delta \theta_i^2}{b_i^2}\right] \exp\left(-\frac{\Delta \theta_i^2}{2b_i^2}\right) \end{aligned} \quad (5.10)$$

Similarly, the linearization of (5.9) with respect to the process noise components yields:

$$\begin{aligned} \frac{\partial F_0}{\partial \omega} &= \delta & \frac{\partial F_1}{\partial \eta} &= 1 & i \in \{P, Q, R, S, T\} \\ \frac{\partial F_0}{\partial \alpha_i} &= \frac{\partial F_0}{\partial b_i} = \frac{\partial F_0}{\partial \theta_i} = \frac{\partial F_0}{\partial \eta} &= 0 \\ \frac{\partial F_1}{\partial \alpha_i} &= -\delta \frac{\omega \Delta \theta_i}{b_i^2} \exp\left(-\frac{\Delta \theta_i^2}{2b_i^2}\right) \\ \frac{\partial F_1}{\partial b_i} &= 2\delta \frac{\alpha_i \omega \Delta \theta_i}{b_i^3} \left[1 - \frac{\Delta \theta_i^2}{2b_i^2}\right] \exp\left(-\frac{\Delta \theta_i^2}{2b_i^2}\right) \\ \frac{\partial F_1}{\partial \theta_i} &= \delta \frac{\alpha_i \omega}{b_i^2} \left[1 - \frac{\Delta \theta_i^2}{b_i^2}\right] \exp\left(-\frac{\Delta \theta_i^2}{2b_i^2}\right) \\ \frac{\partial F_1}{\partial \omega} &= -\sum_i \delta \frac{\alpha_i \Delta \theta_i}{b_i^2} \exp\left(-\frac{\Delta \theta_i^2}{2b_i^2}\right) \end{aligned} \quad (5.11)$$

5.3.2 Observation Equations

The noisy ECG recordings are assumed to be observations for the KF. The relationship between the states and observations of the KF depends on the location of the electrodes and the origin of the measurement noise. For example, motion artifacts, environmental noise or bioelectrical artifacts such as EMG or electrogastric noise, may be assumed as additive measurement noises.

In addition to the noisy ECG observations, the phase θ can also be added as a second observation. In fact, by studying the values of Table 4.1, it is noticed that the R-peak is always assumed to be located at $\theta = 0$ and the ECG contents lying between two consecutive R-peaks are assumed to have a phase between 0 and 2π (or $-\pi$ and π). So by simply detecting the R-peaks an additional observation, namely ϕ_k , is achieved. While R-wave detection is a rather simple and routine procedure in most cases (cf. Appendix D), one may benefit from more sophisticated and robust approaches for very low SNR applications [9]. This additional phase information will also help to synchronize the dynamical KF trajectories with the reference noisy signals, without the need for manual synchronization. This RR-interval *phase warping* technique may be assumed as a generalization of the *external reference* previously used for the synchronization of AFs for event-related signals [143], [113].

Hence the phase observations ϕ_k and the noisy ECG measurements s_k may be related to the state vector as follows:

$$\begin{bmatrix} \phi_k \\ s_k \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} \theta_k \\ z_k \end{bmatrix} + \begin{bmatrix} u_k \\ v_k \end{bmatrix} \quad (5.12)$$

where $R_k = E\{[u_k, v_k]^T [u_k, v_k]\}$ is the observation noise covariance matrix.

In the context of estimation theory, the variance of the observation noise in (5.12) represents the degree of reliability of a single observation. In other words, when a rather precise measurement of the




Figure 5.1: An illustration of the phase assignment approach




Figure 5.2: Several cycles of the ECG phase-wrapped in the state space

states of a system is valid the diagonal entries of R_k are small, and the KF gain is adapted so as to rely on that specific measurement. However, for the epochs where the measurements are too noisy or there are no measurements available, the R_k entries are large and the KF tends to follow its internal dynamics rather than tracking the observations [172]. Recall that the phase state variable θ_k is periodic, starting on $\theta = 0$ at the R-peak and ending on $\theta = 2\pi$ at the next R-peak. Although the only valid phase observation is obtained from the R-peak locations, it is possible to linearly assign a phase value between 0 and 2π to the intermediate samples, as illustrated in Fig. 5.1. The later presented results are all based on this linear phase assignment. However, to indicate the increased uncertainty in the phases assigned to the intermediate samples, the first diagonal entry of R_k corresponding to the time varying variance of the measurement phase noise, may be increased. Another alternative and rather sophisticated approach for the estimation of the intermediate phase values is to directly detect the location of the P, Q, S, and T-waves from the original signal. However, the previous approach is preferred since R-peak detection is far more reliable in high noise scenarios.

5.3.3 Estimation of the Model Parameters

Prior to the implementation of the filter, it is necessary to select the values of the process and measurement noise covariance matrices. Generally, by using m Gaussian kernels in (5.7), the process noise vector defined in (5.8) has $3m + 2$ entries (here 17), leading to a $(3m + 2) \times (3m + 2)$ process noise covariance matrix Q_k . But if the noise sources are assumed to be uncorrelated with each other, a reasonable approximation adopted here, then the matrix is simplified to be diagonal. The measurement noise covariance matrix R_k is similarly considered to be diagonal.

In order to automate the parameter selection procedure for any given ECG, the parameters should be estimated from the signal itself. For this, as described in the previous subsection, any noisy ECG may be transformed to a three-dimensional representation by plotting the noisy ECG versus the periodic phases that are assigned to each sample in polar coordinates on the unit circle ($r = 1$). A typical phase-wrapped ECG with additive noise may be seen in Fig. 5.2. It is now possible to estimate the dynamic model parameters for the given ECG. For this, the mean and variance of the *phase-wrapped ECG* is calculated for all phases between 0 and 2π . This gives the average of the ECG waveform. A typical phase-wrapped ECG with this approach is depicted in Fig. 5.3. The error bar in this figure corresponds with the standard deviation (SD) of different ECG cycles around the mean ECG. Next, the problem is to find the optimal parameters of (5.7) that can best fit the mean ECG. In this stage, many optimization methods may be




Figure 5.3: An average and standard deviation-bar of 30 ECG cycles of a noisy ECG

used. For example as suggested in [38], by using a nonlinear least-squares approach, the best estimate of these parameters in the MMSE sense can be found. A practical means of solving this nonlinear least-squares problem is the *lsqnonlin* function of Matlab[®] that was used to estimate the initial parameters for the later presented results.

The next step is to find an estimate for the covariance values of Q_k . Theoretically, Q_k is a measure of the ECG morphology consistency. Therefore, large entries of Q_k correspond to the non-consistent parameters, e.g. due to abnormal beats. We can practically find rough estimated for Q_k , by using the error values as depicted in Fig. 5.3. In fact, in this step we are attempting to calculate the magnitude of the deviation of the parameters of the Gaussian functions in (4.4) around the estimated mean that best model the acceptable deviations of the ECG around the mean ECG ($\bar{\text{ECG}}(\theta)$). This is again a nonlinear least-squares problem that is solved by finding the optimal parameters that generate the best fit of the mean ECG within the upper and lower ranges of $\bar{\text{ECG}}(\theta) + \sigma_{\text{ECG}}(\theta)$ and $\bar{\text{ECG}}(\theta) - \sigma_{\text{ECG}}(\theta)$. A similar method was used in Section 4.4.3 for fitting the synthetic model parameters to real recordings.

It should be noted that the parameter estimation procedure detailed above is an offline approach that estimates the optimal parameter values for any dataset. It is also possible to develop an online extension of this algorithm that estimates the model and noise parameters from the most recent cycles of the ECG. However, for short ECG recordings we have found that the parameters of the model remain relatively constant and this online process is unnecessary for such signals. Furthermore, as it will be later noted, by monitoring the innovation signal of the KF, it is possible to fine-tune the estimated parameters throughout the filtering process, without the need for their re-estimation. Therefore, for any new dataset, it is possible to start with approximate values for the parameters of the model, which have been calculated from similar data and to modify these values throughout the filtering process.

The angular frequency ω may be set to $\omega = 2\pi/T$, where T is the RR-interval period in each ECG cycle. For short signals with minor RR-interval deviations, a simpler approximation is to use a global ω using the average RR-interval of the whole signal. It should be noted however, that ω can also be considered to vary on an intra-beat basis too, since the PR and QT-intervals are known to change with varying autonomic tone, heart rate and to some extent, with each changing RR-interval. The results presented in this chapter are based on the simpler approximation above using a global angular frequency.

The variance of the process noise η should also be estimated. Noting that η is a parameter that represents the imprecision of the dynamic model, neglecting the other physiological sources that influence the ECG, a simple estimate for this parameter would be a zero mean Gaussian random variable with an appropriate variance. An intuitive value for this variance may be found from the deviations of the isoelectric segment of the ECG between the end of the T-wave and the beginning of the next P-wave, which correspond to the ending segments of the ECG error-bar of Fig. 5.3.

From (5.12) we can observe that u_k is the phase measurement noise. As mentioned before, the phase for each beat is determined from the R-peaks of the signal. A possible noise source for u_k is the sampling error that occurs when the actual R-peak is located between two sample times. Another noise source arises from the additive noise spikes that can cause a misdetection of the R-peak location. The first of these may be easily modeled by assuming that the R-peak is uniformly distributed between two consecutive samples. By considering that each ECG cycle is equivalent to 2π in the phase domain, u_k would be

uniformly distributed in the range of $\pm\omega\delta/2$, where ω is the angular frequency and δ is the sampling period. With this assumption we have: $E\{u_k^2\} = (\omega\delta)^2/12$. Although there have been rather robust R-peak detectors developed to overcome the misdetection of the R-peaks [9], a precise study of this issue requires the amplitude noise to be related to the phase error (or the so called *phase jitter*), a practice that has been well-studied in other contexts [195]. In this work, for the sake of simplicity, it is assumed that the R-peak detector is reliable and the only phase error is due to the imprecision of the sampling time. Moreover, as mentioned in the previous subsection, for the intermediate phases laying between consecutive peaks the variance of the phase noise can be increased to indicate the imprecision of the phase values.

There are also several ways to estimate the variance of the measurement noise, v_k . One method is to estimate the noise power from the deviations of the whole signal around the phase-wrapped ECG, or from the portions of the ECG between two successive T and P-waves. A quantitative study of the accuracy of these estimates is presented later. There are also online approaches for noise power estimation, which have been previously suggested for similar applications [116], and apparently the selection of the method depends on the origin of the expected noise.

5.3.4 Stability and Convergence Issues

The stability and convergence issues of the KF and its extensions have been well-discussed in the literature. In order to ensure numerical stability of the KF equations, and to prevent the covariance matrices from becoming semi-definite, the *Joseph stabilized form* [76, p. 233] of (5.4) is used for the *a posteriori* covariance estimation, to guarantee positive-definite estimates of the covariance matrix.

In practice due to the Gaussian assumption on the noise sources and the initial state vector values, the state estimate entries of $\hat{\mathbf{x}}_k^+$ should lie within the envelope of the square roots of their corresponding diagonal entries in P_k^+ for the majority of the time. Therefore, by monitoring the variance of the filter estimate, it is possible to detect the filter divergence. Moreover it is shown that it is possible to stabilize the KF online, by introducing a forgetting factor in the original filter equations [68].

Another approach known as *serial measurement update* [76, p. 232], uses different observations one-by-one to aid the stability of the KF. This technique requires the observation noise covariance matrix to be diagonal. However, for non-diagonal observation noise covariance matrices, decorrelation methods have been developed to diagonalize this matrix [76].

It is also practically convenient to monitor the covariance matrix of the innovation signal throughout the filtering procedure and to compare it with the innovation covariance matrices estimated by the KF [125], [76]. This provides a means of monitoring the fidelity of the filter and updating the values of Q_k and R_k . Specifically, with a diagonal (or diagonalized) noise covariance matrix of R_k , the following term can be formed for the i -th ECG measurement:

$$\gamma_i = \frac{1}{N} \sum_{k=i-N+1}^i \frac{(r_k^s)^2}{h_k} \quad (5.13)$$

where r_k^s is the second entry of the zero-mean innovation vector of \mathbf{r}_k defined in (5.4), corresponding with the k -th ECG measurement, N is the length of the averaging window, and h_k is the KF estimated variance of r_k^s given by:

$$h_k = E\{(r_k^s)^2\} = \mathbf{c}_k^T P_k^- \mathbf{c}_k + \sigma_{v_k}^2 \quad (5.14)$$

where \mathbf{c}_k is the second row of the C_k matrix defined in (5.3), and $\sigma_{v_k}^2 = E\{v_k^2\}$ is the second diagonal entry of R_k . An identical term can be defined for the phase observation ϕ_k .

In fact, γ_i is an average of the variances of the N recent ECG innovations, normalized by their KF estimated variances h_k . Therefore, as long as the KF is performing correctly, $\gamma_i \approx 1$. Values of γ_i much greater than unity indicate that the innovation signal variance is being underestimated by the KF, while values close to zero indicate that the innovation signal variance is being overestimated. Therefore, by monitoring γ_i it is possible to adaptively modify the KF noise parameters (such as Q_k and R_k), to ensure the filter stability and to achieve a better filtering performance. For example, by using the M most recent samples of the innovation signal, $\sigma_{v_k}^2$ can be adaptively modified as follows:

$$\sigma_{v_k}^2 = \lambda \sigma_{v_{k-1}}^2 + (1 - \lambda) \frac{1}{M} \sum_{j=k-M}^{k-1} (r_j^s)^2 \quad (5.15)$$

where $0 < \lambda < 1$ is the adaptation coefficient³. In [197], a similar means of online modification of the R_k and Q_k entries have been presented.

For the UKF, the algorithm presented in [84] has three parameters, $\tilde{\alpha}$, $\tilde{\beta}$, and $\tilde{\kappa}$. These control the stability of the filter and enable the algorithm to be fine-tuned for systems with different degrees of nonlinearity and non-Gaussian inputs. The parameter $\tilde{\alpha}$ is an indication of the spread of the state variables around their mean and is selected to be a small positive value in the range of $10^{-4} \leq \tilde{\alpha} \leq 1$. The parameter $\tilde{\beta}$ is used to incorporate prior knowledge about the distribution of the state vector, with $\tilde{\beta} = 2$ being optimal for Gaussian distributions [84]. The parameter $\tilde{\kappa}$ is a secondary *ad hoc* scaling parameter that is selected in accordance with the size of the state vector and the higher order statistics of the noise distributions [98], with $\tilde{\kappa} = 0$ being the optimal selection for a state vector of size two. A mathematical study of the effect of these parameters on the UKF accuracy compared with the EKF, can be found in [98, 97].

5.3.5 Practical Filtering Scheme

Before presenting the experiments and results, the scope of the proposed filtering scheme needs to be further clarified. Following the discussions of this section, by using the Bayesian framework we are attempting to utilize *a priori* information about the underlying dynamics of ECG signals to extract the ECG components from background noise. Hence, compared with conventional filtering schemes that perform rather ‘blindly’, Bayesian filters are naturally expected to give superior results as long as we provide them with valid *a priori* information concerning the signal and noise dynamics. This point becomes important when considering that abnormal ECGs can have high inter-beat variations in their wave timings or morphology, meaning that the underlying dynamics of the signals are not valid in pathological beats.

In the presented approach, due to the phase wrapping of the RR-interval to 2π , normal inter-beat variations of the RR-interval (between 10% to 20%), or consistent RR-interval abnormalities such as *bradycardia* or *tachycardia* do not considerably affect the filter performance. However, for abnormalities that only appear in some of the ECG cycles, the phase error of the model can lead to large errors in the Gaussian functions’ locations. A similar case can occur when the R-peak is misdetected. In particular, for morphological abnormalities that appear in some of the ECG cycles, such as the *premature ventricular contraction* (PVC) [127], the filtering performance is not expected to be satisfactory for low input SNRs, since neither the model nor the measurements are reliable for the filter. For such occasional morphologic changes, even temporal adaptation of the filter parameters is not helpful, as the filter does not have sufficient time to adapt itself. However the benefit of the Gaussian mixture representation is that the effect of each Gaussian term vanishes very quickly (in less than an ECG period), meaning that the errors are not propagated to the following ECG cycles⁴. Moreover, by monitoring the state estimates’ covariance matrices and the variations of the innovation signals, it is possible to detect such unexpected abnormalities. Of course, it should be considered that the accurate denoising of abnormal ECGs with high morphologic changes remains an open problem even for conventional filtering methods.

Finally we note that all the results presented in this work, have been implemented offline. However the recursive KF equations are originally designed for online applications. Even for the EKS, considering the quasi-periodic nature of the ECG, fixed lag smoothers with only one or two cycles of ECG lag can be used.

5.4 Evaluation on Simulated Noisy Mixtures

5.4.1 The Dataset

The MIT-BIH *normal sinus rhythm database* [72, 160], is used to evaluate the performance of the proposed methods. This database has been recorded at a sampling rate of 128Hz from 18 subjects with no significant

³For $M = 1$ (single-step update), (5.15) reduces to the autoregressive model suggested in [197], and for $M > 1$ (5.15) represents a moving average filter with the λ parameter changing the slope of the filter’s response. For ECG signals having sharp changes, the moving average model was found to be more robust to the peak changes.

⁴A formal justification of this statement requires mathematical proof of stability and convergence of the EKF, EKS, and UKF for the proposed model, which is beyond the scope of the current study. The reader is referred to [173], for some related discussions.

Table 5.1: Parameters of the process and observation noises

θ_i (rad)	Gaussian kernel center $\pm 0.05\pi$
α_i (mV)	Gaussian kernel peak $\pm 10\%$ of the peak amplitude
b_i (rad)	Gaussian kernel width $\pm 0.05\pi$
η (mV)	1% of maximum ECG peak
ω (rad/s)	Mean beat-to-beat angular frequency ($\bar{\omega}$) \pm SD
u_k (rad)	$0.00 \pm (\bar{\omega}\delta)/\sqrt{12}$
v_k (rad)	Ranges over different SNRs

arrhythmias. From this database 190 low-noise segments of 30 seconds without considerable artifacts were visually selected for the implementation of the proposed filters. These segments were taken from different subjects, recorded from the standard V_I, V_{II}, and V_{III} ECG leads. The heart rate of these ECG segments varied from 55 BPM to 90 BPM (74.1 ± 12.0 BPM on average), with RR-interval deviations between 5% to 25% ($16.7 \pm 6.2\%$ on average).

5.4.2 Noise Generation

Mathematically, white noise is defined to have a flat spectral density function over all frequencies. However, real noise sources have non-flat spectral densities that decrease in power at higher frequencies, making the spectrum *colored* and the noise samples correlated in time. There are different ways of generating colored noise [105], and realistic ECG artifacts (cf. Section 4.5). For the current study, we model the noise color by a single parameter representing the slope of a spectral density function that decreases monotonically with frequency:

$$S(f) \propto \frac{1}{f^\beta}, \quad (5.16)$$

where f is the frequency and β is a measure of noise color. *White noise* ($\beta = 0$), *pink noise* ($\beta = 1$) or *flicker noise*, and *brown noise* ($\beta = 2$) or the *random walk process*, are three of the most commonly referenced noises. The realization of colored noises with spectral densities described by (5.16), generally require nonlinear frequency domain filtering of white noise⁵. For random processes, the expected value of the squared magnitudes of their frequency transforms, or namely the *periodogram*, is known to be an estimate of the spectral density function of the original samples [152]. Therefore, in order to generate colored noise following (5.16), samples of white noise can be generated and transferred into the frequency domain using the discrete Fourier transform (DFT). By altering the frequency components of the DFT according to (5.16), and transferring the reshaped DFT back to the time-domain, typical samples of colored noise are realized. Note that this approach of frequency domain filtering causes transient behavior in the generated noise time series that should be discarded from the samples.

5.4.3 Implementation

Having derived the state equations (5.7), the observation equations (5.12), the linearized state equations of the ECG dynamic model (5.10), (5.11), and the model parameters, the implementation of the EKF, EKS and the UKF are now possible. The procedure for calculating the parameters of the model and the noise covariance matrix entries were explained in section 5.3. Using the explained methods, the mean ECG waveform was extracted for each ECG segment and the parameters of the five Gaussian kernels were calculated using the nonlinear least-squares method for each ECG segment. Due to the variety of the ECG leads and the wide range of studied SNR, the covariance matrices were calculated using a simple approach using the peak values of the Gaussian functions. The parameter selection approach is summarized in Table 5.1. These parameters can be further customized for specific ECG leads and special ranges of the input SNR.

In section 5.3.4, the control parameters of the UKF algorithm were introduced. Throughout the simulations, these parameters were set to $\tilde{\beta} = 2$ and $\tilde{\kappa} = 0$, and are the optimal selections for a system having two state variables with a Gaussian input noise. The parameter $\tilde{\alpha} = 1$ was also empirically derived as a compromise between performance and stability in different SNRs for the studied database.

⁵Except for special cases such as $\beta = 2$ that can also be achieved through linear time-domain filtering of white noise.

Table 5.2: WD parameter combinations tested over the database

Parameter	Values
Mother wavelet	Daubechies 1...8, Coiflets 1...5, Symlets 1...8
Shrinkage rule	SURE, Heuristic SURE, Universal, Minimax
Thresholding strategy	Hard, Soft
Rescaling approach	No scaling, Single level, Multiple level
Decomposition level	1...10

5.4.4 Benchmark Methods

In order to have a comparison between the performance of the proposed methods and conventional ECG denoising schemes, wavelet denoising (WD), adaptive filtering (AF), and conventional finite impulse response (FIR) filtering were also tested on the database. Implementations of these standard methods are now described.

Conventional WD schemes are characterized by different parameters that allow the algorithms to be customized for different mixtures of signal and noise sources. The type of the *mother wavelet*, *shrinkage rule*, *hard-* versus *soft-thresholding*, *noise level rescaling approach*, and *number of decomposition levels* are among the different parameters of common WD algorithms [137]. There are of course, some general rules concerning the selection of these parameters. For example, the mother wavelet is usually selected from families that somehow resemble the shape of the desired signal, or the rescaling approach and shrinkage rules are selected according to the nature (white versus colored) and variance of the noise. However, the literature on the applications of WD for ECG is rather broad and diverse, making it difficult to judge what may be the best combination of these parameters. Consequently, a rather exhaustive search was carried out on the different combinations of the above mentioned parameters to find the best WD scheme for the database being studied. The different parameters that were tested on this database are listed in Table 5.2. Among the different tested combinations, the *Stein's unbiased risk estimate* (SURE) shrinkage rule, together with a single level rescaling and a soft thresholding strategy always gave superior results. Among the tested mother wavelets, *Coiflets2*, *Coiflets3*, *Symlets4*, and *Symlets5* gave superior results. The best decomposition level ranged from 5 to 8, and showed little significant difference within this range. These results were achieved by the comparison of the SNR curves for input SNR ranging from 35dB to -5dB and averaged over 100 Monte Carlo runs with different random noise input vectors. The following reported results are based on the *Coiflets3* mother wavelet with 6 levels of decomposition.

The adaptive filtering approach suggested in [113], was the next filtering approach implemented on the dataset. The original results presented in [113] have been reported on a synthetic ECG, formed by taking a single QRS complex from an ECG and concatenating the same waveform several times to generate a *deterministic* ECG with no timing or morphology variation from beat to beat. An impulse signal, time-synchronized with the R-wave, was also used as the reference channel to enable single channel filtering of the ECG. Although their reported results are rather impressive for this simulated ECG, the method is not expected to give identical results on real ECGs. This is mainly due to the non-stationary behavior of the ECG that causes the ECG shape and RR-duration to change from beat-to-beat. However this AF scheme was also implemented on the dataset as another benchmark. In the original work [113], the number of the AF weights (L) is selected to be equal to the number of samples of the deterministic QRS complexes. For operating on real ECGs with variable RR-intervals we set L to the maximum sample period between the RR-waves of the input ECG. A convergence rate of $\mu = 0.1$, led to rapid adaptation and stable filter outputs for all the SNR range of this study (-5dB - 30dB).

The last filtering approach applied to the dataset was a typical FIR filter, consisting of two cascaded highpass and lowpass filters, with an overall pass-band of 0.4Hz-40Hz, a pass-band ripple of 1dB, and a stop-band attenuation of about 60dB. The main frequency components of typical ECGs lie within this frequency range [11], and the selection of a wider or narrower bandwidth would be a compromise between the attenuation of in-band ECG and noise components.

5.4.5 Results

In order to investigate the performance of the different methods, artificial white and colored Gaussian noise with different variances were generated and added to the ECG segments, and the noisy signals

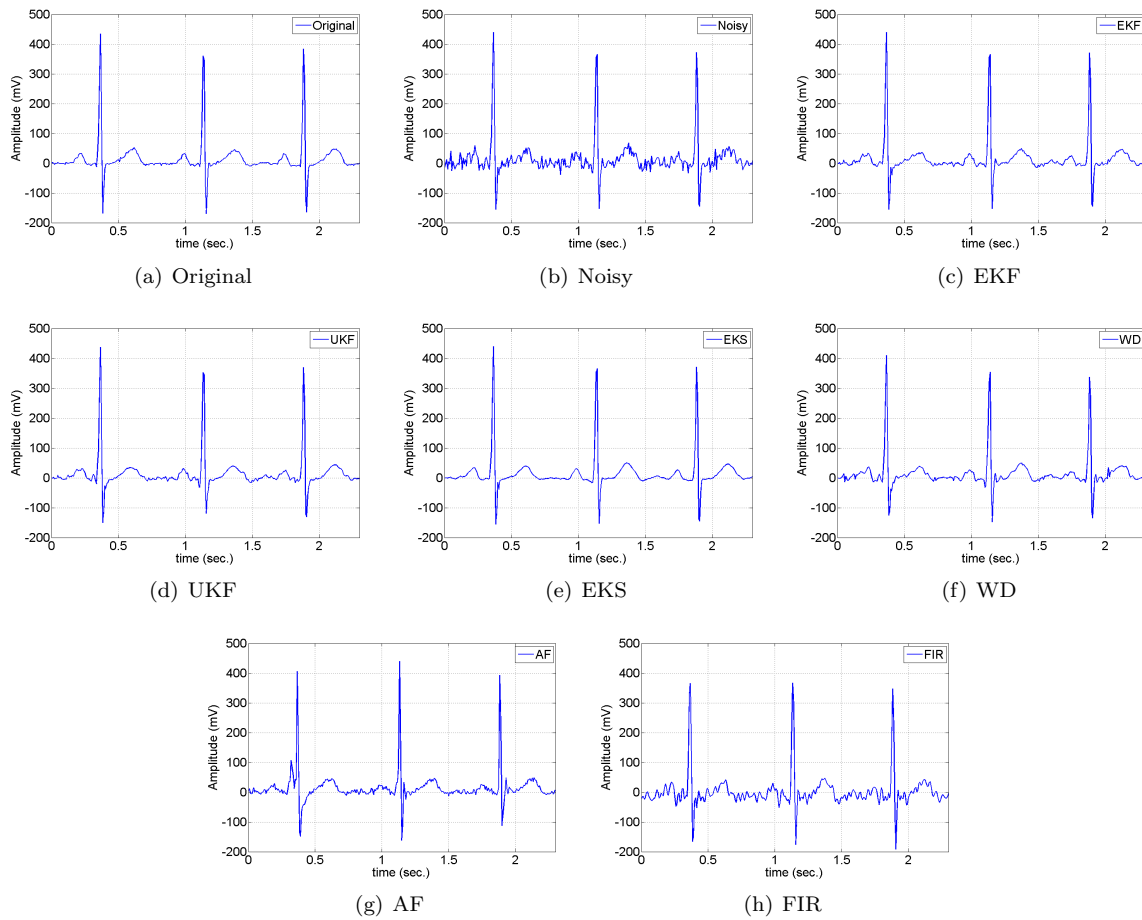


Figure 5.4: Typical filtering results for an input signal of 6dB.

were presented to the proposed filters. To ensure the consistency of the results, the whole procedure was repeated over the 190 ECG segments; each time using a different set of random noise at the input. The filter output SNR calculation was averaged over the whole 190 results. The SNRs were generally calculated over the second half of the filtered segments, to ensure that the transient effects of the filters would not influence the SNR calculations.

In Fig. 5.4, typical results of the FIR, AF, WD, EKF, EKS, and the UKF are presented for an input SNR of 6dB. Visually comparing these results, it can be seen that the proposed methods have effectively tracked the original signal in a rather low input SNR scenario. The EKS demonstrates the smoothest result, while the UKF outperforms the EKF, particularly around the sharp turning points of the signal. In fact, the main difference between the EKF and UKF results are in the QRS complex of the ECG, where the EKF performs slightly less well, since it tends to follow the noisy signal rather than the system dynamics. The reason for this may be seen by observing equations (5.2) to (5.4). According to (5.3), at low sampling rates and in the rapidly changing regions of the ECG, the approximated matrix A_k has large entries due to the differentiation, that in turn causes an increase in the values of the P_{k+1}^- matrix. This means that the KF tends to rely less on the dynamic model. This assumption was validated by limiting the maximum and minimum values of the derivatives calculated from (5.3). This change led to results that demonstrate increased flexibility around the QRS complex of the ECG and is reflective of the fact that derivative-free filters such as the UKF are more robust to severe nonlinearities of the input time series. This also suggests that the EKF and EKS can provide better performance for signals having higher sampling rates (at a cost of higher computation load).

Among the conventional filtering approaches, the WD outperforms the AF and the FIR; but contains some large ripples that do not correspond to the true ECG.

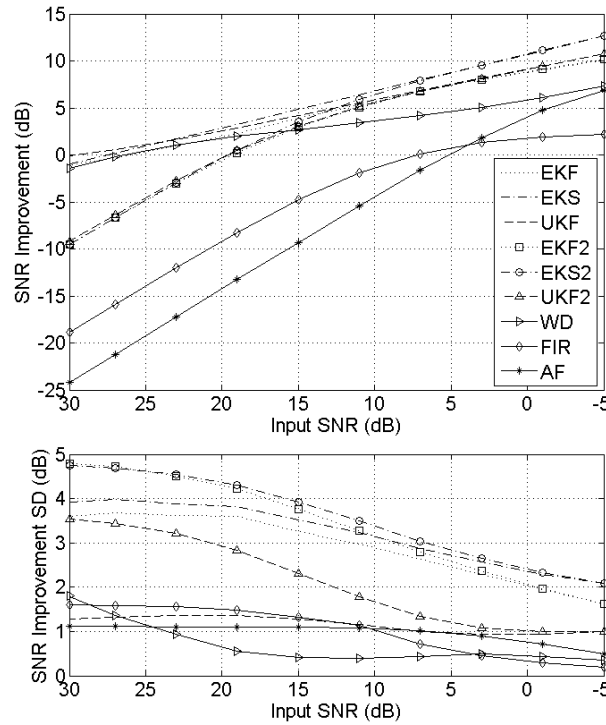


Figure 5.5: The mean (top) and Standard Deviation (bottom) of the filter output SNR improvements versus different input SNRs. In these curves EKF, UKF, and EKS correspond to the results without $\sigma_{v_k}^2$ adaptation, and EKF2, UKF2, and EKS2 refer to the ones with $\sigma_{v_k}^2$ adaptation. Refer to text for further details.

For a quantitative comparison, the mean and SD of the SNR improvements⁶ versus different input SNRs, achieved over 190 ECG segments are plotted in Fig. 5.5. The results of the SNR improvements calculated over the ST-segment, which is extremely sensitive to noise and of great clinical significance, can also be seen in Fig. 5.6. In the presented results of Figs. 5.5 and 5.6, two approaches have been used for the EKF, UKF, and the EKS; the first without the online adaptation of $\sigma_{v_k}^2$, and the second with adaptation. For the first case $\sigma_{v_k}^2$ was fixed to the *a priori* known variance of the additive noise, and for the second case, this initial value was adaptively modified by the filter. For the latter, the adaptation window length was selected to be $M = 13$ and is approximately equivalent to a 100ms window for the sampling rate of 128Hz. This window length is wider than the normal QRS complex, and therefore should be sufficient to prevent sharp variations in $\sigma_{v_k}^2$. The adaptation coefficient in (5.15) was set to $\lambda = 0.6$, a value that was empirically found to provide a compromise between adaptation time and stability of $\sigma_{v_k}^2$. Figs. 5.5 and 5.6 illustrate how the results achieved for a constant $\sigma_{v_k}^2$ (i.e. without its online adaptation), are almost linearly related to the input SNR, and are generally better than the adaptively changing $\sigma_{v_k}^2$ results. In fact for input SNRs below 18dB, the EKF, UKF, and the EKS degrade the input SNR when $\sigma_{v_k}^2$ is adapted; but for SNRs below approximately 10dB, the results with and without adaptation are asymptotically the same. In either case the EKS demonstrates the best average performance, and the UKF performs marginally better than the EKF. Among the conventional filtering methods the WD outperforms both the FIR filter and the AF, both of which are inferior to the other techniques. Furthermore, for input SNRs above 18dB WD outperforms the proposed methods with $\sigma_{v_k}^2$ adaptation, but still underperforms the constant $\sigma_{v_k}^2$ results.

The reason for the asymptotic behavior of the results with and without $\sigma_{v_k}^2$ adaptation can be explained by revisiting (5.14), where we see that the variance of the innovation signal estimated by the KF consists of two parts; the uncertainty of the model parameters (the first term), and the uncertainty of the observations ($\sigma_{v_k}^2$). In high input SNR scenarios, the first term dominates the second. Therefore, when adaptively

⁶The SNR improvement is defined as the output SNR of the filter (in decibels) minus the input SNR (in decibels). Negative SNR improvements apparently indicate a degradation of the input SNR caused by the filtering procedure.

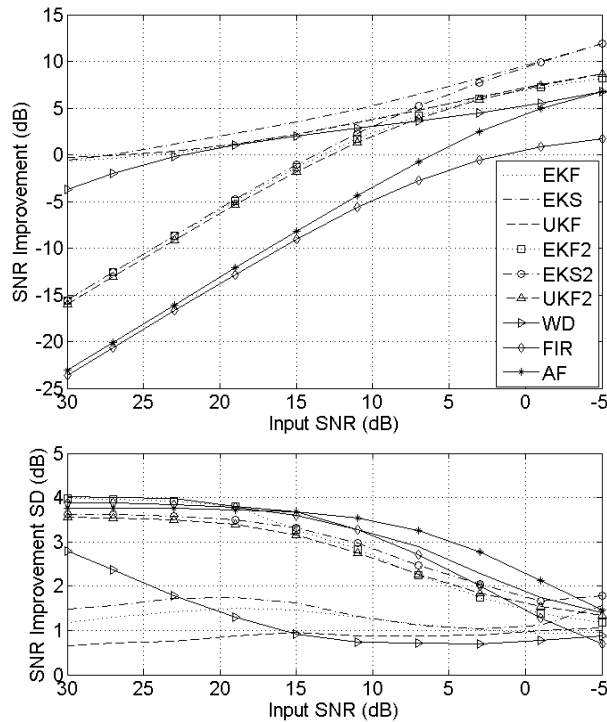


Figure 5.6: The mean (top) and Standard Deviation (bottom) of the filter output SNR improvements over the ST-segment of the ECG, versus different input SNRs. In these curves EKF, UKF, and EKS correspond to the results without $\sigma_{v_k}^2$ adaptation, and EKF2, UKF2, and EKS2 refer to the ones with $\sigma_{v_k}^2$ adaptation. Refer to text for further details.

changing $\sigma_{v_k}^2$ in high input SNRs, we are in fact miscounting the model uncertainties as measurement errors, leading to the overestimation of the measurement noise. Conversely, in low SNR scenarios $\sigma_{v_k}^2$ dominates the first term, and the online adaptation is performed correctly. In different applications with specific ranges of input SNR, this information can help to make online corrections to the estimated noise parameters. These results also suggest that for stationary noise processes, it is preferable to keep the filter noise parameters constant, or alternatively, increase the innovation variance averaging length of M .

From the SD plots in Figs. 5.5 and 5.6, it is seen that WD has the least deviation over the different ECG segments, and among the proposed methods the UKF has the least deviation; meaning that the UKF results are more robust to dataset variations.

In section 5.3.3 it was noted that the initial values of the measurement noise variance may be estimated from the SD of the whole phased-wrapped ECG, or just using the SD of the isoelectric segment of the ECGs between two consecutive T and P waves. To show the accuracy of these estimates, the SNRs estimated from both methods were compared with the true SNR of the signal in presence of additive white noise. The mean and SD of these estimates versus the true input SNR, calculated over the whole database are depicted in Fig. 5.7. As can be seen from this figure, both methods have underestimated the true SNR (especially in high input SNRs); but the SNR estimated from the isoelectric segment is very close to the true values for input SNRs below 20dB. In practice, as we have to estimate the measurement noise variance from the noisy signals, this information may be used as a correction curve for finding accurate estimates of the noise variance.

The result of the noise color study is depicted in Fig. 5.8 for the SNR improvement of the EKS as a function of the input SNR and the input noise color β . This result was achieved using the average of the 190 ECG segments, both with and without the online adaptation of the measurement noise variance $\sigma_{v_k}^2$.

As can be seen in Fig. 5.8, the EKS performance decreases almost linearly as the noise color ranges from white ($\beta = 0$) to pink ($\beta = 1$), while the slope of decrease is larger for lower input SNRs. As with the previous results, it is seen that the constant $\sigma_{v_k}^2$ results (top curve) outperforms the results with $\sigma_{v_k}^2$ adaptation (bottom curve), and both curves merge in input SNRs below approximately 10dB.




Figure 5.7: Estimated signal SNR versus true SNR, from the isoelectric line segment of the ECG between the T-waves and the P-waves (solid), and from the whole ECG cycle (dashed).




Figure 5.8: Effect of noise color on the EKS results. (top) without $\sigma_{v_k}^2$ adaptation, (bottom) with $\sigma_{v_k}^2$ adaptation. Both curves merge for input SNRs below 10dB.

The study of the noise color effect was not extended beyond pink noise, since as the noise becomes more colored in spectrum (β increases), the time-domain samples of the noise will have longer-term correlations. This means that for a valid statistical study of the noise effect longer ECG signals are required. Moreover, colored noise as defined in (5.16) is not generally guaranteed to be a *wide-sense stationary* (WSS) process, meaning that the colored noise samples do not have the same variance⁷. When adding such noise to the ECG, each sample receives a different amount of noise; making the overall SNR criteria rather meaningless. Anyhow, very long-term correlations manifest themselves as slow waves that do not change the critical ECG waveforms and can be effectively removed with baseline wander removal techniques presented in Appendix D.

It should be further noted that although the monotonic shape of the SNR surfaces of Fig. 5.8 proves the consistent behavior of the filtering approaches in different noise colors, the Bayesian framework is originally based on the assumption of white noise sources. In fact, for colored (structured) noise with a known spectral behavior, the systematic approach is to use parametric spectral estimation methods to model the colored noise as the output of a system driven by white noise. Subsequently, the state-space model of this system can be augmented with the original system model. In this way, the dynamics of the noise sources are also considered in the filtering procedure.

5.5 Evaluation on Real Maternal/Fetal Mixtures

In this section we use the proposed method for removing maternal ECG artifacts from fetal recordings. A similar procedure is also used as a post-processing to enhance the fetal ECG.

The well-known DaISy fetal ECG database is used for illustration [46]. The database consists of five abdominal and three thoracic channels recorded from the abdomen and chest of a pregnant woman with a sampling rate of 250Hz. In this example, we will only use the first channel of this dataset for illustration, which is depicted in Fig. 5.9(a). Next, according to the details explained in previous sections,

⁷Specifically for the *random walk process* ($\beta = 2$), the variance of the noise samples increases linearly with time [152].




Figure 5.9: The first channel of the DaISy dataset [46], recorded from a maternal abdominal lead before and after the EKS procedure.

the following steps were applied to remove the maternal ECG: (1) maternal R-peak detection, using one of the thoracic channels of the dataset, (2) calculation of the maternal phase signal ϕ_k , (3) maternal ECG ensemble average calculation, (4) deriving the parameters of the KF (similar to Table 5.1), and (5) applying the EKS to estimate the maternal ECG. By subtracting the maternal ECG estimate from the original signal, the fetal components plus other noises, are left. The maternal ECG estimate and the fetal residual components are depicted in Figs. 5.9(b) and 5.9(c). As a post-processing step, the EKS was applied to the residual fetal components, this time training the filter parameters over the fetal signals. The post-filtered fetal signals are depicted in Fig. 5.9(d).

From the results of Fig. (5.9) we can see that the proposed filtering framework is very effective for the extraction of fetal components from noisy maternal abdominal mixtures, even from a single channel. However, as we can see in Fig. 5.9(d), between $t=6s$ and $t=7s$, the filter has failed to discriminate between the maternal and fetal components when the ECG waves of the mother and fetus fully overlap in time. The reason is that when the maternal and fetal components coincide in time, we have no other *a priori* information for separating these components over the coinciding epochs. This may in fact be considered as an intrinsic limitation of single-channel processing. In later chapters we will present a framework for benefiting from the *mutual information* existing in multichannel recordings to overcome this limitation.




Figure 5.10: Several fetal ECG beats adopted from Fig. 5.9, before and after the post-processing EKS, together with the $\pm\sigma$ and $\pm3\sigma$ confidence envelopes.

As noted in section 5.2.1, an important benefit of Bayesian filtering, is the ability of predicting the accuracy of the estimates. For the Kalman filter, this is readily achieved through the calculation of the error covariance matrices P_k^- and P_k^+ , defined in Section 5.2.1. Specifically, for our application of interest, if the ECG estimation error is Gaussian, the second entry of the state vector \mathbf{x}_k defined in (5.8), corresponding to the ECG, is bounded within the $\pm\sigma$ envelope in 68% of the sample points, where σ^2 is the second diagonal entry of P_k^+ . This is due to the fact that about 68% of the values drawn from a Gaussian distribution are within one standard deviation away from the mean, about 95% of the values are within two standard deviations, and about 99.7% lie within three standard deviations. These probabilities are different for non-Gaussian estimation errors. However, the $\pm\sigma$ envelope can still be used as a rough measure of error spread [209, p. 79]. In Fig. 5.10, several beats of the fetal ECG before and after the post-processing EKS (adopted from Figs. 5.9(c) and 5.9(d)) together with their corresponding $\pm\sigma$ and $\pm3\sigma$ envelopes are plotted. As we can see, this envelope provides the confidence region of the denoised signals.

5.6 Multichannel Extension

Another extension of the proposed filtering framework is for multichannel recordings. Commonly, the ECG is recorded from multiple channels that are highly correlated with each other. This high correlation implies representational *redundancy*, which is not desirable for many applications. However, for signal denoising this redundancy can be exploited to estimate the ECG information embedded in background noise.

This idea may be mathematically expressed as follows: suppose that we use a linear transform to decompose a set of noisy ECG recordings $\mathbf{y}_k \in \mathbb{R}^N$ into N components:

$$\mathbf{y}_k = \mathbf{a}_1 s_k^1 + \mathbf{a}_2 s_k^2 + \dots + \mathbf{a}_N s_k^N \quad (5.17)$$

where k is the time index, $\mathbf{a}_i \in \mathbb{R}^N$ are unit-length columns of the decomposition matrix, and s_k^i ($i = 1 \dots N$) are the decomposing components. Without loss of generality, let's assume that the s_k^i components are sorted in descending order of their contribution in the energy of the noiseless ECG⁸. So if we decide to choose the first M terms as the most dominant ECG components, the desired ECG could be approximated with an M -dimensional subspace and the other $N - M$ components would be considered as the non-ECG or noise components. Therefore, by separating the approximate ECG subspace from the noise subspace, (5.17) may be rewritten as follows:

$$\mathbf{y}_k = H\mathbf{s}_k + \mathbf{v}_k, \quad (5.18)$$

where $H = [\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_M] \in \mathbb{R}^{N \times M}$, $\mathbf{s}_k = [s_k^1, s_k^2, \dots, s_k^M]^T$, and \mathbf{v}_k corresponds to the summation of the last $N - M$ terms in (5.17) having negligible contribution in the desired ECG.

⁸In Chapters 6 and 7, we will present appropriate decomposition methods that can be used for this purpose.

The vector \mathbf{s}_k can be related to the notion of multipole expansion of body surface potentials, presented in Section 3.4. Accordingly, each of the entries of \mathbf{s}_k has a pseudo-periodic rhythm that is synchronous with the heart beat. We can therefore use a multidimensional extension of (5.9) to represent the dynamics of this multipole expansion:

$$\begin{aligned} \theta_{k+1} &= (\theta_k + \omega\delta) \bmod (2\pi) \\ s_{k+1}^1 &= \sum_j \frac{-\delta\alpha_j^1\omega}{(b_j^1)^2} \Delta\theta_j^1 \exp\left[-\frac{(\Delta\theta_j^1)^2}{2(b_j^1)^2}\right] + s_k^1 + \eta_k^1 \\ &\vdots \\ s_{k+1}^M &= \sum_j \frac{-\delta\alpha_j^M\omega}{(b_j^M)^2} \Delta\theta_j^M \exp\left[-\frac{(\Delta\theta_j^M)^2}{2(b_j^M)^2}\right] + s_k^M + \eta_k^M \end{aligned} \quad (5.19)$$

We can now put together the new set of state and observation equations to be used in a Kalman filtering framework. For this, we define the new state and observation vectors as follow:

$$\begin{aligned} \mathbf{X}_k &\doteq [\theta_k, s_k^1, \dots, s_k^M]^T \\ \mathbf{Y}_k &\doteq [\phi_k, \mathbf{y}_k^T]^T \end{aligned} \quad (5.20)$$

where \mathbf{y}_k is the multidimensional observations at time k , and ϕ_k is a coarse estimate of θ_k , generated from the R-peaks, previously defined in (5.3.2). Clearly, \mathbf{X}_k is an $(M+1)$ -dimensional state vector and \mathbf{Y}_k is an $(N+1)$ -dimensional observation vector.

In analogy with (5.8), the new process and observation noise vectors are also defined as follows:

$$\begin{aligned} \mathbf{W}_k &\doteq [\omega, \{\alpha_i^1, \theta_i^1, b_i^1\}, \dots, \{\alpha_i^M, \theta_i^M, b_i^M\}, \eta_k^1, \dots, \eta_k^M]^T \\ \mathbf{V}_k &\doteq [u_k, \mathbf{v}_k^T]^T \end{aligned} \quad (5.21)$$

where \mathbf{W}_k is a large vector containing all the parameters of the Gaussian functions and the other process noised defined in (4.3) and its dimension depends on the number of Gaussian functions used for modeling each multipole term, u_k is the phase noise due to the uncertainty in the R-peak detection, and \mathbf{v}_k is a vector of the ECG observation noises of each channel. All the variables in the definition of \mathbf{W}_k can generally vary with time and the i indexes range over all the number of Gaussian kernels used for modeling the desired waveforms.

With these definitions, the vector representation of the state dynamics and the observation equations of the multidimensional extension are as follows:

$$\begin{aligned} \mathbf{X}_{k+1} &= \mathbf{F}(\mathbf{X}_k, \mathbf{W}_k) \\ \mathbf{Y}_k &= \begin{bmatrix} 1 & \mathbf{0}^T \\ \mathbf{0} & H \end{bmatrix} \cdot \mathbf{X}_k + \mathbf{V}_k \end{aligned} \quad (5.22)$$

where $\mathbf{F}(\cdot)$ is a vector representation of (5.19), and $H \in \mathbb{R}^{N \times M}$ is the transformation matrix defined in (5.18). Due to the nonlinearity of the entries of $\mathbf{F}(\cdot)$, their linearized versions, identical to (5.10)-(5.11), are required for the EKF and EKS. However, the UKF can directly use the nonlinear equations as explained in Section 5.2.3.

We can see that the multidimensional extension is rather straightforward, as long as we can find a suitable decomposing matrix H . This idea was studied in [176] for a dipole expansion ($M=3$) and using the Dower transform to estimate the matrix H . In Chapter 6, we will study different multichannel decomposition techniques that can be used for this purpose.

5.7 Summary and Conclusions

In this chapter, a mathematical framework was proposed for the model-based Bayesian filtering of single and multi-channel noisy ECG recordings. Within this framework several suboptimal filtering schemes were developed and the results were compared with conventional filtering methods. The results demonstrate that the proposed approach can serve as a novel framework for achieving high-resolution ECG of adults or fetuses.

The filtering schemes were based on a modified version of a previously proposed dynamic ECG model. However, the generality and modularity of the proposed methods allow for the improvement of the selected dynamic model. As it was seen, the architecture of the Bayesian filtering methods is such that the filters can work with a rather coarse dynamic model as long as the covariance matrices of the system noise vectors are well estimated. Moreover, the derivative-free characteristics of the UKF also removes the need for an analytical form of the system dynamic model and the statistics of the signals are directly estimated from a finite number of samples, rather than the linearized models required in the EKF or the EKS. This suggests that other morphological models of the ECG can be used instead of mathematical dynamic models. In other words, the nonlinear dynamical model of the ECG may be replaced with other synthetic ECG generators that are capable of generating normal or abnormal ECGs.

The presented results were based on the assumption of additive Gaussian noise sources. With recent developments in Bayesian filtering techniques such as the *particle filter* (PF) [35], other noise distributions may also be considered.

As discussed in Section 5.3.5, in future works, problems related with abnormal ECGs should also be addressed in the model. The extension of these methods to many of the common ECG abnormalities is rather straightforward, since the model parameters may be simply recalculated and used in the filter model. However, for heart defects such as the PVC, where the abnormal wave only appears in certain cycles of the ECG, some revisions are necessary in the filtering process to be able to simultaneously filter the normal and abnormal ECG beats. An intuitive approach to this problem would be to precalculate the filter model parameters for different abnormalities and switch between these values by predicting the existence of an abnormal beat, using a standard ECG classifier, a hidden Markov model [40], or directly using the derived parameters of the model to make a classification itself, as suggested in [38].

Another related issue is the study of the appropriateness of the filtering procedure. In fact, while the Bayesian filtering framework can serve as a powerful tool for the removal of noise from the ECG, one should always be aware of over-filtering of the signals that can lead to the removal of clinically important information from the ECG, or introducing components that do not exist in the true ECG. Therefore the calculation of theoretical lower estimation error bounds, such as the *Cramér Rao* bounds [106], should also be considered in future works. Moreover, some other theoretical and practical aspects such as convergence-time, stability, and estimation bias, also need further studies.

While the results of this chapter were mainly achieved with artificially generated noise, without any specific assumption on the noise origin, the generality of the approach allows the model to be customized for specific applications. In fact, as mentioned in previous sections, for situations where the nature of the contaminating noises are better known, the noise model may be selected according to the particular situation. For example, any spectral information about the noise sources, such as the baseline wander caused by the respiratory system, may be transferred into state-space form by using conventional spectral factorization methods and be augmented with the dynamic model of the system.

In Section 5.6, the single-channel filtering framework was extended to multichannel recordings. In later chapters multichannel decomposition methods are presented, which enable the proposed methods to be linked with multichannel source separation techniques, to simultaneously benefit from the dynamical filtering abilities of KFs and the spatial filtering abilities of source separation methods.

Chapter 6

Linear Multichannel Analysis of Cardiac Signals

6.1 Introduction

Statistical decomposition techniques such as the *principal component analysis* (PCA) [144], *singular value decomposition* (SVD) [102], and *independent component analysis* (ICA) [45, 28] have been widely used as promising methods of multichannel ECG analysis, and noninvasive fetal ECG extraction. However, the research in this area has been mainly carried out from the signal processing perspective and there are few works on the electrophysiological interpretations of the extracted components. Moreover, in previous studies the cardiac source has commonly been treated as a point source or at most a vectorial source [28], while in reality the heart is a distributed source consisting of numerous micro-sources. This is not only a change of viewpoint; but as we will see, it has important influence on the interpretation of the extracted components.

In this chapter, it is shown that by reconsidering the origins of the ECG signals and their relationship with other representations of the cardiac activity such as the *vectorcardiogram* (VCG), it is possible to relate the shapes of the extracted components to their physiological counterparts. Other issues, concerning the dimensionality of ECG signals, necessity of preprocessing, and some proposals for new clinical ECG indexes are also studied.

The main focus of the chapter and the presented case studies are on adult ECG signals without considerable non-cardiac interference. However, in later chapters, we will use these ideas to develop new methods for removing maternal ECG contaminations and fetal ECG filtering.

6.2 Electrophysiology of the Heart: A Signal Processing Perspective

In Chapter 3, the electrophysiology of the heart was studied from an electromagnetic viewpoint. There we noted that an appropriate electromagnetic framework for studying this problem is based on Poisson's equation presented in (3.1). It was further noted that using the Green's theorem, (3.1) may be transformed to an integral equation for calculating the body surface potentials from the current dipole densities. The integral form of Poisson's equation was further used to derive (3.7), a multipole expansion of body surface potentials, which we rewrite here:

$$\phi = \lim_{L \rightarrow \infty} \sum_{l=0}^L \sum_{m=-l}^l a_{lm} s_{lm} \quad (6.1)$$

where s_{lm} and a_{lm} were the *multipole moments* and *multipole coefficients*, respectively, and L was the order of the expansion.

A detailed theoretical discussion on the multipole expansion, from the electromagnetic perspective, can be found in [193, 36, 93]. Here we adopt a signal processing approach to study this problem; but before that, there are several important lessons that should be thought from the electrophysiological

viewpoint. In the following, without going into the electromagnetic details, we briefly refer to the issues that directly influence our signal processing perspective:

1. The heart is a *distributed source* consisted of numerous micro-sources and not a *point source*.
2. The multipole moments are not physically existing sources; they are projections of the distributed source onto a set of basis functions. Therefore, any linear transform of a finite set of body surface electrodes can only retrieve a linear mixture of multipole moments of the cardiac source distribution; but not the source distribution itself.
3. From the expansion (6.1), we can see that the body surface potentials may only be expanded in terms of infinite series of multipole moments of the cardiac source. Depending on the source distribution, volume conduction media, and distance of electrodes from the heart, this infinite expansion may be approximated with a finite number of terms. The *monopole*¹ ($L = 0$), *dipole* ($L = 1$), and *quadrupole* ($L = 2$) expansions are the most common approximations for this expansion [127].
4. Depending on the conductivities of the volume conductor, some sort of *reflections* of the source distributions onto the surfaces of conductivity discontinuity can also appear in the multipole expansion [70, 71]. In this case, what we observe on the body surface is partially due to the distributed source itself and partially due to such reflections. As shown in Section 3.4, the reflection terms can also be factorized as additive terms in (6.1).
5. The multipole expansion is not unique, as we can use any basis function in the expansion [93].

Now, coming back to the signal processing side of the fence, suppose that we have an array of N signals, namely $\mathbf{x}(t) = [x_1(t), x_2(t), \dots, x_N(t)]^T$, recorded from the body thorax². Each entry of $\mathbf{x}(t)$ measures the difference of potentials between two points on the body surface having potentials $\phi_1(t)$ and $\phi_2(t)$ that satisfy Poisson's equation.

Considering that the multipole moments defined in (6.1), are only a function of the source distribution and the multipole coefficients are only a function of electrode locations, $\mathbf{x}(t)$ can be considered as a linear (not necessarily unique) mixture of the multipole moments. However, our linear multichannel analysis algorithms are 'blind' about the aforementioned electrophysiological models³ and using the statistical properties of the observed signals $\mathbf{x}(t)$, they merely seek for linear transforms of the form:

$$\mathbf{y}(t) = B\mathbf{x}(t) \tag{6.2}$$

that maximize some statistical measure, such as uncorrelatedness, independence, or periodicity. On the other hand, the statistical properties of $\mathbf{x}(t)$ depend on both the electrode configuration and the statistics of the distributed sources⁴, and unlike point sources there is no reason to expect the sources estimated from one electrode set to be the same, or even linear mixtures of the same sources extracted from another electrode configuration. In Appendix C, we will show this relationship in a simplified case. In fact, depending on the distributed source statistics and the electrode configuration, each sensor configuration can have a null-space which is different from another electrode configuration. This point has been shown in other contexts for EEG and MEG recordings [165]. For our problem of interest, this implies that *the estimated cardiac subspace is not necessarily invariant under the change of electrode configuration*, which is not a good news from the blind source separation perspective. Nevertheless, depending on the distributed source statistics and the electrode configuration, some specific electrode configurations may be more robust to these variations, the study of which is beyond the scope of the current work.

Another issue which is the direct consequence of the cardiac source distributedness is that the cardiac signal subspace is not finite dimensional. It is of course known that the multipole moments defined in (6.1) are inversely proportional to powers of the distance between the observation point and the distributed source elements, and the terms above the dipole terms ($L > 1$) decay very quickly with the distance

¹As noted in Section 3.4.2, the monopole term is zero in a bounded volume conductor having only internal sources.

²In this chapter, we neglect any other electrical activity except the heart.

³Unless the model information is exploited in the form of a prior, which is not considered here.

⁴As an example of the statistics of the distributed sources, we can refer to second-order statistics such as the cross-correlation of the impressed current dipole moments (Section 3.4), of the different points inside the charge distribution: $R_{\mathbf{J}}(\mathbf{x}, \mathbf{y}, t) = E\{\mathbf{J}(\mathbf{x}, t)\mathbf{J}(\mathbf{y}, t)^T\}$, where $\mathbf{J}(\mathbf{x}, t)$ and $\mathbf{J}(\mathbf{y}, t)$ are the impressed current dipole moments of the points located at \mathbf{x} and \mathbf{y} , respectively (cf. Appendix C).

[93]. However, the higher order terms do exist, and depending on the SNR of the observations and the strength of the non-cardiac components they can be observed in the components extracted from ECG signals. In other words, *theoretically*, no matter how many channels we use, the cardiac components may always manifest as a full-rank signal in the body surface potentials; unless they lie in a null-space of the observations, for some specific electrode configuration. However, *practically*, due to the relatively far distance of the heart from the electrodes, the low SNR of the observations, or the presence of non-cardiac sources, we might only extract a limited number of these cardiac components, while the weaker ones are embedded in the other non-cardiac and noise components⁵ [43].

In Fig. 6.1, a typical segment of fifteen channels of ECG signals, adopted from the Physikalisch-Technische Bundesanstalt diagnostic ECG database (PTBDB) is depicted for illustration [158, 22, 112]. These signals have been recorded from an adult from the twelve standard ECG leads and the three Frank lead VCGs. In this dataset, the electrodes are rather close to the heart. We can visually see that the different channels are all synchronous with the heart beat; but have different shapes. They can be considered as projections of the same distributed source onto the different electrode pairs.

To illustrate the aforementioned properties of the ECG, PCA was used to extract the principal components of the dataset. It is known that due to the electrode configuration used in recording the standard twelve-lead electrodes, only two of the six leads I, II, III, AVR, AVL, and AVF, are linearly independent [127]. Therefore, among the fifteen ECG channel we can only have eleven independent linear mixtures. Using PCA, this fact was used to reduce the dimensions of the ECG to eleven. The resultant components are depicted in Fig. 6.2. Traces of the ECG (specifically the R-peaks), can be seen in most of the extracted components.

As another illustration, the JADE algorithm based on higher-order statistics [31, 30], was also applied to this data for extracting the most independent representation of this dataset. The resultant independent components may be seen in Fig. 6.3. From this figures, we can see traces of the ECG in eight of the extracted components. The other components are mainly noise or due to other sources, such as the respiration. However, a closer look at these components shows that minor traces of the ECG R-peaks exist in them. This is not an observation that has just happened by chance; similar results can be achieved over other ECG dataset, recorded from electrodes close to the heart.

From the discussion of this section, now the question is whether using linear multichannel transforms, the multiple components extracted from the cardiac source convey additional morphological information (other than just the R-peaks) for cardiac studies. This issue will be further discussed in the following section.

6.3 A Study of the Components Extracted from Multichannel Cardiac Recordings

Up to now, we discussed the electrophysiological interpretation of linear mixtures of body surface recordings in terms of multipole expansions. In this section, we will study the problem from a morphological perspective. Some of the hereby presented results were previously reported in [170].

6.3.1 ECG Dimensionality

In the problem of blind or semi-blind source separation, regardless of the origin of the multichannel recordings (cardiac or other), given T samples of an N -channel signal $\mathbf{x}(t)$, we can seek for linear transforms of the form (6.2), which satisfy some desired statistical property. Some of these statistical properties may be visually observed in the scatter plot of the observed signals. To illustrate, in Fig. 6.4 the three-dimensional scatter plot of the ECG signals of Fig. 6.1 are plotted for some different pairs of the original ECG channels. Specifically, the last panel in Fig. 6.4(f) has been constructed from the last three channels of Fig. 6.1, which were recorded from the orthogonal VCG leads. Of course, the original signals in Fig. 6.1 are in a fifteen-dimensional space (or eleven-dimensional due to the electrode redundancy) and we cannot visualize more than three-dimensional pairs. However, even from the few plots in Fig. 6.4, we can notice that due to the pseudo-periodicity of the ECG channels, the scatter plot of any pair of these

⁵For example, using ICA methods based on measures of Kurtosis, ICA will only extract the components that are more ‘kurtotic’ than the other non-cardiac components.

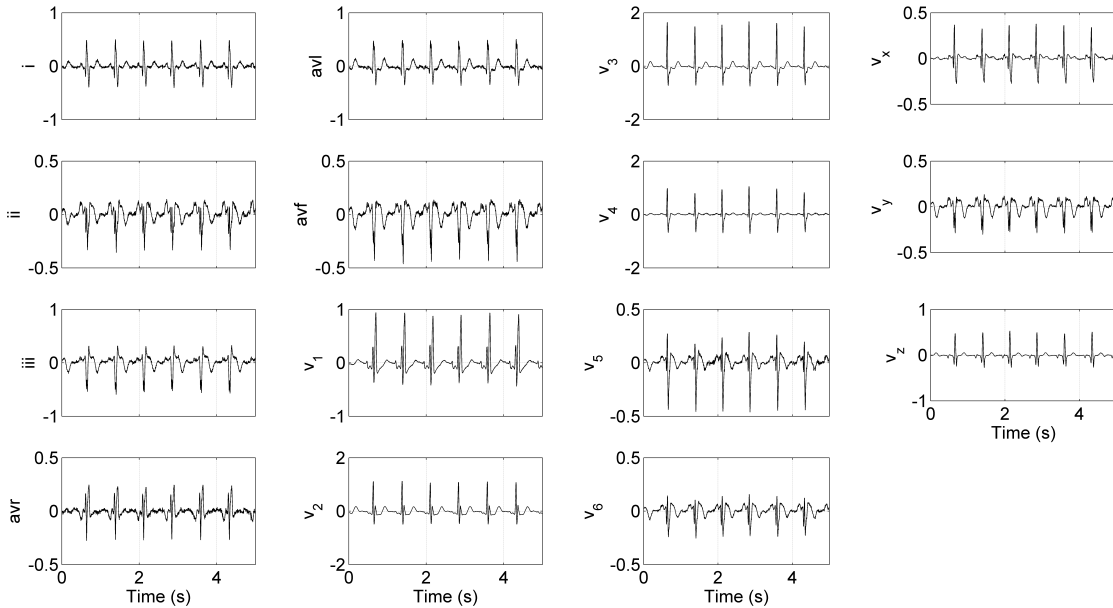


Figure 6.1: A typical segment of the standard 12-lead and the 3 Frank lead ECGs, adopted from the PTB diagnostic database [158, 22, 112].

signals is also pseudo-periodic. Moreover, the region of space swept by the scatter plot is rather *sparse*; meaning that the ECG only occupies a very small region of the N -dimensional space.

Correlation Dimensions

In data analysis, there are several approaches for the estimation of multivariate data dimensions [85]. Here, we do not intend to present the different existing measures of dimensionality; but just to give an idea of how sparse the multichannel ECG representation is, we will calculate a fractal dimension measure known as the *correlation dimension* [85, ch. 9]. For the N -channel signal $\mathbf{x}(t)$, we define

$$C(r) = \frac{1}{N(N-1)} \sum_{i=1}^N \sum_{j=1, j \neq i}^N I_{\|\mathbf{x}_i - \mathbf{x}_j\| < r}, \quad (6.3)$$

where $I_{\|y\| < r} = \{1 : \text{if } \|y\| \leq r, 0 : \text{if } \|y\| > r\}$ and $C(r)$ is the average number of samples of \mathbf{x} that are within the distance of r from each other. $C(r)$ can therefore be considered as the *cumulative distribution function* (CDF) of the sample distances. The correlation dimension is now defined as follows:

$$D_c = \lim_{r \rightarrow 0} \frac{\partial \log C(r)}{\partial \log r} \quad (6.4)$$

For finite number of samples, it is not possible to take the limit in (6.4). Instead, D_c without the limit operator, is plotted versus r and the correlation dimension can be interpolated from this plot for when $r \rightarrow 0$.

This measure was calculated from the data samples of Fig. 6.1. The resultant correlation dimension can be seen in Fig. 6.5(a). For comparison, we have also plotted the correlation dimension of a fifteen-dimensional white Gaussian noise in Fig. 6.5(b). While the D_c of the Gaussian noise is very close to fifteen (the number of its channels), the D_c of the ECG is close to one. However, although the ECG space is very sparse, the regions of space swept by the ECG lie in a rather complicated and curved *manifold*. Therefore, using linear transforms of the array ECG we can not map the ECG onto a lower-dimensional subspace, without losing the ECG components. This is in fact, a geometric restatement of what we mentioned in Section 6.2, concerning the infiniteness of the cardiac signal dimensions.




Figure 6.2: Principal components extracted from the ECG channels of Fig. 6.1 in descending order of their corresponding eigenvalues.




Figure 6.3: Independent components extracted from the ECG channels of Fig. 6.1 using the JADE algorithm in descending order of their energy contribution (explained variance).

Dimensionality Study using Principal and Independent Component Analysis

PCA and ICA can be used to study the sensitivity of the different ECG segments to dimension reduction. As an example, the onset and ending points of each of the P, QRS, and T-segments of the ECG of Fig. 6.1 were manually detected. Next, due to the aforementioned redundancy in the standard ECG lead configuration, a primary PCA was applied to the fifteen channels to reduce the dimensions to eleven. Considering this reduced-dimension data as the reference, PCA was applied to it several times. Each time by eliminating the least significant components corresponding to the smallest eigenvalues, the residual components were back-projected to the original ECG subspace by using the inverse of the decomposing matrix. The average *mean-square error* (MSE) of all channels was calculated in each case over the different ECG segments. In Fig. 6.6, the resultant normalized MSE are plotted for different numbers of dimensions removed. The same procedure was carried out using JADE and each time eliminating the component with the least significance on the total signal energy (explained variance). The normalized MSE results are plotted in Fig. 6.7. As seen in Figs. 6.6 and 6.7, the P and QRS-waves are respectively the most and least sensitive portion of the ECG to dimension reduction. Accordingly, with only five

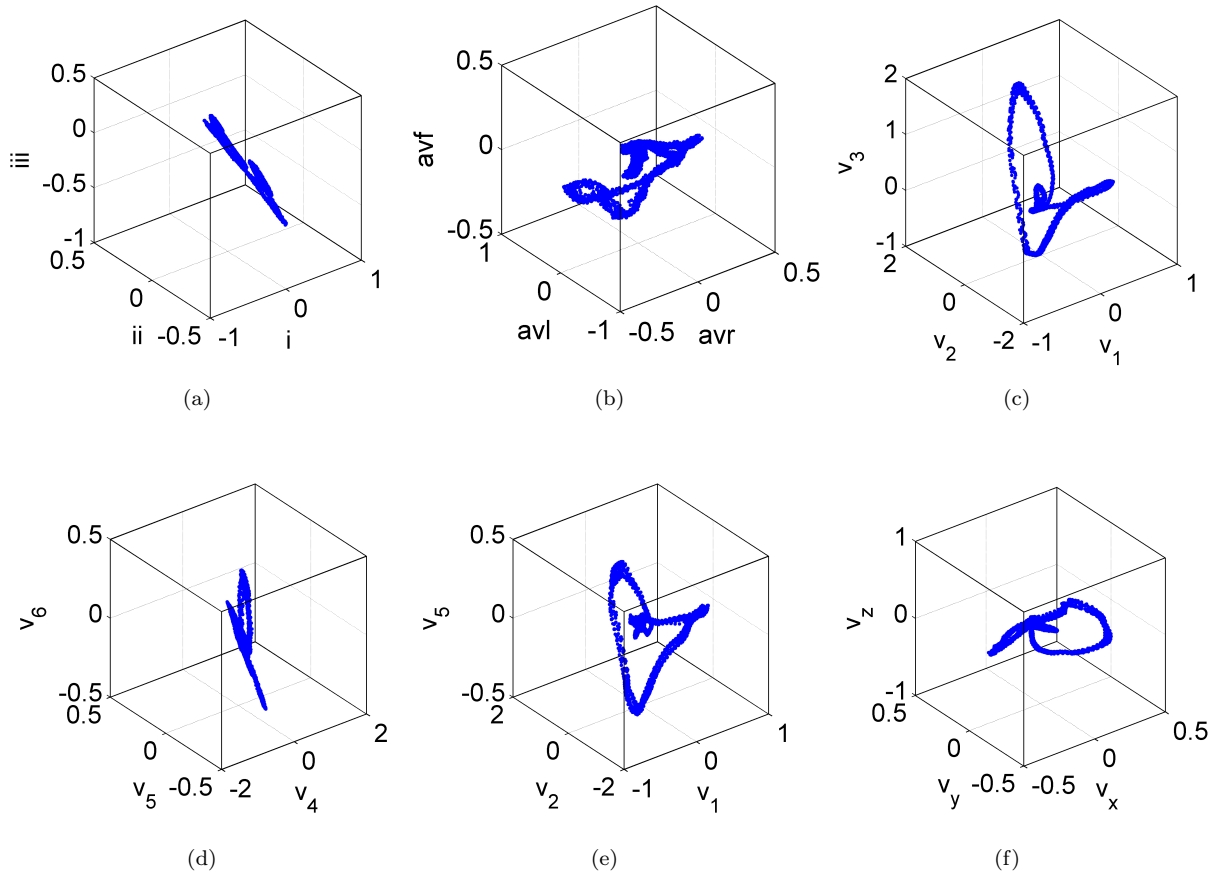


Figure 6.4: Three-dimensional scatter plot of the signals of Fig. 6.1.

dimensions, i.e., having removed ten dimensions, will lead to a MSE of -30dB for the QRS-waves; while at least nine dimensions are required for the same MSE over the P-waves.

Moreover, eight to nine of the independent components extracted by ICA have considerable contribution in the ECG energy, which approves our previous discussions in Section 6.2, concerning the ECG dimensionality. Note that the sensitivity of the different channels to dimension reduction depends on the electrode configuration, the subject, and other signals and noises in the environment. We will discuss about the interpretation of the extracted components in the following.

6.3.2 Interpretation of Independent Components Extracted from Cardiac Recordings

Following our previous discussions, we know that multichannel ECG recordings form a sparse quasi-periodic loop in a curved manifold. By applying linear ICA to such recordings, we are trying to find the directions conveying the most ‘information’ in the N -dimensional space. Intuitively, these directions should correspond to the regions of space in which the data samples are more concentrated. In Fig. 6.8, the VCG representation of the Frank lead electrodes of another sample dataset adopted from the Physikalisch-Technische Bundesanstalt diagnostic ECG database (PTBDB) [158, 22, 112], can be seen. The original data consisted of the twelve standard ECG leads together with the three Frank lead VCG leads, from which we have only plotted the VCG loop for illustration. The baseline wander of this data was removed as preprocessing. In this figure, we have also depicted the column vectors of the mixing matrix estimated by JADE. A segment of the independent components corresponding to these vectors are also seen in Fig. 6.9. As we can see in Fig. 6.8, due to the sparsity of the VCG loop the estimated vectors,




Figure 6.5: The correlation dimensions of (a) a typical ECG, and (b) white Gaussian noise




Figure 6.6: Sensitivity of different ECG components to dimension reduction by applying PCA to eleven-dimensional data

regardless of their sign⁶, give the orientations along which the sample points are more concentrated, i.e., the major VCG loops and planes. This also explains the special shapes and the time-lags that we see in the components in Fig. 6.9. In fact, the different independent components extracted from multichannel ECG, can be considered as projections of the VCG loop (in an N -dimensional hyperspace) onto the directions of ‘significant’ importance. In case of robustness, these components can be used to extract clinical features for normal and pathological cases.

Extraction of Clinical Measures

In previous works, different clinical indexes have been extract from the VCG loop [144, 1]. Following our discussions on the interpretation of the components extracted from multichannel ECG recordings, it is possible to use the information extracted by PCA or ICA as additional clinical measures. In this context, some primary ideas were originally proposed in [43, 146], using singular values and eigenvalues. In the following, we refer to some other possible indexes:

1. The eigenvalues and eigenvectors extracted by PCA convey information about the spread of the data distribution and may be used to find canonical representations of cardiac signals. Specifically, the product of the M largest eigenvalues is a measure of the volume occupied by the VCG loop:

$$c_1 = \prod_{i=1}^M \lambda_i,$$

⁶The ambiguity of sign and variance is an intrinsic limitation of ICA.

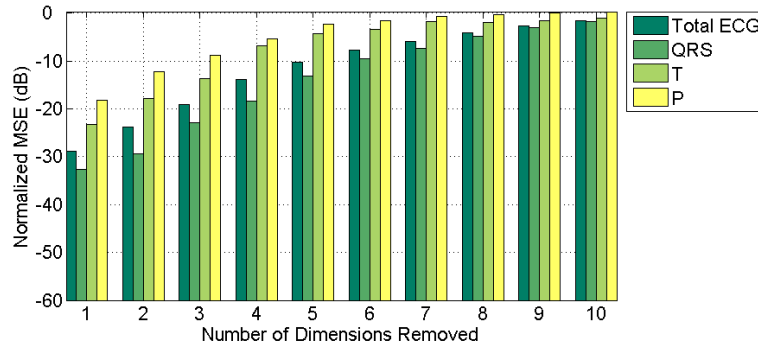


Figure 6.7: Sensitivity of different ECG components to dimension reduction by applying JADE to eleven-dimensional data

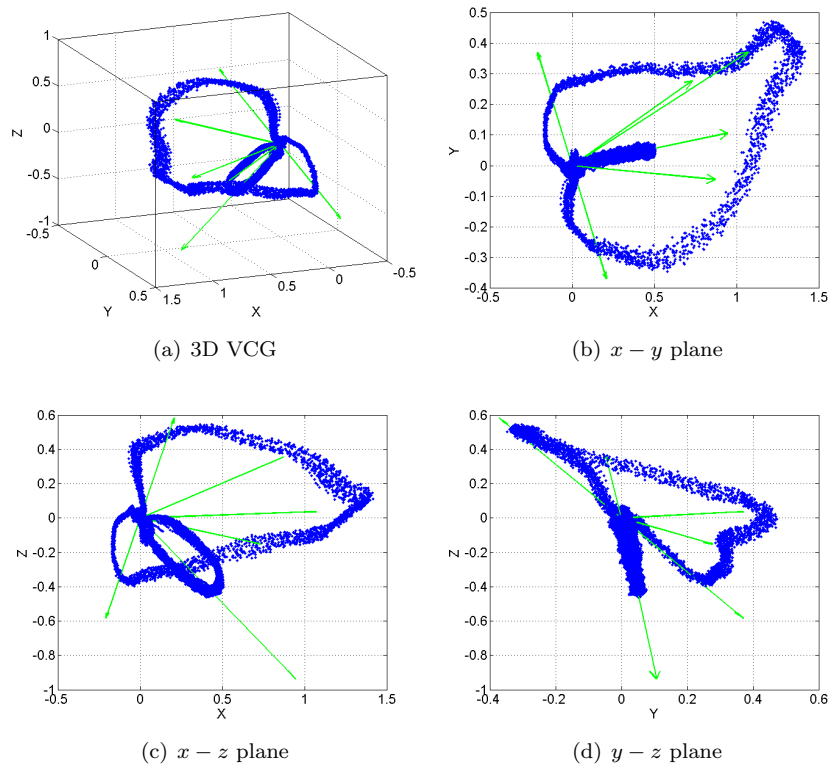


Figure 6.8: Different views of the VCG representation of a typical ECG signal recorded from the Frank lead electrodes, together with column vectors of the mixing matrix estimated by JADE

A similar measure can be derived from the energy of the independent components with the highest explained variance. In this case, we should give the norm of the columns of the mixing vectors estimated by ICA, to their corresponding sources to have unit norm mixing vectors.

2. The angle between the mixing vectors (\mathbf{v}_i) corresponding to the independent components with the highest explained variance:

$$\phi_{ij} = \cos^{-1} \left(\frac{\mathbf{v}_i^T \mathbf{v}_j}{\|\mathbf{v}_i\| \|\mathbf{v}_j\|} \right)$$

3. As mentioned before the VCG loop can lie in a rather complex manifold. However, the visual inspection of each of the P, QRS, and T-loops shows that each of these loops can be rather flat. Interesting measures can be found by detecting the onset and endings of each segment of the ECG




Figure 6.9: A segment of independent components extracted from a typical ECG dataset.

and applying PCA and ICA to the samples of each segment (i.e. *local PCA* and *local ICA*). The upper-mentioned indexes can also be extracted from each segment.

4. The *principal angles* (PA) [73], between the P, QRS, and T-loops also conveys useful information. For this, we apply local PCA or local ICA to find the subspaces corresponding to each segment of the ECG segments. Next, we calculate the PA between each subspace, denoted as follows:

$$\theta_{ij} = \text{PA}(A_i, A_j)$$

where A_i and A_j are the estimated mixing matrices associated with any of the P, QRS, or T-segments. The procedure of calculating the PA is explained in Appendix A.

6.3.3 Necessity of Preprocessing

To this end, we have considered rather clean ECG signals without any non-cardiac interference and noise. However, in presence of noise (especially full-rank noise), the extracted components will not necessarily have the same interpretation. Therefore, in order to use the principal and independent components extracted from multichannel ECG as reliable and robust clinical measures, the data should be preprocessed to remove the non-cardiac interference before applying PCA or ICA.

Baseline Wander Removal

For illustration, JADE is applied to the fifteen-channel ECG signals used in the previous subsection; but this time without prior baseline wander removal. In Fig. 6.10, eight ‘independent’ components extracted from these raw data can be seen⁷. The VCG representation of these signals and the corresponding columns of the mixing matrix estimated by JADE can be seen in Fig. 6.11. In this figure, we can see that the scatter plot of the VCG loop is slightly drifted. Therefore, the baseline wander appears in most of the extracted cardiac components of Fig. 6.10. This is due to the fact that baseline wander is a full-rank signal and it has not been extracted as a single independent component. As a result, the estimated mixing vectors in Fig. 6.11 are tilted toward the direction of the baseline wander.

As we can see, baseline wander removal is an essential preprocessing step that should be done before applying ICA, in order to have meaningful components that purely correspond to the cardiac signals. Additional to baseline wander removal, VCG loop alignment methods such as the one proposed in [188], can also be very effective for removing the minor inter-beat variations of the VCG.

⁷We know that due to the temporal structure of the ECG there exists some dependencies between the components extracted by ICA. Therefore, by ‘independence’ we mean ‘as independent as possible’; which is achieved through the maximization of some measure of independence.




Figure 6.10: A segment of independent components extracted from a typical ECG signal containing baseline wander. The baseline wander can be seen in most of the extracted components.

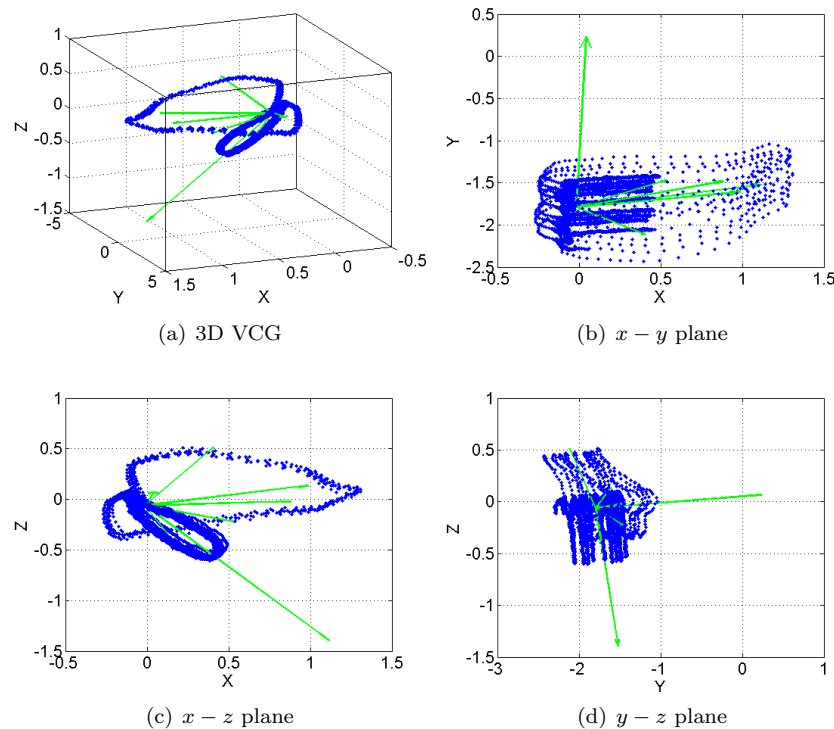


Figure 6.11: Different views of the VCG representation of a typical ECG signal having baseline wander.

Distribution Assymetry

An implicit assumption of common ICA algorithms is that the input signals are symmetric around their mean values; since the mean of the data is initially removed and the sign of the ICs are not determinable [89]. However, ECG data are not symmetric around their mean values, and the *isoelectric point* of cardiac signals slightly differs from the mean value of the data. This means that when applying conventional ICA algorithms to signals such as the ECG (without considering the data asymmetry), the extracted ICs do not exactly correspond to the main planes in the VCG representation.

As an illustration, the scatter plot of a synthetic asymmetric dataset is depicted in Fig. 6.12(a). The columns of the mixing matrix, extracted by JADE have also been depicted. As seen in this figure, due




Figure 6.12: The effect of PDF asymmetry on the extracted ICs for an (a) asymmetric PDF, and the (b) PDF made symmetric using equation (6.5). The vectors represent the mixing vectors estimated by JADE (in positive and negative directions), centered at the mean of the distribution.

to the data asymmetry, ICA does not give the directions of the two lobes of the distribution. Similarly, for the problem of ECG signals we might be more interested in the directions in which the data has the most variations with respect to the isoelectric point, rather than the actual ICs, which are derived with respect to the mean of the data. A simple method that we propose for solving this problem, is to remove the baseline wander of the ECG, which is identical to making zero the isoelectric line of the ECG (i.e. making zero the isoelectric point of the multidimensional VCG), and then to make the dataset virtually zero-mean with respect to the isoelectric point, by augmenting the data samples with their minus values:

$$\mathbf{x}_{n \times 2T} \leftarrow [\mathbf{x}_{n \times T}, -\mathbf{x}_{n \times T}] \quad (6.5)$$

By this way, we are adding the image of the dataset to the distribution space and therefore making the density function of the samples symmetric. This makes the dataset twice as long; but has the benefit that the ICs are found with respect to the isoelectric point of the VCG⁸. The result of applying this idea to the data of Fig. 6.12(a) and the new columns of the mixing matrix extracted by the JADE algorithm are shown in Fig. 6.12(b). This simple idea may also be useful in other applications for exploiting *local* behaviors of complex multivariate distributions around any point other than the actual data mean value.

The effect of data asymmetry becomes much more important for the combination of the fetal and maternal ECGs, where the fetal signals can be an order of magnitude weaker than the maternal signals. Therefore, considering the small fetal VCG loop size, the mean value of the data can be very far from the fetal VCG loop and by applying conventional ICA algorithms that find the components with respect to the mean of the data, the details of the fetal VCG loop are not extracted as independent components. This can explain why in practice only a few fetal components are extracted by applying ICA to maternal abdominal recordings while traces of the fetal signal still exist in the other components. In this case, maternal ECG removal before applying ICA— or before any other linear transform— can be very effective. We will come back to this issue in later chapters.

6.4 Curse of Dimensionality and Over-fitting

A well-known problem in multichannel data analysis is a problem known as the *curse of dimensionality* (CoD), initially coined by R.E. Bellman [16]. CoD describes the problems caused by the exponential increase in volume due to adding extra dimensions to an N -dimensional space. For example, 100 evenly-spaced sample points suffice to sample a unit interval with no more than 0.01 of distance between points. An equivalent sampling of a 10-dimensional unit hypercube with a lattice having a spacing of 0.01 between adjacent points would require 10^{20} sample points. Thus, in some sense, the 10-dimensional hypercube can be said to be a factor of 10^{18} *larger* than the unit interval (example adopted from [222]).

In practical statistical signal processing problems, we work with finite number of samples taken from some random process. In this case, the CoD implies that with an increase of data dimensions, the sampled

⁸Evidently the second-half of the extracted sources will also be redundant in this case.




Figure 6.13: Typical signals extracted from 72 maternal abdominal channels using FastICA. The last three components are typical ‘bumpy’ signals without any physical origin and due to the ‘spikes and bumps’ phenomenon [90, 179, 178].

hyperspace becomes more and more sparse, unless if the number of samples are exponentially increased. In such problems, the CoD can manifest itself in different ways. The most common in ICA applications is that the estimated statistical measures become less relevant and more sensitive to outliers⁹.

As a case study, in some of our early studies we were working on 72-channel signals recorded from an array of maternal abdominal sensors. In order to extract the fetal ECG components, we directly applied FastICA to the 72-channel recordings. Some of the typical signals that we found are depicted in Fig. 6.13. In this study, there were usually about five components that corresponded to the maternal ECG, one to the fetal ECG, few components that corresponded to other known sources, such as the maternal respiration and noise, and about fifty ‘bumpy’ components, such as the ones in the last three panels of Fig. 6.13, for which we could not find any physiological explanations. It was later found that these components were merely due to the high number of dimensions (say above 20 channels), when applying ICA to spectrally colored noisy signals. This phenomenon, known as *spikes and bumps*, was originally introduced in [90, 179, 178]. To approve the fakeness of the bumpy components, similar results can be reproduced by generating 1000 to 5000 samples of Brown noise in 20 to 50 channels. Then, by applying FastICA or JADE to these synthetic signals results in bumpy signals as shown in Fig. 6.14, which despite their misleading appearance do not have any physical origin. In Appendix B, we will draft some ideas for theoretical justification of this phenomenon. The problem is however open to further research in future studies.

To overcome the problem of bumps and spikes or more generally the CoD, the exponential increase

⁹Note that in parametric estimation methods in which a model is assumed as the underlying data model, it is the number of unknown parameters that should be considered as the dimensions. For instance in conventional ICA methods based on the $\mathbf{x} = \mathbf{A}\mathbf{s}$ model, the number of unknown parameters grows with the square of the number of observation channels.




Figure 6.14: Typical bumpy signals extracted from 40 channels of brown noise using FastICA.

of the sample points is not possible, as it would either require very high sampling rates or an increase in the signal registration time. On the other hand, almost all biosignals are nonstationary in time and we usually prefer to process short stationary blocks of data. We can add to this, the higher amount of processing load and memory required for high-dimensional signal processing.

For high-dimensional observations, more practical methods include (1) projection onto lower dimensions, (2) channel selection, (3) processing groups of channels, and (4) nonlinear mapping; which all lead to the processing of a lower-dimensional subspace. In the following, we introduce these methods and their limitations. A survey of different dimension reduction techniques may be found in [61, 119].

1. *Projection onto Lower-Dimensions:* Projection onto lower-dimensions subspaces using PCA is perhaps the most popular means of dimension reduction. It simply consists of eigenvalue decomposition and eliminating the dimensions corresponding to the minor eigenvalues. Despite of its effectiveness, this procedure has also some limitations, especially for noisy data. It is difficult to automatically decide on the number of dimensions to keep. For this, we usually seek for some large gaps between successive eigenvalues or for some sharp changes of slope in the logarithmic plot of the eigenvalues. However, using PCA we can not assure to have removed all the redundant dimensions or guarantee to preserve the important ones.
2. *Channel Selection:* A more effective procedure, which at the same time requires *a priori* information of the signals of interest, is to select the more useful channels for processing. In [177, 213], some electrode selection strategies based on *mutual information* criteria were proposed for the fetal ECG extraction problem.

In practice, it is more effective to have a combination of channel selection and projection. It is also better to eliminate the very noisy channels or the ones that do not convey useful information for the desired signals, before applying a PCA dimension reduction.

3. *Processing Groups of Channels:* Another approach is to process small groups of the channels separately and to merge the information extracted by such processing in later steps of the analysis. For instance, MCG signals are very commonly recorded from several hundreds of sensors. However, the dominant components of the existing bio-sources, such as the maternal and fetal cardiac signals, do not usually exceed about fifteen to twenty dimensions. So it is rather safe to split the sensor array into blocks of, say, twenty electrodes and to extract the fetal components from these smaller subsets. The information found from each group of electrodes can be later merged together for constructing the fetal MCG subspace, or even at higher levels, such as feature extraction.
4. *Nonlinear Mapping:* As mentioned before the desired signal subspace can lie in complex nonlinear manifolds; although each of the signals can be rather simple and structured in the time-domain. We have already shown examples of this phenomenon for the ECG scatter plots in Fig. 6.4. In such cases, an alternative dimension reduction approach is to use nonlinear mappings to ‘unfold’ the complex manifold. This however requires *a priori* information about the manifold structure. Examples of this procedure can be found in [47, 119].

6.5 Optimal Sensor Number and Positioning for Fetal ECG Extraction

A fundamental question in designing fetal ECG or MCG recording protocols is how many channels to use and where to place the sensors on the maternal abdomen. From the discussions of this chapter, we now know that although the cardiac signal subspace can be theoretically infinite dimensional (due to the distributedness of the cardiac source); but in practice we can only see between 4 to 8 dominant components on the body surface of adults. For the fetal signals the number of components is between 1 to 3 and the other components are dominated by the maternal signals and noise. There are also other biosignal sources such as the maternal respiration and uterus contractions signals, which can appear as additional sources. Therefore, empirically between 8 to 16 channels, well-distributed over the maternal abdomen and close to the fetal heart are rather sufficient for extracting the fetal components. Note that although the fetal R-peaks may also be detected using fewer number of electrodes; but the morphological details of these signals require more electrodes.

The nonstationarity of the fetal cardiac signals and the possible movements of the fetus should also be considered in the electrode positioning. While some researchers perform ultrasound measurements for detecting the fetal position and choosing the appropriate electrode configuration prior to signal recording, it is practically more convenient to have a fixed electrode configuration, with even more electrodes than required, and to dynamically re-select the most important ones as the fetus or fetuses move.

Concerning the optimal electrode configuration for retrieving the full fetal ECG morphology, it is better to have the electrodes rather distributed over the maternal abdomen and not to localize them all close to one another. This helps us to have a full view of the fetal heart from different sides and prevents the ECG or MCG signals from becoming degenerate mixtures of the cardiac signals.

6.6 Summary and Conclusions

In this chapter the components extracted by linear transforms of multichannel ECG recordings, such as ICA, were studied from an electrophysiological and signal processing viewpoint. Considering the distributedness of the heart, we found some electrophysiological interpretation for the extracted components and explained the multiple components that are typically extractable from multichannel ECG. We also studied the close relationship between such components and the vectorcardiogram representation of the ECG and proposed some quantitative measures that can be derived from such signals and be used as clinical measures. Issues concerning the dimensionality of these signals were also studied.

The main focus of the chapter and the presented case studies were on adult ECG signals recorded from the thorax. Nevertheless, the results and interpretations are more or less applicable to the fetal ECG, except that the fetal components are much weaker in amplitude and some times become electromagnetically shielded by low-conductivity layers such as the *vernix caseosa* (cf. Section 3.3).

Although the results of applying ICA to multichannel ECG recordings are rather subjective and database dependent, the hereby presented results are believed to be rather insightful from both the signal processing and physiological viewpoints. In later chapters we will use these insights to develop novel algorithms for adult and fetal ECG analysis.

In the light of the hereby presented results, there are many other issues which need to be studied in future works, including:

- A general framework for the study of distributed sources and the theoretical relationship between the statistical properties of a distributed source and the principal or independent components extracted from body surface electrodes.
- The study of ECG manifolds, their dimensionality, complexity, and other geometric and statistical properties of these manifolds. In this context, the Bayesian filtering framework presented in Chapter 5 might also be useful for nonlinearly unfolding the ECG manifolds.
- The impact of cardiac defects on the ECG dimensionality and the clinical measures proposed in Section 6.3.2, especially for fetal ECG studies.

Chapter 7

Multichannel Electrocardiogram Decomposition using Periodic Component Analysis

7.1 Introduction

Independent Subspace Analysis (ISA) has been introduced as an alternative to Independent Component Analysis (ICA), for problems in which we deal with groups of signals having inter-group independence and intra-group dependencies. ISA was first introduced in [41] and mathematically developed in [28], where the notion of ICA was generalized to the notion of multidimensional ICA. As noted there, ISA relies on the idea of *vector-valued* components rather than *scalar* source signals. The first, and yet most commonly studied, application of ISA is the extraction of fetal ECG from a set of maternal abdominal mixtures. From the discussions of Chapter 6, the cardiac signals of either the mother or the fetus are generally multidimensional signals. Therefore, by applying *any* linear transform to the observation signals we can extract several maternal and fetal components. From the signal processing point of view, this means that the maternal and fetal ECG components, individually, form a signal subspace with internal dependencies, rather than being single dimensional components; while the components of the maternal and fetal subspaces are independent from each other.

ISA may be realized by applying an initial ICA step on multichannel observations and then empirically regrouping the independent components that belong to the same subspace from prior knowledge of the subspace structures, and to achieve a canonical representation of the desired subspace. In fact, there is an intrinsic ambiguity in retrieving the components inside the subspaces, which may not be resolved with the same measure of independence used for extracting the subspaces. Therefore, from the ISA viewpoint, no representation of components inside a signal subspace can be considered to be better than another. This is the reason for which the components that belong to the same subspace are regrouped after the initial ICA step.

The currently challenging issues in the ISA context are (1) how to find the dimensions of each subspace [28], (2) how to automatically regroup the components [219, 10, 191], and (3) the impact of subspace distances and noise on the stability of the extracted subspaces [135, 82]. For the fetal ECG extraction problem, previous studies have focused on the feasibility of extracting the independent subspaces [28, 45] and regrouping strategies [10].

Although the ISA framework has answered many questions concerning vectorial components, it has some limitations for being used for ECG decomposition. Following the explanations in Chapter 6, we know that cardiac signals (or any other distributed source) do not generally have a limited number of dimensions and any low-rank subspace of the body surface recorded signals is only an approximation of the true cardiac subspace. Therefore, for ECG signals, rather than seeking for subspaces of fixed finite dimensions, it is more reasonable to somehow rank the subspace components according to their contribution in the surface ECG and to threshold the least relevant dimensions.

Another limitation of ISA (and ICA) is that they are usually based on the maximization of some measure of component independence. However, for pseudo-periodic signals such as the ECG, the temporal

$$\mathbf{x} \xrightarrow[\text{(rotation)}]{\text{PCA}} \mathbf{v} \xrightarrow[\text{(scaling)}]{\text{Sphering}} \mathbf{z} \xrightarrow[\text{(rotation)}]{\text{ICA}} \mathbf{y}$$

Figure 7.1: General scheme of ICA algorithms with spatial whitening

structure of the signal is rich in information. For such signals, a measure of periodicity of the extracted signals, both clinically and mathematically, is a more appropriate criterion as compared to independence. In some recent works, the temporal periodicity of ECG signals has been exploited in source separation algorithms [94]. These methods are however of the same order of complexity as conventional ICA algorithms, and are basically designed for signals having a constant fundamental period; an assumption that is not true for real ECG recordings.

In this chapter, we present a linear source separation method that is specifically customized for pseudo-periodic signals, such as the ECG. The proposed method is partially based on the notion of *periodic component analysis* (π CA) introduced in [180] and *generalized eigenvalue decomposition* [154]. The relationship between this method and well-known source separation techniques such as AMUSE [201], JADE [31], and SOBI [17] are also discussed. As it is shown, for ECG signals, the method has several benefits over conventional ICA and is applicable to both adult and fetal ECG embedded in noise.

7.2 Background

7.2.1 Generalized Eigenvalue Decomposition

For symmetric matrices $A, B \in \mathbb{R}^{N \times N}$, the problem of generalized eigenvalue decomposition (GEVD) [192], of the matrix pair (A, B) , consists of finding the matrices W and D , such that:

$$W^T A W = D \quad , \quad W^T B W = I \quad (7.1)$$

where D is the diagonal generalized eigenvalue matrix corresponding to the eigenmatrix $W = [\mathbf{w}_1, \dots, \mathbf{w}_N]$, with real eigenvalues sorted in ascending order on its diagonal¹. As it is seen from (7.1), W is a transformation that simultaneously diagonalizes A and B . Moreover, the first eigenvector \mathbf{w}_1 , corresponding to the largest generalized eigenvalue, also maximizes the following ratio, known as the *Rayleigh quotient* [192]:

$$J(\mathbf{w}) = \frac{\mathbf{w}^T A \mathbf{w}}{\mathbf{w}^T B \mathbf{w}} \quad (7.2)$$

This property is used in the derivation of the methods in Chapter 8.

7.2.2 ICA versus Generalized Eigenvalue Decomposition

In the context of ICA, in practice, we have a finite number of samples of an N -dimensional observation vector $\mathbf{x}(t) = [x_1(t), \dots, x_N(t)]^T$, and seek for a linear mixture of these observations that maximizes some measure of independence, or namely a *contrast function*. Under some general assumptions the estimated components are solutions of a BSS problem with a linear latent variable model. Moreover, most ICA algorithms perform spatial whitening on the dataset, which as shown in Fig. (7.1) only leaves a rotation matrix to be estimated by maximizing the contrast function of ICA.

An algebraic approach to ICA is to diagonalize a set of matrices containing second or higher-order statistics of the dataset [31, 17]. For signals with temporal structure, there are various algorithms that use this algebraic approach. For example, for a wide-sense stationary or cyclostationary real observation vector $\mathbf{x}(t)$, if we define the covariance matrix as:

$$C_x(\tau) = E_t\{\mathbf{x}(t + \tau)\mathbf{x}(t)^T\} \quad (7.3)$$

where $E_t\{\cdot\}$ indicates averaging over t , the AMUSE algorithm provides a linear transform that jointly whitens the data and diagonalizes $C_x(\tau)$ for some arbitrary τ , i.e., the solution of the GEVD problem

¹ In the problem of interest, A and B are symmetric positive definite matrices, therefore the eigenvalues of D are real and positive [192].

for the matrix pair $(C_x(\tau), C_x(0))$ [201, 154]. However, it is known that no more than two matrices may be *exactly* diagonalized by a quadratic transform of the form (7.1), except when they belong to the same eigenspace [17]. Due to this fact, methods have been proposed which approximately diagonalize a set of desired matrices, and are more robust to data outliers and computational errors as compared with AMUSE [29]. In this context, the SOBI algorithm is an example of a time-domain algorithm that whitens the data and approximately diagonalizes $C_x(\tau)$ for several time-lags τ [17]. In [156], joint approximate diagonalization was further used for the separation of nonstationary and (possibly) Gaussian sources. Similar time-domain methods have also been proposed for cyclostationary sources, in which the data is again pre-whitened and matrices representing cyclostationary statistics of the dataset are (approximately) diagonalized [59]. However, the final output of any of these approximate diagonalization methods is a linear transformation of the form:

$$\begin{aligned} \mathbf{y}(t) &= W^T \mathbf{x}(t) \\ W &= Q\Lambda^{-1/2}R \end{aligned} \quad (7.4)$$

where Q and Λ are respectively the eigenmatrix and eigenvalue matrix of $C_x(0)$, such that $C_x(0) = Q\Lambda Q^T$, and R is the rotation matrix found by maximizing the contrast function of the corresponding ICA method. In (7.4), the terms Q , $\Lambda^{-1/2}$, and R respectively correspond with the PCA, sphering, and ICA steps indicated in Fig. 7.1. It is easy to show that W , diagonalizes the following symmetric matrix:

$$\Gamma = Q\Lambda^{1/2}RDR^T\Lambda^{1/2}Q^T \quad (7.5)$$

where D is an arbitrary diagonal matrix ². It therefore follows that the matrix pair $(\Gamma, C_x(0))$ satisfies (7.1).

Equation (7.5) implies that for pre-whitened data and for a rotation matrix R found by maximizing *any* ICA contrast function, there exists a set of matrices Γ that are exactly diagonalizable via GEVD. The matrix Γ , by itself is not necessarily any of the second or higher-order statistics matrices of the data; but contains some overall statistics of the data that is *exactly* diagonalized through the linear transform W . From this point of view, any ICA algorithm of the form of Fig. 7.1 may be transformed into a GEVD problem for diagonalizing a single matrix containing some cross statistical measure of the multi-channel dataset. This suggests that in some applications, instead of forming a set of matrices containing conventional statistics of the data, such as the second or higher-order statistics, and approximately diagonalizing these matrices, we can directly use the *a priori* information about the desired signals (e.g. their periodicity) to form an *ad hoc* symmetric matrix Γ , and seek for linear transforms that diagonalize this matrix, i.e., decorrelate the *ad hoc* statistics. The matrix Γ may for example be a *covariance-like* matrix as defined in (7.3); but using a time-varying time lag τ that is derived from our prior knowledge of the dataset. This idea is investigated for ECG signals in the following sections.

7.2.3 Periodic Component Analysis

Here we restate the problem of π CA, adapted from [180], which is merely a restatement of the AMUSE algorithm derived from a measure of periodicity.

Given an N -dimensional observation vector $\mathbf{x}(t) = [x_1(t), \dots, x_N(t)]^T$, we seek for the linear mixture $s(t) = \mathbf{w}^T \mathbf{x}(t)$ with a maximal periodic structure that minimizes the following *measure of periodicity*:

$$\epsilon(\mathbf{w}, \tau) = \frac{\sum_t |s(t+\tau) - s(t)|^2}{\sum_t |s(t)|^2} \quad (7.6)$$

where $\mathbf{w} = [w_1, \dots, w_N]^T$ and τ is the period of interest. It can be shown that (7.6) may be rearranged as follows:

$$\epsilon(\mathbf{w}, \tau) = \frac{\mathbf{w}^T A_x(\tau) \mathbf{w}}{\mathbf{w}^T C_x(0) \mathbf{w}} = 2 \left[1 - \frac{\mathbf{w}^T C_x(\tau) \mathbf{w}}{\mathbf{w}^T C_x(0) \mathbf{w}} \right] \quad (7.7)$$

where

$$\begin{aligned} A_x(\tau) &= E_t \{ [\mathbf{x}(t+\tau) - \mathbf{x}(t)][\mathbf{x}(t+\tau) - \mathbf{x}(t)]^T \} \\ &= 2C_x(0) - 2C_x(\tau). \end{aligned} \quad (7.8)$$

In the second part of (7.8) we have used the symmetry of the matrix $C_x(\tau)$, and assumed that the time averages asymptotically converge to statistical averaging as $t \rightarrow \infty$.

²In fact, in (7.5), the term RDR^T represents an arbitrary matrix from the eigenspace of R and $\Gamma = W^{-T}DW^{-1}$.

Then using the Rayleigh-Ritz theorem from linear algebra [192], it follows that the weight vector \mathbf{w} minimizing (7.7) is given by the eigenvector corresponding to the smallest generalized eigenvalue of the matrix pair $(A_x(\tau), C_x(0))$, or equivalently the largest generalized eigenvalue of $(C_x(\tau), C_x(0))$. As with (7.1), by assuming D as the diagonal generalized eigenvalue matrix corresponding to the eigenmatrix W , with real eigenvalues sorted in descending order on its diagonal, the transformation $W^T \mathbf{x}(t)$ gives the *most periodic* components in descending order of τ -periodicity.

7.3 Modifications for Time-varying Periods

ECG signals have a pseudo-periodic structure that is repeated in every cycle of the ECG beat. However, normal ECGs can have RR-interval deviations of up to 20% (cf. [160]), which means that a constant period τ as defined in the previous section, does not fully describe the periodicity of the ECG.

For such signals, we propose to use a time-varying period that is updated on a beat-to-beat basis. For this, as shown in Fig. 7.2, by detecting the R-peaks of the ECG, a linear phase $\phi(t)$ ranging from $-\pi$ to π is assigned to each ECG sample, with the R-peak being fixed at $\phi(t) = 0$. The linear phase $\phi(t)$, provides a means for phase-wrapping the RR-interval onto the $[-\pi, \pi]$ interval. Therefore, the ECG – regardless of its RR-interval deviations – may be converted to a polar representation in which the ECG components in different beats, such as the P, Q, R, S, and T-waves, are more or less phase-aligned with each other, especially over the QRS segment (Fig. 7.3).

On the other hand, minimizing (7.7), requires the calculation of the cross-correlation between the samples having a time-lag τ in different channels. So in order to apply π CA to ECG signals, we can replace the constant time-lag τ in (7.8), with a variable τ_t that is calculated from $\phi(t)$ from beat to beat. Therefore, in each ECG cycle the sample at the time instant t is compared with the sample $t + \tau_t$, namely its *dual sample*, which is the sample with the same phase value in the successive ECG beat. The time-varying period τ_t may be mathematically defined as follows:

$$\tau_t = \min\{\tau | \phi(t + \tau) = \phi(t), \tau > 0\} \quad (7.9)$$




Figure 7.2: Illustration of the phase assignment procedure used for calculating τ_t in each ECG beat.




Figure 7.3: Polar representation of a noisy ECG using the ECG phase $\phi(t)$.

Using this definition, the covariance matrix defined in (7.3) is redefined as follows:

$$\tilde{C}_x = E_t\{\mathbf{x}(t + \tau_t)\mathbf{x}(t)^T\} \quad (7.10)$$

Moreover, in order to assure the symmetry of \tilde{C}_x and the realness of its eigenvalues, the following step is required in practice:

$$\tilde{C}_x \leftarrow (\tilde{C}_x + \tilde{C}_x^T)/2 \quad (7.11)$$

In fact, if we consider the polar representations of each of the channels of $\mathbf{x}(t)$ (as in Fig. 7.3), the matrix \tilde{C}_x represents the covariance of the polar representation around its mean value. Following the discussions in Section 7.2.2, \tilde{C}_x is the matrix containing the *ad hoc* statistics of the ECG, i.e., a measure of ECG periodicity extracted from the ECG R-peak information.

Next, we find W , the GEVD solution of the $(\tilde{C}_x, C_x(0))$ pair with the eigenvectors ranked in descending order of their corresponding generalized eigenvalues. The desired signal vector $\mathbf{y}(t) = [y_1(t), \dots, y_N(t)]^T$ is then found from (7.4). The components of $\mathbf{y}(t)$ are sorted according to the amount of their periodicity, relative to the heart beat. In other words, $y_1(t)$ is the most periodic component and $y_N(t)$ is the least periodic, with respect to the R-peaks of the ECG. Hereinafter, we refer to the algorithm presented in this section, which is based on a time-varying lag, as π CA.

The proposed method is rather flexible, and may be extended to other *ad hoc* statistics extracted from ECG recordings. For instance, for the problem of fetal ECG extraction, if we define $\phi_m(t)$ and $\phi_f(t)$ as the maternal and fetal ECG phases found from the maternal and fetal R-peaks, \tilde{C}_x^m and \tilde{C}_x^f representing the covariance matrices of the mother and fetus, are found by averaging (7.10) over the maternal and fetal periods, respectively. Then the matrix \tilde{C}_x used in the GEVD may be set to any of the following matrices:

$$\begin{aligned} \tilde{C}_x &= \tilde{C}_x^m \\ \tilde{C}_x &= \tilde{C}_x^f \\ \tilde{C}_x &= \tilde{C}_x^m - \tilde{C}_x^f \end{aligned} \quad (7.12)$$

If we assume the data to be pre-whitened, the diagonalization of the matrices defined in (7.12) is respectively equivalent to finding (a) the most periodic components with respect to the maternal ECG, (b) the most periodic components with respect to the fetal ECG, and (c) the most periodic components with respect to the maternal ECG while being the least periodic components with respect to the fetal ECG. In this latter case the extracted components should gradually change from the maternal ECG to the fetal ECG, from the first to the last component³. It should of course be noted that the last two cases are difficult to implement in practice, as they require the prior extraction of the fetal R-peaks to form the \tilde{C}_x^f matrix.

7.4 Results

As our first illustration, we apply π CA on the ECG signals used in Fig. 6.1 (page 58). The R-peak detection, time-varying lag calculation, and covariance matrix calculations were carried out according to the procedures described in the previous section. The resultant periodic components are depicted in Fig. 7.4. We can see in this figure, how the components are ranked according to their ECG contribution. These components may be compared with the principal and independent components in Fig. 6.2 and 6.3. The sensitivity study of the P, QRS, and the T-segments of the ECG, introduced in Section 6.3.1, was also carried out on the extracted periodic components. The MSE results, similar to Figs. 6.6 and 6.7 may be seen in Fig. 7.5.

The DaISy fetal ECG database is used as our second illustration [46]. As explained in previous chapters, the database consists of five abdominal and three thoracic channels recorded from the abdomen and chest of a pregnant woman with a sampling rate of 250Hz. The eight channels of the dataset may be seen in Fig. 7.6. The independent subspace decomposition approach proposed in [28], using the JADE algorithm [31], is used as the benchmark ICA method. The eight independent components extracted by this algorithm are depicted in Fig. 7.7.

By performing R-wave detection on one of the maternal thoracic channels, the maternal ECG phase $\phi_m(t)$ was calculated according to the explanations of the previous section. Next, the time-varying

³Note that in the last definition of (7.12), \tilde{C}_x is not necessarily positive definite. However, due to its symmetry the generalized eigenvalues are real and may be ranked in ascending/descending order.




Figure 7.4: Periodic components extracted from the ECG channels of Fig. 6.1 in descending order of periodicity.

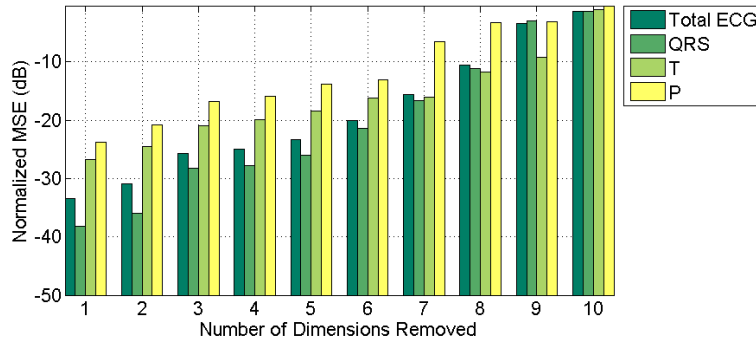


Figure 7.5: Sensitivity of different ECG components to dimension reduction by applying π CA to eleven-dimensional data of Fig. 6.1




Figure 7.6: The DaISy dataset consisting of five maternal abdominal and three thoracic channels [46].

maternal ECG period τ_t^m was calculated, from which the matrix \tilde{C}_x and W , the generalized eigenmatrix of the $(\tilde{C}_x, C_x(0))$ pair, were found and sorted in descending order of the eigenvalues. The resultant




Figure 7.7: Independent components extracted from the dataset of Fig. 7.6, using JADE. Notice that components 1, 2, 3, and 5 correspond to the maternal subspace and components 4 and 8 to the fetal subspace.




Figure 7.8: Periodic components extracted from the dataset with **maternal** ECG beat synchronization. The maternal ECG contribution reduces from top to bottom.

periodic components calculated from (7.4) are depicted in Fig. 7.8. As it is seen in this figure, the first component, corresponding to the largest eigenvalue, has the most resemblance with the maternal ECG, while as the eigenvalues decrease, the signals become less similar to the maternal ECG. Interestingly, two of the last extracted components (components six and seven) resemble the fetal ECG. This can be explained by considering that π CA is ranking the extracted components according to their resemblance with the maternal ECG period, while the fetal components do not resemble the maternal ECG, when averaged synchronously with respect to the maternal R-peaks. The fetal components are therefore not extracted among the first components.

As explained in Section 7.3, it is also possible to consider the fetal ECG periodicity in the matrix \tilde{C}_x , which requires the fetal R-peaks for extracting the time-varying fetal period τ_t^f . For this, the fetal ECG component extracted by JADE in the fourth channel of Fig. 7.7 is used for fetal R-peak detection and phase calculation. Having calculated the fetal ECG phase $\phi_f(t)$, the previously explained procedures are repeated to extract the periodic components of the fetal ECG. The resultant periodic components are depicted in Fig. 7.9. This time we see that the extracted components are ranked according to their resemblance with the fetal ECG.

The last results correspond to the last type of covariance matrix defined in (7.12). The covariance matrix \tilde{C}_x for this part is calculated from the difference of \tilde{C}_x^m and \tilde{C}_x^f . After performing the previously




Figure 7.9: Periodic components extracted from the dataset with **fetal** ECG beat synchronization. Notice that the fetal ECG contribution reduces from top to bottom.




Figure 7.10: Periodic components extracted from the dataset with **maternal & fetal** ECG beat synchronization. The maternal (fetal) ECG contribution reduces (increases) from top to bottom.

explained GEVD stages, the periodic components are found from (7.4). These components are depicted in Fig. 7.10. As expected before, the first component has the most resemblance with the maternal ECG, while the last component mostly resembles the fetal ECG and the intermediate components are mixtures of maternal and fetal components and noise.

7.5 Multichannel Fetal MCG in Multiple Pregnancy: A Case Study

The methods proposed up to now have been principally designed for ECG signals. Nevertheless, as mentioned before, due to the morphological similarity of the ECG and the MCG, the same methods and tools are applicable to MCG recordings. In what follows, we present a subjective study for the extraction of twin fetal cardiac signals from an array of MCG recordings using a combination of methods, including ICA and π CA. In this case study, we can see how the methods are adapted in practice to extract and track fetal signals from long datasets.

The dataset has been recorded in the Biomagnetic Center of the Department of Neurology, in Friedrich Schiller University, Jena, Germany⁴. It consists of several sets of MCG and other signals, in arrays of

⁴This dataset has been provided by Dr. Dirk Hoyer, from the Biomagnetic Center of the Department of Neurology, in

208 channels recorded over 30 minutes, with a sampling rate of 1025Hz. The description of the recorded channels are summarized in Table 7.1. This data has been recorded by a SQUID Biomagnetometer system. The pregnant women were positioned *supine*, i.e., with a slight twist to either side, to prevent compression of the inferior *vena cava* by the pregnant uterus. The dewar was positioned with its curvature above the fetuses after sonographic localization as close to the maternal abdominal wall without contact as possible [77].

The current results have been achieved on one of the available datasets, namely the q00002252 dataset. By visual inspection of the dataset, some of the channels were noticed to be zero or at the order of the quantization error of the analog-to-digital converter of the recording device. These channels (listed in Table 7.1) were excluded from further analysis to prevent noisy results.

In the following, we explain the procedures that were carried out over the dataset to extract the twin MCGs. Due to the large size of the data the mentioned steps have been applied on segments of 10000 samples (about 10s), and as later explained, the procedures are repeated over all such windows throughout the 30 minute data, each time with 10%-50% of overlap with previous windows, for consistency.

7.5.1 Baseline Wander Removal and Preprocessing

Visual inspection of the original data shows that the channels are highly contaminated with baseline wander (BW). From our previous discussions on ECG signals (Section 6.3.3), we know that the results of source separation methods are very susceptible to BW and it is necessary to remove it before applying any source separation. For this, a two-step moving window median filter with 100ms and 200ms window lengths, was used (cf. Appendix D). The BW removal was followed by a lowpass filter with a cut-off frequency of 150Hz. This filtering improves the SNR of the source separation procedures and reduces their susceptibility to noise. Typical channels achieved after the preprocessing procedure are depicted in Fig. 7.11.

7.5.2 Maternal MCG Removal

The maternal MCG is the most considerable artifact for the fetal MCG. We therefore require an effective method for removing it. For this we use a combination of π CA and the other methods discussed in previous chapters. In order to apply π CA for maternal MCG extraction (and removal), the locations of the maternal R-peaks are required, which may be detected from any ECG or MCG signal that is synchronously recorded with the dataset. For this, we used the maternal reference ECG channel (channel number 196, cf. Table 7.1), with a simple peak-detection method that seeks for the maximum points of the data within a moving window (cf. Appendix D). The maternal R-peaks were detected for the whole 30 minute data and saved for later processing. Using the maternal R-peaks, π CA was applied to the preprocessed data to sort the components according to their resemblance with the maternal MCG. The first ten components extracted by π CA, from a typical segment of the preprocessed data may be seen in Fig. 7.12.

The next step is to remove the maternal MCGs. The simplest (but not necessarily the best) way is to set the first K components found by π CA to zero and to project the remaining components back to the sensor space. This is a common practice that is carried out in conventional subspace cancellation techniques using PCA or ICA [99]. On the other hand, its drawback is that by nulling K components,

Table 7.1: Description of the recorded channels

Channels	Description
1-168	magnetic channels
169-195	magnetic reference channels
196-199	electric channels (mother's ECG)
200-208	others
Excluded channels	
6,7,8,27,28,29,99,100,101,102,108,109,110,129,130,131,180,181,182,183, 186,187,188,189,193,194,195,199,200,201,202,203,204,205,206,207,208	




Figure 7.11: A typical segment of five channels after the preprocessing step

the dimension of the multichannel dataset is reduced. Moreover, if the separation procedure (ICA, π CA, etc.) does not fully separate the subspaces of the maternal and fetal MCGs, i.e., traces of the fetal MCG exist in the maternal components, these components are eliminated with the maternal MCG. However, for the current dataset having many dimensions, it was noticed that by simply nulling the first 5-6 components extracted by π CA, the fetal MCG were still detectable from the remaining channels. We therefore adopted this simple method and as later explained, the quality of the fetal MCG was improved in a post-processing step. In Chapter 8, a more robust method is presented that can be used for maternal ECG/MCG cancellation without loss of dimensions.

7.5.3 Fetal MCG Extraction - First Run

Up to this step, we have preprocessed the data and removed the maternal MCG. Considering the excluded channels in Table 7.1 and the 5-6 dimensions lost in the maternal MCG cancellation, the data has about 165 dimensions. Since we are processing windows of 10000 samples, the ratio between the number of channels to the number of samples is too high and we have the problem of *curse of dimensionality* described in Section 6.4. To remedy this effect, we use PCA to reduce the dimensions of the data to 30, which was found from the slope changes of the normalized logarithmic plot of the data's covariance matrix eigenvalues, in Fig. 7.13. This selection was also practically found to be a compromise between twin fetal MCG extractability and erroneous components (and noise) extraction.

The next step was to apply ICA on the remaining components to extract the fetal MCGs as independent components. For this we used FastICA [67], due to its higher speed for high-dimensional data (as compared with JADE based on higher-order statistics), and also the ability of initializing its initial values from window to window. For this, we used FastICA with the *symmetric* component extraction method with a cubic polynomial nonlinearity ($g(u) = u^3$) in its fixed-point algorithm, which is related to a kurtosis-based approximation of *negentropy* [89]. With these choices, the twin MCGs were empirically found to be always extractable⁵.

As mentioned before, the explained procedures were carried out over blocks of 10000 samples and to assure the consistency of the results, the blocks were chosen to have 10%-50% of overlap with the previous blocks. Moreover, in order to preserve the order of the extracted components (which would otherwise be arbitrary in ICA), the estimate of the mixing matrix of the FastICA algorithm for each block was selected as the initial value of the mixing matrix for the next block. This property, besides preserving the order of the components, speeds up the calculations too; since the statistical properties of adjacent blocks are close to each other, and the initial estimate of the mixing matrix will very soon converge to

⁵Note that in this example we are more interested in the extraction and tracking of fetal MCGs from long recordings, without any claims on the optimality of the parameters used in the FastICA algorithm.




Figure 7.12: The first ten periodic components extracted from a typical segment of the data corresponding to the maternal MCG

its actual value in few iterations. It was later found that a more conservative (and more robust) choice of the initial mixing matrix is a weighted average of the mixing matrices of previous blocks, rather than the mixing matrix of the last block. This choice of the mixing matrix still converges much faster than a blind search (without initialization), and also reduces the risk of losing track of the components due to short noisy segments of the data.

The ambiguity of sign and variance is a known intrinsic limitation of ICA. In order to have a normalized representation, the variances of the independent components extracted by FastICA were always normalized to one and the variance was given to their corresponding columns in the mixing matrix. This method worked well and due to the overlap of the blocks the signs of the components were rather consistent. However, in some noisy segments of the data the sign was changed. This problem was overcome when we increased the overlap of the windows to 90% for fine-tuning. The components found in this extreme case (90% overlap) were very robust without abrupt changes from one block to another.

Note that in the overlapping segments of the data, the extracted components were averaged over the corresponding samples found from the different block ICAs. For example, for a block length of 10000 samples and an overlap of 90%, we will have 9 estimates for each sample (except for the beginning and ending blocks) that correspond to the number of overlapping blocks that contain that specific sample. This averaging generally improves the SNR of components since the signals of the components (e.g. one specific fetal MCG complex) are correlated with each other, while the noises of the components




Figure 7.13: The logarithm of the eigenvalues normalized by the first eigenvalue (cf. [120])

are uncorrelated. Theoretically, by averaging N trials, we can achieve an SNR improvement of up to $10 \log_{10} N$ dB. This can be simply verified by assuming that for each sample we are estimating a random variable of the form $x_i = s + n_i$, where s is the deterministic part of interest that is common in all trials (e.g. a fetal MCG complex), while n_i is the noise part that is uncorrelated from one trial to another.

Typical examples of the first fifteen components found by FastICA are plotted in Fig. 7.14. Among these components, four of them (components 4, 8, 9, and 10), are persistent in all the different blocks. Two of these correspond with one of the fetuses (components 4 and 10) and the other one (component 9) corresponds to the other fetus. This may be verified by plotting these channels in one single plot shown in Fig. 7.15, where we notice that the temporal cardiac rhythm of two of the components are the same, while the rhythm of the other one is different. The fact that we have extracted more than one component from one of the fetuses, corresponds to the multidimensionality of the cardiac signal subspace described in Chapter 6. The fourth persistent component (component 8) in Fig. 7.14, which is a noisy and spiky signal with high frequency and irregular spikes, resembles the magnetic gradient noise that is commonly seen in MCG recordings.

7.5.4 Fetal Peak Detection and HRV Calculation - First Run

The fetal MCG found by FastICA contains some noise. Therefore, the simple peak detection procedure that was used for the maternal R-peak detection in Section 7.5.2 is not very efficient for the fetus and is very likely to lose or miscount some of the peaks. This was verified over one of the fetal components. The well-known Pan-Tompkins R-peak detection method was also applied on the dataset [151]. The results were however rather similar to the previous approach. A more effective method that we developed was to use a combination of a matched filter and a peak-detector for R-peak detection. This method is further described in Appendix D. By this way we achieved to extract the fetal R-peaks for both of the fetuses with much less erroneous peaks as compared with previous classical methods.

The HRV of the fetuses were next calculated from the estimated fetal peaks. The HRV results may be seen in Fig. 7.16 for the whole 30 minute of the dataset. From this figure we notice some erroneous peaks, which are due to the noisiness of the fetal MCG in the corresponding epochs leading to the misdetection of the fetal R-peaks. The negative peaks in the HRV sequence having almost half the average HRV values indicate a missed R-peaks and the ones having one third of the average HRV indicate two successive missed peaks. For automatic detection and analysis, it is further possible to make some general rules for making corrections over an estimated HRV sequence; an issue which is not considered in the current study.

A close look at the HRV extracted of the second fetus in Fig. 7.16 shows an interesting phenomenon. We have zoomed over a small segment of the HRV of the second fetus in Fig. 7.17. As we can see from this figure, there are some sharp variations in the HRV over some epochs. This can be explained by considering that the second fetal MCG in Fig. 7.15 is not in the conventional QRS form, i.e., we have another morphological projection of the vectorcardiogram loop (cf. Chapter 6). Therefore the R-peaks have not been precisely detected from this representation.




Figure 7.14: Independent components extracted from a typical segment of the data after preprocessing and maternal MCG cancellation using π CA




Figure 7.15: Overlaid and zoomed plot of components 4, 9, and 10 from Fig. 7.14. We can see that the green and red plots correspond to one of the fetuses while the blue plot corresponds to the other fetus




Figure 7.16: A primary estimate of the fetal HRVs




Figure 7.17: A zoomed view of the second fetal HRV plot of Fig. 7.16

7.5.5 Fetal MCG Extraction - Second Run

We now have a coarse estimate of the fetal MCGs, the corresponding fetal R-peaks, and the cardiac HRV. These coarse results may be further used to improve the performance of the previous steps of the algorithm. In this post-processing stage there are broad ranges of possibilities that we can test. One of the ideas is to use the initial fetal cardiac R-peaks in a π CA procedure, over the output of the preprocessing step presented in Section 7.5.1. This is similar to what we did for the maternal MCG extraction using π CA. In fact, having the fetal R-peaks and by using π CA we can seek linear mixtures of the preprocessed data (BW removed and lowpass filtered data) that is maximally periodic with respect to the fetal R-peaks. While this procedure can be very effective for extracting multiple dimensions of the fetal MCG (more than one single component for each fetus), we noticed that in the segments in which we have misdetections of the R-peaks it can lead to erroneous components. We can nevertheless keep the idea for when we have more accurate estimates of the fetal R-peaks, for instance at the output of this second post-processing step. Note however that π CA is not very sensitive to single R-peak misdetections, as the covariance estimates required in the π CA procedure are not very likely to be biased by just a few misdetections of R-peaks.

However, with our current estimate of the fetal R-peaks, one of the possible and effective post-processings is to consider the fetal R-peaks in the maternal MCG cancellation step. In equation (7.12) (page 73), several possible covariance matrices were proposed for diagonalization. The procedure that we used in Section 7.5.2 was based on the first cost function that only seeks for the components that are maximally synchronous with the maternal R-peaks. This cost function however does not guarantee to preserve the fetal components. In other words, there is no guarantee that none of the fetal components exist in the first components extracted by applying π CA over the data using the maternal R-peaks. In order to guarantee this issue we can use a modified version of (7.12) to assure that the fetal components are not among the first components extracted by π CA. For this, suppose that we have the preprocessed

multidimensional data $\mathbf{x}(t)$, and the time-varying periods of the maternal and the two fetal MCGs, namely τ_t^m , τ_t^{f1} and τ_t^{f2} have been estimated from the maternal and fetal R-peaks. We propose the following measure:

$$\tilde{C}_x = \tilde{C}_x^m - (\tilde{C}_x^{f1} + \tilde{C}_x^{f2}) \quad (7.13)$$

where $\tilde{C}_x^m = E_t\{\mathbf{x}(t + \tau_t^m)\mathbf{x}(t)^T\}$, $\tilde{C}_x^{f1} = E_t\{\mathbf{x}(t + \tau_t^{f1})\mathbf{x}(t)^T\}$ and $\tilde{C}_x^{f2} = E_t\{\mathbf{x}(t + \tau_t^{f2})\mathbf{x}(t)^T\}$ are the maternal and the two fetal covariance matrices, respectively. Now if we jointly diagonalize \tilde{C}_x and $C_x = E_t\{\mathbf{x}(t)\mathbf{x}(t)^T\}$, the extracted components will be uncorrelated and ranked according to their similarity with the maternal MCG (appearing in the first components) and the two fetal MCGs (appearing in the last components). This approach is rather similar to the ideas used in the classification context using linear classifiers for multiple classes or clusters [65].

In Figs. 7.18 and 7.19 we have depicted the first and last ten components extracted by this method from a 10s segment of the 171-channel preprocessed data. As it is seen, the first components resemble the maternal MCG (least similar to the fetal MCGs) while the last components are mixtures of the fetal MCGs with the least maternal contribution. We can therefore, safely eliminate the first K components that have minimal contribution in the twin MCGs. In this step, since we are localizing the fetal contribution in the last π CA components, we can eliminate up to the first 80-100 (or even more) dimensions of the 171-dimensional preprocessed data, without losing the fetal components. Therefore, this step also replaces the dimension reduction that was used in Section 7.5.3. Of course, with the 60-80 remaining dimensions we are still likely to have the problem of curse of dimensionality. We can therefore additionally use PCA to reduce the dimensions to about 30 components, as suggested in Section 7.5.3, and re-run FastICA for re-extracting the fetal components.

In fact, the key difference between single-step dimension reduction (using PCA) and our two-step dimension reduction (first by π CA and next by PCA) is to better target the components that we want to preserve, i.e., the fetal MCG. Typical segments of the fetal MCG components extracted by this post-processing step are depicted in Fig. 7.20.

7.5.6 Fetal Movement Tracking

Thirty minute recordings are rather long and we do not expect two healthy twins to keep silent and rested during this period. Looking more carefully at the components extracted by the methods proposed in Sections 7.5.3 and 7.5.5, it was noticed that the fetal MCG morphologies are changing in time. Therefore, if we assume that the mother has not considerably moved, the MCG electrode configuration can be considered to be fixed with respect to the mother's body axis. Therefore, the morphological change of the fetal MCGs should be due to the fetal movements.

An interesting framework for studying the fetal movements is the idea of *subspace tracking*. In other words, in every segment of the data being processed, we can estimate and track the rotations of specific subspaces of interest from frame to frame. This tracking can be in terms of tracking the *principal angles* of the subspaces (cf. Appendix A), or tracking the components themselves. As a tracking problem, Kalman filtering can also be very effective and due to the rather slow (and limited) possible movements of the fetuses, a low-dimensional and accurate state-space model can be proposed for the required Kalman filter. This issue is left for future works; in this study we only used a rather simple least squares estimator for tracking the fetal movements.

Suppose that $A, B \in \mathbb{R}^{N \times P}$ are the mixing matrices of two segments of the data estimated by ICA. We seek a rotation matrix Q that minimizes the following measure:

$$d = \|A - BQ\|^2, \quad \text{subject to } Q^T Q = I \quad (7.14)$$

This is a classical optimization problem in linear algebra and the optimal Q is found by *singular value decomposition* (SVD) as follows [73, p. 601]:

- Define $C \doteq B^T A$,
- Compute the SVD of C : $U^T C V = \Sigma$,
- $Q = UV^T$.




Figure 7.18: The first ten periodic components extracted by using a mixed measure of maternal/fetal MCGs, proposed in Section 7.5.5. Note that these components are not as clean as the maternal components in Fig. 7.12, which were based on the maternal periodicity alone.

Using this method, we calculated the rotation matrices (Q) for 10 second blocks with 97.5% of overlap over the total 30 minute record. The mixing matrix of the first extracted block was used as the reference matrix A in (7.14), for these calculations. We next calculated the angles between the columns of each Q with its corresponding column in the Q calculated from the first two blocks (used as a reference). Assuming that $Q \in \mathbb{R}^{P \times P}$, we find P angles corresponding to the number of components extracted by FastICA. Each angle can therefore be considered as a measure of the rotation of a specific component in time. The calculated angles for four of the components extracted by FastICA (one component from the first fetus, two from the other fetus, and one for the spiky component that probably corresponds with the magnetic noise) are plotted in Fig. 7.21. Interestingly, from these results we notice some clear changes in the rotation angles, which should principally correspond to the movements of the fetuses in these specific epochs. Another interesting observation is that the spiky component, which does not have a physiological source, has rotated once around the fifth minute of the data recording. This abrupt change has also influenced the other rotation angles. This can possibly be due to some specific situational change during that specific epoch, such as the movement of the mother with respect to the SQUID system.

Note that the matrix Q^T may also be used as a rotation matrix for compensating the fetal rotations




Figure 7.19: The last ten periodic components extracted by using a mixed measure of maternal/fetal MCGs, proposed in Section 7.5.5

and to derive canonical representations of the fetal signals.

7.6 Summary and Conclusions

In this chapter a linear transform was presented that is able to extract the ‘most periodic’ components corresponding to a desired ECG signal from a set of multichannel recordings by forming some covariance-like matrix and jointly diagonalizing this matrix and the actual covariance matrix of the dataset. The intuition behind this method is to find any periodic structure that is synchronous with the reference ECG R-peaks extracted from a suitably clean ECG reference. As it was explained in Section 7.2.2, the bases of the proposed method are rather similar to PCA and ICA. However, due to the proper use of the temporal pseudo-periodicity of the ECG, the method has some interesting benefits over conventional source separation techniques, which are noted in the following:

- From the physiological point of view, the independence criterion of conventional ICA has been replaced with a periodic temporal structure criterion, which is a more reasonable assumption for ECG signals and in accordance with the clinical intuition about the ECG. In fact, cardiologists




Figure 7.20: Independent components extracted from a typical segment of the data in the post-processing step (Section 7.5.5)




Figure 7.21: Rotation angles (in degrees) calculated according to the procedure presented in Section 7.5.6

are not familiar with the interpretation of independent components extracted from multichannel recordings; but they are interested in periodic structures that are repeated in each ECG beat.

- From the mathematical point of view, the temporal information of the ECG is gathered in the matrix \hat{C}_x . Therefore, the conventional iterative ICA algorithm is replaced by an algorithm consisting of an initial R-wave detection step, covariance matrix calculation, and a single step of GEVD. The method is therefore more time-efficient.
- The proposed method can also be evaluated from the sparsity viewpoint. ECG signals are rather sparse in time (cf. Fig. 7.6). For such signals, the second or higher-order statistics estimated in conventional ICA techniques can be rather susceptible to noise. The measure of periodicity that was considered in the π CA method is a way of using *a priori* information to have better estimates of the required statistics.
- The extracted components are ranked according to their degree of synchronization (periodicity)

with the R-peaks, while in conventional ICA it is not possible to predict the order of the extracted components. This feature is very helpful, especially for automating the removal of the maternal ECG from fetal ECG recordings, or generally for removing cardiac interference from multichannel biosignals.

- Due to the *mean square error compression* property of eigenvalue decomposition [89, ch. 6], the ECG signal is concentrated within the least number of components. In other words, by using the proposed method, we achieve a minimal representation of the ECG signal. This feature can be used in ECG compression, as we can define some threshold on the extracted eigenvalues and remove the components corresponding to the smaller eigenvalues by simple thresholding or through a *hypothesis test*. In this way, since the components are ranked according to their resemblance with the ECG, we are sure that only the ‘least important’ components have been removed. This is however not the case for ICA or PCA compression techniques when applied to the ECG, since the low variance components extracted by these methods can convey important parts of the ECG morphology and may not be removed by thresholding.
- The eigenvalues found from (7.1) are also a measure of the noisiness of the extracted component and may be used for thresholding the minor components. In the general BSS problem, if we consider the observed data to follow the latent variable model $\mathbf{x} = H\mathbf{s} + \mathbf{n}$, any linear combination of the data \mathbf{x} will contain some noise [44]. However, depending on the criterion that is used, the noise may be accumulated in a few undesired components while keeping the desired ones such as the ECG cleaner than the rest of the components. In fact, if we consider any source of aperiodicity as noise, the proposed method may be interpreted as a transformation that is distributing the noise variance in the less important components. This is an important issue that is not achieved with conventional ICA, as they seek for the most independent components and not the most periodic (least noisy) ones.
- The solutions of GEVD problems are generally more susceptible to noise, as compared with joint approximate diagonalization methods [17]. However, in the proposed method although only two matrices are jointly diagonalized, the results are still robust to deviations of heartbeat and noise. The reason is that the time lags τ_t required in the calculation of \tilde{C}_x , are extracted from the beat-to-beat information of the ECG. The robustness may however be improved in future works, if we split the overall information of the RR-interval that is carried by \tilde{C}_x , into several matrices that contain local information of the ECG cycle, such as the P, QRS, and T-segments. These matrices can then be approximately diagonalized using joint approximate diagonalization methods [29].
- On the other hand, we should also consider that conventional joint approximate diagonalization methods require some criterion for measuring the amount of diagonalization of the set of matrices; for instance by minimizing a weighted sum of the off-diagonal entries [17, 156]. Although it has been shown that the maximization of such criteria is equivalent to the maximization of other measures such as the *likelihood function* [156], the cost function is otherwise rather empirical. Therefore, in applying approximate diagonalization methods, we have a sort of ‘double-blind’ problem, in the sense that the matrices used in the diagonalization procedure and the performance measure are both *ad hoc*, which makes the interpretation of the results more difficult. However, in GEVD-based methods we no longer have problems for measuring the amount of diagonalization; but at a cost of the more sophisticated way of forming the original matrices and at a risk of higher sensitivity.

Although the method is a logical extension of conventional source separation techniques customized for ECG signals, more quantitative studies are required for its evaluation and comparison with classical methods. Therefore, in future works, simulated multichannel ECG signals as proposed in Chapter 4, together with quantitative measures of performance can be used to evaluate the performance of this method. Issues concerning the choice of the number of effective ECG dimensions, using independence and periodicity measures, also need further consideration.

In the fetal MCG extraction case study, the proposed method was used for ranking and eliminating the less-informative components. This procedure also reduces the rank of the observations and is appropriate for MCG signals, which are commonly recorded from hundreds of channels. However, for ECG signals recorded from a dozen of electrodes, it is not efficient to reduce the data dimensions by this way. For such

signals, we prefer to separate the desired and undesired subspaces without losing the dimensionality of the recordings. This issue is studied in the next chapter.

Chapter 8

A Deflation Procedure for Subspace Decomposition

8.1 Introduction

In this chapter, we present a general deflation framework for the separation of a desired signal subspace of arbitrary dimensions from noisy multichannel observations. The proposed method is an iterative procedure that is repeatedly applied to the input signal until all the dimensions of the desired subspace are extracted. This method simultaneously uses single and multichannel priors to split the desired and undesired subspaces, even for coplanar (intersecting) subspaces. We will see that by appropriate use of such priors we can even extract the signals from degenerate mixtures of signals and noise in low SNR scenarios. The proposed method is fairly general and may be applied to various applications.

A special case for the hereby proposed method is the problem of separating fetal cardiac signals from interferences and noise. It will be shown that the Bayesian filtering framework developed in Chapter 5 and the multichannel periodic component analysis of Chapter 7 can be effectively used within this framework to recover the weak fetal ECG components embedded in noise.

8.2 Data Model

We consider zero-mean N -dimensional observations $\mathbf{x}(t) \in \mathbb{R}^N$, which follow an additive model as follows:

$$\mathbf{x}(t) = \mathbf{x}_s(t) + \mathbf{x}_n(t) \quad (8.1)$$

where $\mathbf{x}_s(t)$ is the *desired* part of the observations considered as the signal part, and $\mathbf{x}_n(t)$ is the *undesired* part considered as noise and/or interference. The covariance matrix of the observations is defined as follows:

$$C_{\mathbf{x}} \doteq E_t\{\mathbf{x}(t)\mathbf{x}(t)^T\} \quad (8.2)$$

where $E_t\{\cdot\}$ represents averaging over time. We further assume that the desired and undesired parts of the observation be uncorrelated; but the desired parts of the different channels be dependent. They can for example be a linear or nonlinear transform of a signal subspace projected onto the observation space:

$$\mathbf{x}_s(t) = f[\mathbf{s}(t)] \quad (8.3)$$

where $\mathbf{s}(t) \in \mathbb{R}^M$ and $f(\cdot) : \mathbb{R}^M \mapsto \mathbb{R}^N$. The undesired part of the signal is also assumed to be a mixture of low-rank and full-rank noise and/or interferences. The objective is to denoise the observations, i.e. to retrieve $\mathbf{x}_s(t)$.

The data model is very general and appears in many applications. For example, the problem of blind source (or subspace) extraction with a latent variable model is a special case for this model:

$$\begin{aligned} \mathbf{x}(t) &= H\mathbf{u}(t) + \mathbf{n}(t) \\ &= \underbrace{u_1(t)\mathbf{h}_1 + \dots + u_M(t)\mathbf{h}_M}_{\mathbf{x}_s(t)} + \underbrace{\dots + u_P(t)\mathbf{h}_P + \mathbf{n}(t)}_{\mathbf{x}_n(t)} \end{aligned} \quad (8.4)$$

where $H \in \mathbb{R}^{N \times P}$, $\mathbf{u}(t) \in \mathbb{R}^P$, $\mathbf{n}(t) \in \mathbb{R}^N$, and $\mathbf{h}_i \in \text{span}\{H\}$. In this case, we have assumed that the signal of interest ($\mathbf{x}_s(t)$) is a subspace of the column space of H , while the other subspaces and observation noise form the undesired part $\mathbf{x}_n(t)$. Under the assumption of independent sources $\mathbf{u}(t)$ with $P \leq N$, and a moderately small noise $\mathbf{x}_n(t)$, the latent variable model can be solved with second and higher-order statistical ICA methods [31, 17]. However, the performance of these solutions is limited, since:

- a. the estimation of the required demixing matrix is sensitive to the data model and observation noise;
- b. they can not solve the degenerate (under-determined) case ($P > N$);
- c. the dimensions of the desired subspace is to be known;
- d. non-independent, coplanar subspaces may not be separated;
- e. the full-rank noise is not separated and may even be amplified in the extracted components;
- f. they commonly reduce the rank of the observation signals, if used for denoising.

We therefore seek for a method for retrieving $\mathbf{x}_s(t)$, without these limitations based on some *a priori* knowledge of the signal/noise subspace structure.

8.3 Method

8.3.1 Single Channel Denoising

Regardless of the multichannel structure of the data model (8.1), in each of the channels the desired and undesired parts of the observation may be separable (although approximately) using some linear or nonlinear denoising scheme, applied in the time or transform domain. The optimal denoising scheme is of course application dependent and its performance depends on the SNR of each channel. Simple frequency domain filters, wavelet denoisers, optimal Wiener filters, linear or nonlinear Kalman filters, or any other model-based Bayesian filter, are amongst the possible denoising schemes that can be used for single channel denoising. However, due to the commonly low SNR of the observations, single channel denoising is not very efficient for our applications of interest. We therefore seek for a method to benefit from the mutual spatial information of the different channels to improve the signal quality before denoising.

8.3.2 Linear Decomposition using GEVD

Any linear transform of the observation signals $\mathbf{x}(t)$, defined in (8.1), can be represented as follows:

$$y(t) = \mathbf{w}^T \mathbf{x}(t) = \mathbf{w}^T \mathbf{x}_s(t) + \mathbf{w}^T \mathbf{x}_n(t) \doteq y_s(t) + y_n(t) \quad (8.5)$$

Apparently, depending on the vector \mathbf{w} , different mixtures of signal and noise are achieved. The performance of the later proposed method highly relies on the efficient use of the prior information of the signal and noise subspaces in such linear transforms. In the following, we review several different cost functions that can appear in different applications, each leading to a different linear transform.

SNR Maximization

Suppose that we want to find a linear mixture of the input with a maximal SNR. The SNR of the linear mixture $y(t)$ defined in (8.5), can be calculated as follows:

$$\text{SNR}(\mathbf{w}) \doteq \frac{E_t\{y_s(t)^2\}}{E_t\{y_n(t)^2\}} = \frac{\mathbf{w}^T C_{\mathbf{x}_s} \mathbf{w}}{\mathbf{w}^T C_{\mathbf{x}_n} \mathbf{w}} = \frac{\mathbf{w}^T C_{\mathbf{x}} \mathbf{w}}{\mathbf{w}^T C_{\mathbf{x}_n} \mathbf{w}} - 1 \quad (8.6)$$

where $C_{\mathbf{x}_s} \doteq E_t\{\mathbf{x}_s(t)\mathbf{x}_s(t)^T\}$ and $C_{\mathbf{x}_n} \doteq E_t\{\mathbf{x}_n(t)\mathbf{x}_n(t)^T\}$ are the covariance matrices of the signal and noise parts and we have used the uncorrelatedness of the signal and noise parts for the second part of this equation.

Following the explanations in Section 7.2.1, the maximum value of the SNR is achieved by GEVD of the matrix pair $(C_{\mathbf{x}_s}, C_{\mathbf{x}_n})$, or equivalently, the GEVD of the matrix pair $(C_{\mathbf{x}}, C_{\mathbf{x}_n})$. This approach is therefore applicable for the case that $C_{\mathbf{x}_n}$ or $C_{\mathbf{x}_s}$ are known or can be estimated.

Nonstationarity Maximization

Suppose that the signal (or noise) parts of the observations have a ‘burst’ or ‘sparse’ behavior, i.e. only appear over a certain period of the time, and the active times of the desired signals are known or could be estimated from the dataset. This is the case for many practical applications such as mixtures of electroencephalogram (EEG) signals corrupted with eye movement or blink artifacts, EEG evoked responses to ocular or audio stimuli, or EEG signals corrupted by switching magnetic resonance artifacts in fMRI experiments. In all of these applications, the observed signals are nonstationary mixtures of signal and noise. We can therefore define a cost function that accounts for this nonstationarity:

$$\zeta(\mathbf{w}) \doteq \frac{E_{\theta}\{y(\theta)^2\}}{E_t\{y(t)^2\}} = \frac{\mathbf{w}^T \tilde{C}_{\mathbf{x}} \mathbf{w}}{\mathbf{w}^T C_{\mathbf{x}} \mathbf{w}} \quad (8.7)$$

where θ is the active time epochs of the burst signal (or more generally the time instants with a known nonstationarity), and $\tilde{C}_{\mathbf{x}} \doteq E_{\theta}\{\mathbf{x}(\theta)\mathbf{x}(\theta)^T\}$ is the covariance matrix of the observations over the nonstationary periods. Next, by maximizing (8.7), we can find the components that have a maximal contribution in the energy of the burst epochs (the numerator), while having the least contribution in the overall signal energy (the denominator). Again, following (7.2), the maximum value of $\zeta(\mathbf{w})$ is achieved by GEVD of the matrix pair $(\tilde{C}_{\mathbf{x}}, C_{\mathbf{x}})$.

Spectral Contrast Maximization

The signal/noise separability may be achieved in domains other than the time domain. For instance, suppose that we are interested in extracting band-limited signals from the observation signals. For example the extraction of *alpha* or *beta* (or both) rhythms from noisy EEG recordings. In this case, the linear mixture defined in (8.5), can be transformed into the frequency domain:

$$Y(f) \doteq \mathcal{F}\{y(t)\} = \mathbf{w}^T \mathcal{F}\{\mathbf{x}(t)\} = \mathbf{w}^T \mathbf{X}(f)$$

where $\mathcal{F}\{\cdot\}$ represents the Fourier transform. We can now define the contrast function as follows:

$$\sigma(\mathbf{w}) \doteq \frac{E_{\nu}\{|Y(\nu)|^2\}}{E_f\{|Y(f)|^2\}} = \frac{\mathbf{w}^T S_{\mathbf{x}} \mathbf{w}}{\mathbf{w}^T C_{\mathbf{x}} \mathbf{w}} \quad (8.8)$$

where ν is the frequency band of interest, f is the whole frequency axis (or the Nyquist band for the discrete case), and $S_{\mathbf{x}} \doteq E_{\nu}\{\mathbf{X}(\nu)\mathbf{X}(\nu)^H\}$ is the cross-spectrum of the observation vector averaged over the bandwidth of interest. In the denominator of the last part of (8.8), we have used the Parseval’s relation. In order to have a real symmetric matrix $S_{\mathbf{x}}$, the summation of $E_{\nu}\{\cdot\}$ is taken over both positive and negative values of ν (or symmetric frequencies with respect to the Nyquist frequency for the discrete case), over the frequency bands of interest. This condition guarantees the realness of $\sigma(\mathbf{w})$. For this cost function, the maximal spectral contrast is achieved by GEVD of the matrix pair $(S_{\mathbf{x}}, C_{\mathbf{x}})$.

Note that here the objective is to transform the signals into a domain in which the signal and noise are better separated. Therefore, any other transform that preserves the linearity, such as wavelet transforms, may also be used in this procedure.

Periodicity Maximization

The desired or undesired parts of the observations can have a periodic or pseudo-periodic structure. This is also the case in applications in which a periodic signal is corrupted by noise, such as mixtures of ECG or MCG signals and noise. In this case we may seek for linear transformations that maximize some measure of periodicity, while keeping the signal energy bounded:

$$\epsilon(\mathbf{w}) \doteq \frac{E_t\{y(t)y(t+\tau_t)\}}{E_t\{y(t)^2\}} = \frac{\mathbf{w}^T P_{\mathbf{x}} \mathbf{w}}{\mathbf{w}^T C_{\mathbf{x}} \mathbf{w}} \quad (8.9)$$

where τ_t is a (possibly) time-varying period of the periodic or pseudo-periodic signal and $P_{\mathbf{x}} \doteq E_t\{\mathbf{x}(t)\mathbf{x}(t+\tau_t)^T\}$. The AMUSE algorithm with a constant τ_t [201], and periodic component analysis with a time-varying time-lag τ_t presented in Chapter 7, are examples of algorithms that maximize such a cost function. In either case, the maximum value of $\epsilon(\mathbf{w})$ is again achieved by GEVD of the matrix pair $(P_{\mathbf{x}}, C_{\mathbf{x}})$.




Figure 8.1: The general iterative subspace decomposition scheme

The cost functions defined in (8.6)-(8.9), were all defined such that the problem of finding the optimal linear transform would lead to GEVD of the covariance matrix $C_{\mathbf{x}}$ and another positive definite matrix¹ that is designed using our prior knowledge of the desired signals. For our problem of interest, linear transforms based on GEVD have several advantages over other possible linear decompositions:

- The components are uncorrelated and whitened, which means that there is no redundancy in the extracted components up to the second-order statistics, as in PCA.
- They use all the N^2 possible degrees of freedom of an $N \times N$ linear transform; unlike PCA or whitening transforms that only use $N(N-1)/2$ and $N(N+1)/2$ degrees of freedom, respectively.
- The components are ranked according to the statistical measure used in defining the second diagonalized matrix; unlike conventional ICA that can not rank the components.

8.3.3 Iterative Subspace Decomposition

We are now at a point, where we can merge linear multichannel decomposition with single channel denoising schemes. The proposed method is presented in Fig. 8.1. The method consists of an iterative procedure that repeatedly applies a sequence of *linear decomposition (projection)*, *denoising*, and *recomposition (back-projection)* to the input data. Following the explanations in previous subsections, the linear decomposition unit is a GEVD procedure based on our prior knowledge of the signal and noise subspaces. The output components of this unit are assumed to be ranked in descending (ascending) order of ‘similarity’ with the signal (noise) subspace. The objective of this block is to concentrate the components of the desired subspace in the first few components of its output. This unit is followed by a linear or nonlinear denoising filter that is applied to the first L components ($L \ll N$) of the previous block. This filter can be a single channel filter applied to each channel separately, or a multichannel filter applied to the L components together. As explained before, such denoising can indeed be directly applied to the original $\mathbf{x}(t)$ signals too; but by applying it after the linear decomposition in Fig. 8.1, we benefit from the improved SNR of the first few components extracted by the linear decomposition block. Finally, the residual signals of the L denoised components and the other $N-L$ unchanged components are back-projected to the observation space, using the inverse of the linear decomposition matrix. In each iteration of the algorithm some portion of the signal and noise subspaces are separated and the procedure is repeated until the output signals satisfy some predefined measure of signal/noise separability.

The iterative algorithm depicted in Fig. 8.1 can be expressed as follows:

¹The positive definiteness of the matrices, guarantees that the cost functions defined in (8.6)-(8.9) are always positive, which makes their maximization sensible. Among the defined matrices, $P_{\mathbf{x}}$ is the only matrix that depending on the choice of τ_t can become non-positive definite. However, for quasi-periodic signals, such as the ECG, a proper choice of τ_t usually leads to dominant positive eigenvalues of $P_{\mathbf{x}}$.

The Subspace Decomposition by Deflation Algorithm

- 1: $\mathbf{x}^{(0)}(t) = \mathbf{x}(t)$, $k = 0$
- 2: **repeat**
- 3: Calculate $C_{\mathbf{x}}^{(k)}$, the covariance matrix of $\mathbf{x}^{(k)}(t)$
- 4: Calculate $Q_{\mathbf{x}}^{(k)}$, the matrix containing the desired statistics of $\mathbf{x}^{(k)}(t)$
- 5: $W^{(k)} \leftarrow \text{GEVD}(C_{\mathbf{x}}^{(k)}, Q_{\mathbf{x}}^{(k)})^T$
- 6: $\mathbf{y}^{(k)}(t) \leftarrow W^{(k)}\mathbf{x}^{(k)}(t)$
- 7: $A^{(k)} \doteq [\mathbf{a}_1^{(k)}, \dots, \mathbf{a}_N^{(k)}] \leftarrow W^{(k)-1}$
- 8: $\mathbf{s}^{(k)}(t) \doteq [s_1^{(k)}(t), \dots, s_L^{(k)}(t)]^T \leftarrow \mathbf{G}(\mathbf{y}^{(k)}(t))$
- 9: $\mathbf{x}^{(k+1)}(t) \leftarrow \mathbf{x}^{(k)}(t) - \sum_{j=1}^L \mathbf{a}_j^{(k)} s_j^{(k)}(t)$
- 10: $c \leftarrow \zeta(\mathbf{x}^{(k+1)}(t))$
- 11: $k \leftarrow k + 1$
- 12: **until** $c \leq th$

In this algorithm the index k represents the k -th iteration, $W^{(k)}$ is the transpose of the decomposition matrix found by GEVD, $\mathbf{x}^{(k+1)}(t)$ is the output of each iteration, $\mathbf{G}(\cdot)$ is the denoising function for removing (keeping) the undesired (desired) components applied to the first L channels of $\mathbf{y}^{(k)}(t)$, $s_j^{(k)}(t)$ is the output of the denoising block in channel j , $\mathbf{a}_j^{(k)}$ is the j -th column vectors of $A^{(k)}$, $\zeta(\cdot)$ is a measure of the desired subspace removal used as a stopping criterion, and th is a predefined threshold. The output of each iteration of the algorithm can also be represented in the following compact form:

$$\mathbf{x}^{(k)}(t) = \mathbf{x}(t) - \sum_{i=0}^{k-1} \sum_{j=1}^L \mathbf{a}_j^{(i)} s_j^{(i)}(t) \quad (8.10)$$

It should be noted that the covariance matrices required for the GEVD procedure of the proposed method, are recalculated in every iteration from the output signals of the previous iteration. Moreover, the separation criterion used for designing the linear projector and also the denoising block of Fig. 8.1 do not need to be the same in all iterations. For instance, for EEG signals, if we use the spectral contrast function defined in (8.8), one iteration of the algorithm can be designed to extract the alpha rhythms, while another iteration extracts the beta rhythms of the EEG. Similar examples are presented in the following for the extraction of maternal and fetal ECG mixtures.

One of the most important advantages of the proposed method is that we can remove the desired (or undesired) subspaces without losing the dimensionality of the recordings. This is, in fact, due to the intermediate denoising step that is ‘breaking’ the linearity of the transform, which is an important issue for real noisy signal mixtures recorded from a few number of channels.

8.4 Application in Fetal ECG Extraction

As discussed in previous chapters, the fetal ECG recorded from the maternal abdomen are heavily contaminated with maternal ECG, which depending on the gestational age and electrode locations, can be up to ten to thirty times stronger than the fetal components. Independent subspace analysis (ISA) and other ICA-based methods are the most common approaches for fetal ECG extraction. However, these methods have limitations that were explored in previous chapters. In fact, although the maternal and fetal ECGs form two independent subspaces, and ISA tends to separate these two subspaces, but in the presence of noise or with special lead configurations, the subspaces of the fetal and maternal cardiac signals are not fully separated by linear ICA. This usually results in fetal signals that are still contaminated by the maternal components, or in fetal components removed with the maternal signals. We should add to this, the problem of distributedness of the cardiac sources and the ambiguity in the number of dominant dimensions corresponding to the maternal and fetal subspaces².

In the following, the proposed deflation procedure is used for the removal of maternal ECG interference from fetal recordings. Due to the pseudo-periodic structure of the ECG, the linear step of the algorithm is designed according to the periodic component analysis procedure developed in Chapter 7, and for the

²The problem of dimensionality is more dominant for the maternal ECG where the far-field approximation of the cardiac potentials is less relevant.




Figure 8.2: Maternal phase calculation procedure





Figure 8.3: The overall iterative procedure for maternal ECG cancellation

Figure 8.4: The L -channel Kalman filter for maternal ECG removal

denoising step, we use the Kalman filtering framework presented in Chapter 5. The method is applied to both simulated and real ECG recordings. The detailed procedure of maternal ECG removal based on the deflation method is depicted in Figs. 8.2-8.4.

Note that after removing the maternal interference, a post ICA, π CA, or the deflation procedure may be applied to the residual signals, to separate the fetal components from the other background noises.

8.4.1 Simulated Data

The Dataset

Due to the multi-dimensional nature of the ECG and our processing procedure, a realistic ECG model with multiple dimensions is required. Here we use a multichannel ECG and noise generator similar to the one proposed in Section 4.6.2 for generating simulated ECGs. The model is as follows:

$$\begin{aligned} \mathbf{x}(t) &= \alpha H_m \mathbf{s}_m(t) + H_f \mathbf{s}_f(t) + \beta H_\eta \mathbf{v}(t) + \gamma \mathbf{n}(t) \\ &\doteq \alpha \mathbf{x}_m(t) + \mathbf{x}_f(t) + \beta \boldsymbol{\eta}(t) + \gamma \mathbf{n}(t) \end{aligned} \quad (8.11)$$

where $\mathbf{s}_m(t)$ and $\mathbf{s}_f(t)$ are three-dimensional sources representing the maternal and fetal cardiac components³, $\mathbf{v}(t)$ is structured low-rank noise, $\mathbf{n}(t)$ is full-rank noise, and H_m , H_f , and H_η are the volume conduction transfer matrices for the mother, fetus, and structured noises, respectively. In this model, the maternal signal $\mathbf{x}_m(t)$ is assumed as *interference*, while $\boldsymbol{\eta}(t)$ and $\mathbf{n}(t)$ are assumed as *noises* for the fetal signal $\mathbf{x}_f(t)$. Therefore, the parameters α , β and γ control the signal-to-interference ratio (SIR) and signal-to-noise ratio (SNR) of the fetal components.

For this simulation, three independent leads from 24 subjects of the Physikalisch-Technische Bundesanstalt diagnostic ECG database (PTBDB) [158], were used. This database has a sampling rate of 1000Hz. The signals were further preprocessed for baseline wander removal and bandpass filtered between 0.7Hz and 150Hz.

For each sample trial, three channels of the preprocessed data were randomly selected among the 24 subjects to represent $\mathbf{s}_m(t)$ and $\mathbf{s}_f(t)$. Moreover, considering that the fetal heart beat is almost twice as

³the assumption of three dimensions for the maternal and fetal components is based on the far-field dipole approximation of cardiac potentials [127].

fast as the maternal heart beat, the signals representing $\mathbf{s}_f(t)$ were resampled to 500Hz to mimic fetal ECGs with higher heart beats. All simulations were performed over 10s signals.

The full-rank noise $\mathbf{n}(t)$ was considered as white Gaussian noise, while for the structured noise $\mathbf{v}(t)$, real baseline wander, muscle artifacts, and electrode movement signals from the NSTDB [138, 139] were randomly selected and resampled to $f_s=1000\text{Hz}$.

Similar to the example in Section 4.6.2, H_m , H_f , and H_η are random 8×3 matrices, which result in 8 channel observation signals $\mathbf{x}(t)$. Moreover, to control the angle between the maternal and fetal subspaces, we calculate the *principal angles* (PA) between these two subspaces, denoted as follows:

$$\theta_{mf} = \text{PA}(H_m, H_f)$$

The procedure of calculating the PA is explained in Appendix A.

Identical to (4.15)-(4.17), we calculate the following SIR, SNR_n , SNR_η , and the signal to interference plus noise ratio (SINR), as measures of fetal signal quality *before denoising*:

$$\text{SIR} \doteq \frac{E_{t,i}\{x_{f_i}(t)^2\}}{\alpha^2 E_{t,i}\{x_{m_i}(t)^2\}} \quad (8.12)$$

$$\text{SNR}_\eta \doteq \frac{E_{t,i}\{x_{f_i}(t)^2\}}{\beta^2 E_{t,i}\{\eta_i(t)^2\}} \quad (8.13)$$

$$\text{SNR}_n \doteq \frac{E_{t,i}\{x_{f_i}(t)^2\}}{\gamma^2 E_{t,i}\{n_i(t)^2\}} \quad (8.14)$$

$$\text{SINR} \doteq \frac{E_{t,i}\{x_{f_i}(t)^2\}}{E_{t,i}\{[\alpha x_{m_i}(t) + \beta \eta_i(t) + \gamma n_i(t)]^2\}} \quad (8.15)$$

where $x_{f_i}(t)$, $x_{m_i}(t)$, $\eta_i(t)$, and $n_i(t)$, are respectively the entries of $\mathbf{x}_f(t)$, $\mathbf{x}_m(t)$, $\boldsymbol{\eta}(t)$, and $\mathbf{n}(t)$ in channel i , and $E_{t,i}\{\cdot\}$ represents summation over time and channels. Moreover, for this simulation, the parameters α , β and γ were selected such that

$$\begin{aligned} \text{SNR}_\eta &= \text{SIR} + 10\text{dB}, \\ \text{SNR}_n &= \text{SIR} + 20\text{dB}, \end{aligned}$$

and the overall fetal input SINR was swept in the range of -25dB to -5dB, which are close to their actual values. With these choices, the maternal ECG is the dominant artifact (as it is so in reality).

Evaluation

For evaluation, three algorithms were used:

1. *ICA*: As a benchmark method, the JADE algorithm was applied to the simulated data $\mathbf{x}(t)$ to achieve independent components $\mathbf{y}(t)$. Then, in order to find and regroup the dominant components that correspond to the fetal subspace, the optimal linear transform was found such that the following error be minimized:

$$e_i = E_t\{(y_i(t) - \mathbf{w}_i^T \mathbf{s}_f(t))^2\} \quad (8.16)$$

where $y_i(t)$ is the i -th entry of $\mathbf{y}(t)$ and $\mathbf{w}_i \in \mathbb{R}^{3 \times 1}$ is a projection vector to be found. The optimal solution of this problem, is known to be:

$$\mathbf{w}_i^* = E_t\{\mathbf{s}(t)\mathbf{s}_f(t)^T\}^{-1} E_t\{y_i(t)\mathbf{s}_f(t)\} \quad (8.17)$$

Using this optimal transform, the corresponding errors were calculated from (8.16) and ranked in ascending order. The first three channels of $\mathbf{y}(t)$ with the smallest value of e_i were selected as the fetal subspace⁴. These three components were back-projected to the sensor space using the inverse of the demixing matrix estimated by JADE.

⁴Note that this procedure is somehow unrealistic for real ISA problems; since in reality we do not have access to the actual fetal subspace $\mathbf{s}_f(t)$. Instead the fetal components are selected and regrouped by visual inspection or other classification criteria.

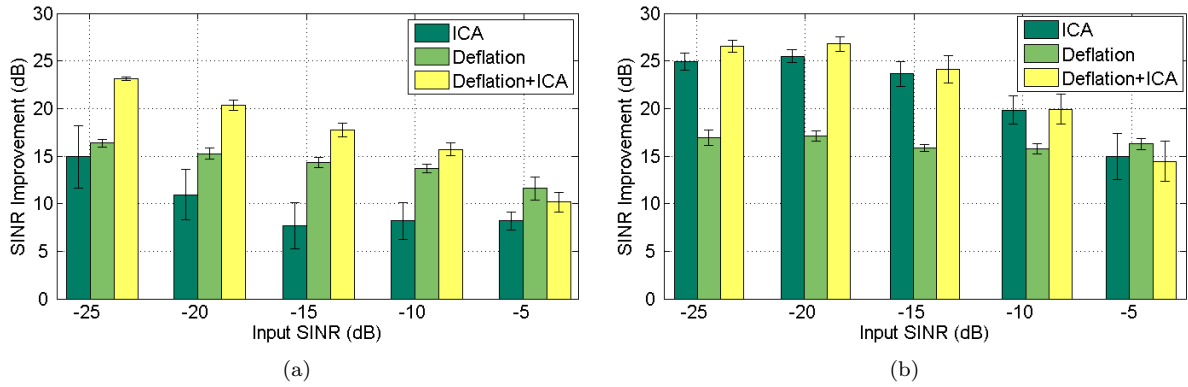


Figure 8.5: The mean and standard deviation bar of SINR improvements achieved in different input SINRs for (a) $\theta_{mf} < 10^\circ$ (b) $\theta_{mf} > 60^\circ$

2. *Deflation*: According to the procedure developed in previous sections, the deflation method was also applied to the synthetic signals for removing the maternal ECG interferences. The R-peaks of the maternal ECG were detected from an arbitrary channel which had dominant maternal ECG artifacts. The deflation procedure was repeated in four iterations for each dataset, which was empirically found to be sufficient for removing the three-dimensional maternal subspace used in the simulations.
3. *Deflation + ICA*: As mentioned before, the described deflation procedure is only able to remove the maternal subspace. The deflation results may be improved by applying a post-filtering, for instance another ICA, to the signals after maternal ECG removal. This ICA step is identical to the one explained in item (1), except that the mixture is ‘less-degenerate’, in the sense that the maternal ECG has already been removed and there are no more dimensions occupied by the maternal ECG components at the ICA output. Note that in this post-processing step, we again take the first three independent components corresponding to the fetal subspace.

For each value of the input SINR, the simulations were repeated 10 times, each time with different signal and noise samples and for different values of θ_{mf} . Following the discussions in previous chapters, multichannel source separation techniques are believed to be rather sensitive to the angles between the mixing subsets. Therefore, all experiments were performed for two extreme cases: (1) all principle angles smaller than 10° with an average $\theta_{mf} = 5.9^\circ$, which we denote by $\theta_{mf} < 10^\circ$, and (2) all principal values greater than 60° with an average $\theta_{mf} = 73.1^\circ$, which we denote by $\theta_{mf} > 60^\circ$. These two cases represent very close and very far maternal-fetal subspaces, respectively. They were achieved by generating random matrices H_m and H_f , and then altering them with a Givens rotation [73], to satisfy the mentioned conditions on their principle angles.

After applying the filtering procedure, the SINR improvement⁵ achieved by the filtering procedures were calculated and compared for the different denoising methods. In Figs. 8.5(a) and 8.5(b), the results of this study can be seen for $\theta_{mf} < 10^\circ$ and $\theta_{mf} > 60^\circ$, respectively.

We can see from Fig. 8.5(a) that the SINR improvement is significantly lower as compared with Fig. 8.5(b), indicating the fact that subspace separation becomes more difficult as the two subspaces become close to each other. In this case, the deflation procedure outperforms ICA and the combination of deflation and ICA (denoted by deflation+ICA) outperforms both methods, except in very high SINR where the deflation method alone has outperformed the ICA-based methods. For $\theta_{mf} < 10^\circ$, the standard deviation of the ICA method is higher than the other methods, indicating its lower robustness.

On the other hand, in Fig. 8.5(b) we can see that for $\theta_{mf} > 60^\circ$ ICA alone and the combination of ICA and deflation have close performances, while the latter has been slightly more effective. Interestingly, in the highest input SINR (-5dB), the deflation method has again outperformed the ICA-based methods.

⁵As before, the SNR improvement is the output SNR, in decibels (dB), minus the input SNR, in dB.

8.4.2 Real Data

For real data, it is not possible to calculate the SNR improvement as a measure of performance. However, in order to have a quantitative measure, we propose to compare the amount of ‘periodicity’ of the signal before and after the filtering procedure. Here, the idea is that an artifact-free signal should not contain any waveform that is synchronous with the ECG signal. Based on this idea, we first find the ECG R-peaks and the phase ϕ_k , as explained in Chapter 5. Then by using ϕ_k , the time-lag τ_k is calculated as the time distance between the sample x_k and its *dual sample* $x_{k+\tau_k}$ ⁶. From this, the following correlation coefficient is proposed as an overall *periodicity measure* (PM) for a signal x_k :

$$\text{PM} = |\text{corrcoef}(x_k, x_{k+\tau_k})| = \frac{|E\{x_k x_{k+\tau_k}\}|}{|E\{x_k^2\}E\{x_{k+\tau_k}^2\}|^{1/2}} \quad (8.18)$$

where $E\{\cdot\}$ represents averaging over the time index k . From this definition $0 \leq \text{PM} \leq 1$, where $\text{PM} = 0$ indicates an aperiodic signal, and $\text{PM} = 1$ indicates a fully periodic one (with respect to the target ECG signal). An effective filter should be able to remove any component that is temporally synchronous with the heartbeat and the samples x_k and $x_{k+\tau_k}$ should become uncorrelated; resulting in a PM close to zero. It should of course be noted that the reduction of PM is a *necessary* but not a *sufficient* measure of the filtering performance. In fact, the PM might be reduced, e.g. by an increase of the overall noise, without an improvement of the signal quality. Therefore, other evidence such as the visual inspection of the resultant waveforms or a comparison of the signal spectra before and after filtering (especially for EEG signals) is always required besides this measure.

First Dataset

We first apply the proposed method to the DaISy fetal ECG database [46], introduced in previous chapters, for removing the maternal ECG. This dataset has been replotted in Fig. 8.6(a). As a preprocessing step, the baseline wander of the data was removed, the maternal R-peaks were detected and the proposed method was applied on it in five iterations. The results of these iterations are seen in Figs. 8.6(b)-8.6(f). Note that for this example we have fixed the number of iterations of the recursive algorithm to five, and we do not use a threshold on the PM for stopping the algorithm.

The PM of all channels in different iterations can also be seen in Table 8.1. It is seen that the proposed method has effectively reduced the maternal ECG contaminants from the first to the last iteration. This result is also quantitatively approved as we see that the average PM of all channels has also decreased in each iteration. These values can be used for defining the appropriate threshold (*th*) for stopping the iterative procedure.

In order to compare the performance of the deflation procedure with ICA methods applied directly on the original data, π CA was applied to the fetal ECG signals of Fig. 8.6(f), i.e., to the results of the last iteration of the deflation procedure; but this time using the fetal R-peaks as the periodic signal of interest. Therefore, in the output, the components were ranked in descending order of contribution in the fetal ECG. The ensemble averages of these periodic components are compared with the independent components extracted by ICA in Fig. 8.7. In this figure we can see that there are three dominant components extracted from the fetal ECG by the combination of the deflation procedure and π CA with some small trace of the fetal P-wave, while ICA has only extracted two dominant components.

Next, by nulling the last six components, i.e. keeping only two dimensions for the fetal ECG, the components were back-projected to the signal subspace using the inverse of the π CA decomposing matrix. A similar procedure was applied on the independent components of Fig. 7.7 (page 75), by back-projecting the fourth and eighth components (that correspond to the fetal ECG) onto the sensor subspace using the estimated mixing matrix. The resultant components of both methods are depicted in Fig. 8.8. As seen in this figure, while the fetal contributions in the first five abdominal channels are rather the same for both methods; but the deflation procedure has additionally found the fetal components that existed in the last three thoracic channels. This means that there are considerable fetal components even in the electrodes recorded close to the maternal heart, which are not extractable using simple ICA. We should add to this that denoising using ICA requires the manual selection of the fetal channels, while it can be done automatically in the deflation procedure.

⁶ τ_k was previously defined in equation (7.9).

Table 8.1: fetal ECG data PM in different iterations for the DaISy database

channel	1	2	3	4	5	6	7	8	mean
iteration									
original	0.837	0.959	0.923	0.617	0.963	0.977	0.981	0.978	0.904
1	0.704	0.801	0.231	0.828	0.403	0.697	0.874	0.954	0.687
2	0.306	0.424	0.223	0.152	0.359	0.501	0.586	0.447	0.375
3	0.281	0.173	0.063	0.151	0.150	0.459	0.268	0.316	0.233
4	0.276	0.169	0.062	0.151	0.152	0.284	0.277	0.316	0.211
5	0.078	0.061	0.058	0.012	0.135	0.281	0.270	0.305	0.150

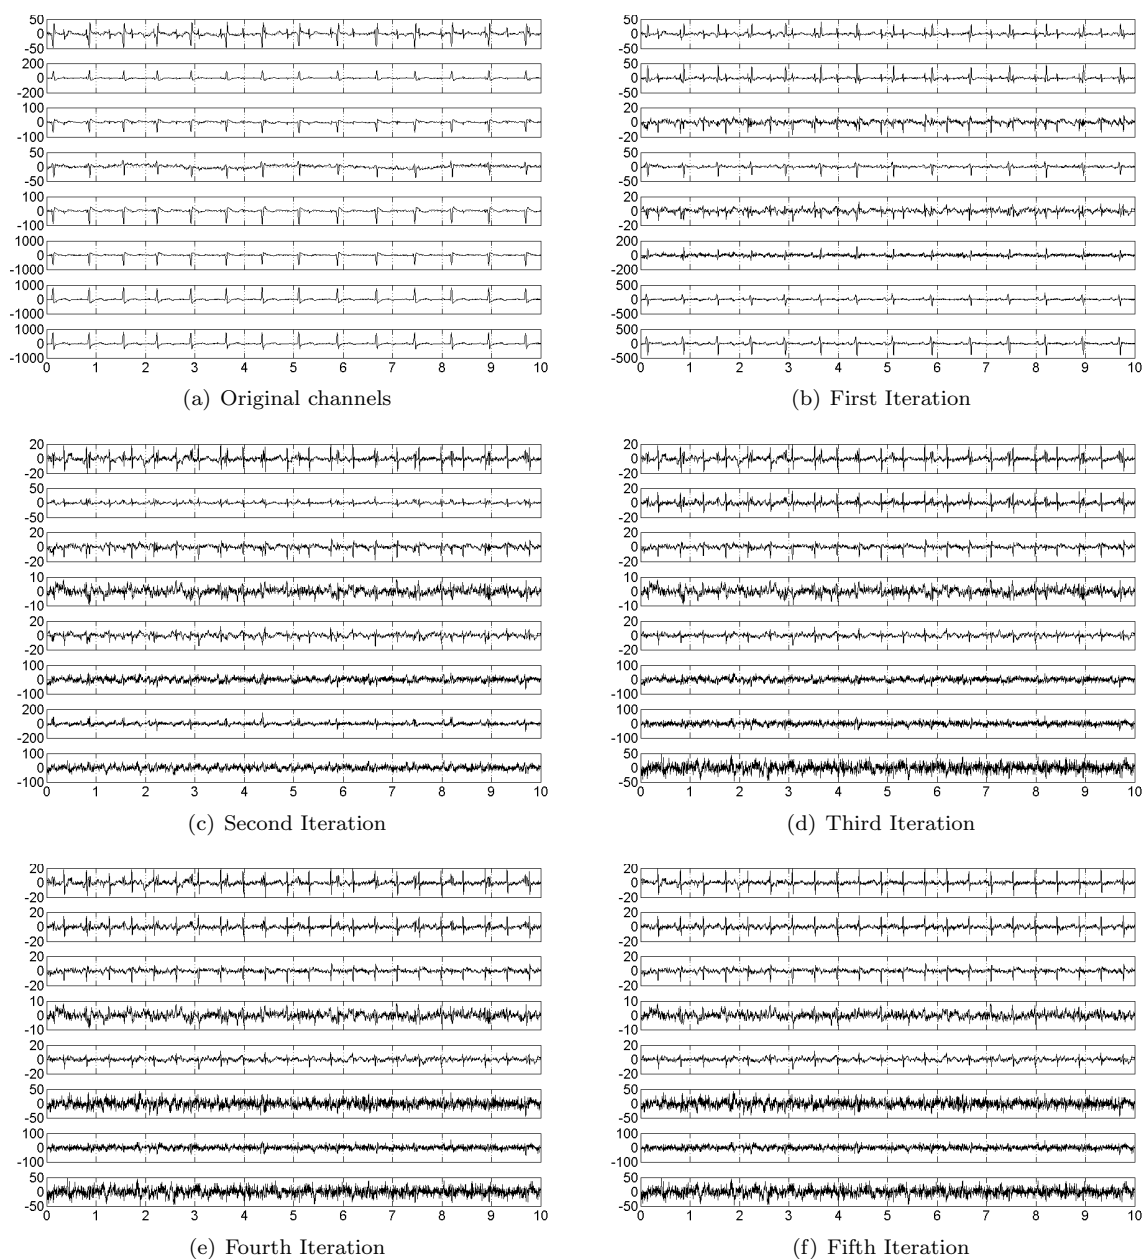


Figure 8.6: The results of the proposed method on the DaISy fetal dataset in five iterations




Figure 8.7: Comparison of the fetal ensemble average extracted by (a) ICA and (b) the deflation method followed by π CA from the DaISy database




Figure 8.8: Residual fetal components extracted by ICA (in red) and the deflation procedure (in blue) from the DaISy database

Table 8.2: fetal ECG data PM in different iterations for the second database

channel iteration	1	2	3	4	5	6	mean
original	0.291	0.623	0.809	0.790	0.791	0.963	0.705
1	0.515	0.527	0.245	0.733	0.556	0.875	0.575
2	0.471	0.527	0.198	0.680	0.477	0.510	0.477
3	0.266	0.492	0.156	0.019	0.033	0.499	0.244
4	0.112	0.378	0.076	0.019	0.032	0.321	0.156
5	0.056	0.026	0.082	0.006	0.021	0.242	0.072

Second Dataset

The second dataset was recorded in the Biological Signal Analyses group, München, Germany⁷ [181]. It consists of five abdominal channels and one thoracic channel for recording the maternal ECG. The data has a sampling rate of 1600Hz. This dataset was a special case in which, due to the specific electrode configuration, none of the common linear transforms including PCA, ICA using FastICA, JADE, SOBI, and π CA achieved in separating the maternal and fetal subspaces. It was therefore a suitable dataset for evaluating the deflation procedure.

The original dataset and the fetal signals after maternal ECG removal extracted by five iterations of the deflation method are seen in Fig. 8.9. In Table 8.2 the PM after each iteration of the deflation method is listed. We can again see that the maternal ECG has been effectively removed in five iterations.

8.5 Summary and Conclusion

In this chapter, a general deflation framework was presented for the separation of a desired signal subspace of arbitrary dimensions from noisy multichannel (and possibly degenerate) mixtures of signal and noise.

While the method is general and may be adapted to different applications, we specifically focused on the extraction of fetal ECG signals from maternal abdominal recordings. For this application the developed method is partially based on the Bayesian filtering framework developed in Chapter 5 and the multichannel periodic component analysis of Chapter 7. The results of the proposed method were applied on simulated and real data.

Following our discussions in Chapter 6, ECG signals of either the mother or the fetus, have several dimensions. Other biosignals and structured noises can also occupy several dimensions. Therefore, multichannel recordings having six to eight channels can often become degenerate mixtures of the desired and undesired sources, which means that we can not retrieve all dimensions of the fetal cardiac signals. In this case, the proposed method, which is able to remove the maternal contaminants without reducing the rank of the observations, can be very effective for retrieving the weak fetal components.

As compared with simple denoising schemes, the deflation procedure is more effective; but computationally more expensive. However, due to the simplicity of GEVD as compared to ICA methods based on approximate joint diagonalization, the proposed method can be more efficient. The cost of this improved performance is the need for *a priori* information about the signal and noise subspaces.

The performance of the deflation method highly depends on the performance of its decomposition and denoising steps. From (8.10), it is seen that the output of each iteration of this algorithm is equal to the original input vector minus a summation of terms corresponding to the undesired signals. However, an open question that requires further study is whether the deflation procedure is stable and under which (theoretical) conditions does it converge to the desired subspace.

⁷ This dataset has been recorded by Dr. Evelyn Huhn and provided by Dr. Raphael Schneider from the Biological Signal Analyses group, München, Germany.

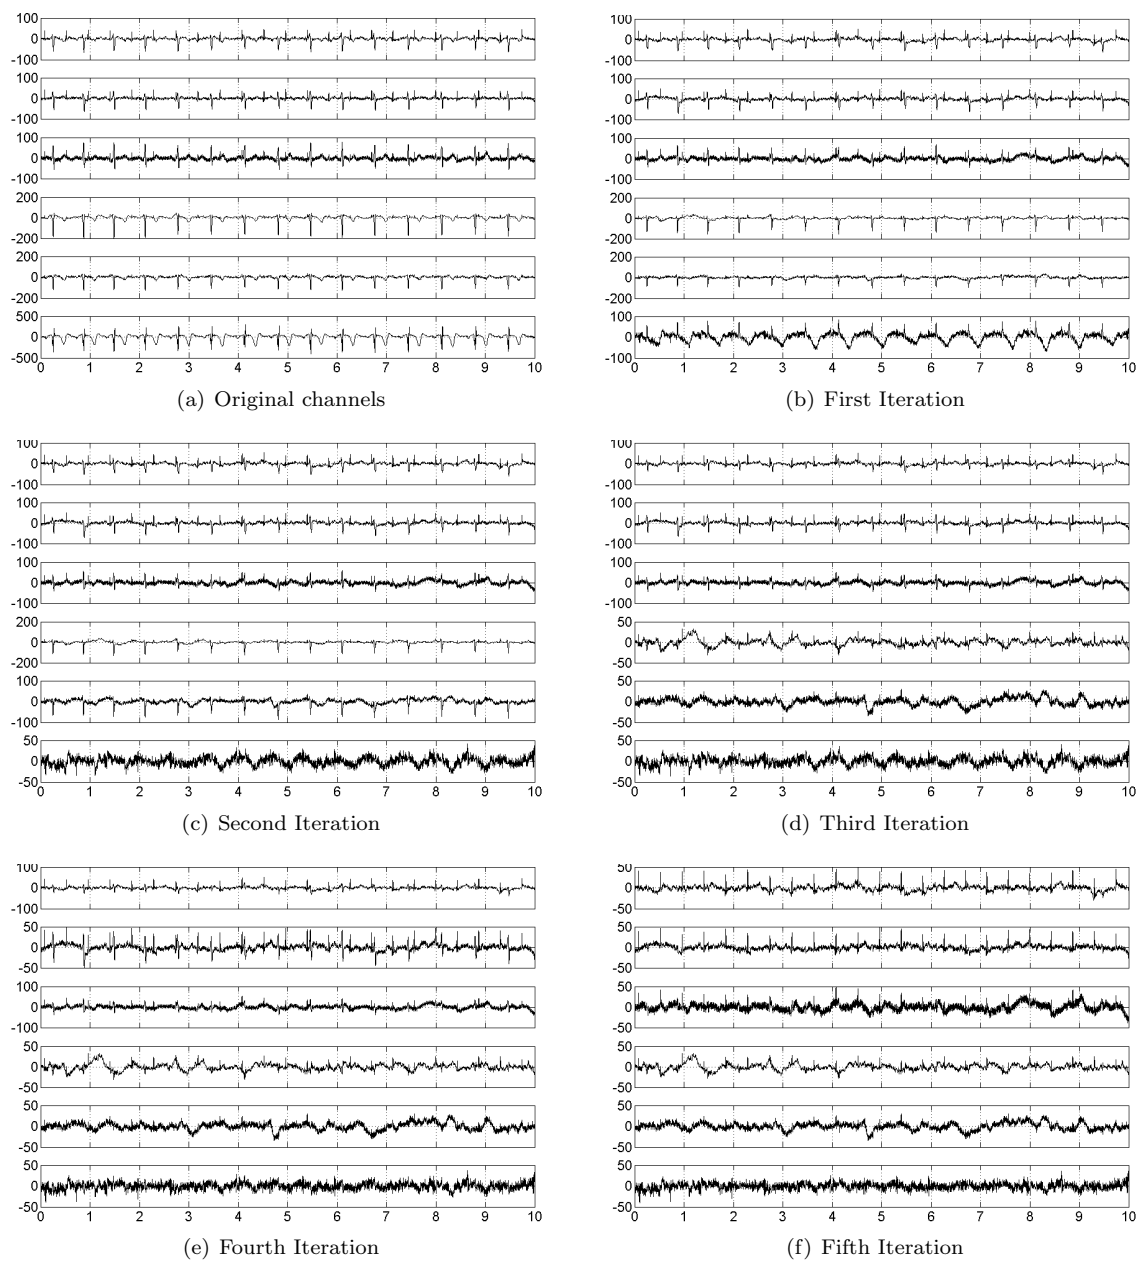


Figure 8.9: The results of the proposed method on the second fetal dataset in five iterations

Chapter 9

Model-Based Bayesian Filtering of Cardiac Contaminants from Biomedical Recordings

9.1 Introduction

The Bayesian filtering framework presented in Chapter 5 was originally developed for ECG denoising. However, as we showed in Section 5.4 and Chapter 8 the same procedure may be used for eliminating undesired ECG signals, such as the maternal ECG, from the desired background signals, such as the fetal ECG. It was later noticed that the same idea could be used as a general framework for removing *cardiac contaminants* (CC) such as the ECG, ballistocardiogram¹ (BCG), or MCG, from other biomedical recordings. This idea has been since implemented and tested over several databases from rather diverse applications, in which conventional CC removal techniques had not led to satisfying results.

In this chapter, the details of this general framework are presented and simulated data are used to quantify its performance. Moreover, several examples from real world applications are presented, on which we have evaluated the method up to now. The proposed method is a filtering framework with several tunable parameters, which enables it to be simply adapted to different applications. Some general rules are presented throughout the chapter, by which the filter parameters can be manipulated and tuned for different situations.

For the sake of brevity, we only present the EKS results that were shown to outperform the EKF and UKF in a wide range of signal SNRs in Chapter 5.

9.2 Cardiac Artifact Removal

In (9.1) and (9.2), the nonlinear state-space model of the ECG, presented in Chapter 5 are rewritten:

Process equation:

$$\begin{cases} \theta_{k+1} = (\theta_k + \omega\delta) \bmod (2\pi) \\ s_{k+1} = -\sum_{i=1}^N \delta \frac{\alpha_i \omega}{b_i^2} \Delta\theta_i \exp\left(-\frac{\Delta\theta_i^2}{2b_i^2}\right) + s_k + \eta, \end{cases} \quad (9.1)$$

Observation equation:

$$\begin{cases} \phi_k = \theta_k + u_k \\ y_k = s_k + v_k, \end{cases} \quad (9.2)$$

where δ is the sampling period, $\Delta\theta_i = (\theta_k - \theta_i) \bmod (2\pi)$, $\omega = 2\pi f$, f is the beat-to-beat heart rate, and N is the number of Gaussian functions used for modeling the shape of the desired ECG. In (9.1) and (9.2), θ_k and s_k are assumed as the state variables, and ω , α_i , θ_i , b_i and η are assumed as i.i.d Gaussian random variables considered as process noises.

¹ The ballistocardiogram is a record of the body's recoil caused by the cardiac contraction.




Figure 9.1: The overall denoising scheme. As shown in this figure the R-peaks of the contaminating signals (CC) may either be detected from an arbitrary reference ECG or from the noisy biosignal after baseline wander (BW) removal.

The basic difference between our approach here and the method in Chapter 5, is that we are interested in the residual signal v_k rather than s_k . In other words, s_k represents the CC, while v_k is the signal of interest. Therefore, we can model the CC with the dynamic model in (9.1) and apply the KF to find an estimate of the artifacts, namely \hat{s}_k . The background signal may then be found as follows:

$$\hat{v}_k = y_k - \hat{s}_k, \quad (9.3)$$

which is the *innovation* signal of the KF.

The overall filtering procedure is illustrated in Fig. 9.1 and may be summarized as follows:

1. *Baseline wander removal.* For the reliable extraction of the average CC templates, the baseline wander of the noisy records should be removed beforehand.
2. *R-peak detection.* These peaks are required for constructing the phase signal ϕ_k , which is in terms needed for synchronizing the noisy ECG with the dynamic model in (9.1). They are also used for extracting the mean CC by synchronous averaging over the heart beats. Depending on the power of the CC, as compared with the background signals and noise, the R-peaks may be detectable from the noisy recordings or from any arbitrary ECG channel synchronously recorded with the noisy dataset.
3. *CC template extraction.* Using the R-peaks, the *ensemble average* (EA) and standard deviation of the CC are extracted through synchronous averaging. Several methods have been proposed in the literature for synchronous averaging. One of the most effective approaches in this area is the *Robust Weighted Averaging* method [121], which outperforms conventional EA extraction methods and is useful for noisy nonstationary mixtures.
4. *Model fitting.* As proposed in Chapter 5, by using a nonlinear least square estimation, the parameters of the Gaussian kernel defined in (9.1) are found, such that the model will best fit the mean CC waveform.
5. *Covariance matrix calculations.* The standard deviation of the average CC is used to find the entries of the covariance matrices of the dynamic model noise Q_k , and measurement noise R_k .
6. *Filtering.* Having extracted all the required parameters, the CC may be estimated by the KF framework and the desired background signal is found from (9.3).

A rather general algorithm for selecting the filter parameters was presented in Chapter 5. Some other rules-of-thumb are summarized in Table 9.1, by which the original filter parameters can be adapted to different situations.

Note further that for online applications or denoising long nonstationary datasets, all of the model parameters and the covariance matrices may be updated in time, by recalculating them from the most recent cardiac beats.

Table 9.1: Rules chart for the manipulation of the filter parameters

Situation	Comments
Strong baseline wanders	Remove before Kalman filtering
Low SNR (strong CC)	Decrease the corresponding R_k entries
High SNR (weak CC)	Increase the corresponding R_k entries; additional ECG reference may be required for R-wave detection
Desired signal with highly colored spectrum	Remove the baseline wander as much as possible
Desired signal highly nonstationary	Allow the filter to adaptively change R_k using the KF innovation signal according to the procedure proposed in Chapter 5
Irregular CC (high inter-beat variations of CC)	Increase the corresponding Q_k entries
Desired signal highly non-Gaussian	Better performance may be achieved using UKF
Batch offline processing	EKS outperforms EKF
Online processing	Use EKF or fixed-lag EKS with a single heart beat delay for parameter estimation

9.3 Bayesian Filtering versus Conventional Techniques

Cardiac signal processing is a highly developed and competitive area that might raise suspicions about the effectiveness of the proposed method. However, we can argue that this method is in fact a generalization of conventional denoising techniques.

In fact, from a filtering point of view, Kalman filters can be assumed as *adaptive filters* that continuously move the poles and zeros of their transfer functions, according to the signal/noise contents of the input signal and the prior model of the signal dynamics. This feature allows the filter to adapt with different spectral shapes and temporal non-stationarities. The method is therefore more general than other frequency domain filters such as FIR, IIR, and Wiener filters that use fixed pole-zero configurations in their transfer function.

A similar argument holds about *wavelet denoising* (WD) [107], as it performs a sort of ‘blind thresholding’ between the signal and noise, without using the temporal structure and pseudo-periodicity of the cardiac contaminants. The proposed method may nevertheless be linked with the known WD techniques, by applying the Bayesian filter in the wavelet domain rather than the original time domain.

As explained in the previous section, since the ensemble average (EA) of the CC is required for training the KF parameters, the proposed method is also comparable with conventional EA removal techniques that simply remove the EA of the contaminating signals from the noisy recordings [5]. The proposed method is again more general in the sense that it adaptively changes its ensemble average template, accounting for the inter-beat variations of the CC. In fact, in the extreme case that the desired signal is very clean (with low cardiac contamination), since the KF is seeking for the CC that is highly defected by the background signal, the observation signal of the KF is not ‘trusted’, and the filter will follow its dynamic model that was trained by the EA of the CC. Mathematically speaking, in this case, the entries of the observation noise covariance matrix R_k are large and the Kalman gain is small [105]. The proposed method therefore reduces to simple EA subtraction. On the other hand, for extremely noisy signals (with high cardiac contamination), the R_k entries will be rather small and the filter will ‘trust’ the observations for CC estimation. Therefore, the filter fully tracks the CC within the input signal.

Hence, the Bayesian approach theoretically outperforms many of the conventional techniques, in the same manner that model-based parametric methods typically outperform nonparametric methods, as long as their underlying model is consistent [25]. However, it does not seem to be possible to prove this claim in the general case, and case-by-case quantitative studies are required for different applications. In what follows, we present the results of applying the proposed method on simulated data and typical results that were achieved in several diverse applications, for which, conventional denoising methods did not yield satisfying results.

9.4 Applications

9.4.1 Simulated Data

The application domain of the proposed method is rather vast, ranging from low amplitude EEG to nonstationary EMG. In these applications, the SNR and spectral color of the *target signal* are very different. Therefore, in order to show the applicability of the proposed method in these applications, we can use spectrally colored signals to simulate arbitrary signals representing the EEG or EMG. Then by diluting this target signal with ECGs (or simulated ECGs) in different strengths, noisy signals are achieved that may be used for performance studies. For illustration, in Fig. 9.2(a) a thirty second segment of an arbitrary signal with a frequency range below 25Hz and a $1/f$ spectral shape is depicted. This signal was diluted with a rather clean ECG, such that its SNR reduced to 7.6dB. The resultant signal, depicted in Fig. 9.2(b), was then denoised by the EKS. The denoised signal can be seen in Fig. 9.2(c). In this example, the output SNR was increased to 15.1dB using the proposed method. The periodicity measure (PM), defined in Section 8.4.2, for the original, noisy, and denoised signals were 0.095, 0.179, and 0.074, respectively.

For a consistent quantitative study, *Monte Carlo* simulations were carried out with different input signals and by sweeping the SNR and spectral color of the target signal over the entire range of their practical values. There are different ways of generating spectrally colored signals. Here we adopt the approach that we used in Section 5.4.2, i.e., we model the signal color by a single parameter β representing the slope of a spectral density function that decreases monotonically with frequency.

On the other hand, for simulating the cardiac artifacts, the MIT-BIH normal sinus rhythm database was used [72, 160]. This database has a sampling rate of 128Hz. From this database, 20 low-noise segments of 30 seconds ECG without considerable artifacts were visually selected from different channels. Next,

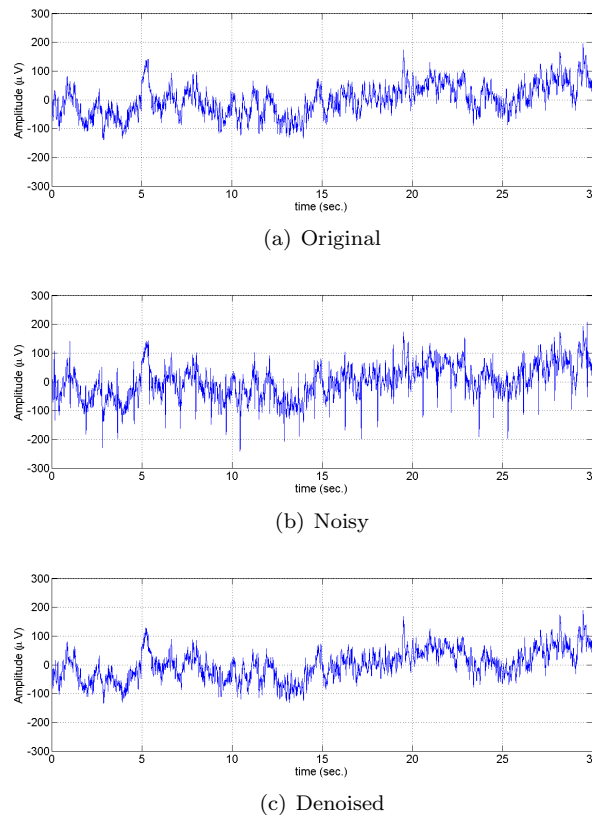


Figure 9.2: Results of the EKS on a mixture of simulated signal plus ECG artifacts. The (a) original, (b) noisy, and (c) denoised signals can be seen in this figure. The SNR of the noisy signal was improved from 7.6dB to 15.1dB using the EKS.

normalized portions of these segments were added to the simulated target signals to achieve simulated signals in different SNRs. These signals were filtered by the proposed method. In high SNRs, in which the ECG peaks were not detectable from the noisy signal, the required R-peak detection was performed on the original ECG, which is equivalent to using an arbitrary reference ECG channel, as shown in Fig. 9.1.

The overall simulation procedure was repeated 10 times for each of the 20 ECG segments, in different SNRs and with different spectrally colored target signals. The input SNR was swept in the range of -20dB to 20dB, in 5dB steps, and the signal spectral parameter β was swept in the range of 0 (white noise) to 2.5 (beyond brown noise), in 0.2 steps. This range of parameters is believed to be sufficient for simulating practical scenarios².

The SNR improvement achieved by the proposed filtering procedure was finally averaged over the 200 results (20 ECG segments \times 10 trials). The result of this study is depicted in Fig. 9.3. This result shows an SNR improvement surface, continuously changing according to the input SNR and spectral color. It is seen that the SNR improvement is generally better in low SNRs than in high SNRs. This was already expected, since in high input SNRs there is not very much of CC to be removed by the filter. The SNR improvement also generally increases as the signal spectrum becomes more colored. This is partially due to the fact that a high spectral color corresponds to a high β value in (5.16), representing a narrow-band signal. According to the filtering scheme in Fig. 9.1, such components are not affected by the KF, as they are bypassed by the BW removal block. However, in Fig. 9.3 it is interesting to see that in input SNRs above 5dB, a minimum SNR improvement is achieved for $\beta \approx 1$, i.e., for a target signal having a pink spectrum. This can be explained by considering the KF as an adaptive filter that adaptively changes its input-output transfer function in the frequency domain, to separate the signal from noise. This separation becomes more difficult, when the target signal and the CC have similar spectral shapes. Therefore, the minimum performance around $\beta = 1$ indicates that the CC used in our Monte Carlo simulations were spectrally closer (on average) to a pink spectrum than other spectral shapes. Therefore, the filtering of the CC has been most difficult for $\beta \approx 1$, leading to smaller SNR improvements at the filter output.

Practical target signals do not usually have a monotonically decreasing spectrum. However, from the Monte Carlo simulations in Fig. 9.3, the approximate performance of the proposed method may be roughly predicted for different applications, depending on the SNR and spectral color of the target signal.

The PM, defined in (8.18), was also calculated for the BW removed simulated data, before and after the filtering procedure. As explained before, an effective CC removal should decrease PM (although not being a sufficient measure of performance). This can be seen in the results shown in Fig. 9.4. Note however that from Fig. 9.4(a) it is seen that in high input SNRs and low spectral colors (target signal close to white noise), the input PM is already rather low and the filter does not decrease it more. Nevertheless, the output PM is always less than the input PM.

It should also be noted that results of Figs. 9.3 and 9.4 are an underestimate of the actual performance, since for these simulations the KF parameters were all selected in an automated way from the noisy signals, while in real applications (especially in offline processing), we can always fine-tune the parameters for the specific signal.




Figure 9.3: SNR improvement results achieved on simulated data in different input SNRs and different spectral colors.

²Refer to [127] for some typical spectral curves.




Figure 9.4: Periodicity measures (PM) achieved on simulated data in different input SNRs and different spectral colors, (a) before and (b) after denoising. The decrease of PM is used as a necessary, but not sufficient, measure of the filtering performance.

9.4.2 Single-Channel Fetal ECG Extraction and HRV Analysis

The first example is the six channel maternal abdominal signals used in the second example of Section 8.4.2. As mentioned there, this dataset was a degenerate mixture of maternal and fetal signals, and the fetal components were not separable from the maternal components with typical linear multichannel analysis methods. In Section 8.4.2 we used the deflation method for retrieving the fetal components. However, looking at the first channel of this data depicted in Fig. 9.5, we notice that the fetal SINR is rather high. In this figure, the R-peaks can be clearly seen in this channel with positive peaks while the maternal signal appears as negative peaks. Therefore, for HRV analysis that only requires the RR-interval sequence (and not the ECG morphology), we might achieve in extracting the fetal signals by using the CC cancellation method, for maternal ECG estimation and removal. After that due to the high SNR of the fetal components, the R-peaks can be detected from the resultant signal. As a preliminary step, the baseline wander of the signal was removed by using a two-stage moving median filter with the window lengths of 200ms and 600ms. In Appendix D, we will show that this method of baseline wander removal is rather robust and effective. Next, the CC cancellation procedure of Section 9.2, was applied on 100s of the preprocessed data to remove the maternal ECG. The fetal ECG found by this method can be seen in Fig. 9.6. We can see that the maternal ECG has been effectively removed, leaving behind the fetal ECG, which may be used for HRV analysis.

In order to extract the HRV of the fetus, the fetal R-peaks were detected with a robust R-peak detector based on *matched filtering* (cf. Appendix D). The original 100s data, the fetal signals after maternal ECG removal using the EKS, and the fetal HRV are depicted in Fig. 9.7. From this figure, we can notice some misdetections of the fetal R-peaks, which have led to errors in the fetal HRV sequence. These cases correspond to the very noisy epochs of the data and also to the time instants in which the maternal and fetal waves have coincided with each other and the fetal components have been removed as part of the maternal ECG. It is rather difficult to prevent these minor cases in a long set of single-channel recording; but we can do so using multichannel recordings with multichannel methods proposed




Figure 9.5: A typical segment of channel 1 of the dataset, including the fetal and maternal ECG




Figure 9.6: The signal of Fig. 9.5, after maternal ECG removal using the EKS

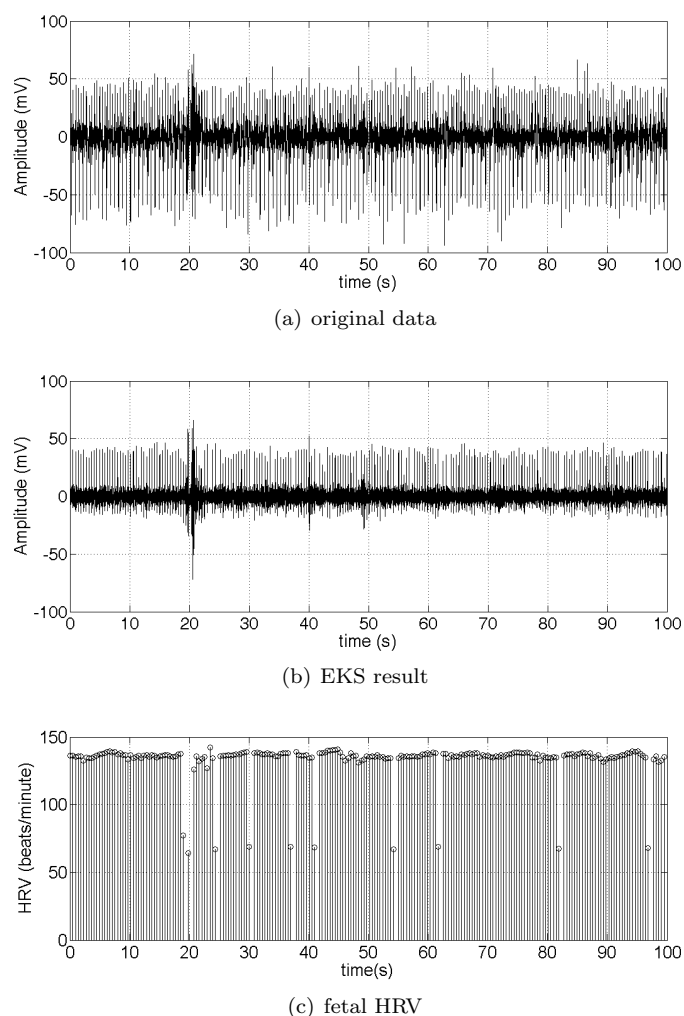


Figure 9.7: (a) A segment of the original data, (b) the fetal signals after mECG removal, and (c) the estimated fetal HRV. We can see that the fetal R-peaks have been misdetections at some points and have caused errors in the HRV sequence.

in previous chapters. In this example, the average of the fetal heartbeat, discarding the misdetections, is about 136 beats/minute.

9.4.3 EEG Denoising

The effective removal of cardiac artifacts such as the ECG, MCG, and the BCG artifacts from EEG and MEG recordings, remains a challenging issue. The conventional method of CC cancellation, is

to subtract the *ensemble average* (EA) of the CC, either directly or adaptively from the noisy brain recordings [194, 5, 4]. However, EA subtraction is not enough because it does not account for inter-beat variations of the cardiac waveforms. More effective means of EEG denoising are KF-based approaches that use additional channels such as the electrooculogram (EOG) [92], or a motion channel attached to the temporal artery [21], for recording the BCG and motion artifacts. In these works, the CC of the EEG has been approximated by a linear mixture of the EOG or an artery channel, and the weights of the mixture have been estimated by a KF. The drawback of these approaches is that they require additional EOG or reference artery channels. Moreover, due to the three-dimensional propagation of the cardiac potentials, the CC recorded from different leads can not be reconstructed from a single reference channel, unless the reference channel is recorded from a location close to the distorted channel, or alternatively, multiple reference channels are used. Both of these methods are practically limiting and do not fully remove the artifacts. A comparative study of the different methods may be found in [78]. Other methods, based on wavelet transforms and nonlinear noise reduction [216], and ICA [189, 128], have also been proposed.

With this background, the method explained in previous sections was used for removing CC from EEG signals recorded during an fMRI experiment. In this case, the CC of the EEG can be considered as a mixture of ECG and BCG signals. The presented method may be considered as an extension of the work presented in [194, 92, 21]. The advantage of the proposed method over these previous approaches is that the overall shape of the contaminating ECG plus BCG is extracted from the noisy EEG channel itself, without the need of any reference BCG channels or lead projections. Next, by using the extracted waveform, the artifacts are removed by an EKF or EKS. Note that in the proposed method, a single ECG channel is also required for R-peak detection, which can be simultaneously recorded from any arbitrary ECG lead.

This idea was tested on a dataset consisting of an EEG channel recorded from the standard AF₈ lead of an awake subject in rest condition with closed eyes, during an fMRI experiment³. An arbitrary ECG channel was also simultaneously recorded from the subject. The sampling rate of the dataset was 250Hz. Visually inspecting the data, the EEG appeared to have considerable artifacts that were temporally synchronous with the ECG beats. These artifacts are clearly seen in the EEG channel depicted in Fig. 9.9(b). In order to remove the CC, the single ECG channel was used to localize the artifacts within the EEG recordings. For this, the baseline wander of the raw data was removed and the R-peaks of the ECG were detected from the ECG channel. In this example, a two-stage median filter, with 200ms and 600ms window lengths, was used for baseline wander removal. Next, by using the R-peak information, the baseline removed EEG channels were synchronously averaged over the ECG period and the average and standard deviation of the ECG beat template were extracted. This gave the average shape of the CC (ECG plus BCG) that were contaminated over the EEG. The resultant mean and standard deviation bars can be seen in Fig. 9.8, for the reference ECG channel and the EEG before and after filtering. As we see, the CC of the EEG in Fig. 9.8(b), seem to be smoothed by the fMRI magnetic field and do not resemble the reference ECG in Fig. 9.8(a). This means that conventional adaptive noise cancellation ideas that remove linear mixtures of three orthogonal ECG leads from the EEG channels are not applicable to this data.

Next, using the average CC template, the parameters of the Gaussian mixtures required for the EKS were extracted according to the steps described in Chapter 5. Finally, by applying the EKS to the noisy EEG recording, the CC was estimated from the background EEG. The noise free EEG signal was achieved by subtracting the estimated CC from the original noisy recording. In Fig. 9.9, a typical segment of the resultant EEG is plotted versus the original AF₈ channel. As it can be seen, the EKS has effectively removed the CC from the EEG channel. The PM defined in (8.18), was 0.42 for the original AF₈ channel, 0.07 using conventional EA subtraction, and 0.04 after applying the proposed method. Even though the PM values of conventional EA subtraction and the proposed method are rather close, the proposed method was found to be more robust to inter-beat deviations of the cardiac signals.

The spectral density functions of the original and denoised EEG signals are depicted in Fig. 9.10. As it is seen, the filtering procedure has changed the spectrum in the *theta* and *alpha* bands, which are important frequency bands of the EEG spectrum. This implies that spectral analysis of the EEG without CC removal can be rather misleading.

Another example is presented in Fig. 9.11, for an EEG segment of the MIT-BIH polysomnographic

³ The EEG recordings used in this section have been provided by Dr. Christophe Phillips from the Cyclotron Research Centre, Liege, Belgium.

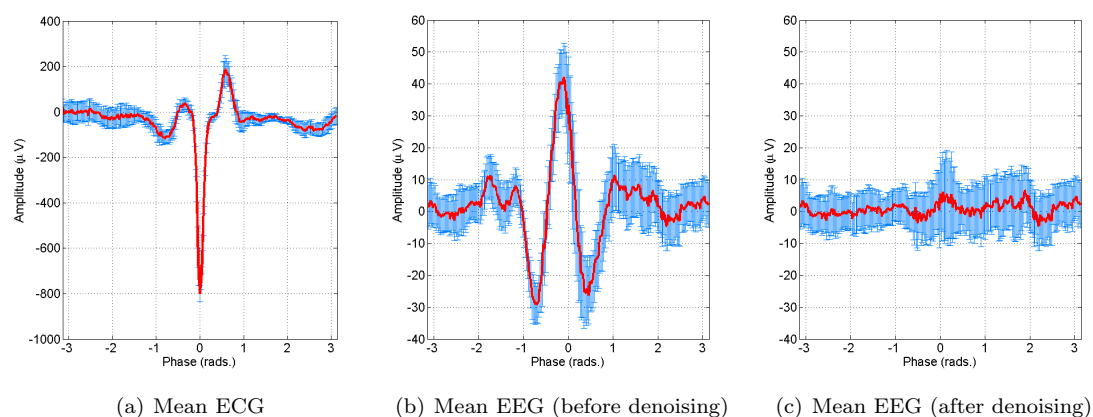


Figure 9.8: The average and SD-bars of the (a) ECG, (b) EEG before denoising, and (c) EEG after denoising, by synchronous averaging of the data using the ECG R-peaks.

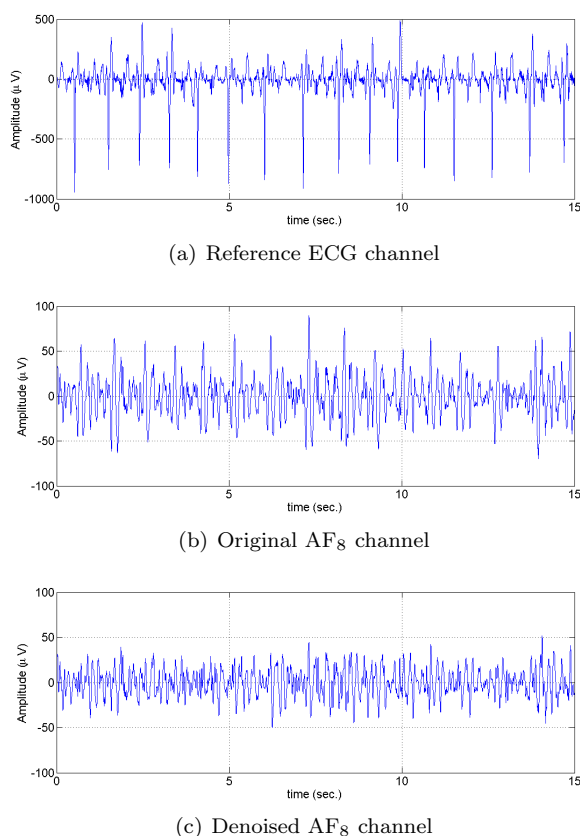


Figure 9.9: Results of the proposed method on a segment of EEG signals recorded during an fMRI experiment. (a) Reference ECG channel, (b) Noisy AF₈ EEG channel, (c) AF₈ channel after ECG removal.

database, recorded during sleep from the C₄-A₁ channel, with a sampling rate of 250Hz[161]. From this figure, we can again notice that the CC peaks that were synchronous with the ECG R-peaks have been effectively removed, while the non-ECG contents and the alpha rhythms have been preserved. The PM of the noisy and denoised signals were 0.15 and 0.001, respectively. The noisy and denoised signal spectra are also compared in Fig. 9.12, where we can see that the signal spectrum has been considerably changed after CC removal, especially in the beta band and higher frequencies.







Figure 9.10: The spectral estimate of the original and denoised EEG signals, from the fMRI experiment data. The difference of the two spectra is plotted at the bottom.



(a) Reference ECG channel



(b) Original EEG channel



(c) Denoised EEG channel

Figure 9.11: Results of the proposed method on a segment of EEG signals from the MIT-BIH Polysomnographic Database. (a) Reference ECG channel, (b) Noisy C_4-A_1 EEG channel, (c) C_4-A_1 channel after ECG removal.

9.4.4 EMG Signal Denoising

Another tested application for the proposed method was the removal of ECG contaminants from diaphragmatic EMG signals recorded from an *intraoesophageal* electrode⁴. In this application we were interested in the extraction of diaphragmatic EMG bursts synchronous with the respiration. The exact detection of the beginning and ending points of the EMG burst are widely used in respiratory studies. However, the recorded EMG are usually highly contaminated with ECG. The conventional method for

⁴ The EMG recordings used in this section have been provided by Dr. Vincent Vigneron from the Laboratory of Informatique, Biologie Intégrative et Systèmes Complexes (IBISC), CNRS FRE 2873, Evry, France.




Figure 9.12: The spectra of the original and denoised EEG signals, from the MIT-BIH Polysomnographic Database. The difference of the two spectra is plotted at the bottom.

removing such artifacts is to detect the QRS-segments of the ECG and to remove their EA or to directly set them to zero. This however removes considerable portions of the EMG signal and causes spurs in the EMG spectrum.

For this example, the procedure was similar to the previous example. Since the CC was strong, the R-peaks of the ECG were directly estimated from the noisy recording, from which the ECG EA was calculated. Next, the dynamic model parameters were selected according to the EA of the ECG artifacts and the EKS was applied to the noisy recordings. After estimating the ECG and removing it from the original recordings, the residual signal contains the desired diaphragmatic EMG. For this application, due to the nonstationary nature of EMG bursts, the adaptive modification of the measurement noise variance proposed in Section 5.3.4, was used. This procedure was based on the monitoring of the KF innovation signal, allowing the filter to adapt the KF measurement noise variance R_k , according to the noise contents of the signal, such that the innovation signal becomes spectrally white⁵.

An example of an EMG burst signal denoised with this approach is depicted in Fig. 9.13. In this example, the PM before and after ECG cancellation were 0.68 and 0.002, respectively. The resultant signal may be further used for EMG burst analysis and depending on the quality of the recorded signals, additional post-filtering may be required to reject the out-of-band noise.

9.5 Summary and Conclusions

In this chapter, a model-based Bayesian filtering framework was presented for removing cardiac contaminants from various biomedical signals. It was shown that the method is applicable to as few as single channel recordings with an arbitrary reference ECG/MCG channel for R-peak detection.

In summary, the general idea is to use *a priori* knowledge of the pseudo-periodic structure of ECG, MCG, or BCG signals to separate the desired and undesired parts of the observations. This method is hence using the temporal and frequency domain information about the CC. Nevertheless, when multi-channel recordings are available, one can use the additional information provided by the spatial diversity of the sensors, which is the essence of spatial filtering methods presented in previous chapters. From this point of view, this filtering framework can be integrated within the denoising block of the deflation procedure developed in Chapter 8.

For the sake of brevity, the presented results were only based on the EKS. However, other types of Bayesian filters such as the UKF and the *particle filter* (PF) can be used in the same manner for highly nonlinear and non-Gaussian noise scenarios [84].

Due to the recursive structure of the KF, the proposed method is also computationally efficient and of special interest for real-time applications. Generally, the computation time of this method is linearly proportional to the signal length in samples. For the currently developed Matlab[®] source codes (available at [167]) the computation time is already close to real-time on a 3GHz CPU for signals with a sampling rate of up-to 1kHz, except for the CC template fitting step of the algorithm that is carried out by the user

⁵Note that in the KF context, the whiteness of the innovation signal implies that all the information concerning the estimated process, up to the second order statistics, have been extracted from the observation signal.

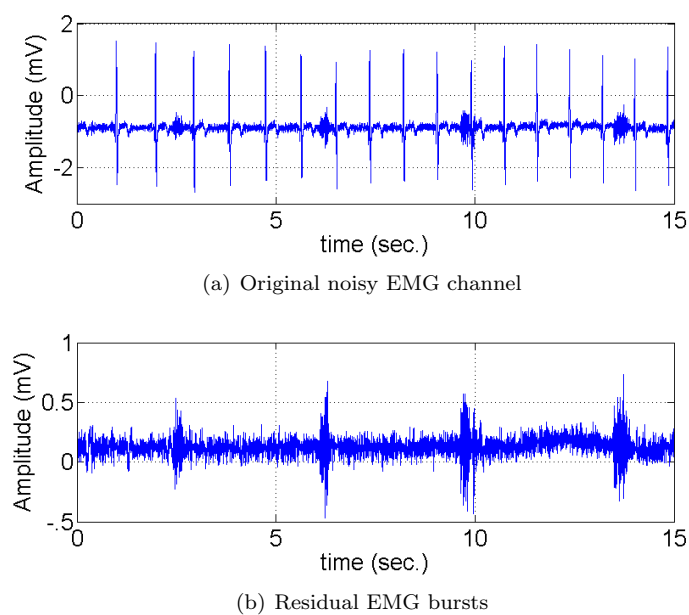


Figure 9.13: Results of the proposed method on EMG recordings highly contaminated with ECG.

through an interactive graphical user interface, which allows the user to adjust the number and locations of the Gaussian kernels used for fitting the EA of the cardiac waveforms. However, these Matlab[®] codes may be further optimized and converted into low-level languages for being used in pre-processing units of clinical monitoring systems.

Chapter 10

Conclusions and Future Works

In this research the problem of extracting fetal cardiac signals from an array of maternal abdominal recordings was studied. Various methods were proposed for this problem and were evaluated on real and synthetic signals using realistic fetal ECG models. Further issues concerning the interpretation and dimensionality of these signals were also studied. Considering the very low SNR of fetal cardiac signals, the main contribution of this work was to develop single and multi-channel signal processing tools that use *a priori* information of different natures, about the fetal and contaminating signals, to improve the quality of fetal cardiac signal extraction.

More specifically, the method developed in Chapter 5 is a novel approach that utilizes the temporal dynamics of cardiac signals within a Bayesian filtering framework. Therefore, considering the advances in Bayesian methods, we believe that this framework can serve as a basis for utilizing Bayesian techniques for cardiac signal processing and analysis.

It is generally known that semi-blind source separation methods based on *a priori* information, are more effective than totally ‘blind’ methods. Therefore, a challenging (and *ad hoc*) issue in the source separation context, is how to use the existing *a priori* information to solve specific problems. The periodic component analysis technique presented in Chapters 7 and 9, is one such method that uses the pseudo-periodicity of cardiac signals for their separation.

The notion of distributed sources, as compared with point or vectorial sources, presented in Chapter 6, and the interpretation of components extracted by linear decomposition methods, is also an important contribution of this study. In this context, we should also consider the necessity of preprocessing, e.g., baseline wander removal and its impact on the extracted components.

Although linear decomposition techniques are commonly preferred over the nonlinear ones (mainly due to their relative simplicity), issues such as signal mixture degeneracy and noise, limit the performance of these methods. The deflation procedure developed in Chapter 8, was a technique that combined the benefits of linear and nonlinear filtering for dealing with the problem of data degeneracy (through appropriate use of the signal priors), and without losing the data dimensions.

Here we add that the methods proposed in previous chapters should not be envisaged as a replacement, but rather as complements for the existing methods. In fact, due to the various objectives, measurement setups, fetal conditions and gestational ages, SNR, etc., it is not reasonable (nor possible) to present a one-for-all universal filtering solution. What is more feasible is to focus on specific applications, such as fetal R-peak detection using a fixed electrode configuration and perhaps for specific ranges of gestational age. This approach results in methods and tools that can be used for developing monitoring systems with specific applications. An alternative approach, which was adopted in this project, was to try to develop a toolbox by which a set of methods can be rearranged for doing a certain job, e.g. maternal ECG removal or fetal ECG enhancement, with a minimal interaction with an expert operator. In this approach, for each case, depending on the aforementioned variable factors, a signal processing scenario can be setup that benefits from a set of algorithms for achieving the desired objective. For instance, if we are only interested in fetal heart rate variability analysis and the recorded data is rather clean, a simple maternal ECG canceller followed by a peak-detector may be satisfying. However, for morphologic studies in low SNRs a combination of decomposition and filtering methods may be required. It is therefore important to know the strength and weak-points of each method. In Table 10.1 (page 118), a general comparison is made between the previously existing and the proposed methods. Accordingly, each method has its own

benefits and limitations and is applicable for specific scenarios.

We believe that the proposed methods are specifically powerful in the following cases:

- Maternal cardiac interference cancellation using any electrode configuration from single or multiple channels with minimal influence on fetal components; as long as the maternal R-peaks are detectable either from the noisy data itself or any arbitrary reference channel.
- Fetal cardiac signal enhancement using the fetal cardiac peaks.
- Decomposition of a desired subspace of cardiac signals from undesired signal subspaces. The subspace dimensions need not be known *a priori* nor fixed, and the signal mixture can be degenerate.
- More generally, the denoising and enhancement of any class of signals characterized by a pseudo-periodic structure.
- A general framework for the decomposition of multichannel mixtures of signals and noises, based on *a priori* information concerning the signal and noise structure.

As we see the methods that we developed for cardiac signals are highly based on the effective use of prior information, especially the pseudo-periodicity of cardiac signals. Although this is an important factor that can improve the signal quality, it can also be considered as a point of weakness for the cases that:

- The R-peaks (either for the mother or the fetus) are not well-detected.
- The cases in which the assumption of pseudo-periodicity does not fully hold. For example, for some pathologic cases such as Premature Ventricular Contractions, where the ECG/MCG morphology are not consistent, there can be considerable morphologic variations from beat-to-beat. Therefore, methods such as the π CA or decomposition by deflation will not be very effective in these cases. We can nevertheless use these methods for the cases in which only one of the maternal or fetal ECG/MCG are abnormal. We have a similar limitation for the Bayesian filtering framework, except that in this case the statistical variations of the filter's *innovation* signal can be used as an indication of signal abnormality.

Despite these limitations, it may however be argued that no filtering technique is fully blind. In fact, any other filtering procedure is also using some *a priori* model or other prior information from the time, frequency, space, or feature domain, and up to now no general filtering procedure has been developed which can be universally applied to normal and abnormal cardiac signals.

10.1 Future Works

Eternally, there are infinite numbers of questions to answer and works to do.... In the following, some of the questions and possible future fields of study concerning the extraction and analysis of fetal cardiac signals are noted.

- All practical blind source separation problems contain portions of full-rank or non-full rank noise. Therefore, from the estimation theoretical perspective, there is always a theoretical limit on the accuracy of the estimated mixing (demixing) matrices and the estimated signals. In previous research, this issue has been studied in the context of ICA, by using Cramér-Rao bounds [32, 157, 27, 225, 109, 200, 50]. These ideas can be extended to the calculation of similar theoretical bounds for other decomposition techniques and also for the estimated signals. Specifically, for ECG signals, we believe that the dynamic model presented in Chapter 5 can be used as an appropriate tool for this purpose. The calculation of such theoretical bounds can also be helpful for finding theoretical benchmarks for the performance of different ECG filtering techniques.
- Linear decomposition techniques are very common for blind and semi-blind source separation problems. In these methods, some measure of signal separability, such as non-gaussianity, second or higher-order statistics, periodicity, spectral contrast, or etc. is maximized to achieve the desired transform. However, using many of these methods we frequently find very similar solutions regardless of the cost-function that was maximized. These similarities are rather common between

a broad range of linear transforms and seem to be due to an intrinsic behavior of linear models themselves, which impose a very strong constraint on the data model¹. It is therefore interesting to study such common characteristics to find intrinsic properties of linear and nonlinear decomposition techniques, regardless of the type of criteria that are maximized.

- The modeling of abnormal ECG signals (either for adults or fetuses) is another issue of interest. The dynamic models presented in Chapter 4 can effectively generate multichannel normal or abnormal beats with consistent morphology. However, cardiac abnormalities usually appear as occasional odd-beats between a set of normal ones. Therefore, a smooth transition between normal and abnormal beats should also be considered in our morphologic models. We believe that the Hidden-Markov Model (HMM) is an effective means for this purpose. A preliminary idea of this extension was recently presented in [40]. The HMM-based extension of the ECG generator can also be used within the Bayesian filtering framework of Chapter 5 for, at the same time, denoising and detecting abnormal ECGs.
- The clinical validation of the presented methods should also be considered in future works. In fact, the proposed methods were presented as general ECG processing tools and were validated on discrete databases, each having a different recording protocol. The methods were also validated over realistic simulated data. However, due to the lack of a unique fetal database recorded at different gestational ages and from various subjects, the presented methods have not yet passed standard clinical validation procedures. This is an essential step that should be taken before using any of the proposed methods in clinical monitoring systems.
- In Chapter 7, a primary attempt was made to track fetal cardiac signals in long recordings and to compensate for the fetal movements. This is a topic that requires further studies in future work. One of the byproducts of this study would be the development of a canonical fetal ECG representation with respect to the fetal or maternal body axis. Such a representation is believed to be of significant importance for clinical morphologic studies.

¹This idea originated from one of our early discussions with Dr. Vincent Vigneron from the Laboratory of Informatique, Biologie Intégrative et Systèmes Complexes (IBISC), CNRS FRE 2873, Evry, France.

Table 10.1: Comparison of the different methods and their applications

Method	Overall Performance	SNR improvement	Computational Cost	Realtime*	Implementation Complexity	Operator interaction	Applications/Comments
FIR/IIR Filtering	Low	Low	Low	Yes	Simple	Not required	Simple frequency domain denoising
Wiener Filtering	Medium	Medium	Medium	Possible	Medium	Not required	Denoising using morphologic or spectral priors; realtime implementation is only possible for the causal Wiener filter
Wavelet Denoising	Medium	Medium	Low	Yes	Medium	Not required	More robust denoising for signal-noise mixtures having different scales
Adaptive Filtering	Medium	Medium	Median	Yes	Medium	Sometimes required	Low quality maternal ECG cancellation or fetal ECG enhancement for fetal HRV analysis or average morphologic studies; extendable to multichannels
Nonlinear Filtering	Medium	Medium/High	High	No	Complex	Required	Maternal ECG cancellation or fetal ECG enhancement in single channel recordings or degenerate multichannel mixtures
Bayesian Filtering	High	High	Median	Yes	Complex	Sometimes required	ECG denoising, maternal ECG removal or fetal ECG enhancement having the R-peaks of the signal and with rather consistent morphologic shape; extendable to multichannels
PCA/SVD/ Factor Analysis	Low	Low	Low	No	Simple	Not required	Dimension reduction for high-dimensional data and noise removal
ICA	Medium	Medium	Medium/High	No	Medium/Complex	Required	Blind or semi-blind decomposition of multichannel maternal-fetal mixtures without considerable <i>a priori</i> information and in relatively high SNRs SNR improvement not guaranteed
π CA	Medium	High	Low	No	Simple	Not required	Decomposition of multichannel ECG mixtures having the R-peaks of the signal; dimension reduction with minimal ECG information loss
Deflation	High	High	Medium/High	No	Complex	Sometimes required	Decomposition of (possibly) degenerate multichannel mixtures of desired and undesired subspaces having <i>a priori</i> information about the signal-noise structures

* Here we have not considered the possibility of block delayed realtime analysis.
 ** The indicated performances and complexities of multichannel methods are generally higher than single-channel methods; the qualitative descriptions should only be compared within each group separately.

List of Related Publications

Journals

1. R. Sameni, G.D. Clifford, C. Jutten, and M.B. Shamsollahi, Multichannel ECG and Noise Modeling: Application to Maternal and Fetal ECG Signals, *EURASIP Journal on Advances in Signal Processing*, Volume 2007, Article ID 43407, doi:10.1155/2007/43407. (**Chapter 4**)
2. R. Sameni, M.B. Shamsollahi, C. Jutten, and G.D. Clifford, A Nonlinear Bayesian Filtering Framework for ECG Denoising, *IEEE Transactions on Biomedical Engineering*, Vol. 54, No. 12, pp. 2172–85, Dec. 2007. (**Chapter 5**)
3. R. Sameni, C. Jutten, and M.B. Shamsollahi, Multichannel Electrocardiogram Decomposition using Periodic Component Analysis, *IEEE Transactions on Biomedical Engineering*, Vol. 55, No. 8, pp. 1935–1940, Aug. 2008. (**Chapter 7**)
4. T. Tsalaila, R. Sameni, S. Sanei, C. Jutten, and J. Chambers, Sequential Blind Source Extraction for Quasi-Periodic Signals with Time-Varying Period, [to appear in] *IEEE Transactions on Biomedical Engineering*, 2009. (**Chapter 7**)
5. R. Sameni, M.B. Shamsollahi, and C. Jutten, Model-Based Bayesian Filtering of Cardiac Contaminants from Biomedical Recordings, *Physiological Measurements*, Vol. 29, No. 5, pp. 595–613, May 2008. (**Chapter 9**)

Conferences

1. R. Sameni, M.B. Shamsollahi, and C. Jutten, Filtering Electrocardiogram Signals Using the Extended Kalman Filter, *Proc. of the 27th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBS), Shanghai, China*, Sep. 1-4, 2005, pp. 5639–5642. (**Chapter 5**)
2. R. Sameni, M.B. Shamsollahi, and C. Jutten, Filtering Noisy ECG Signals Using the Extended Kalman Filter Based on a Dynamic ECG Model, *Proc. of the 32nd Annual International Conference on Computers in Cardiology, Lyon, France*, Sep. 25-28, 2005, vol. 32, pp. 1017–1020. (**Chapter 5**)
3. R. Sameni, M.B. Shamsollahi, and C. Jutten, Multi-Channel Electrocardiogram Denoising Using a Bayesian Filtering Framework, *Proc. of the 33rd Annual International Conference on Computers in Cardiology, Valencia, Spain*, Sep. 17-20, 2006., vol. 33, pp. 185–188. (**Chapter 5**)
4. R. Sameni, C. Jutten, and M. B. Shamsollahi, What ICA Provides for ECG Processing: Application to Noninvasive Fetal ECG Extraction, *Proc. of the International Symposium on Signal Processing and Information Technology (ISSPIT'06), Vancouver, Canada*, Aug. 27-30, 2006, pp. 656–661. (**Chapter 6**)
5. R. Sameni, F. Vrins, F. Parmentier, C. Hérail, V. Vigneron, M. Verleysen, C. Jutten, and M.B. Shamsollahi, Electrode Selection for Noninvasive Fetal Electrocardiogram Extraction using Mutual Information Criteria, *Proc. of the 26th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering (MaxEnt 2006), CNRS, Paris, France*, Jul. 8-13, 2006, vol. 872, pp. 97–104. (**Chapter 6**)

6. C. Jutten, R. Sameni, and H. Hauksdóttir. On the Relevance of Independent Components. *In Proc. of the ICA Research Network International Workshop (ICArn 2006), Liverpool, UK*, September 18-19 2006, pp. 1-8, [Invited paper]. **(Chapter 6)**
7. L. Amini, R. Sameni, C. Jutten, G.A. Hossein-Zadeh and H. Soltanian-Zadeh, MR Artifact Reduction in the Simultaneous Acquisition of EEG and fMRI of Epileptic Patients, *Proc. of the 16th European Signal Processing Conference (EUSIPCO2008), Lausanne, Switzerland*, Aug. 25-29, 2008 [to appear]. **(Chapter 8)**
8. G.D. Clifford and R. Sameni, An Artificial Multi-Channel Model for Generating Abnormal Electrocardiographic Rhythms, *The 35th Annual International Conference on Computers in Cardiology, Bologna, Italy*, Sep. 14-17, 2008 [abstract accepted]. **(Chapter 4)**

Appendix A

Angles Between Subspaces

Let F and G be subspaces in \mathbb{R}^m whose dimensions satisfy¹:

$$p = \dim(F) \geq \dim(G) = q \geq 1 \quad (\text{A.1})$$

The *principle angles* $\theta_1, \dots, \theta_q \in [0, \pi/2]$ between F and G are defined recursively as follows:

$$\cos(\theta_k) = \max_{u \in F} \max_{v \in G} u^T v = u_k^T v_k \quad (\text{A.2})$$

subject to:

$$\begin{aligned} \|u\| &= \|v\| = 1 \\ u^T u_i &= 0 \quad , \quad i = 1 : k-1 \\ v^T v_j &= 0 \quad , \quad j = 1 : k-1 \end{aligned} \quad (\text{A.3})$$

The principal angles satisfy $0 \leq \theta_1 \leq \dots \leq \theta_q \leq \pi/2$ and the vectors $\{u_1, \dots, u_q\}$ and $\{v_1, \dots, v_q\}$ are called the *principal vectors* between F and G .

Principal angles and vectors arise in many statistical applications. The largest principal angle is related to the notion of *distance* between equidimensional subspaces. Specifically, if $p = q$ then

$$\text{dist}(F, G) \doteq \|F - G\|_2 = \sqrt{1 - \cos(\theta_p)^2} = \sin(\theta_p) \quad (\text{A.4})$$

With this definition, given two matrices $A \in \mathbb{R}^{m \times p}$ and $B \in \mathbb{R}^{m \times q}$ ($p \geq q$) each with linearly independent columns, the following algorithm computes the orthogonal matrices $U = [u_1, \dots, u_q]$ and $V = [v_1, \dots, v_q]$ and $\cos(\theta_1), \dots, \cos(\theta_q)$ such that the θ_k are the principal angles between $\text{ran}(A)$ and $\text{ran}(B)$ and u_k and v_k are the associated principal vectors:

1. Compute the QR factorizations of the matrices A and B ([73] Chapter 5):

$$\begin{aligned} A &= Q_A R_A & Q_A^T Q_A &= I_p \quad , \quad R_A \in \mathbb{R}^{p \times p} \\ B &= Q_B R_B & Q_B^T Q_B &= I_q \quad , \quad R_B \in \mathbb{R}^{q \times q} \end{aligned}$$

2. $C = Q_A^T Q_B$
3. Compute the singular value decomposition (SVD) of C : $Y^T C Z = \text{diag}(\cos(\theta_k))$
4. $Q_A Y(:, 1 : q) = [u_1, \dots, u_q]$
5. $Q_B Z = [v_1, \dots, v_q]$

¹ This appendix has been adopted from [73, ch. 12].

Appendix B

On the Spikes and Bumps Phenomenon

In Chapter 6, an over-fitting phenomenon known as *spikes and bumps* was introduced, which practically occurs when applying many of the conventional ICA methods to multichannel data having many noisy channels. This phenomenon was originally introduced in [90, 179, 178], and it was justified that in the extreme case of having as many samples and sensors, sparse spiky signals can maximize a measure of kurtosis; which explains their extraction by marginal distribution-based ICA methods. In [179], some qualitative explanations have also been presented for bumpy components that appear as independent components, when applying ICA to high-dimensional signals with $1/f$ spectral shape.

In our experiments, we further noticed that this phenomenon can also happen when we apply the aforementioned ICA algorithms to spectrally colored noise (even without having any mixture of true physical sources), and the shapes and widths of the spikes and bumps highly depend on the spectral form of the input signals (or noises) and the number of data channels. We made some attempts to formulate these practical findings and to relate the exact shape of the bumps with the spectral shape and number of channels. In the following, we present some of our findings in this area. The problem is however open to further research in future works.

Consider that we have T samples of an N -channel real random process $\mathbf{x}[n] = [x_1[n], \dots, x_N[n]]^T$. Also assume that all the channels of this process have the same spectral shape $S_x(e^{j\omega})$. We can therefore represent each channel $x_i[n]$, as follows:

$$x_i[n] = w_i[n] * h[n] \quad (\text{B.1})$$

where $w_i[n]$ is a unit-variance white noise, $*$ represents convolution, and $h[n]$ is the *innovations filter* for the process $x_i[n]$ [152, ch. 12], such that

$$S_x(e^{j\omega}) = |H(e^{j\omega})|^2 \quad (\text{B.2})$$

where $H(e^{j\omega})$ is the discrete-time Fourier transform of $h[n]$. Following (B.2), we also have

$$h[n] * h[-n] = R_x[n] \quad (\text{B.3})$$

where $R_x[n]$ is the autocorrelation function of each channel of $\mathbf{x}[n]$. We are now interested in linear transforms of the form

$$y[n] = \sum_{i=1}^N b_i x_i[n] = \mathbf{b}^T \mathbf{x}[n] \quad (\text{B.4})$$

which maximize the following criterion:

$$C_0 = \frac{1}{T} \sum_{n=1}^T f(y[n]) \quad \text{subject to } \|\mathbf{b}\| = 1 \quad (\text{B.5})$$

where $\mathbf{b} = [b_1, \dots, b_N]^T$ and $f(\cdot)$ is some nonlinear *contrast function* (CF) that we want to maximize. Using a Lagrange multiplier λ , this constraint optimization problem may be represented as follows:

$$C = \frac{1}{T} \sum_{n=1}^T f(y[n]) + \lambda \left(\sum_{i=1}^N b_i^2 - 1 \right) \quad (\text{B.6})$$

By setting to zero the derivative of (B.6) with respect to b_p ($1 < p < N$), we have:

$$\frac{\partial C}{\partial b_p} = \frac{1}{T} \sum_{n=1}^T f'(y[n])x_p[n] + 2\lambda b_p = 0 \quad (\text{B.7})$$

where $f'(y) = \partial f / \partial y$. We further multiply both sides of (B.7) by b_p and take a summation over all $p = 1, \dots, N$, to find λ :

$$\lambda = \frac{-1}{2T} \sum_{n=1}^T f'(y[n])y[n] \quad (\text{B.8})$$

Combining (B.7) and (B.8) we find the b_p as follows:

$$b_p = \frac{\sum_{n=1}^T f'(y[n])x_p[n]}{\sum_{n=1}^T f'(y[n])y[n]} = \frac{\langle f'(y[n]), x_p[n] \rangle}{\langle f'(y[n]), y[n] \rangle} \quad (\text{B.9})$$

By multiplying both sides of (B.9) by $x_p[m]$ and taking the summation over $p = 1, \dots, N$, we can find a recursive function for calculating $y[n]$:

$$y[m] = \alpha \sum_{n=1}^T \sum_{p=1}^N f'(y[n])x_p[m]x_p[n] \quad (\text{B.10})$$

where $\alpha = \langle f'(y[n]), y[n] \rangle^{-1}$ is a scaling factor. Up to now, we have not used the spectral form of $\mathbf{x}[n]$. Following (B.1) we have:

$$\sum_{p=1}^N x_p[m]x_p[n] = \sum_{m'} \sum_{n'} h[m']h[n'] \sum_{p=1}^N w_p[m-m']w_p[n-n'] \quad (\text{B.11})$$

Considering that the $w_p[n]$ are unit variance white noise, the last term in (B.11) may be approximated by:

$$\sum_{p=1}^N w_p[m-m']w_p[n-n'] \approx \delta[m-n+n'-m'] \quad (\text{B.12})$$

where $\delta[\cdot]$ is the Dirac delta function. Therefore, (B.11) simplifies to:

$$\sum_{p=1}^N x_p[m]x_p[n] = \sum_{n'} h[m-n+n']h[n'] = h[u] * h[-u] \Big|_{u=m-n} = R_x[m-n] \quad (\text{B.13})$$

This can be used to further simplify (B.10):

$$y[m] = \alpha \sum_{n=1}^T f'(y[n])R_x[m-n] \quad (\text{B.14})$$

or

$$y[n] = \alpha f'(y[n]) * R_x[n] \quad (\text{B.15})$$

Equation (B.15), may also be represented in the frequency domain:

$$S_x(e^{j\omega}) = \frac{\mathcal{F}\{y[n]\}}{\alpha \mathcal{F}\{f'(y[n])\}} \quad (\text{B.16})$$

where $\mathcal{F}\{\cdot\}$ represents the discrete-time Fourier transform. As we see, (B.15) and (B.16) provide a compact form for relating the spectral information of the input signal $\mathbf{x}[n]$ with the output signals $y[n]$. An additional information that is given by (B.16) is that since $S_x(e^{j\omega})$ is real, $\mathcal{F}\{y[n]\}$ and $\mathcal{F}\{f'(y[n])\}$ have the same phase. However, to our knowledge, there is no closed form solution for these equations, unless perhaps for some special $f(\cdot)$ and $R_x[n]$. For instance, if we use $f(u) = u^4$, which leads to a kurtosis-based CF, (B.15) reduces to:

$$y[n] = \alpha' y[n]^3 * R_x[n] \quad (\text{B.17})$$

which belongs to the class of Volterra nonlinear equations with power nonlinearity kernels, which have some general asymptotic properties [140].

The hereby presented formulations are believed to provide a basis for finding better insights of the practical problem of spikes and bumps. In further studies, the problems with finite number of samples, the inaccuracies in the approximation in (B.12), and the necessity of adding additional orthogonality constraints for the different channels of $y[n]$, should also be considered.

Appendix C

Multichannel Decomposition of Distributed Sources

In Chapter 6, the relationship between distributed sources and principal and independent components extracted from multichannel observations of such sources was studied from a practical viewpoint. In this appendix, we will show this relationship in more details for a simplified case of a finite distributed source within a homogeneous volume conductor.

Consider a distributed source with a time varying stochastic density $\rho(\mathbf{x}, t)$ within a finite homogeneous volume conductor (Fig. C.1). The field potentials due to this distribution can be represented as follows:

$$\Phi(\mathbf{x}, t) = \int \frac{\rho(\mathbf{x}', t)}{|\mathbf{x} - \mathbf{x}'|^p} \mathbf{d}\mathbf{x}' \quad (\text{C.1})$$

where $\mathbf{x} = [x, y, z]^T$ is the Cartesian coordinates of the observation point¹, p is the decay exponent which depends on the nature of $\rho(\mathbf{x}, t)$, and the integral is taken over the volume containing the distributed source. We further assume that we make n such measurements (observations) from different locations, each time with respect to the reference point \mathbf{x}_r . Therefore, the potential difference of each point with respect to the reference can be written as follows:

$$\Delta\Phi(\mathbf{x}_i, t) = \Phi(\mathbf{x}_i, t) - \Phi(\mathbf{x}_r, t) = \int \rho(\mathbf{x}', t) \left[\frac{1}{|\mathbf{x}_i - \mathbf{x}'|^p} - \frac{1}{|\mathbf{x}_r - \mathbf{x}'|^p} \right] \mathbf{d}\mathbf{x}' \doteq \int \rho(\mathbf{x}', t) D(\mathbf{x}_i, \mathbf{x}') \mathbf{d}\mathbf{x}' \quad (\text{C.2})$$

By defining $\phi(t) = [\Delta\Phi(\mathbf{x}_1, t), \dots, \Delta\Phi(\mathbf{x}_n, t)]^T$ and $\mathbf{G}(\mathbf{x}') = [D(\mathbf{x}_1, \mathbf{x}'), \dots, D(\mathbf{x}_n, \mathbf{x}')]^T$, (C.2) can be represented in the following compact form:

$$\phi(t) = \int \rho(\mathbf{x}', t) \mathbf{G}(\mathbf{x}') \mathbf{d}\mathbf{x}' \quad (\text{C.3})$$

where $\phi(t)$ can be considered as a vector of differential observations (e.g. body surface potentials). Without loss of generality, we consider these observations to be zero-mean and calculate their covariance matrix as follows:

$$R_\phi = E_t\{\phi(t)\phi(t)^T\} = \iint \mathbf{G}(\mathbf{x}') E_t\{\rho(\mathbf{x}', t)\rho(\mathbf{y}', t)\} \mathbf{G}(\mathbf{y}')^T \mathbf{d}\mathbf{x}' \mathbf{d}\mathbf{y}' \quad (\text{C.4})$$

where $E_t\{\cdot\}$ represents averaging over time. We define $R_\rho(\mathbf{x}', \mathbf{y}') = E_t\{\rho(\mathbf{x}', t)\rho(\mathbf{y}', t)\}$ as the cross-correlations of the source densities of different points inside the distribution. Next, using the Karhunen-Loève transform [152, ch. 12], $R_\rho(\mathbf{x}', \mathbf{y}')$ may be expanded in terms of a unique set of orthogonal basis functions, namely *eigen-functions*:

$$\lambda_i f_i(\mathbf{u}) = \langle R_\rho(\mathbf{v}, \mathbf{u}), f_i(\mathbf{v}) \rangle, \quad \langle f_i(\mathbf{v}), f_j(\mathbf{v}) \rangle = \delta_{i-j} \quad (\text{C.5})$$

where the λ_i are the *eigenvalues* and δ_{i-j} is the Dirac delta function. Next, following Mercer's theorem, $R_\rho(\mathbf{x}', \mathbf{y}')$ can be expanded as follows:

$$R_\rho(\mathbf{x}', \mathbf{y}') = \sum_{i=1}^{\infty} \lambda_i f_i(\mathbf{x}') f_i(\mathbf{y}') \quad (\text{C.6})$$

¹For simplicity, we have normalized the equation by the conductivity of the medium, as it does not affect our calculations.




Figure C.1: Distributed source model and the electrode positions

Therefore, by defining the vector $\mathbf{E}_i \doteq \int f_i(\mathbf{x}')\mathbf{G}(\mathbf{x}')d\mathbf{x}'$ and combining (C.4) and (C.6), we find:

$$R_\phi = \sum_{i=1}^{\infty} \lambda_i \iint \mathbf{G}(\mathbf{x}')f_i(\mathbf{x}')f_i(\mathbf{y}')\mathbf{G}(\mathbf{y}')^T d\mathbf{x}'d\mathbf{y}' = \sum_{i=1}^{\infty} \lambda_i \mathbf{E}_i \mathbf{E}_i^T \quad (\text{C.7})$$

The $\mathbf{E}_i \mathbf{E}_i^T \in \mathbb{R}^{n \times n}$ are symmetric rank-one matrices with an eigenvalue $\|\mathbf{E}_i\|^2$. R_ϕ is therefore an infinite weighted sum of such positive definite matrices. Note that although the eigen-functions $f_i(\cdot)$ are orthogonal, the \mathbf{E}_i are not necessarily orthogonal. Therefore, (C.7) is a ‘redundant’ decomposition of R_ϕ in the n -dimensional space.

What does this equation imply? R_ϕ is a matrix that we directly calculate from the observed potentials, while the λ_i and \mathbf{E}_i are derived from the physical properties of the distributed source and the geometry of the electrodes. Equation (C.7), shows how this matrix can be related to the stochastic properties of the source density and measures of the electrode positions and distances. In fact, depending on the form of the cross-correlation function $R_\rho(\cdot, \cdot)$, the infinite summation in (C.6) and (C.7) can be approximated by a finite number of terms. This is the idea behind the *multipole expansion* discussed in Chapters 3 and 6.

On the other hand, the matrix R_ϕ is at the heart of all multichannel decomposition techniques. For instance in PCA, we are interested in the following decomposition:

$$R_\phi = QDQ^T = \sum_{i=1}^n d_i \mathbf{q}_i \mathbf{q}_i^T \quad (\text{C.8})$$

where $Q = [\mathbf{q}_1, \dots, \mathbf{q}_n]$ is the orthogonal eigenmatrix, $D = \text{diag}(d_1, \dots, d_n)$ is the eigenvalue matrix, and the $\mathbf{q}_i \mathbf{q}_i^T$ are orthogonal projectors onto $\text{range}(Q)$. Therefore, comparing (C.7) and (C.8), we can see how the eigenvectors and eigenvalues of a set of observations $\phi(t)$, are related to the source distribution and electrode geometry. The concept of infinite dimensions of the observations, discussed in Chapter 6, can also be seen from the derived formulations.

Similar (but more complicated) equations can be derived for *cumulant tensors* [89, ch. 11], of $\phi(t)$. These tensors can be used to relate ICA solutions based on higher-order statistics to the above-mentioned physical properties of the distributed source.

Note that in the derived equations, we made no specific assumptions on the nature of the source density $\rho(\mathbf{x}, t)$. For the case of ECG signals, taking $p = 1$, this density function can represent $\nabla \cdot \mathbf{J}$, the divergence of the cardiac impressed current dipole moment density defined in Section 3.4.

In future studies, a more realistic formulation would be to consider non-homogeneous propagation media, which as mentioned in Section 3.4 leads to some sort of ‘reflection’ sources due to the surfaces of conductivity discontinuity.

Appendix D

The Open Source ECG Toolbox (OSET)

Throughout the manuscript the main methods and algorithms that were specifically developed for fetal ECG/MCG extraction were presented. Some of these methods resulted in byproducts that were not restricted to fetal cardiac signals and were applicable to a broader range of biosignals. In this appendix we present general explanations about some of these byproducts which had methodological novelty and may be considered as additional contributions of this work.

Most of the proposed methods have been implemented in Matlab[®]. A collection of these implementations have been provided in the *open source ECG toolbox* (OSET), online available at [167]. This toolbox contains the synthetic ECG/MCG generator proposed in Chapter 4 (for both adults and fetuses), the Kalman filtering framework presented in Chapter 5, and several preprocessing functions which may be used for ECG/MCG or even other biosignals. Snapshots of the graphical user interface of this toolbox are shown in Figs. D.3-D.7 (at the end of the appendix). In what follows the details of some of these functions are presented.

D.1 Baseline Wander Removal

Several baseline wander (BW) removal functions are provided in OSET, including (1) single pole lowpass filters using forward-backward filtering to attain zero-lag, (2) bandpass DFT filters, (3) single and two-stage moving average (MA) filters, and (4) single and two-stage moving window median (MWM) filters.

Specifically, the idea of the MWM filter is to slide a window over the signal and to assign the median of the window to the sample at the center of each window. It is therefore rather similar to a MA filter, except that the median value is more robust to outliers than the mean value; therefore spikes such as the ECG (MCG) R-peaks are less likely to pass the MWM filter¹. A theoretical study of the MWM filter and its extensions may be found in [8]. The intuition behind using a two stage MWM is that the first filter (with a shorter window length) removes the narrow QRS-wave, while the second one removes the P and T-waves. Therefore, by subtracting the output of the second MWM filter from the input signal, we can assure that the complete ECG (MCG) morphology is preserved and the BW components are effectively removed. Despite the effectiveness of the median filter, its implementation requires the sorting of the samples within the moving windows. The function is therefore slower than a MA filter. In order to overcome this issue, we implemented the MWM filter in C++ using the QuickSort sorting algorithm [163]. A mex-file dynamically linked library (dll) of this filter was also developed, which is directly executable from Matlab.

In a comparative study, we performed a Monte Carlo simulation for evaluating the performance of different BW removal methods. For this simulation, the twelve standard leads from 24 subjects of the Physikalisch-Technische Bundesanstalt diagnostic ECG database (PTBDB) [158], were used to train synthetic ECG signals and to generate thirty second segment ECG signals, as proposed in Chapter 4. Next, real segments of BW noise were randomly selected from the MIT-BIH noise stress test database [139, 159], and added to the synthetic ECG in different portions to achieve noisy ECG in various SNRs.

¹It should be noted that the MWM filter is a nonlinear filter.

The noisy signals were given to the different BW removal methods and the input-output SNR improvement of each filter was compared as a performance measure. For each input SNR, the simulation was repeated 10 times with different BW noise segments for each of the 24 subjects and over 12 channels. The average and standard deviation of the SNR improvement over 2880 trials (10 repetitions \times 24 subjects \times 12 channels) can be seen in Fig. D.1. We can see from these results that the two-stage MWM has outperformed




Figure D.1: The mean and standard deviation bars of SNR improvements achieved in different input SNRs for BW removal based on (MD2) two-stage MWM filter, (MD1) single-stage MWM filter, (MN2) two-stage MA filter, (LP) zero-lag lowpass filter, and (BP) bandpass DFT filter

the other methods. It was also experienced that the MWM filter results can be further smoothed by applying a lowpass filter, with a cutoff frequency of 30Hz-50Hz, to the BW estimated by the MWM filter. Other combinations of these filters, such as the *alpha-trimmed mean filter* [15, 8], can also be of interest for BW removal.

D.2 R-Peak Detectors

Three R-peak detectors are provided in OSET: (1) a simple peak detector that seeks for the local peaks that occur in a window of certain width (corresponding to the approximate heart rate), (2) the classical Pan-Tompkins R-peak detector [151], and (3) an *ad hoc* matched filter. The intuition behind the latter method is that the ECG (or MCG) morphology is rather consistent in time. Therefore, if we use one of the ECG complexes as a reference of a matched filter and convolve the signal with this reference, the filter output will have a peak at the points that the signal maximally match the template, i.e., under each of the ECG complexes. Then by thresholding the matched filter output (e.g. at 0.2 of the maximum absolute value of the filter output) and detecting the peaks that occur in a sliding window of a certain length (shorter than the smallest expected RR-interval), we can find the actual ECG peaks. This method was especially very effective for the detection of the twin fetal MCG peaks presented in Section 7.5, for which the fetal MCGs were highly noisy.

Note that the matched filter peak-detector finds the time instants in which the signal maximally matches the template segment. This time instant is not necessarily the same as the signal's peak location in that segment, which are found by classical peak detection methods. This leaves us with a more fundamental question:

Where should the ECG R-peak be defined? The point that the ECG complex has a local peak, or the center of windows having a maximal similarity (in the mean square error sense) with other ECG complexes?

The clarification of this definition can also influence HRV studies for both adults and fetuses.

D.3 Power-Line Noise Notch Filter

Besides the classical lowpass and bandpass filters available in OSET, a specific notch filter has also been designed for power-line (PL) cancellation. This filter is a Kalman notch filter, which can adapt itself to mixtures of nonstationary signals and noise. Up to now, several types of PL cancellers have been proposed in the literature, ranging from simple second order notch filters to sophisticated adaptive filters [58, 83, 130]. Second order notch filters are rather simple with two parameters for altering the notch frequency and the filter's Q-factor; but can be rather unstable or can have long transition times. More sophisticated notch filters have also been proposed, which are able to track the frequency, phase, and amplitude of the PL in a broad range of variations; but at the expense of higher complexity. Looking at the problem of PL interference more realistically, we notice that PL characteristics are highly controlled in power systems for both economic and safety reasons. The standards are even tighter, when it comes to hospitals and clinical conditions that are in direct contact with patients. Therefore, a PL canceller that tracks the PL frequency in the range of several Hertz or with high phase drifts would only be of theoretical interest. With this background we proposed a rather simple linear Kalman notch filter for PL cancellation.

We can consider the PL as a sinusoidal signal with an arbitrary amplitude and phase:

$$x_n = B \cos(\omega n + \phi), \quad (\text{D.1})$$

where $\omega = 2\pi f/f_s$, f is the 50Hz or 60Hz PL frequency, and f_s is the sampling frequency. Using trigonometric identities, (D.1) can be recursively written as follows:

$$x_{n+1} + x_{n-1} = 2 \cos(\omega) x_n \quad (\text{D.2})$$

A more conservative form of this model, which is also more suitable for a KF, is to consider an additive zero-mean random term w_n , to represent the model errors, including minor frequency, phase or amplitude deviations. The modified dynamic model is represented as follows:

$$x_{n+1} + x_{n-1} = 2 \cos(\omega) x_n + w_n \quad (\text{D.3})$$

We further consider the observation signals as a superposition of the PL signal and the ECG (or any other target signal of interest), i.e.,

$$y_n = x_n + v_n, \quad (\text{D.4})$$

where x_n is the PL and v_n is a zero-mean random term representing all the non-PL signals and noises.

In order to set up a KF for estimating the PL, and eventually removing it from the desired biosignal, the dynamic equation in (D.3) needs to be converted to a matrix state-space form. There are several standard ways for representing a dynamic model in a state-space form, which are known as *canonical representations*. Two of the most common canonical representations, which are of great importance in control theory, are the *controllable* and *observable* canonical forms [101]. The resulting models of these forms are guaranteed to be *controllable* and *observable*, respectively².

On the other hand, (D.3) is equivalent with a system with w_n as its input and x_n as its output. This transfer function is *strictly proper*, i.e., the degree of the numerator is less than the degree of the denominator. It can therefore be represented in a controllable canonical form. By defining $x'_n \doteq x_{n-1}$, we can rewrite (D.3) and (D.4) as follows:

$$\begin{bmatrix} x_{n+1} \\ x'_{n+1} \end{bmatrix} = \begin{bmatrix} 2 \cos(\omega) & -1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} x_n \\ x'_n \end{bmatrix} + \begin{bmatrix} 1 \\ 0 \end{bmatrix} w_n \quad (\text{D.5})$$

$$y_n = [1 \quad 0] \begin{bmatrix} x_n \\ x'_n \end{bmatrix} + v_n, \quad (\text{D.6})$$

and by defining $\mathbf{x}_n = [x_n \quad x'_n]^T$, $A = \begin{bmatrix} 2 \cos(\omega) & -1 \\ 1 & 0 \end{bmatrix}$, $\mathbf{b} = [1 \quad 0]^T$, and $\mathbf{h} = [1 \quad 0]^T$, we come to the following matrix representation:

$$\begin{cases} \mathbf{x}_{n+1} = A\mathbf{x}_n + \mathbf{b}w_n \\ y_n = \mathbf{h}^T \mathbf{x}_n + v_n \end{cases} \quad (\text{D.7})$$

² In the context of control theory, *controllability* is a property of a system, describing the ability to drive the system states to arbitrary values through the input signal or noise in finite time. Its dual notion of *observability* describes the ability to infer the system states given output measurements [101, 202].

This dynamic model is only *marginally stable*, since the eigenvalues of the matrix A are $\lambda(A) = e^{\pm j}$ and they exactly lie on the unit circle. This is clearly due to the oscillatory nature of the PL. However, since the matrices $M = [\mathbf{b} \ \mathbf{b}A]$ and $N = [\mathbf{h} \ A^T\mathbf{h}]^T$ are full-rank, the system is both controllable and observable [101]. This practically means that despite the marginal stability of the model, the PL may indeed be estimated by this KF and the filter converges to its steady state value, i.e., the value with the least minimum mean square error, in a finite time from any initial state. It is also proved that for stationary signals, the filter converges to the optimal Wiener filter in its steady state. The time of convergence depends on the covariances of the model and observation noises, namely $q = E\{w_n^2\}$ and $r = E\{v_n^2\}$.

Moreover, it is seen from (D.5) that the amplitude and phase of the PL defined in (D.1), do not appear in the dynamic model of the system. This implies that the filter will internally estimate these values regardless of their initial values.

Note that although we run the proposed KF to estimate \hat{x}_n (i.e. the PL), we are finally interested in the innovation signal of the KF, i.e.

$$\hat{v}_n = y_n - \hat{x}_n, \quad (\text{D.8})$$

which represents the PL-removed signal.

To complete the discussion, the values of q and r should also be selected. q is the tunable parameter of the filter, having a rather small value as compared with the peak amplitude of the expected PL noise. r is the variance of the non-PL components. For moderate values of input PL noise, we can set $r = E\{y_n^2\}$, since r is merely the variance of the input signal without the single frequency component f . However, in high PL noise values we can estimate r by calculating the DFT of y_n , removing the 50Hz or 60Hz PL frequency component, and calculating the variance of the remaining components.

The filter is now ready to be applied to ECG signals using the classical KF equations [68, 106]. Apparently, multiple stages of the proposed filter may be used for removing the different harmonics of the PL noise. For illustration, in Fig. D.2 PL noise has been synthetically added to clean ECG signals and the Kalman notch filter results are compared with the original and noisy ECG. A detailed study of the theoretical properties of this filter, including its steady state performance was reported in [168].




Figure D.2: A segment of diluted, denoised, and original ECG signal (from top to bottom)

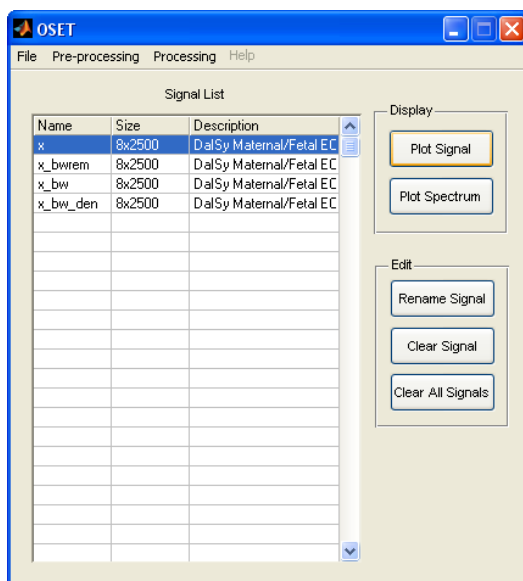
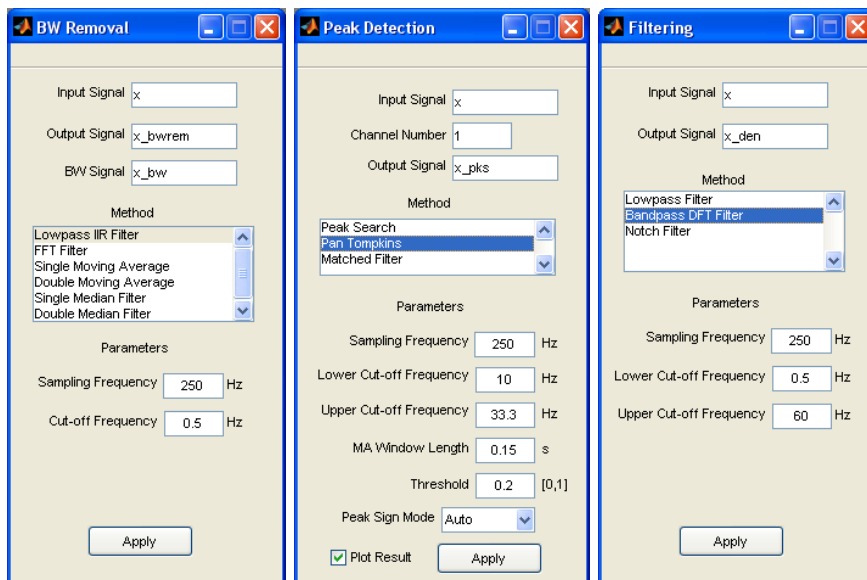


Figure D.3: Snapshot of the main interface and workspace of the OSET GUI



(a) BW removal

(b) Peak detection

(c) Filtering

Figure D.4: Snapshot of some of the functions of the OSET GUI

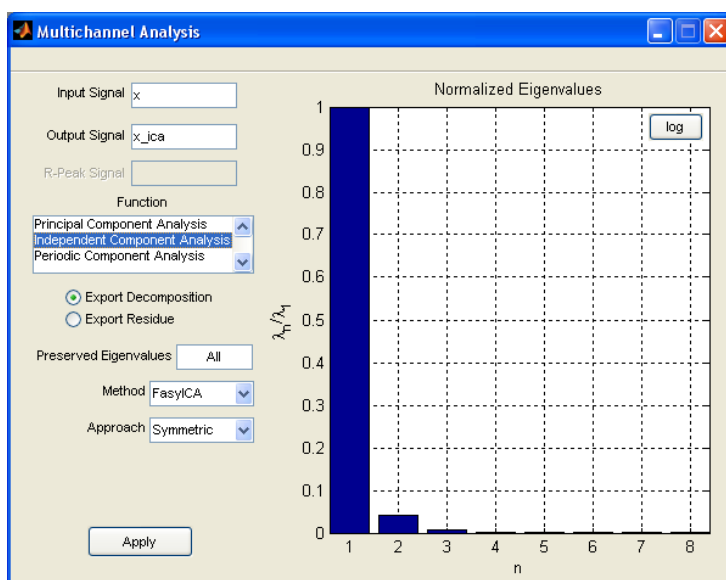


Figure D.5: Multichannel analysis panel containing some of the implemented procedures




Figure D.6: Signal plot panel




Figure D.7: Kalman filter parameter optimization panel

Bibliography

- [1] R. Abeysekera, “Some physiologically meaningful features obtained from Fourier descriptors of vectorcardiograph,” *IEEE Eng. Med. Biol. Mag.*, vol. 10, pp. 58–63, 1991.
- [2] P. M. Agante and J. P. M. de Sá, “ECG Noise Filtering Using Wavelets with Soft-thresholding Methods,” in *Proc. Computers in Cardiology’99*, 1999, pp. 535–542.
- [3] M. Akay, M. Akay, and E. Mulder, “Examining fetal heart-rate variability using matching pursuits,” *IEEE Eng. Med. Biol. Mag.*, vol. 15, no. 5, pp. 64–67, 1996.
- [4] P. J. Allen, O. Josephs, , and R. Turner, “A Method for Removing Imaging Artifact from Continuous EEG Recorded during Functional MRI,” *Neuroimage*, vol. 12, no. 2, pp. 230–239, August 2000.
- [5] P. J. Allen, G. Polizzi, K. Krakow, D. R. Fish, and L. Lemieux, “Identification of EEG Events in the MR Scanner: The Problem of Pulse Artifact and a Method for Its Subtraction,” *Neuroimage*, vol. 8, no. 3, pp. 229–239, October 1998.
- [6] *Congenital Heart Defects in Children Fact Sheet*, American Heart Association, 2008. [Online]. Available: <http://www.americanheart.org/children>
- [7] L. Amini, R. Sameni, C. Jutten, G. Hossein-Zadeh, and H. Soltanian-Zadeh, “MR Artifact Reduction in the Simultaneous Acquisition of EEG and fMRI of Epileptic Patients,” in *EUSIPCO2008 - 16th European Signal Processing Conference*, Lausanne, Switzerland, August 25-29 2008, [to appear].
- [8] G. R. Arce, *Nonlinear Signal Processing: A Statistical Approach*. New York: John Wiley & Sons Inc., 2004.
- [9] B-U Köhler, C. Hennig, and R. Orglmeister, “The principles of software QRS detection. Review and comparing algorithms for detecting this important ECG waveform,” *IEEE Eng. Med. Biol. Mag.*, vol. 21, pp. 42–57, Jan/Feb 2002.
- [10] F. Bach and M. Jordan, “Beyond independent components: trees and clusters,” *Journal of Machine Learning Research*, vol. 4, pp. 1205–1233, 2003. [Online]. Available: <http://cmm.ensmp.fr/~bach/bach03a.pdf>
- [11] J. J. Bailey, A. S. Berson, and A. J. Garson, “Recommendations for standardization and specifications in automated electrocardiography: bandwidth and digital signal processing.” *Circulation*, vol. 81, pp. 730–739, 1990.
- [12] P. N. Baker, I. R. Johnson, P. A. Gowland, J. Hykin, V. Adams, P. Mansfield, and B. S. Worthington, “Measurement of fetal liver, brain and placental volumes with echo-planar magnetic resonance imaging.” *Br J Obstet Gynaecol*, vol. 102, no. 1, pp. 35–39, Jan 1995.
- [13] R. Barr, “Genesis of the electrocardiogram,” pp. 129–151, 1989.
- [14] A. K. Barros and A. Cichocki, “Extraction of specific signals with temporal structure.” *Neural Comput*, vol. 13, no. 9, pp. 1995–2003, Sep 2001. [Online]. Available: <http://dx.doi.org/10.1162/089976601750399272>
- [15] J. Bednar and T. Watt, “Alpha-trimmed means and their relationship to median filters,” *Acoustics, Speech, and Signal Processing, IEEE Transactions on*, vol. 32, no. 1, pp. 145–153, February 1984.

- [16] R. Bellman, *Dynamic Programming*. Princeton University Press, Princeton, NJ., 1957, republished 2003: Dover.
- [17] A. Belouchrani, K. Abed-Meraim, J.-F. Cardoso, and E. Moulines, "A Blind Source Separation Technique Using Second-Order Statistics," *IEEE Trans. Signal Processing*, vol. 45, pp. 434–444, Feb. 1997.
- [18] J. Ben-Arie and K. Rao, "Nonorthogonal representation of signals by Gaussians and Gabor functions," *IEEE Trans. Circuits Syst. II*, vol. 42, no. 6, pp. 402–413, June 1995.
- [19] P. Bergveld and W. J. H. Meijer, "A New Technique for the Supression of the MECG," *IEEE Trans. Biomed. Eng.*, vol. BME-28, pp. 348–354, Apr. 1981.
- [20] C. Bishop, *Neural Networks for Pattern Recognition*. New York: Oxford University Press, 1995.
- [21] G. Bonmassar, P. L. Purdon, I. P. Jaaskelainen, K. Chiappa, V. Solo, E. N. Brown, and J. W. Belliveau, "Motion and ballistocardiogram artifact removal for interleaved recording of EEG and EPs during MRI," *NeuroImage*, vol. 16, no. 4, pp. 1127–1141, 2002.
- [22] R. Boussejot, D. Kreiseler, and A. Schnabel, "Nutzung der EKG-Signaldatenbank CARDIODAT der PTB uber das Internet," *Biomedizinische Technik*, vol. 40, no. 1, pp. S317–S318, 1995.
- [23] R. A. Brace and E. J. Wolf, "Normal amniotic fluid volume changes throughout pregnancy." *Am J Obstet Gynecol*, vol. 161, no. 2, pp. 382–388, Aug 1989.
- [24] J. Bronzino, Ed., *The Biomedical Engineering Handbook*, 2nd ed. CRC Press, Inc., 2000.
- [25] J. V. Candy, *Model-Based Signal Processing*. Wiley-IEEE Press, 2005.
- [26] *Physiology*, Cardiovascular Consultants, 2006. [Online]. Available: <http://www.cardioconsult.com/Physiology/>
- [27] J.-F. Cardoso, "Blind signal separation: statistical principles," *Proc. IEEE*, vol. 86, no. 10, pp. 2009–2025, 1998.
- [28] —, "Multidimensional independent component analysis," in *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP '98)*, vol. 4, May 1998, pp. 1941–1944.
- [29] —, "High-order contrasts for independent component analysis," *Neural Comput.*, vol. 11, no. 1, pp. 157–192, 1999.
- [30] —, *Blind Source Separation and Independent Component Analysis*, 2005. [Online]. Available: <http://www.tsi.enst.fr/~cardoso/guidesepsou.html>
- [31] J.-F. Cardoso and A. Souloumiac, "Blind beamforming for non Gaussian signals," *IEE - Proceedings -F*, vol. 140, pp. 362–370, 1993.
- [32] J. Cardoso, "On the performance of orthogonal source separation algorithms," in *Proc. EUSIPCO*, 1994, pp. 776–779.
- [33] *CancerWEB*, Centre for Cancer Education, University of Newcastle upon Tyne, 1998. [Online]. Available: <http://cancerweb.ncl.ac.uk/>
- [34] M. Chen, R. T. Wakai, and B. V. Veen, "Eigenvector based spatial filtering of fetal biomagnetic signals." *J Perinat Med*, vol. 29, no. 6, pp. 486–496, 2001.
- [35] Z. Chen, "Bayesian filtering: from kalman filters to particle filters, and beyond," 2003. [Online]. Available: http://soma.crl.mcmaster.ca/~zhechen/ieee_bayes.ps
- [36] D. K. Cheng, *Field and Wave Electromagnetics*, 2nd ed. Addison-Wesley, Reading, Mass., 1989.
- [37] I. I. Christov and I. K. Daskalov, "Filtering of electromyogram artifacts from the electrocardiogram," *Medical Engineering and Physics*, vol. 21, pp. 731–736, 1999.

- [38] G. D. Clifford, A. Shoeb, P. E. McSharry, and B. A. Janz, "Model-based Filtering, Compression and Classification of the ECG," *International Journal of Bioelectromagnetism*, vol. 7, no. 1, pp. 158–161, 2005.
- [39] G. Clifford, "A novel framework for signal representation and source separation," *Journal of Biological Systems*, vol. 14, no. 2, pp. 169–183, June 2006.
- [40] G. Clifford and R. Sameni, "An Artificial Multi-Channel Model for Generating Abnormal Electrocardiographic Rhythms," in *Computers in Cardiology, 2008*, Bologna, Italy, 14-17 September 2008, [accepted].
- [41] P. Comon, "Supervised classification, a probabilistic approach," in *ESANN-European Symposium on Artificial Neural Networks*, Verleysen, Ed. Brussels: D facta Publ., Apr 19-21 1995, pp. 111–128, [invited paper].
- [42] M. Cremer, "Über die Direkte Ableitung der Aktionströme des Menschlichen Herzens vom Oesophagus und Über das Elektrokardiogramm des Fetus," *Münchener Medizinische Wochenschrift*, vol. 53, pp. 811–813, April 1906.
- [43] A. A. Damen and J. Van Der Kam, "The use of the singular value decomposition in electrocardiography," *Med Biol Eng Comput*, vol. 20, no. 4, pp. 473–82, July 1982.
- [44] M. Davies, "Identifiability Issues in Noisy ICA," *IEEE Signal Processing Lett.*, vol. 11, no. 5, pp. 470–473, 2004.
- [45] L. De Lathauwer, B. De Moor, and J. Vandewalle, "Fetal electrocardiogram extraction by blind source subspace separation," *IEEE Trans. Biomed. Eng.*, vol. 47, pp. 567–572, May 2000.
- [46] B. De Moor, *Database for the Identification of Systems (DaISy)*, 1997. [Online]. Available: <http://homes.esat.kuleuven.be/~smc/daisy/>
- [47] P. Demartines and J. Hérault, "Curvilinear component analysis: a self-organizing neural network for nonlinear mapping of data sets," *Neural Networks, IEEE Transactions on*, vol. 8, no. 1, pp. 148–154, Jan 1997.
- [48] D. Devedeux, C. Marque, S. Mansour, G. Germain, and J. Duchêne, "Uterine electromyography: a critical review." *Am J Obstet Gynecol*, vol. 169, no. 6, pp. 1636–1653, Dec 1993.
- [49] D. L. Donoho, "De-noising by soft-thresholding," *IEEE Trans. Inform. Theory*, vol. 41, pp. 613–627, 1995.
- [50] E. Doron, E. Doron, A. Yeredor, and P. Tichavsky, "Cramér-Rao-Induced Bound for Blind Separation of Stationary Parametric Gaussian Sources," *IEEE Signal Processing Lett.*, vol. 14, no. 6, pp. 417–420, 2007.
- [51] O. Dössel, "Inverse problem of electro- and magnetocardiography: Review and recent progress," *Int. J. Bioelectromagnetism*, vol. 2, no. 2, 2000. [Online]. Available: <http://www.ijbem.org/volume2/number2/doessel/paper.htm>
- [52] G. E. Dower, H. B. Machado, and J. A. Osborne, "On deriving the electrocardiogram from vectorcardiographic leads," *Clin. Cardiol.*, vol. 3, p. 87, 1980.
- [53] J. A. Drose, *Fetal Echocardiography*. Saunders W B Co, 1998.
- [54] L. Edenbrandt and O. Pahlm, "Vectorcardiogram synthesized from a 12-lead ECG: Superiority of the inverse Dower matrix," *J. Electrocardiol.*, vol. 21, p. 361, 1988.
- [55] W. Einthoven, "Über die Form des menschlichen electrocardiogramms," *Arch Gesamte Physiol.*, vol. 60, pp. 101–123, 1895.
- [56] *phonocardiography*, Encyclopaedia Britannica, Inc., 2008. [Online]. Available: <http://www.britannica.com/>

- [57] A. G. Farvet, "Computer Matched Filter Location of Fetal R-Waves," *Medical & Biological Engineering*, vol. 6, no. 5, pp. 467–475, September 1968.
- [58] M. Ferdjallah and R. Barr, "Adaptive digital notch filter design on the unit circle for the removal of powerline noise from biomedical signals," *Biomedical Engineering, IEEE Transactions on*, vol. 41, no. 6, pp. 529–536, June 1994.
- [59] A. Ferreol and P. Chevalier, "On the Behavior of Current Second and Higher Order Blind Source Separation Methods for Cyclostationary Sources," *IEEE Trans. Signal Processing*, vol. 48, pp. 1712–1725, June 2000.
- [60] G. F. Fletcher, G. Balady, V. F. Froelicher, L. H. Hartley, W. L. Haskell, and M. L. Pollock, "Exercise Standards : A Statement for Healthcare Professionals From the American Heart Association," *Circulation*, vol. 91, no. 2, pp. 580–615, 2001. [Online]. Available: <http://circ.ahajournals.org>
- [61] I. K. Fodor, "A survey of dimension reduction techniques," U.S. Department of Energy, Tech. Rep., June 2002. [Online]. Available: <https://e-reports-ext.llnl.gov/pdf/240921.pdf>
- [62] E. Frank, "An Accurate, Clinically Practical System For Spatial Vectorcardiography," *Circulation*, vol. 13, no. 5, pp. 737–749, 1956. [Online]. Available: <http://circ.ahajournals.org/cgi/content/abstract/13/5/737>
- [63] L. Frenkel and M. Feder, "Recursive Expectation-Maximization (EM) Algorithms for Time-varying Parameters with Applications to Multiple Target Tracking," *IEEE Transactions on Signal Processing*, vol. 47, pp. 306–320, Feb. 1999.
- [64] G. M. Friesen, T. C. Jannett, M. A. Jadallah, S. L. Yates, S. R. Quint, and H. T. Nagle, "A comparison of the noise sensitivity of nine QRS detection algorithms," *IEEE Trans. Biomed. Eng.*, vol. 37, no. 1, pp. 85–98, 1990.
- [65] K. Fukunaga, *Introduction to Statistical Pattern Recognition*, 2nd ed. Academic Press, New York, 1990.
- [66] D. Gabor and C. V. Nelson, "Determination of the Resultant Dipole of the Heart from Measurements on the Body Surface," *J. Appl. Phys.*, vol. 25, no. 4, pp. 413–416, 1954.
- [67] H. Gävert, J. Hurri, J. Särelä, and A. Hyvärinen, *The FastICA package for MATLAB - Version 2.1*, 2001. [Online]. Available: <http://www.cis.hut.fi/projects/ica/fastica/>
- [68] A. Gelb, Ed., *Applied Optimal Estimation*. MIT Press, 1974.
- [69] E. Geneviev, A. Deans, M. Carter, and P. Steer, "Separation of fetal and maternal ECG complexes from a mixed signal using an algorithm based on linear regression," *Medical Engineering & Physics*, vol. 17, no. 7, pp. 514–522, October 1995.
- [70] D. B. Geselowitz, "On bioelectric potentials in an inhomogeneous volume conductor," *Biophys J*, vol. 7, pp. 1–11, 1967.
- [71] —, "On the Theory of the Electrocardiogram," *Proc. IEEE*, vol. 77, pp. 857–876, Jun. 1989.
- [72] A. L. Goldberger, L. A. N. Amaral, L. Glass, J. M. Hausdorff, P. C. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C.-K. Peng, and H. E. Stanley, "PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals," *Circulation*, vol. 101, no. 23, pp. e215–e220, June 2000, circulation Electronic Pages: <http://circ.ahajournals.org/cgi/content/full/101/23/e215>.
- [73] G. Golub and C. van Loan, *Matrix Computations*, 3rd ed. The Johns Hopkins University Press, 1996.

- [74] A. Gotchev, N. Nikolaev, and K. Egiazarian, "Improving the transform domain ECG denoising performance by applying interbeat and intra-beat decorrelating transforms," in *Proc. of the 2001 IEEE International Symposium on Circuits and Systems; IEEE Circuits and Systems Society 2001*, Sydney, Australia, 2001, pp. 17–20.
- [75] C. Gouy-Pailler and R. Sameni, "Iterative Subspace Decomposition for Ocular Artifact Removal from EEG Recordings," GIPSA-LAB, INP-Grenoble, Tech. Rep., 2008.
- [76] M. S. Grewal, L. R. Weill, and A. P. Andrews, Eds., *Global Positioning Systems, Inertial Navigation, and Integration*. John Wiley & Sons Inc., 2001.
- [77] B. Grimm, J. Haueisen, M. Huotilainen, S. Lange, P. V. Leeuwen, T. Menendez, M. J. Peters, E. Schleussner, and U. Schneider, "Recommended standards for fetal magnetocardiography," *Pacing Clin. Electrophysiol.*, vol. 26, pp. 2121–2126, 2003.
- [78] F. Grouiller, L. Vercueil, A. Krainik, C. Segebarth, P. Kahane, and O. David, "A comparative study of different artefact removal algorithms for EEG signals acquired during functional MRI," *NeuroImage*, vol. 38, no. 1, pp. 124–37, 2007.
- [79] R. Gulrajani, "The forward and inverse problems of electrocardiography," *Engineering in Medicine and Biology Magazine, IEEE*, vol. 17, no. 5, pp. 84–101, 122, Sep/Oct 1998.
- [80] E. J. Gumbel, *Statistics of Extremes*. Columbia University Press, 1958.
- [81] L. Hadzievski, B. Bojovic, V. Vukcevic, P. Belicev, S. Pavlovic, Z. Vasiljevic-Pokrajcic, and M. Ostojic, "A novel mobile transtelephonic system with synthesized 12-lead ECG," *IEEE Trans. Inform. Technol. Biomed.*, vol. 8, pp. 428–438, Dec. 2004.
- [82] S. Hamerling, F. Meinecke, and K.-R. Müller, "Analysing ICA components by injecting noise," in *Proceedings of the 4th Int. Symp. on Independent Component Analysis and Blind Source Separation (ICA2003)*, Nara, Japan, April 1-4 2003, pp. 149–154. [Online]. Available: http://www.lis.inpg.fr/pages_perso/bliss/deliverables/d19.html
- [83] P. Hamilton, "A comparison of adaptive and nonadaptive filters for reduction of power line interference in the ECG," *Biomedical Engineering, IEEE Transactions on*, vol. 43, no. 1, pp. 105–109, January 1996.
- [84] S. Haykin, Ed., *Kalman Filtering and Neural Networks*. John Wiley & Sons Inc., 2001.
- [85] R. C. Hilborn, *Chaos and Nonlinear Dynamics*, 2nd ed. Oxford University Press, 2000.
- [86] L. K. Hornberger and D. J. Sahn, "Rhythm abnormalities of the fetus." *Heart*, vol. 93, no. 10, pp. 1294–1300, Oct 2007. [Online]. Available: <http://dx.doi.org/10.1136/hrt.2005.069369>
- [87] X. Hu and V. Nenov, "A single-lead ECG enhancement algorithm using a regularized data-driven filter," *IEEE Trans. Biomed. Eng.*, vol. 53, pp. 347–351, Feb. 2006.
- [88] J. W. Hurst, "Naming of the waves in the ECG, with a brief account of their genesis," *Circulation*, vol. 98, no. 18, pp. 1937–1942, Nov 1998.
- [89] A. Hyvarinen, J. Karhunen, and E. Oja, *Independent Component Analysis*. Wiley-Interscience, 2001.
- [90] A. Hyvärinen, J. Särelä, and R. Vigário, "Spikes and bumps: Artefacts generated by independent component analysis with insufficient sample size," in *Int. Workshop on Independent Component Analysis and Signal Separation (ICA'99)*, Aussois, France, 1999, pp. 425–429. [Online]. Available: <http://citeseer.ist.psu.edu/513307.html>
- [91] Z. Ihara, A. van Oosterom, and R. Hoekema, "Atrial repolarization as observable during the pq interval," *J Electrocardiol*, vol. 39, pp. 290–297, 2006.

- [92] M. H. In, S. Y. Lee, T. S. Park, T.-S. Kim, M. H. Cho, and Y. B. Ahn, "Ballistocardiogram artifact removal from eeg signals using adaptive filtering of eeg signals," *Physiological Measurement*, vol. 27, no. 11, pp. 1227–1240, 2006. [Online]. Available: <http://stacks.iop.org/0967-3334/27/1227>
- [93] J. D. Jackson, *Classical Electrodynamics*, 3rd ed. John Wiley & Sons Inc., 1999.
- [94] M. G. Jafari, W. Wang, J. A. Chambers, T. Hoya, and A. Cichocki, "Sequential blind source separation based exclusively on second-order statistics developed for a class of periodic signals," *IEEE Trans. Signal Processing*, vol. 54, pp. 1028–1040, March 2006.
- [95] M. G. Jafari and J. A. Chambers, "Fetal electrocardiogram extraction by sequential source separation in the wavelet domain." *IEEE Trans Biomed Eng*, vol. 52, no. 3, pp. 390–400, Mar 2005. [Online]. Available: <http://dx.doi.org/10.1109/TBME.2004.842958>
- [96] L. Jana, *Early Fetal Heart Development: 0-9 Weeks*, The Dr. Spock Company, 2004. [Online]. Available: <http://www.drspock.com/article/0,1510,5287,00.html>
- [97] S. Julier, J. Uhlmann, and H. Durrant-Whyte, "A new method for the nonlinear transformation of means and covariances in filters and estimators," *IEEE Trans. Automat. Contr.*, vol. 45, pp. 477–482, Mar. 2000.
- [98] S. J. Julier, J. K. Uhlmann, and H. F. Durrant-Whyte, "A New Approach for Filtering Nonlinear Systems," in *Proceedings of the American Control Conference*, vol. 3, 21-23 June 1995, pp. 1628–1632.
- [99] T. Jung, C. Humphries, T. Lee, M. McKeown, V. Iragui, S. Makeig, and T. Sejnowski, "Removing electroencephalographic artifacts by blind source separation," *Journal of Psychophysiology*, vol. 37, pp. 163–178, 2000.
- [100] C. Jutten, R. Sameni, and H. Hauksdóttir, "On the Relevance of Independent Components," in *Proc. of the ICA Research Network International Workshop (ICArn 2006)*, Liverpool, UK, September 18-19 2006, pp. 1–8.
- [101] T. Kailath, *Linear Systems*. Prentice Hall, 1980.
- [102] P. P. Kanjilal, S. Palit, and G. Saha, "Fetal ECG extraction from single-channel maternal ECG using singular value decomposition," *IEEE Trans. Biomed. Eng.*, vol. 44, pp. 51–59, Jan. 1997.
- [103] H. Kantz and T. Schreiber, "Human ECG: nonlinear deterministic versus stochastic aspects," vol. 145, no. 6, Nov 1998, pp. 279–284.
- [104] V. Kariniemi and K. Hukkinen, "Quantification of fetal heart rate variability by magnetocardiography and direct electrocardiography," *Am J Obstet Gynecol*, vol. 128, no. 5, pp. 526–30, July 1977.
- [105] S. M. Kay, "Efficient Generation of Colored Noise," *Proc. IEEE*, vol. 69, pp. 480–481, Apr. 1981.
- [106] ———, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Prentice Hall PTR, 1993.
- [107] H. A. Kestler, M. Haschka, W. Kratz, F. Schwenker, G. Palm, V. Hombach, and M. Höher, "Denosing of High-Resolution ECG-Signals by Combining the Discrete Wavelet Transform with the Wiener Filter," in *Proceedings IEEE Conference on Computers in Cardiology*, September 13-16, 1998 Cleveland, Ohio, USA, 1998, pp. 233–236.
- [108] A. Khamene, A. Khamene, and S. Negahdaripour, "A new method for the extraction of fetal ECG from the composite abdominal signal," *IEEE Trans. Biomed. Eng.*, vol. 47, no. 4, pp. 507–516, 2000.
- [109] Z. Koldovsky and P. Tichavsky, "Methods of Fair Comparison of Performance of Linear ICA Techniques in Presence of Additive Noise," in *Acoustics, Speech and Signal Processing, 2006. ICASSP 2006 Proceedings. 2006 IEEE International Conference on*, vol. 5, May 2006, pp. 873–876.

- [110] M. Kotas, "Projective filtering of time-aligned ECG beats," *IEEE Trans. Biomed. Eng.*, vol. 51, no. 7, pp. 1129–1139, 2004.
- [111] F. Kovacs, M. Torok, and I. Habermajer, "A rule-based phonocardiographic method for long-term fetal heart rate monitoring," *IEEE Trans. Biomed. Eng.*, vol. 47, no. 1, pp. 124–130, 2000.
- [112] D. Kreiseler and R. Bousseljot, "Automatisierte EKG-Auswertung mit Hilfe der EKG-Signaldatenbank CARDIODAT der PTB," *Biomedizinische Technik*, vol. 40, no. 1, pp. S319–S320, 1995.
- [113] P. Laguna, R. Jane, O. Meste, P. W. Poon, P. Caminal, H. Rix, and N. V. Thakor, "Adaptive filter for event-related bioelectric signals using an impulse correlated reference input," *IEEE Trans. Biomed. Eng.*, vol. 39, pp. 1032–1044, 1992.
- [114] K.-C. Lai and J. Shynk, "A successive cancellation algorithm for fetal heart-rate estimation using an intrauterine ECG signal," *IEEE Trans. Biomed. Eng.*, vol. 49, no. 9, pp. 943–954, 2002.
- [115] P. Lander and E. J. Berbari, "Time frequency plane Wiener filtering of the high resolution ECG: background and time frequency representations," *IEEE Trans. Biomed. Eng.*, vol. 44, pp. 247–255, 1997.
- [116] —, "Time frequency plane Wiener filtering of the high resolution ECG: development and applications," *IEEE Trans. Biomed. Eng.*, vol. 44, pp. 256–265, 1997.
- [117] S. D. Larks, "Present status of fetal electrocardiography," *Bio-Medical Electronics, IRE Transactions on*, vol. 9, no. 3, pp. 176–180, July 1962.
- [118] C. D. Lawrence, *Fetal Heart Development*, 1995. [Online]. Available: <http://user.gru.net/clawrence/vccl/chpt1/embryo.htm>
- [119] J. A. Lee and M. Verleysen, *Nonlinear Dimensionality Reduction*. Springer, 2007.
- [120] J. A. Lee, "Analysis of high-dimensional numerical data : from principal component analysis to nonlinear dimensionality reduction and blind source separation," Ph.D. dissertation, Louvain-la-Neuve, UCL, Belgium, 2003.
- [121] J. Leski, "Robust weighted averaging [of biomedical signals]," *IEEE Trans. Biomed. Eng.*, vol. 49, no. 8, pp. 796–804, 2002.
- [122] C. Li, C. Zheng, and C. Tai, "Detection of ECG characteristic points using wavelet transforms," *Biomedical Engineering, IEEE Transactions on*, vol. 42, no. 1, pp. 21–28, January 1995.
- [123] Y. Li and Z. Yi, "An algorithm for extracting fetal electrocardiogram," *Neurocomput.*, vol. 71, no. 7-9, pp. 1538–1542, 2008.
- [124] D. B. Lindsley, "Heart and Brain Potentials of Human Fetuses in Utero," *The American Journal of Psychology*, vol. 55, no. 3, pp. 412–416, July 1942.
- [125] M. A. Maasoumnia, "Estimation Theory and Optimal Filtering, [Lecture Notes]," 2003, Sharif University of Technology, Tehran, Iran.
- [126] E. F. Magann, J. D. Bass, S. P. Chauhan, R. A. Young, N. S. Whitworth, and J. C. Morrison, "Amniotic fluid volume in normal singleton pregnancies." *Obstet Gynecol*, vol. 90, no. 4 Pt 1, pp. 524–528, Oct 1997.
- [127] J. A. Malmivuo and R. Plonsey, *Bioelectromagnetism, Principles and Applications of Bioelectric and Biomagnetic Fields*. Oxford University Press, 1995. [Online]. Available: <http://butler.cc.tut.fi/~malmivuo/bem/bembook>
- [128] D. Mantini, M. Perrucci, S. Cugini, A. Ferretti, G. Romani, and C. D. Gratta, "Complete artifact removal for EEG recorded during continuous fMRI using independent component analysis," *Neuroimage*, vol. 34, no. 2, pp. 598–607, Jan 2007.

- [129] *Congenital Heart Defects*, March of Dimes, 2005. [Online]. Available: http://www.marchofdimes.com/professionals/14332_1212.asp
- [130] S. Martens, M. Mischi, S. Oei, and J. Bergmans, "An Improved Adaptive Power Line Interference Canceller for Electrocardiography," *Biomedical Engineering, IEEE Transactions on*, vol. 53, no. 11, pp. 2220–2231, November 2006.
- [131] S. M. M. Martens, C. Rabotti, M. Mischi, and R. J. Sluijter, "A robust fetal ECG detection method for abdominal recordings," *Physiol Meas*, vol. 28, no. 4, pp. 373–388, Apr 2007. [Online]. Available: <http://dx.doi.org/10.1088/0967-3334/28/4/004>
- [132] P. E. McSharry and G. D. Clifford, *ECGSYN - A realistic ECG waveform generator*. [Online]. Available: <http://www.physionet.org/physiotools/ecgsyn/>
- [133] P. E. McSharry, G. D. Clifford, L. Tarassenko, and L. A. Smith, "A Dynamic Model for Generating Synthetic Electrocardiogram Signals," *IEEE Trans. Biomed. Eng.*, vol. 50, pp. 289–294, March 2003.
- [134] W. J. H. Meijer and P. Bergveld, "The Simulation of the Abdominal MEEG," *IEEE Trans. Biomed. Eng.*, vol. BME-28, pp. 354–357, Apr. 1981.
- [135] F. Meinecke, A. Ziehe, M. Kawanabe, and K.-R. Müller, "A resampling approach to estimate the stability of one-dimensional or multidimensional independent components." *IEEE Trans. Biomed. Eng.*, vol. 49, no. 12 Pt 2, pp. 1514–25, 2002.
- [136] A. M. Minino, M. P. Heron, S. L. Murphy, and K. D. Kochanek, "Deaths: Final data for 2004," *National Vital Statistics Reports*, vol. 55, no. 19, August 2007.
- [137] M. Misiti, Y. Misiti, G. Oppenheim, and J.-M. Poggi, *Matlab® Wavelet Toolbox User's guide version 3*, 2005. [Online]. Available: <http://www.mathworks.com/access/helpdesk/help/toolbox/wavelet/>
- [138] G. Moody, W. Muldrow, and R. Mark, "The MIT-BIH Noise Stress Test Database," <http://www.physionet.org/physiobank/database/nstdb/>.
- [139] —, "A noise stress test for arrhythmia detectors," in *Computers in Cardiology*, 1984, pp. 381–384.
- [140] W. Mydlarczyk, "A Volterra inequality with the power type nonlinear kernel," *Journal of Inequalities and Applications*, vol. 6, no. 6, pp. 625–631, 2001.
- [141] N. Nikolaev and A. Gotchev, "ECG signal denoising using wavelet domain Wiener filtering," in *Proc. European Signal Processing Conf. EUSIPCO-2000*, Tampere, Finland, September 2000, pp. 51–54.
- [142] M. Nørgaard, *The Kalman Filter Toolbox*. [Online]. Available: <http://www.iau.dtu.dk/research/control/kalmtree.html>
- [143] S. Olmos, L. Sörnmo, and P. Laguna, "Block adaptive filters with deterministic reference inputs for event related signals; BLMS and BRLS," *IEEE Trans. Signal Processing*, vol. 50, pp. 1102–1112, 2002.
- [144] T. Oostendorp, *Modeling the Fetal ECG*. Ph.D. dissertation, K. U. Nijmegen, The Netherlands, 1989.
- [145] T. F. Oostendorp, A. van Oosterom, and H. W. Jongasma, "The effect of changes in the conductive medium on the fetal ECG throughout gestation," *Clin. Phys. physiol Meas.*, vol. 10 Sup. B, pp. 11–20, 1989.
- [146] —, "The fetal ECG throughout the second half of gestation," *Clin Phys physiol Meas.*, vol. 10, no. 2, pp. 147–160, 1989.
- [147] T. Oostendorp, A. van Oosterom, and H. Jongasma, "Electrical properties of tissues involved in the conduction of fetal ECG," *Biomed Eng Comput*, vol. 27, pp. 322–324, 1989.

- [148] E. K. Osei and K. Faulkner, "Fetal position and size data for dose estimation." *Br J Radiol*, vol. 72, no. 856, pp. 363–370, Apr 1999.
- [149] N. J. Outram, E. C. Ifeachor, P. W. J. V. Eetvelt, and J. S. H. Curnow, "Techniques for optimal enhancement and feature extraction of fetal electrocardiogram," in *IEE Proc.-Sci. Meas. Technol.*, vol. 142, no. 6, November 1995, pp. 482–489.
- [150] E. Pajkrt, B. Weisz, H. V. Firth, and L. S. Chitty, "Fetal cardiac anomalies and genetic syndromes." *Prenat Diagn*, vol. 24, no. 13, pp. 1104–1115, Dec 2004. [Online]. Available: <http://dx.doi.org/10.1002/pd.1067>
- [151] J. Pan and W. J. Tompkins, "A Real-Time QRS Detection Algorithm," *IEEE Trans. Biomed. Eng.*, vol. BME-32, no. 3, pp. 230–236, 1985.
- [152] A. Papoulis, *Probability, random variables, and stochastic processes*, 3rd ed. McGraw-Hill, 1991.
- [153] Y. Park, K. Lee, D. Youn, N. Kim, W. Kim, and S. Park, "On detecting the presence of fetal R-wave using the moving averaged magnitude difference algorithm," *IEEE Trans. Biomed. Eng.*, vol. 39, no. 8, pp. 868–871, August 1992.
- [154] L. Parra and P. Sajda, "Blind Source Separation via Generalized Eigenvalue Decomposition," *Journal of Machine Learning Research*, vol. 4, pp. 1261–1269, 2003.
- [155] M. Peters, J. Crowe, J.-F. Piéri, H. Quartero, B. Hayes-Gill, D. James, J. Stinstra, and S. Shakespeare, "Monitoring the fetal heart non-invasively: a review of methods," *J. Perinat. Med.*, vol. 29, pp. 408–416, 2001.
- [156] D.-T. Pham and J.-F. Cardoso, "Blind separation of instantaneous mixtures of nonstationary sources," *Signal Processing, IEEE Transactions on [see also Acoustics, Speech, and Signal Processing, IEEE Transactions on]*, vol. 49, no. 9, pp. 1837–1848, Sep 2001.
- [157] D. T. Pham and P. Garat, "Blind separation of mixture of independent sources through a quasi-maximum likelihood approach," *IEEE Trans. Signal Processing*, vol. 45, no. 7, pp. 1712–1725, 1997.
- [158] *The PTB Diagnostic ECG Database*, Physikalisch-Technische Bundesanstalt (PTB). [Online]. Available: <http://www.physionet.org/physiobank/database/ptbdb/>
- [159] PhysioNet, *The MIT-BIH Noise Stress Test Database*, National Institutes of Health. [Online]. Available: <http://www.physionet.org/physiobank/database/nsttdb/>
- [160] —, *The MIT-BIH Normal Sinus Rhythm Database*, National Institutes of Health, 1991. [Online]. Available: <http://www.physionet.org/physiobank/database/nsrdb/>
- [161] —, *MIT-BIH Polysomnographic Database*, National Institutes of Health, 1999. [Online]. Available: <http://www.physionet.org/physiobank/database/slpdb/>
- [162] M. Popescu, P. Cristea, and A. Bezerianos, "High Resolution ECG Filtering Using Adaptive Bayesian Wavelet Shrinkage," in *Proc. Computers in Cardiology 1998 Conference*, Cleveland, Ohio, USA, September 13-16 1998, pp. 401–404.
- [163] W. H. Press, S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes in C: The Art of Scientific Computing*, 2nd ed. Cambridge University Press, 1992.
- [164] M. Richter, T. Schreiber, and D. T. Kaplan, "Fetal ECG Extraction with Nonlinear State-Space Projections," *IEEE Trans. Biomed. Eng.*, vol. 45, no. 1, pp. 133–137, January 1998.
- [165] J. J. Riera, P. A. Valdes, K. Tanabe, and R. Kawashima, "A theoretical formulation of the electrophysiological inverse problem on the sphere," *Physics in Medicine and Biology*, vol. 51, no. 7, pp. 1737–1758, 2006. [Online]. Available: <http://stacks.iop.org/0031-9155/51/1737>
- [166] J. B. Roche and E. H. Hon, "The fetal electrocardiogram. V. Comparison of lead systems," *Am J Obstet Gynecol*, vol. 92, pp. 1149–1159, Aug 1965.

- [167] R. Sameni, *Open Source ECG Toolbox (OSET)*, 2006. [Online]. Available: <http://ecg.sharif.ir/>
- [168] —, “A Kalman Notch Filter for Removing Power-Line Noise from Biomedical Signals,” GIPSA-LAB, INP-Grenoble, Tech. Rep., October 2007, Technical Report.
- [169] R. Sameni, G. D. Clifford, C. Jutten, and M. B. Shamsollahi, “Multichannel ECG and Noise Modeling: Application to Maternal and Fetal ECG Signals,” *EURASIP Journal on Advances in Signal Processing*, vol. 2007, pp. Article ID 43407, 14 pages, 2007, ISSN 1687-6172, doi:10.1155/2007/43407. [Online]. Available: <http://www.hindawi.com/GetArticle.aspx?doi=10.1155/2007/43407>
- [170] R. Sameni, C. Jutten, and M. B. Shamsollahi, “What ICA Provides for ECG Processing: Application to Noninvasive Fetal ECG Extraction,” in *Proc. of the International Symposium on Signal Processing and Information Technology (ISSPIT'06)*, Vancouver, Canada, August 2006, pp. 656–661.
- [171] —, “Multichannel electrocardiogram decomposition using periodic component analysis,” *Biomedical Engineering, IEEE Transactions on*, vol. 55, no. 8, pp. 1935–1940, Aug. 2008.
- [172] R. Sameni, M. B. Shamsollahi, and C. Jutten, “Filtering Electrocardiogram Signals Using the Extended Kalman Filter,” in *Proceedings of the 27th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBS)*, Shanghai, China, September 1-4 2005, pp. 5639–5642.
- [173] —, “Model-based Bayesian filtering of cardiac contaminants from biomedical recordings,” *Physiological Measurement*, vol. 29, no. 5, pp. 595–613, May 2008.
- [174] R. Sameni, M. B. Shamsollahi, C. Jutten, and M. Babaie-Zadeh, “Filtering Noisy ECG Signals Using the Extended Kalman Filter Based on a Modified Dynamic ECG Model,” in *Proceedings of the 32nd Annual International Conference on Computers in Cardiology*, Lyon, France, September 25-28 2005, pp. 1017–1020.
- [175] R. Sameni, M. B. Shamsollahi, C. Jutten, and G. D. Clifford, “A Nonlinear Bayesian Filtering Framework for ECG Denoising,” *IEEE Trans. Biomed. Eng.*, vol. 54, no. 12, pp. 2172–2185, Dec 2007.
- [176] R. Sameni, M. Shamsollahi, and C. Jutten, “Multi-Channel Electrocardiogram Denoising Using a Bayesian Filtering Framework,” in *Proc. of the 33rd Annual International Conference on Computers in Cardiology*, Valencia, Spain, September 17-20 2006, pp. 185–188. [Online]. Available: <http://cinc.mit.edu/archives/2006/>
- [177] R. Sameni, F. Vrins, F. Parmentier, C. Hérial, V. Vigneron, M. Verleysen, C. Jutten, and M. Shamsollahi, “Electrode Selection for Noninvasive Fetal Electrocardiogram Extraction using Mutual Information Criteria,” in *Proc. of the 26th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering (MaxEnt 2006)*, vol. 872, CNRS, Paris, France, July 8-13 2006, pp. 97–104.
- [178] J. Särelä, *Exploratory source separation in biomedical systems*, October 2004, PhD Thesis, Helsinki University of Technology. [Online]. Available: <http://lib.tkk.fi/Diss/2004/isbn9512273438/>
- [179] J. Särelä and R. Vigário, “Overlearning in marginal distribution-based ICA: analysis and solutions,” *Journal of machine learning research*, vol. 4, pp. 1447–1469, 2003. [Online]. Available: <http://jmlr.csail.mit.edu/papers/volume4/sarela03a/sarela03a.pdf>
- [180] L. K. Saul and J. B. Allen, “Periodic component analysis: An eigenvalue method for representing periodic structure in speech.” in *NIPS*, 2000, pp. 807–813. [Online]. Available: <http://www.cs.cmu.edu/Groups/NIPS/00papers-pub-on-web/SaulAllen.pdf>
- [181] R. Schneider, A. Bauer, P. Barthel, and G. Schmidt, “libRASCH: a programming framework for signal handling,” in *Computers in Cardiology, 2004*, 19-22 Sept. 2004, pp. 53–56.

- [182] T. Schreiber and D. T. Kaplan, "Nonlinear noise reduction for electrocardiograms," *Chaos*, vol. 6, no. 1, pp. 87–92, 1996.
- [183] ———, "Signal separation by nonlinear projections: The fetal electrocardiogram," *Phys. Rev. E.*, vol. 53, pp. R4326–R4329, 1996.
- [184] M. J. Sexton and L. A. Latson, *Fetal blood flow*, Healthwise, Inc., October 2005. [Online]. Available: <http://www.memorialhermann.org/>
- [185] M. Shao, K. Barner, and M. Goodman, "An interference cancellation algorithm for noninvasive extraction of transabdominal fetal electroencephalogram (TaFEEG)," *IEEE Trans. Biomed. Eng.*, vol. 51, no. 3, pp. 471–483, 2004.
- [186] M. G. Signorini, G. Magenes, S. Cerutti, and D. Arduini, "Linear and nonlinear parameters for the analysis of fetal heart rate signal from cardiotocographic recordings." *IEEE Trans Biomed Eng*, vol. 50, no. 3, pp. 365–374, Mar 2003. [Online]. Available: <http://dx.doi.org/10.1109/TBME.2003.808824>
- [187] S. Snowden, N. A. Simpson, and J. J. Walker, "A digital system for recording the electrical activity of the uterus." *Physiol Meas*, vol. 22, no. 4, pp. 673–679, Nov 2001.
- [188] L. Sornmo, "Vectorcardiographic loop alignment and morphologic beat-to-beat variability," *Biomedical Engineering, IEEE Transactions on*, vol. 45, no. 12, pp. 1401–1413, December 1998.
- [189] G. Srivastava, S. C.-H. K. Lau, G. Glover, and V. Menon, "ICA-based procedures for removing ballistocardiogram artifacts from EEG data acquired in the MRI scanner," *Neuroimage*, vol. 24, no. 1, pp. 50–60, Jan 2005.
- [190] J. Stinstra, "Reliability of the fetal magnetocardiogram," Ph.D. dissertation, University of Twente, Enschede, The Netherlands, 2001. [Online]. Available: <http://doc.utwente.nl/35964/>
- [191] H. Stogbauer, A. Kraskov, S. A. Astakhov, and P. Grassberger, "Least dependent component analysis based on mutual information," *Physical Review E*, vol. 70, p. 066123, 2004. [Online]. Available: [doi:10.1103/PhysRevE.70.066123](https://doi.org/10.1103/PhysRevE.70.066123)
- [192] G. Strang, *Linear Algebra and Its Applications*, 3rd ed. Brooks/Cole, 1988.
- [193] J. Stratton, *Electromagnetic Theory*. McGraw-Hill Book Company Inc., 1941.
- [194] P. Strobach, K. Abraham-Fuchs, and W. Härer, "Event-synchronous cancellation of the heart interference in biomedical signals," *IEEE Trans. Biomed. Eng.*, vol. 41, no. 4, pp. 343–350, 1994.
- [195] Y. Surrel, "Additive noise effect in digital phase detection," *Applied Optics*, vol. 36, pp. 271–276, 1997.
- [196] A. Swami, G. Giannakis, and S. Shamsunder, "Multichannel ARMA processes," *Signal Processing, IEEE Transactions on [see also Acoustics, Speech, and Signal Processing, IEEE Transactions on]*, vol. 42, no. 4, pp. 898–913, April 1994.
- [197] M. P. Tarvainen, S. D. Georgiadis, P. O. Ranta-aho, and P. A. Karjalainen, "Time-varying analysis of heart rate variability signals with a Kalman smoother algorithm," *Physiol. Meas.*, vol. 27, pp. 225–239, Mar. 2006.
- [198] H. ter Brake, A. Rijpma, J. Stinstra, J. Borgmann, H. Holland, H. Krooshoop, M. Peters, J. Flokstra, H. Quartero, and H. Rogalla, "Fetal magnetocardiography: clinical relevance and feasibility," *Physica C*, vol. 368, pp. 10–17, March 2002. [Online]. Available: <http://www.ingentaconnect.com/content/els/09214534/2002/00000368/00000001/art01132>
- [199] N. V. Thakor and Y. S. Zhu, "Application of adaptive filtering to ECG analysis: noise cancellation and arrhythmia detection," *IEEE Trans. Biomed. Eng.*, vol. 38, pp. 785–794, 1991.

- [200] P. Tichavsky, Z. Koldovsky, and E. Oja, "Performance analysis of the fastica algorithm and cramer-rao bounds for linear independent component analysis," *IEEE Trans. Signal Processing*, vol. 54, no. 4, pp. 1189–1203, 2006.
- [201] L. Tong, R.-W. Liu, V. Soon, and Y.-F. Huang, "Indeterminacy and identifiability of blind identification," *IEEE Trans. Circuits Syst.*, vol. 38, pp. 499–509, May 1991.
- [202] K. S. Tsakalis, "Stability, controllability, observability," June 2001, Lecture Notes. [Online]. Available: <http://www.fulton.asu.edu/~tsakalis/notes/sco.pdf>
- [203] T. Tsalaile, R. Sameni, S. Sanei, C. Jutten, and J. Chambers, "Sequential Blind Source Extraction For Quasi-Periodic Signals With Time-Varying Period," *IEEE Trans. Biomed. Eng.*, 2009, [to appear].
- [204] *PubMed*, The U.S. National Library of Medicine and the National Institutes of Health, 2008, A service of the U.S. National Library of Medicine and the National Institutes of Health. [Online]. Available: <http://www.pubmed.gov/>
- [205] P. Van Leeuwen, S. Lange, H. Bettermann, D. Grönemeyer, and W. Hatzmann, "Fetal heart rate variability and complexity in the course of pregnancy," *Early Hum Dev*, vol. 54, no. 3, pp. 259–269, Apr 1999.
- [206] P. van Leeuwen, S. Lange, A. Klein, D. Geue, and D. H. Grönemeyer, "Dependency of magnetocardiographically determined fetal cardiac time intervals on gestational age, gender and postnatal biometrics in healthy pregnancies," *BMC Pregnancy Childbirth*, vol. 4, no. 6, 2004.
- [207] A. van Oosterom, "Spatial filtering of the fetal electrocardiogram," *J. Perinat Med.*, vol. 14, no. 6, pp. 411–419, 1986.
- [208] —, "Beyond the dipole; modeling the genesis of the electrocardiogram," *100 years Einthoven*, pp. 7–15, 2002, the Einthoven Foundation, Leiden.
- [209] H. Van-Trees, *Detection, Estimation, and Modulation Theory. Part I*. John Wiley & Sons, 2001.
- [210] 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 Trans Biomed Eng*, vol. BME-34, pp. 233–243, 1987.
- [211] P. Várady, L. Wildt, Z. Benyó, and A. Hein, "An advanced method in fetal phonocardiography," *Computer Methods and Programs in Biomedicine*, vol. 71, no. 3, pp. 283–296, 2003.
- [212] V. Vigneron, A. Paraschiv-Ionescu, A. Azancot, O. Sibony, and C. Jutten, "Fetal electrocardiogram extraction based on non-stationary ICA and wavelet denoising," *Signal Processing and Its Applications, 2003. Proceedings. Seventh International Symposium on*, vol. 2, pp. 69–72 vol.2, 1-4 July 2003.
- [213] F. Vrins, C. Jutten, and M. Verleysen, "Sensor array and electrode selection for non-invasive fetal electrocardiogram extraction by independent component analysis," in *Independent Component Analysis and Blind Signal Separation*, ser. Lecture Notes in Computer Science (LNCS 3195), A. P. C.G. Puntonet, Ed. Springer, 2004, pp. 1017–1024.
- [214] R. T. Wakai, "Assessment of fetal neurodevelopment via fetal magnetocardiography." *Exp Neurol*, vol. 190 Suppl 1, pp. S65–S71, Nov 2004. [Online]. Available: <http://dx.doi.org/10.1016/j.expneurol.2004.04.019>
- [215] R. Wakai and W. Lutter, "Matched-filter template generation via spatial filtering: application to fetal biomagnetic recordings," *IEEE Trans. Biomed. Eng.*, vol. 49, no. 10, pp. 1214–1217, 2002.
- [216] X. Wan, K. Iwata, J. Riera, M. Kitamura, and R. Kawashima, "Artifact reduction for simultaneous EEG/fMRI recording: Adaptive fir reduction of imaging artifacts," *Clinical Neurophysiology*, vol. 117, no. 3, pp. 681–692, March 2006. [Online]. Available: <http://dx.doi.org/10.1016/j.clinph.2005.07.025>

- [217] E. Wan *et al.*, *ReBEL: Recursive Bayesian Estimation Library*. [Online]. Available: <http://choosh.ece.ogi.edu/rebel/>
- [218] J. G. Webster, Ed., *Medical Instrumentation: Application and Design*, 3rd ed. John Wiley & Sons, 1998.
- [219] Y. Weiss, "Segmentation using eigenvectors: A unifying view," in *Proc. IEEE Int. Conf. Computer Vision (2)*, 1999, pp. 975–982. [Online]. Available: citeseer.ist.psu.edu/weiss99segmentation.html
- [220] B. Widrow, J. Glover, J. McCool, J. Kaunitz, C. Williams, H. Hearn, J. Zeidler, E. Dong, and R. Goodlin, "Adaptive noise cancelling: principles and applications," *Proc. IEEE*, vol. 63, no. 12, pp. 1692–1716, 1975.
- [221] *Cardiotocograph*, Wikipedia. [Online]. Available: <http://en.wikipedia.org/wiki/Cardiotocograph>
- [222] *Curse of dimensionality*, Wikipedia. [Online]. Available: http://en.wikipedia.org/wiki/Curse_of_dimensionality
- [223] *Superconducting Quantum Interference Devices (SQUID)*, Wikipedia. [Online]. Available: <http://en.wikipedia.org/wiki/SQUID>
- [224] J. Wladimiroff and G. Pilu, Eds., *Ultrasound and the Fetal Heart*. New York, NY: Parthenon Publishing, 1996.
- [225] D. Yellin and B. Friedlander, "Multichannel system identification and deconvolution: performance bounds," *IEEE Trans. Signal Processing*, vol. 47, no. 5, pp. 1410–1414, 1999.
- [226] V. Zarzoso and A. K. Nandi, "Comparison between blind separation and adaptive noise cancellation techniques for fetal electrocardiogram extraction," in *IEE Colloq. Medical Applications for Signal Processing*, 1999, pp. 1/1–1/6.
- [227] V. Zarzoso, A. K. Nandi, and E. Bacharakis, "Maternal and fetal ECG separation using blind source separation methods," *IMA J. Math. Appl. Med. Biol.*, vol. 14, pp. 207–225, 1997.
- [228] V. Zarzoso and A. Nandi, "Noninvasive fetal electrocardiogram extraction: blind separation versus adaptive noise cancellation," *Biomedical Engineering, IEEE Transactions on*, vol. 48, no. 1, pp. 12–18, January 2001.
- [229] Z.-L. Zhang and Z. Yi, "Extraction of temporally correlated sources with its application to non-invasive fetal electrocardiogram extraction," *Neurocomputing*, vol. 69, no. 7-9, pp. 894–899, 2006.
- [230] A. Zuckerwar, R. Pretlow, J. Stoughton, and D. Baker, "Development of a piezopolymer pressure sensor for a portable fetal heart rate monitor," *IEEE Trans. Biomed. Eng.*, vol. 40, no. 9, pp. 963–969, 1993.